

SIPREC
Internet-Draft
Intended status: Informational
Expires: March 13, 2013

A. Hutton, Ed.
Siemens Enterprise
Communications
L. Portman, Ed.
NICE Systems
R. Jain
IPC Systems
K. Rehor
Cisco Systems, Inc.
September 9, 2012

An Architecture for Media Recording using the Session Initiation
Protocol
draft-ietf-siprec-architecture-06

Abstract

Session recording is a critical requirement in many communications environments such as call centers and financial trading. In some of these environments, all calls must be recorded for regulatory, compliance, and consumer protection reasons. Recording of a session is typically performed by sending a copy of a media stream to a recording device. This document describes architectures for deploying session recording solutions in an environment which is based on the Session Initiation Protocol (SIP).

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 March 13, 2013.

Copyright Notice

Copyright (c) 2012 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	3
2. Definitions	4
3. Session Recording Architecture	5
3.1. Location of the Session Recording Client	5
3.1.1. B2BUA acts as a Session Recording Client	5
3.1.2. Endpoint acts as Session Recording Client	7
3.1.3. A SIP Proxy cannot be a Session Recording Client	7
3.1.4. Interaction with MEDIACTRL	8
3.1.5. Interaction with Conference Focus	9
3.2. Establishing the Recording Session	10
3.2.1. Session Recording Client Initiated Recording	11
3.2.2. Session Recording Server Initiated Recording	11
3.2.3. Pause/Resume Recording Session	12
3.2.4. Media Stream Mixing	12
3.2.5. Media Transcoding	12
3.3. Recording Metadata	12
3.3.1. Contents of recording metadata	13
3.3.2. Mechanisms for delivery of metadata to Session Recording Server	13
3.4. Notifications to the Recorded User Agents	13
3.5. Preventing the recording of a SIP session	13
4. IANA considerations	13
5. Security considerations	14
6. Acknowledgements	14
7. Informative References	15
Authors' Addresses	16

1. Introduction

Session recording is a critical requirement in many communications environments such as call centers and financial trading. In some of these environments, all calls must be recorded for regulatory, compliance, and consumer protection reasons. Recording of a session is typically performed by sending a copy of a media stream to a recording device. This document describes architectures for deploying session recording solutions in an environment which is based on the Session Initiation Protocol (SIP) the requirements for which are described in [RFC6341].

This document focuses on how sessions are established between a Session Recording Client (SRC) and the Session Recording Server (SRS) for the purpose of conveying the Replicated Media and Recording Metadata (e.g. Identity of parties involved) relating to the Communication Session.

Once the Replicated Media and Recording Metadata have been received by the Session Recording Server they will typically be archived for retrieval at a later time. The procedures relating to the archiving and retrieval of this information is outside the scope of this document.

This document only considers active recording, where the Session Recording Client purposefully streams media to a Session Recording Server. Passive recording, where a recording device detects media directly from the network (E.g. using port mirroring techniques), is outside the scope of this document. In addition, lawful intercept is outside the scope of this document which takes account of the IETF policy on wiretapping [RFC2804].

The Recording Session that is established between the Session Recording Client and the Session Recording Server uses the normal procedures for establishing INVITE initiated dialogs as specified in [RFC3261] and uses SDP for describing the media to be used during the session as specified in [RFC4566]. However it is intended that some extensions to SIP (E.g. Headers, Option Tags, Etc.) will be defined to support the requirements for media recording. The Replicated Media is required to be sent in real-time to the Session Recording Server and is not buffered by the Session Recording Client to allow for real-time analysis of the media by the Session Recording Server.

2. Definitions

Session Recording Server (SRS): A Session Recording Server (SRS) is a SIP User Agent (UA) that is a specialized media server or collector that acts as the sink of the recorded media. An SRS is typically implemented as a multi-port device that is capable of receiving media from multiple sources simultaneously. An SRS is the sink of the recorded session metadata.

Session Recording Client (SRC): A Session Recording Client (SRC) is a SIP User Agent (UA) that acts as the source of the recorded media, sending it to the SRS. An SRC is a logical function. Its capabilities may be implemented across one or more physical devices. In practice, an SRC could be a personal device (such as a SIP phone), a SIP Media Gateway (MG), a Session Border Controller (SBC) or a SIP Media Server (MS) integrated with an Application Server (AS). This specification defines the term SRC such that all such SIP entities can be generically addressed under one definition. The SRC provides metadata to the SRS.

Communication Session (CS): A session created between two or more SIP User Agents (UAs) that is the subject of recording.

Recording Session (RS): The SIP session created between an SRC and SRS for the purpose of recording a Communication Session.

Recording aware User Agent (UA): A SIP User Agent that is aware of SIP extensions associated with the Communication Session. Such extensions may be used to notify the Recording aware UA that a session is being recorded, or by a Recording aware UA to express preferences as to whether a recording should be started, paused, resumed or stopped.

Recording unaware User Agent (UA): A SIP User Agent that is unaware of SIP extensions associated with the Communication Session. Such Recording unaware UA will be notified that a session is being recorded or express preferences as to whether a recording should be started, paused, resumed or stopped via some other means that is out of scope of SIPREC.

Recording Metadata: The metadata describing the communication session that is required by the Session Recording Server. This will include for example the identity of users that participate in the Communication Session and dialog state. Typically this metadata is archived with the replicated media at the Session Recording Server. The recording metadata is delivered in real-time to the Session Recording Server.

Replicated Media: A copy of the media associated with the Communication Session created by the Session Recording Client and sent to the Session Recording Server. It may contain all the media associated with the communication session (E.g. Audio and Video) or just a subset (E.g. Audio). Replicated Media is part of Recording Session.

3. Session Recording Architecture

3.1. Location of the Session Recording Client

This section contains some example session recording architectures showing how the Session Recording Client is a logical function that can be located in or split between various physical components.

3.1.1. B2BUA acts as a Session Recording Client

A SIP Back to Back User Agent (B2BUA) which has access to the media that is to be recorded may act as a Session Recording Client. The B2BUA may already be aware that a session needs to be recorded before the initial establishment of the communication session or the decision to record the session may occur after the session has been established.

If the B2BUA/SRC makes the decision to initiate the Recording Session (RS) then it will initiate the establishment of a SIP Session by sending an INVITE to the Session Recording Server.

If the Session Recording Server makes the decision to initiate the recording session then it will initiate the establishment of a SIP Session by sending an INVITE to the B2BUA/Session Recording Client.

The RS INVITE contains information which identifies the session as being established for the purposes of recording and prevents the session from being accidentally rerouted to a UA which is not a SRS.

The B2BUA/SRC is responsible for notifying the UAs involved in the communication session that the session is being recorded.

The B2BUA/SRC is responsible for complying with requests from recording aware UAs or through some configured policies indicating that the communication session should not be recorded.

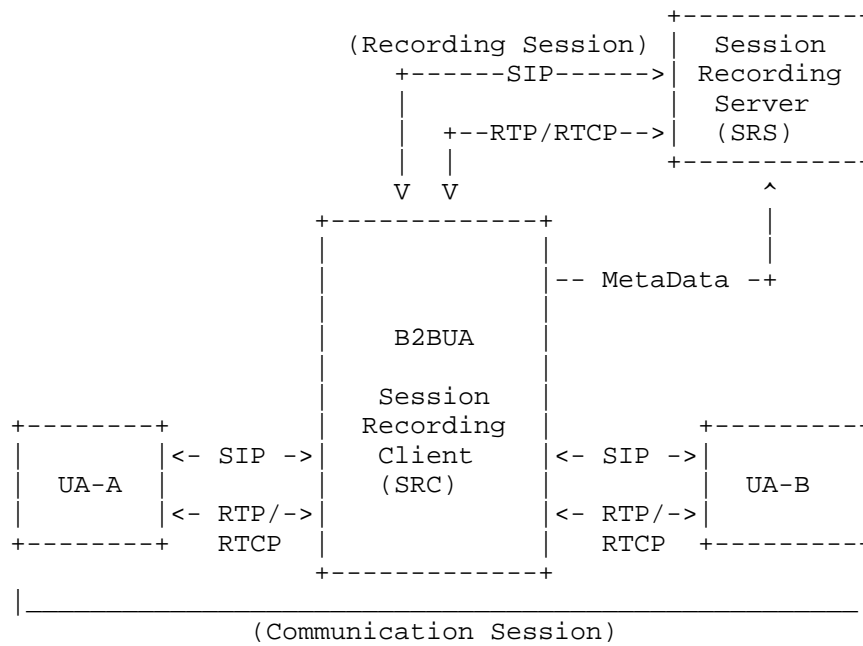


Figure 1: B2BUA Acts as the Session Recording Client.

3.1.2. Endpoint acts as Session Recording Client

A SIP Endpoint / UA may act as a Session Recording Client in which case the endpoint sends the Replicated Media to the Session Recording Server

If the endpoint makes the decision to initiate the Recording Session then it will initiate the establishment of a SIP Session by sending an INVITE to the Session Recording Server.

If the Session Recording Server makes the decision to initiate the Recording Session then it will initiate the establishment of a SIP Session by sending an INVITE to the endpoint. The actual decision mechanism is out of scope of SIPREC.

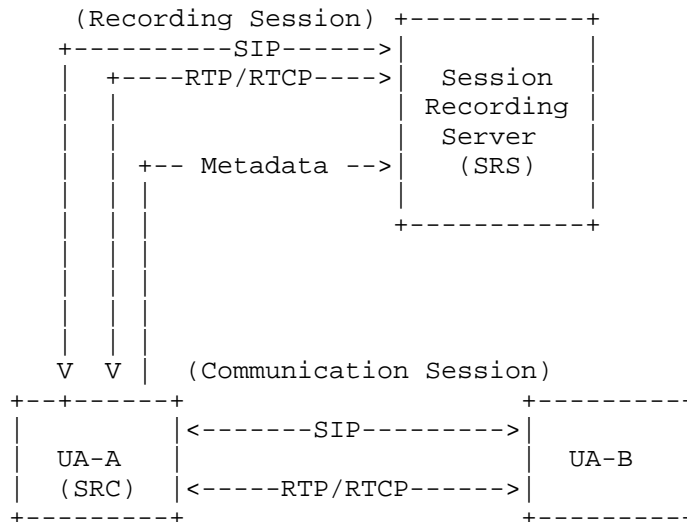


Figure 2: SIP Endpoint acts as the Session Recording Client

3.1.3. A SIP Proxy cannot be a Session Recording Client

A SIP Proxy is unable to act as an SRC because it does not have access to the media and therefore has no way of enabling the delivery of the replicated media to the SRS

3.1.4. Interaction with MEDIACTRL

The mediactrl architecture [RFC5567] describes an architecture in which an Application Server (AS) controls a Media Server (MS) which may be used for purposes such as conferencing and recording media streams. In the [RFC5567] architecture the AS typically uses SIP Third Party Call Control (3PCC) to instruct the SIP UAs to direct their media to the Media Server.

The Session Recording Client and Session Recording Server described in this document may act as an application server as described in [RFC5567]; and therefore, when further decomposed, may be made up of an application server which uses a mediactrl interface to control a media server for the purpose of recording the media streams. However, this interface is considered outside the scope of this document.

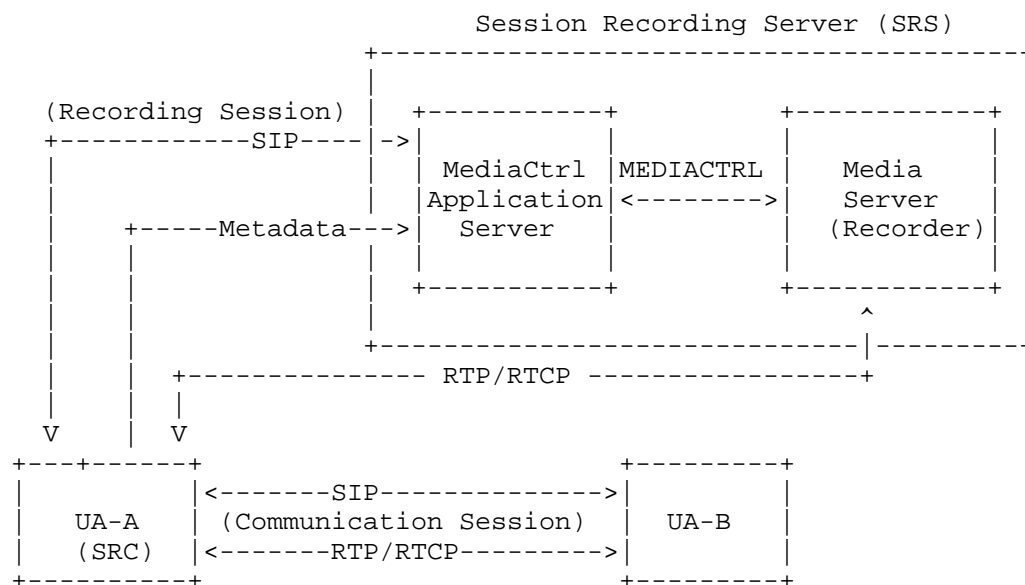


Figure 3: Example of Session Recording Server using MEDIACTRL

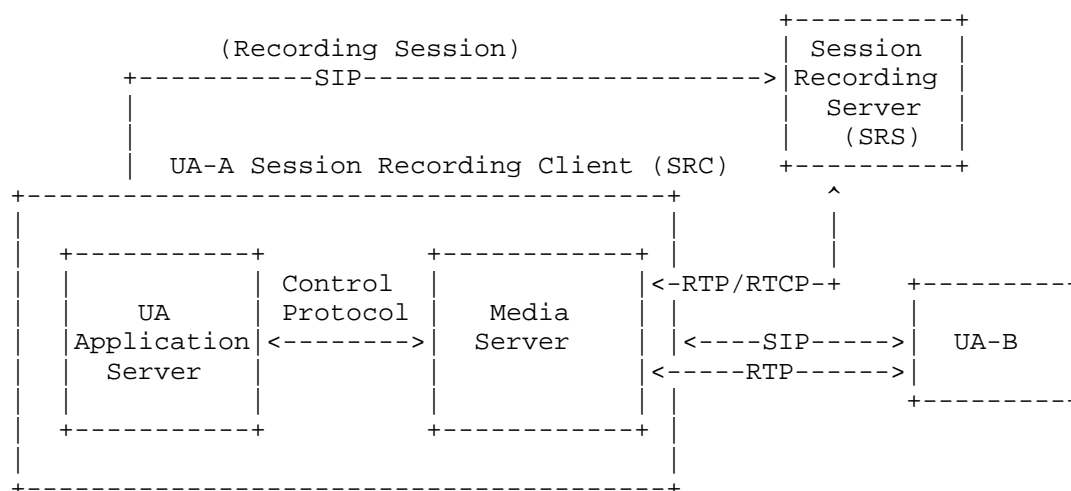


Figure 4: Example of Session Recording Client decomposition

3.1.5. Interaction with Conference Focus

In the case of a centralised conference a combination of the conference focus and mixer [RFC4353] may act as a SRC and therefore provide the SRS with the replicated media and associated recording metadata. In this arrangement the SRC is able to provide media and metadata relating to each of the participants, including, for example, any side conversations where the media passes through the mixer.

Conference Focus can either provide mixed replicated media or separate streams per conference participant (as depicted in the Figure 5).

The conference focus may also act as a Recording Aware UA in the case when one of the participants acts as a SRC.

In an alternative arrangement a SIP endpoint which is a conference participant can act as an SRC. The SRC will in this case have access to the media and metadata relating to that particular participant and may be able to obtain additional metadata from the conference focus. The SRC may for example use the conference event package as described in [RFC4575] to obtain information about other participants which it provides to the SRS within the recording metadata.

The SRC may be involved in the conference from the very beginning or may join at some later point of time.

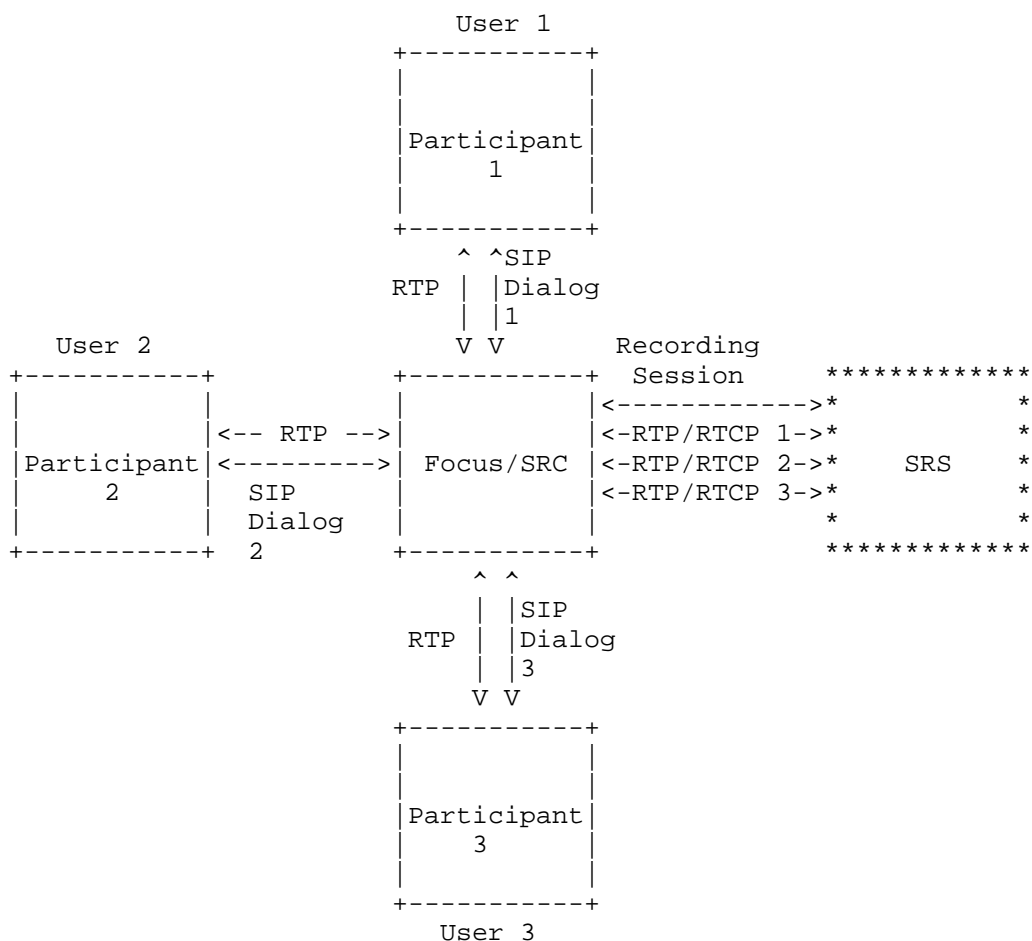


Figure 5: Conference Focus acting as an SRC.

3.2. Establishing the Recording Session

The Session Recording Client or the Session Recording Server may initiate the Recording Session.

It should be noted that the Recording Session is independent from the Communication Session that is being recorded at both the SIP dialog level and at the session level.

Concerning media negotiation, regular SIP/SDP capabilities should be used, and existing transcoding capabilities and media encryption should not be precluded.

3.2.1. Session Recording Client Initiated Recording

When the Session Recording Client initiates the Recording Session for the purpose of conveying media to the Session Recording Server it performs the following actions.

- o The SRC is provisioned with a Unified Resource Identifier (URI) for the SRS, which is resolved through normal [RFC3263] procedures.
- o Initiates the dialog by sending an INVITE request to the Session Recording Server. The dialog is established according to the normal procedures for establishing an INVITE initiated dialog as specified in [RFC3261].
- o Include in the INVITE an indication that the session is established for the purpose of recording the associated media.
- o If the Replicated Media is to be started immediately then the Session Recording Client will include an SDP attribute of "a=sendonly" for each media line or "a=inactive" if it is not ready to transmit the media.
- o The Recording Session may replicate all media associated with the Communication Session or only a subset.
- o Replicate the media streams that are to be recorded and transmit the media to the Session Recording Server.

3.2.2. Session Recording Server Initiated Recording

When the Session Recording Server initiates the media recording session with the Session Recording Client it performs the following actions.

- o The SRS is provisioned with a Unified Resource Identifier (URI) for the SRC, which is resolved through normal [RFC3263] procedures.
- o Send an INVITE request to the Session Recording Client
- o Include in the INVITE an indication that the session is established for the purpose of recording the associated media. Possible mechanisms for this include using the Require header or a media feature tag as defined in [RFC3840].
- o Identify the sessions that is to be recorded. The actual mechanism of the identification depends on SRC policy.

- o If the Recording Session is to be started immediately then the Session Recording Client will include an SDP attribute of "a=recvonly" for each media line or "a=inactive" if it is not ready to receive the media

If the Session Recording Server does not have prior knowledge of what media streams are available to be recorded it can make use of an offerless INVITE which allows the Session Recording Client to make the initial Session Description Protocol (SDP) offer.

3.2.3. Pause/Resume Recording Session

The Session Recording Server or the Session Recording Client may pause the recording by changing the SDP direction attribute to "inactive" and resume the recording by changing the direction back to "recvonly" or "sendonly".

3.2.4. Media Stream Mixing

In a basic session involving only audio there are typically two audio/RTP streams between the two UAs involved transporting media in each direction. When recording this media the two streams may be mixed at the SRC before being transmitted to the SRS or it may be a requirement of the recording server that the media streams are not mixed and are sent to the SRS as two separate streams. The case when media is mixed at the SRC is simple as only a single media stream is required to be sent to the SRS. However in the case when the media streams are not mixed then the SDP offer sent to the SRS must describe two separate media streams.

3.2.5. Media Transcoding

The communication session (CS) and the recording session (RS) are negotiated separately using a standard SDP offer/answer exchange which may result in the SRC having to perform media transcoding between the two sessions. If the SRC is not capable of performing media transcoding it may limit the media formats in the offer to the SRS depending on what media is negotiated on the CS or may limit what it includes in the offer on the CS if it has prior knowledge of the media formats supported by the SRS. However typically the SRS will be a more capable device which can provide a wide range of media format options to the SRC and may also be able to make use of a media transcoder as detailed in [RFC5369].

3.3. Recording Metadata

3.3.1. Contents of recording metadata

The metadata model is defined in [I-D.ietf-siprec-metadata].

3.3.2. Mechanisms for delivery of metadata to Session Recording Server

The SRS obtains session recording metadata from the SRC. The metadata is transported via SIP based mechanisms as specified in [I-D.ietf-siprec-protocol]

It is also possible that metadata is transported via non SIP based mechanisms but these are considered out of scope.

It is also possible to have RS session without the metadata, in such case SRS will be receiving it by some other means or not at all.

3.4. Notifications to the Recorded User Agents

Typically a user that is involved in a session that is to be recorded is notified by an announcement at the beginning of the session or may receive some warning tones within the media. However the standardization of media recording protocols when using SIP enable an indication that the call is being recorded to be included in the SIP requests and responses associated with that communication session.

It is the Session Recording Client that provides the notification to all SIP UAs for which it is replicating received media for the purpose of recording including the local user if the Session Recording Client is a SIP endpoint.

3.5. Preventing the recording of a SIP session

A Recording Aware UA may during the initial session establishment or during an established session provide an indication of their preference with regard to recording the media in the communication session. The mechanism for this are specified in [I-D.ietf-siprec-protocol]

4. IANA considerations

This draft mentions SIP/SDP extensions. The associated IANA considerations are addressed in [I-D.ietf-siprec-protocol] that defines them.

5. Security considerations

The Recording Session is fundamentally a standard SIP dialog and media session and therefore makes use of existing SIP security mechanisms for securing the Recording Session and Recording Metadata.

The intended use of this architecture is only for the case where the users are aware that they are being recorded, and the architecture provides the means for the Session Recording Client to notify users that they are being recorded.

This architectural solution is not intended to support lawful intercept which in contrast requires that users are not informed.

It is the responsibility of the Session Recording Server to protect the Replicated Media and Recording Metadata once it has been received and archived. The mechanism for protecting the storage and retrieval from the SRS is out of scope of this work.

Protection of the RS should be not weaker than protection of the CS, and may need to be stronger because the media is retransmitted (allowing more possibility for interception). This applies to both the signaling and media paths.

It is essential that the SRC will authenticate the SRS because the client must be certain that it is recording on the right recording system. It is less important that the SRS authenticate the SRC, but implementations must have the ability to perform mutual authentication.

In some environments, it is desirable to not decrypt and re-encrypt the media. This means the same key is used on the CS and RS. This only works when the media is encrypted, not the transport. When the media must be decrypted on the CS, and is re-encrypted on the RS, a new key must be used.

The retrieval mechanism for media recorded by this protocol is out of scope. Implementations of retrieval mechanisms should consider the security implications carefully as the retriever is not usually a party to the call that was recorded. Retrievers should be authenticated carefully. The crypto suites on the retrieval should be no less strong than used on the RS, and may need to be stronger.

6. Acknowledgements

Thanks to John Elwell, Brian Rosen, Alan Johnson, Cullen Jennings, Hadriel Kaplan, Henry Lum, Paul Kyzivat, Parthasarathi R, Ram Mohan

R, Charles Eckel and Friso Feenstra for their significant contributions and assistance with this document and Working Group, and to all the members of SIPREC WG mailing list for providing valuable input to this work.

7. Informative References

- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3263] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, June 2002.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC6341] Rehor, K., Portman, L., Hutton, A., and R. Jain, "Use Cases and Requirements for SIP-Based Media Recording (SIPREC)", RFC 6341, August 2011.
- [I-D.ietf-siprec-metadata]
R, R., Ravindran, P., and P. Kyzivat, "Session Initiation Protocol (SIP) Recording Metadata",
draft-ietf-siprec-metadata-06 (work in progress),
March 2012.
- [I-D.ietf-siprec-protocol]
Portman, L., Lum, H., Eckel, C., Johnston, A., and A. Hutton, "Session Recording Protocol",
draft-ietf-siprec-protocol-04 (work in progress),
May 2012.
- [RFC3265] Roach, A., "Session Initiation Protocol (SIP)-Specific Event Notification", RFC 3265, June 2002.
- [RFC3840] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", RFC 3840, August 2004.
- [RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, February 2006.
- [RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, "A Session

Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, August 2006.

[RFC5567] Melanchuk, T., "An Architectural Framework for Media Server Control", RFC 5567, June 2009.

[RFC5369] Camarillo, G., "Framework for Transcoding with the Session Initiation Protocol (SIP)", RFC 5369, October 2008.

[RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.

Authors' Addresses

Andrew Hutton (editor)
Siemens Enterprise Communications
Hofmannstrasse 51
Munich D-81379
Germany

Email: andrew.hutton@siemens-enterprise.com

Leon Portman (editor)
NICE Systems
8 Hapnina
Ra'anana 43017
Israel

Email: leon.portman@nice.com

Rajnish Jain
IPC Systems
777 Commerce Drive
Fairfield, CT 06825
USA

Email: rajnish.jain@ipc.com

Ken Rehor
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

Email: krehor@cisco.com

SIPREC
Internet-Draft
Intended status: Standards Track
Expires: April 20, 2013

Ram Mohan. Ravindranath
Cisco Systems, Inc.
Parthasarathi. Ravindran
Nokia Siemens Networks
Paul. Kyzivat
Huawei
October 17, 2012

Session Initiation Protocol (SIP) Recording Metadata
draft-ietf-siprec-metadata-08

Abstract

Session recording is a critical requirement in many communications environments such as call centers and financial trading. In some of these environments, all calls must be recorded for regulatory, compliance, and consumer protection reasons. Recording of a session is typically performed by sending a copy of a media stream to a recording device. This document describes the metadata model as viewed by Session Recording Server(SRS) and the Recording metadata format.

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 20, 2013.

Copyright Notice

Copyright (c) 2012 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	4
2. Terminology	4
3. Definitions	4
4. Metadata Model	5
5. Recording Metadata Format	6
5.1. XML data format	6
5.1.1. Namespace	7
5.1.2. recording	7
6. Recording Metadata classes	7
6.1. Recording Session	7
6.1.1. Attributes	8
6.1.2. Linkages	8
6.1.3. XML element	8
6.2. Communication Session Group	9
6.2.1. Attributes	9
6.2.2. Linkages	9
6.2.3. XML element	10
6.3. Communication Session	10
6.3.1. Attributes	11
6.3.2. Linkages	11
6.3.3. XML element	12
6.4. CSRSAssociation	12
6.4.1. Attributes	13
6.4.2. Linkages	13
6.4.3. XML element	13
6.5. Participant	13
6.5.1. Attributes	14
6.5.2. Linkages	14
6.5.3. XML element	14
6.6. ParticipantCSAssociation	15
6.6.1. Attributes	15
6.6.2. Linkages	16
6.6.3. XML element	16
6.7. Media Stream	16
6.7.1. Attributes	17
6.7.2. Linkages	17
6.7.3. XML element	17
6.8. ParticipantStream Association	17

6.8.1. Attributes	18
6.8.2. Linkages	18
6.8.3. XML element	18
6.9. associate-time/disassociate-time	18
6.10. Unique ID format	19
7. SIP Recording Metadata Example	19
7.1. Complete SIP Recording Metadata Example	19
7.2. Partial Update of Recording metadata XML body	21
8. XML Schema definition for Recording metadata	21
9. Security Considerations	25
9.1. Connection Security	25
10. IANA Considerations	25
10.1. SIP recording metadata Schema Registration	25
11. Acknowledgement	26
12. Appendix A: Metadata Model Object Instances	26
12.1. Use case 1: Basic Call	26
12.2. Use case 2: Hold/Resume	27
12.3. Use case 3: Basic call with Transfer	29
12.4. Conference Use Cases	30
12.4.1. Case 1:	31
12.4.2. Case 2:	33
12.4.3. Case 3:	35
12.4.4. Case 4:	36
13. Appendix B: Metadata XML schema Instances	37
13.1. Use case 1: Basic Call	37
13.2. Use case 2: Hold/resume	39
13.3. Use case 3: Basic Call with transfer	41
13.4. Use Case 4: Call disconnect	45
14. References	46
14.1. Normative References	46
14.2. Informative References	47
Authors' Addresses	47

1. Introduction

Session recording is a critical requirement in many communications environments such as call centers and financial trading. In some of these environments, all calls must be recorded for regulatory, compliance, and consumer protection reasons. Recording of a session is typically performed by sending a copy of a media stream to a recording device. This document focuses on the Recording metadata which describes the communication session. The document describes a metadata model as viewed by Session Recording Server and the Recording metadata format, the requirements for which are described in [RFC6341] and the architecture for which is described in [I-D.ietf-siprec-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 [RFC2119]. This document only uses these key words when referencing normative statements in existing RFCs."

3. Definitions

Metadata Model: An abstract representation of metadata using a Unified Modelling Language(UML) class diagram.

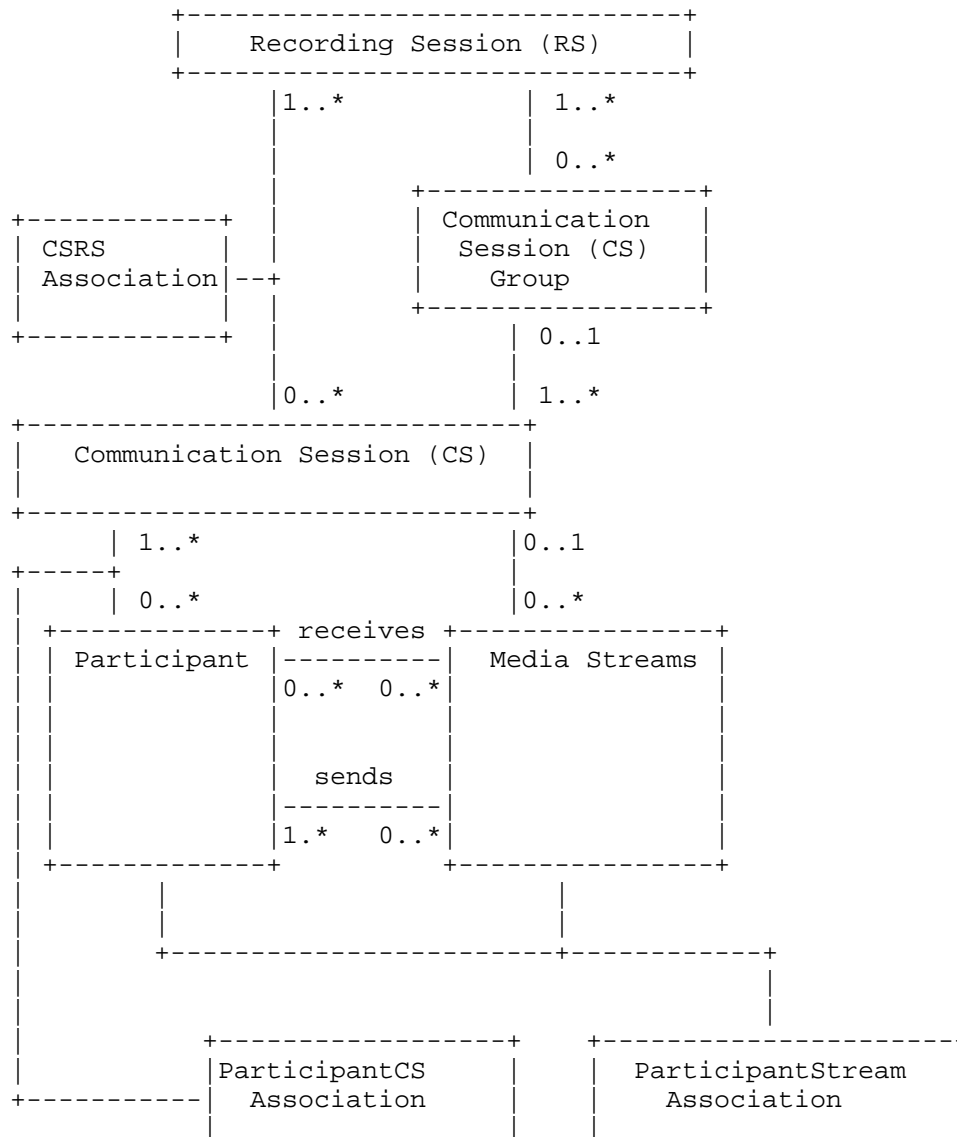
Metadata classes: Each block in the model represents a class. A class is a construct that is used as a blueprint to create instances(called objects) of itself. The description of each class also has representation of its attributes in a second compartment below the class name.

Attributes: Attributes represents the attributes listed in each of the classes. The attributes of a class are listed in the second compartment below the class name. Each instance of class conveys values for these attributes which adds to the recording's Metadata.

Linkages: Linkages represents the relationship between the classes in the model. It represents the logical connections between classes(or objects) in class diagrams/ object diagrams. The linkages used in the Metadata model of this document are associations.

4. Metadata Model

Metadata is the information that describes recorded media and the CS to which they relate. Below diagram shows a model for Metadata as viewed by Session Recording Server (SRS).



+-----+ +-----+

The Metadata model is a class diagram in Unified Modelling Language(UML). The model describes the structure of a metadata in general by showing the classes, their attributes, and the relationships among the classes. Each block in the model above represents a class. The linkages between the classes represents the relationships which can be associations or Composition. The metadata is conveyed from SRC to SRS.

The model allows the capture of a snapshot of a recording's Metadata at a given instant in time. Metadata changes to reflect changes in what is being recorded. For example, if in a conference a participant joins SRC sends a snapshot of metadata having that participant information (with attributes like name/AoR pair and associate-time) to the SRS.

Some of the metadata is not required to be conveyed explicitly from the SRC to the SRS, if it can be obtained contextually by the SRS(e.g., from SIP or SDP signalling).

5. Recording Metadata Format

This section gives an overview of Recording Metadata Format. Some data from the metadata model is assumed to be made available to the SRS through Session Description Protocol (SDP)[RFC4566], and therefore this data is not represented in the XML document format specified in this document. SDP attributes describes about different media formats like audio, video. The other metadata attributes like participant details are represented in a new Recording specific XML document namely application/rs-metadata+xml. The SDP label attribute [RFC4574] provides an identifier by which a metadata XML document can refer to a specific media description in the SDP sent from the SRC to the SRS.

The XML document format can be used to represent either the complete metadata or a partial update to the metadata. The latter includes only elements that have changed compared to the previously reported metadata.

5.1. XML data format

Recording Metadata document is an XML document. recording element MUST be present in all recording metadata XML document. recording acts as container for all other elements in this XML document.

Recording object is a XML document. It MUST have the XML declaration and it SHOULD contain an encoding declaration in the XML declaration, e.g., "<?xml version='1.0' encoding='UTF-8'?>". If the charset parameter of the MIME content type declaration is present and it is different from the encoding declaration, the charset parameter takes precedence.

Every application conforming to this specification MUST accept the UTF-8 character encoding to ensure the minimal interoperability.

Syntax and semantics error in recording XML document has to be informed to the originator using application specific mechanism.

5.1.1. Namespace

The namespace URI for elements defined by this specification is a Uniform Resource Namespace (URN) [RFC2141], using the namespace identifier 'ietf' defined by [RFC2648] and extended by [RFC3688].

The URN is as follows: urn:ietf:params:xml:ns:recording

5.1.2. recording


recording element MUST contain an xmlns namespace attribute with value as urn:ietf:params:xml:ns:recording. One recording element MUST be present in the all recording metadata XML document.

dataMode element shows whether the XML document is complete document or partial update. The default value is complete.

6. Recording Metadata classes

This section describes each class of the metadata model, and the attributes of each class. This section also describes how different classes are linked and the XML element for each of them.

6.1. Recording Session



Each instance of a Recording Session class (namely the Recording Session Object) represents a SIP session created between an SRC and SRS for the purpose of recording a Communication Session.

6.1.1. Attributes

A Recording Session class has the following attributes:

- o Start/End Time - Represents the Start/End time of a Recording Session object.

6.1.2. Linkages

Each instance of Recording Session has:


- o Zero or more instances of Communication Session Group. CSG may be zero because it is optional metadata object. Also the allowance of zero instances is to accommodate persistent recording, where there may be none.
- o Zero or more instances of Communication Session objects.

6.1.3. XML element

Recording Session object is represented by recording XML element. That in turn relies on the SIP/SDP session with which the XML document is associated to provide some of the attributes of the Recording Session element.

Start and End time value are derivable from Date header(if present in SIP message) in RS. In cases where Date header is not present, Start/End time are derivable from the time at which SRS receives the notification of SIP message to setup RS / disconnect RS.

6.2. Communication Session Group



One instance of a Communication Session Group class (namely the Communication Session Group object) provides association or linking of Communication Sessions.

6.2.1. Attributes

A CS Group has the following attributes:

- o Unique-ID - This Unique-ID is to group different CSs that are related. SRC (or SRS) is responsible for ensuring the uniqueness of Unique-ID in case multiple SRC interacts with the same SRS. The mechanism by which SRC groups the CS is outside the scope of SIPREC.
- o Associate-time - Associate-time for CS-Group shall be calculated by SRC as the time when a grouping is formed. The rules that determine how a grouping of different Communication Session objects is done by SRC is outside the scope of SIPREC.
- o Disassociate-time - Disassociate-time for CS-Group shall be calculated by SRC as the time when the grouping ends

6.2.2. Linkages

The linkages between Communication Session Group class and other classes is association. A communication Session Group is associated with RS and CS in the following manner:


- o There is one or more Recording Session objects per Communication Session Group.
- o Each Communication Session Group object has to be associated with one or more RS [Here each RS can be setup by the potentially different SRCs]
- o There is one or more Communication Sessions per CS Group [e.g. Consult Transfer]

6.2.3. XML element

Group element is an optional element provides the information about the communication session group

Each communication session group (CSG) object is represented using one group element. Each group element has unique Base 64 URN UUID attribute which helps to uniquely identify CSG.

6.3. Communication Session



A Communication Session class and its object in the metadata model represents Communication Session and its properties needed as seen by SRC.

6.3.1. Attributes

A communication Session class has the following attributes:

- o Termination Reason - This represents the reason why a CS was terminated. The communication session MAY contain a Call Termination Reason. This MAY be derived from SIP Reason header [RFC3326] of CS.
- o CS Identifier - This attribute is used to uniquely identify a CS.
- o Start-time - This optional attribute represents start time of CS as seen by SRC
- o Stop-time - This optional attribute represents stop time of CS as seen by SRC

This document does not specify attributes relating to what should happen to a recording of a CS after it has been delivered to the SRS, e.g., how long to retain the recording, what access controls to apply. The SRS is assumed to behave in accordance with policy. The ability for the SRC to influence this policy is outside the scope of this document. However if there are implementations where SRC has enough information, this could be sent as Extension Data attached to CS

6.3.2. Linkages

A Communication Session is linked to CS-Group, Participant, Media Stream and Recording Session classes using the association relationship. Association between CS and Participant allows:

- o CS to have atleast zero or more participants
- o Participant is associated with zero or more CSs. This includes participants who are not directly part of any CS. An example of such a case is participants in a premixed media stream. The SRC may have knowledge of such Participants, yet not have any signaling relationship with them. This might arise if one participant in CS is a conf focus. To summarize even if SRC does not have direct signalling relationships with all participants in a CS, it should nevertheless create a Participant object for each participant that it knows about.
- o The model also allows participants in CS that are not participants in the media. An example is the identity of a 3pcc controller that has initiated a CS to two or more participants of the CS. Another example is the identity of a conference focus. Of course a focus is probably in the media, but since it may only be there as a mixer, it may not report itself as a participant in any of the media streams.

Association between CS and Media Stream allows:

- o A CS to have zero or more Streams
- o A stream can be associated with at most one CS. Stream in persistent RS is not required to be associated with any CS before CS is created and hence the zero association is allowed.

Association between CS and RS allows:

- o Each instance of RS has Zero or more instances of Communication Session objects.
- o Each CS has to be associated with one more RS [Here each RS can be potentially setup by different SRCs]


6.3.3. XML element

Session element provides the information about the communication session

Each communication session(CS) object is represented by one session element. Each session element has unique Base 64 URN UUID attribute which helps to uniquely identify CS.

Reason element MAY be included to represent the Termination Reason attribute. group-ref element MAY exist to indicate the group where the mentioned session belongs.

6.4. CSRSAssociation



A CSRS Association class and its objects has attributes of CS object which are attributes of association of a session to a RS.

6.4.1. Attributes

CSRS association class has the following attributes:

- o Associate-time - associate-time is calculated by SRC as the time it sees a CS is associated to a RS
 - o Disassociate-time- Disassociate-time is calculated by SRC as the time it see a CS disassociate from a RS.
- It is possible that a given CS can have multiple associate/disassociate times within given RS.


6.4.2. Linkages

CSRS association class is linked to CS and RS classes. There are no cardinalties for this linkage.

6.4.3. XML element

sessionrecordingassoc is the XML element to represent CSRS association object. session URN UUID is used to uniquely identify this element and link with the specific session.

6.5. Participant



A Participant class and its objects has information about a device that is part of a CS and/or contributes/consumes media stream(s) belonging to a CS.

6.5.1. Attributes

Participant has attributes like:

- o AoR / Name pair list - This attribute is a list of Name/AoR tuple. An AoR MAY be SIP/SIPS/TEL URI. Name represents Participant name(SIP display name) or DN number (in case it is known). There are cases where a participant can have more than one AoR [e.g. P-Asserted-identity header [RFC3325] which can have both SIP and TEL URIs]

This document does not specify other attributes relating to participant e.g. Participant Role, Participant type. An SRC which has information of these attributes can indicate the same as part of extension data to Participant from SRC to SRS.

6.5.2. Linkages

The participant class is linked to MS and CS class using association relationship. The association between participant and Media Stream allows:

- o Participant to receives zero or more media streams
- o Participant to send zero or more media streams. (Same participant provides multiple streams e.g. audio and video)
- o Media stream to be received by zero or more participants. Its possible, though perhaps unlikely, that a stream is generated but sent only to the SRC and SRS, not to any participant. E.g. In conferencing where all participants are on hold and the SRC is collocated with the focus. Also a media stream may be received by multiple participants (e.g. Whisper calls, side conversations).
- o Media stream to be sent by one or more participants (pre-mixed streams).

Example of a case where a participant receives Zero or more streams - a Supervisor may have side conversation with Agent, while Agent converses with customer.

6.5.3. XML element

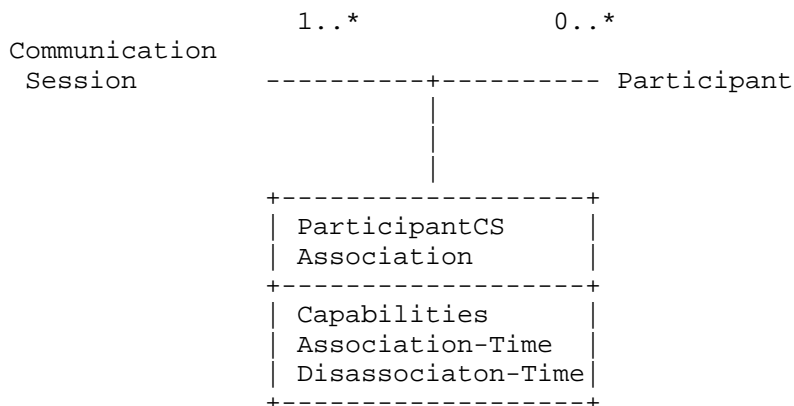
A participant element represents a Participant object.

Participant MUST have a NameID complex element which contains AoR as attribute and Name as element. AOR element is SIP/SIPS URI FQDN or IP address which represents the user. name is an optional element to represent display name.

Each participant element has unique ID (Base 64 URN UUID) attribute

which helps to uniquely identify participant and session Base 64 URN UUID to associate participant with specific session element. Base 64 URN UUID of participant MUST used in the scope of CSG and no new Base 64 URN UUID has to be created for the same element (participant, stream) between different CS in the same CSG. In case Base 64 URN UUID has to be used permanent, careful usage of Base 64 URN UUID to original AoR has to be decided by the implementers and it is implementer's choice.

6.6. ParticipantCSAssociation



A participantCS Association class and its objects has attributes of participant object which are attributes of association of a participant to a Session.

6.6.1. Attributes

ParticipantCS association class has the following attributes:

- o Associate-time - associate-time is calculated by SRC as the time it sees a participant is associated to CS
- o Disassociate-time- Disassociate-time is calculated by SRC as the time it see a participant disassociate from a CS. It is possible that a given participant can have multiple associate/disassociate times within given communication session.
- o Capabilities - A participant capabilities as defined in [RFC3840] which is an optional attribute that includes the capabilities of a participant in a CS. Each participant shall have Zero or more capabilities. A participant may use different capabilities depending on the role it plays at a particular instance. IOW if a

participants moves across different CSs (due to transfer e.t.c)
OR is simultaneously present in different CSs its role may be
different and hence the capability used.

- o "send" or "recv" element in each participant is associating SDP m-lines with the participant. send element indicates that participant is sending the stream of media with the mentioned media description. recv element indicates that participant is receiving the stream and by default all participant will receive the stream. recv element has relevance in case whisper call scenario wherein few of the participant in the session receives the stream and not others.

6.6.2. Linkages


The participantCS association class is linked to participant and CS classes. There are no cardinalities for this linkage.

6.6.3. XML element

participantsessionassoc XML element represent participantCS association object. participant and session id is used to uniquely identify this element

NOTE: RFC 4235 encoding shall be used to represent capabilities attribute in XML.

6.7. Media Stream



A Media Stream class (and its objects) has the properties of media as seen by SRC and sent to SRS. Different snapshots of media stream object may be sent whenever there is a change in media (e.g. dir

change like pause/resume and/or codec change and/or participant change.)).

6.7.1. Attributes

A Media Stream class has the the following attributes:

- o Media Stream Reference - In implementations this can reference to m-line
- o Content - The content of an MS element will be described in terms of value from the [RFC4796] registry.

The metadata model should include media streams that are not being delivered to the SRS. Examples include cases where SRC offered certain media types but SRS chooses to accept only a subset of them OR an SRC may not even offer a certain media type due it its restrictions to record

6.7.2. Linkages

A Media Stream is linked to participant and CS classes using the association relationship. The details of association with the Participant are described in the Participant class section. The details of association with CS is mentioned in the CS section.

6.7.3. XML element

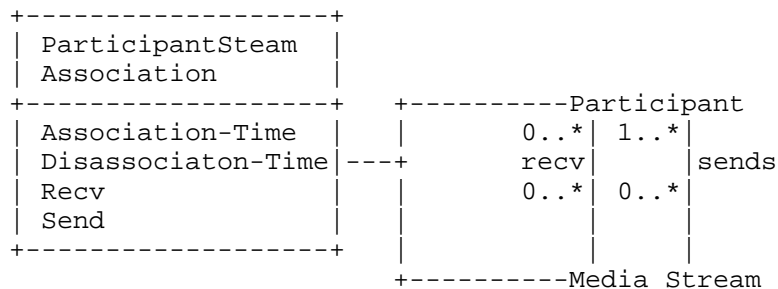
stream element represents a Media Stream object. Stream element indicates SDP media lines associated with the session and participants.

This element indicates the SDP m-line properties like label attributes. Label attribute is used to link m-line SDP body using label attribute in SDP m-line.

Each stream element has unique Base 64 URN UUID attribute which helps to uniquely identify stream and session Base 64 URN UUID to associate stream with specific session element.

The content attribute if an SRC wishes to send is conveyed in RS SDP.

6.8. ParticipantStream Association



A ParticipantStream association class and its object has attributes that are attributes of association of a Participant to a Stream.

6.8.1. Attributes

A participantStream association class has the following attributes:

- o Associate-Time: This attributes indicates the time a Participant started contributing to a Media Stream
- o Disassociate-Time: This attribute indicates the time a Participant stopped contributing to a Media Stream

6.8.2. Linkages

The participantStream association class is linked to participant and Stream classes. There are no cardinalities for this linkage.

6.8.3. XML element

ParticipantStreamAssoc XML element represents participant to stream association object. participant element is used to uniquely identify this element and related with stream using stream unique URN id..

6.9. associate-time/disassociate-time

associate-time/disassociate-time contains a string indicating the date and time of the status change of this tuple. The value of this element MUST follow the IMPP datetime format [RFC3339]. Timestamps that contain 'T' or 'Z' MUST use the capitalized forms. At a time, any of the time tuple associate-time or disassociate-time MAY exist in the element namely group, session, participant and not both timestamp at the same time.

As a security measure, the timestamp element SHOULD be included in all tuples unless the exact time of the status change cannot be

determined.

6.10. Unique ID format

Unique id is generated in two steps:

- o UUID is created using [RFC4122])
- o UUID is encoded using base64 as defined in [RFC4648]

The above mentioned unique-id mechanism SHOULD be used for each metadata element.

7. SIP Recording Metadata Example

7.1. Complete SIP Recording Metadata Example

The following example provides all the tuples involved in Recording Metadata XML body.

```
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>complete</dataMode>
    <group group_id="7+OTCyoxTmqmqyA/1weDag==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
      <!-- Standardized extension -->
      <call-center xmlns='urn:ietf:params:xml:ns:callcenter'>
        <supervisor>sip:alice@atlanta.com</supervisor>
        <uniqueid>FaXHlc+3WruaroDaNE87am==</uniqueid>
      </call-center>
      <mydata xmlns='http://example.com/my'>
        <structure>FOO!</structure>
        <whatever>bar</whatever>
      </mydata>
    </group>
    <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <group-ref>7+OTCyoxTmqmqyA/1weDag==
      </group-ref>

      <!-- Standardized extension -->
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </session>
    <sessionrecordingassoc session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </sessionrecordingassoc>
    <participant
      participant_id="srfBEImCRp2QB23b7Mpk0w==">
```

```
<nameID aor=sip:bob@biloxi.com>
  <name xml:lang="it">Bob B</name>
</nameID>
<!-- Standardized extension -->
<structure>FOO!</structure>
<whatever>bar</whatever>
</participant>
<participantsessionassoc
  participant_id="srfBEImCRp2QB23b7Mpk0w=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="srfBEImCRp2QB23b7Mpk0w==">
  <send>ilPz3to5hGk8fuXl+PbwCw==</send>
  <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
  <recv>8zc6e0lYTlWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<participant
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  <nameID aor=sip:Paul@biloxi.com>
    <name xml:lang="it">Paul</name>
  </nameID>
  <!-- Standardized extension -->
  <structure>FOO!</structure>
  <whatever>bar</whatever>
</participant>
<participantsessionassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
  <send>8zc6e0lYTlWIINA6GR+3ag==</send>
  <send>EiXGlc+4TruqqoDaNE76ag==</send>
  <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
  <recv>ilPz3to5hGk8fuXl+PbwCw==</recv>
</participantstreamassoc>
<stream stream_id="UAAMm5GRQKSCMVvLyl4rFw=="
  session session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>96</label>
</stream>
<stream stream_id="ilPz3to5hGk8fuXl+PbwCw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>97</label>
</stream>
```

```

    <stream stream_id="8zc6e0lYTlWIINA6GR+3ag=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <label>98</label>
    </stream>
    <stream stream_id="EiXGlc+4TruqqoDaNE76ag=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <label>99</label>
    </stream>
  </recording>

```

SIP Recording Metadata Example XML body

7.2. Partial Update of Recording metadata XML body

The following example provides partial update in Recording Metadata XML body for the above example. The example has a snapshot that carries the disassociate-time for a participant from a session.

```

<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <participant
      participant_id="srfBElmCRp2QB23b7Mpk0w==">
      <name ID=sip:bob@biloxi.com>
        <name xml:lang="it">Bob R</name>
      </nameID>
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </participant>
    <participantsessionassoc
      participant_id="srfBElmCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ=="
      <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
    </participantsessionassoc>
  </recording>

```

Partial update of SIP Recording Example XML body

8. XML Schema definition for Recording metadata

This section defines XML schema for Recording metadata document

```

<?xml version="1.0" encoding="UTF-8"?>

```

```
<xs:schema targetNamespace="urn:ietf:params:xml:ns:recording"
  xmlns:tns="urn:ietf:params:xml:ns:recording"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">
  <!-- This import brings in the XML language attribute xml:lang-->
  <xs:import namespace="http://www.w3.org/XML/1998/namespace"/>
  <xs:element name="recording" type="recording"/>
  <xs:complexType name="recording">
    <xs:sequence>
      <xs:element name="datamode" type="dataMode"
        minOccurs="0"/>
      <xs:element name="group" type="group"
        minOccurs="0" maxOccurs="unbounded"/>
      <xs:element name="session" type="session"
        minOccurs="0" maxOccurs="unbounded"/>
      <xs:element name="participant" type="participant"
        minOccurs="0" maxOccurs="unbounded"/>
      <xs:element name="stream" type="stream"
        minOccurs="0" maxOccurs="unbounded"/>
      <xs:any namespace='##other'
        minOccurs='0'
        maxOccurs='unbounded'
        processContents='lax'/>
    </xs:sequence>
  </xs:complexType>
  <xs:complexType name="group">
    <xs:sequence>
      <xs:element name="associate-time" type="xs:dateTime"
        minOccurs="0"/>
      <xs:element name="disassociate-time" type="xs:dateTime"
        minOccurs="0"/>
      <xs:any namespace='##other'
        minOccurs='0'
        maxOccurs='unbounded'
        processContents='lax'/>
    </xs:sequence>
    <xs:attribute name="group_id" type="xs:base64Binary"
      use="required"/>
  </xs:complexType>
  <xs:complexType name="session">
    <xs:sequence>
      <xs:element name="reason" type="xs:string"
        minOccurs="0"/>
      <xs:element name="group-ref" type="xs:base64Binary"
        minOccurs="0" maxOccurs="1"/>
      <xs:any namespace='##other'
        minOccurs='0'
        maxOccurs='unbounded'
        processContents='lax'/>
    </xs:sequence>
  </xs:complexType>
```



```
        maxOccurs='unbounded'
        processContents='lax' />
    </xs:sequence>
    <xs:attribute name="session_id" type="xs:base64Binary"
        use="required" />
</xs:complexType>
<xs:complexType name="sessionrecordingassoc">
    <xs:sequence>
        <xs:element name="associate-time" type="dateTime"
            minOccurs="0" />
        <xs:element name="disassociate-time" type="dateTime"
            minOccurs="0" />
        <xs:any namespace='##other'
            minOccurs='0'
            maxOccurs='unbounded'
            processContents='lax' />
    </xs:sequence>
    <xs:attribute name="session_id" type="xs:base64Binary"
        use="required" />
</xs:complexType>
<xs:complexType name="participant">
    <xs:sequence>
        <xs:element name="nameID" type="nameID"
            maxOccurs="1" />
        <xs:any namespace='##other'
            minOccurs='0'
            maxOccurs='unbounded'
            processContents='lax' />
    </xs:sequence>
    <xs:attribute name="participant_id" type="xs:base64Binary"
        use="required" />
</xs:complexType>
<xs:complexType name="participantsessionassoc">
    <xs:sequence>
        <xs:element name="associate-time" type="xs:dateTime"
            minOccurs="0" />
        <xs:element name="disassociate-time" type="xs:dateTime"
            minOccurs="0" />
        <xs:any namespace='##other'
            minOccurs='0'
            maxOccurs='unbounded'
            processContents='lax' />
    </xs:sequence>
    <xs:attribute name="participant_id" type="xs:base64Binary"
        use="required" />
    <xs:attribute name="session_id" type="xs:base64Binary"
        use="required" />
</xs:complexType>
```

```
<xs:complexType name="participantstreamassoc">
  <xs:sequence>
    <xs:element name="send" type="xs:base64Binary"
      minOccurs="0" maxOccurs="unbounded"/>
    <xs:element name="recv" type="xs:base64Binary"
      minOccurs="0" maxOccurs="unbounded"/>
    <xs:element name="associate-time" type="xs:dateTime"
      minOccurs="0"/>
    <xs:element name="disassociate-time" type="xs:dateTime"
      minOccurs="0"/>
    <xs:any namespace='##other'
      minOccurs='0'
      maxOccurs='unbounded'
      processContents='lax'/>
  </xs:sequence>
  <xs:attribute name="participant_id" type="xs:base64Binary"
    use="required"/>
</xs:complexType>
<xs:complexType name="stream">
  <xs:sequence>
    <xs:element name="label" type="xs:string"
      minOccurs="0" maxOccurs="1"/>
    <xs:any namespace='##other'
      minOccurs='0'
      maxOccurs='unbounded'
      processContents='lax'/>
  </xs:sequence>
  <xs:attribute name="stream_id" type="xs:base64Binary"
    use="required"/>
  <xs:attribute name="session_id" type="xs:base64Binary"/>
</xs:complexType>
<xs:simpleType name="dataMode">
  <xs:restriction base="xs:string">
    <xs:enumeration value="complete"/>
    <xs:enumeration value="partial"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="nameID">
  <xs:sequence>
    <xs:element name="name" type="name" minOccurs="0" maxOccurs="1"/>
  </xs:sequence>
  <xs:attribute name="aor" type="xs:anyURI" use="required"/>
</xs:complexType>

<xs:complexType name="name">
  <xs:simpleContent>
    <xs:extension base="xs:string">
```

```
        <xs:attribute ref="xml:lang" use="optional"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>

</xs:schema>
```

9. Security Considerations

The metadata information sent from SRC to SRS MAY reveal sensitive information about different participants in a session. For this reason, it is RECOMMENDED that a SRC use a strong means for authentication and metadata information protection and that it apply comprehensive authorization rules when using the metadata format defined in this document. The following sections will discuss each of these aspects in more detail.

9.1. Connection Security

It is RECOMMENDED that a SRC authenticate SRS using the normal SIP authentication mechanisms, such as Digest as defined in Section 22 of [RFC3261]. The mechanism used for conveying the metadata information MUST ensure integrity and SHOULD ensure confidentiality of the information. In order to achieve these, an end-to-end SIP encryption mechanism, such as S/MIME described in [RFC3261], SHOULD be used.

If a strong end-to-end security means (such as above) is not available, it is RECOMMENDED that a SRC use mutual hop-by-hop Transport Layer Security (TLS) authentication and encryption mechanisms described in "SIPS URI Scheme" and "Interdomain Requests" of [RFC3261].

10. IANA Considerations

This specification registers a new XML namespace, and a new XML schema.

10.1. SIP recording metadata Schema Registration

URI: urn:ietf:params:xml:ns:recording

Registrant Contact: IETF SIPREC working group, Ram mohan
R(rmohanr@cisco.com)

XML: the XML schema to be registered is contained in Section 6.

Its first line is `<?xml version="1.0" encoding="UTF-8"?>` and its last line is `</xs:schema>`

11. Acknowledgement

We wish to thank John Elwell(Siemens-Enterprise), Henry Lum(Alcatel-Lucent), Leon Portman(Nice), De Villers, Andrew Hutton(Siemens-Enterprise), Deepanshu Gautam(Huawei), Charles Eckel(Cisco), Muthu Arul(Cisco), Michael Benenson(Cisco), Hadriel Kaplan (ACME), Brian Rosen(Neustar), Scott Orton(Broadsoft), Ofir Roth for their valuable comments and inputs.

We wish to thank Joe Hildebrand(Cisco), Peter Saint-Andre(Cisco) for the valuable XML related guidance and Martin Thompson for validating the XML schema and providing comments on the same.


12. Appendix A: Metadata Model Object Instances

Note: This Appendix has to be moved to callflow document after the discussion in the mailing alias

This section describes the metadata model object instances for different use cases of SIPREC. For the sake of simplicity as the media streams sent by each of the participants is received by every other participant in these use cases, it is NOT shown in the object instance diagrams below. Also for the sake of ease not all attributes of each object are shown in these instance diagrams.


12.1. Use case 1: Basic Call

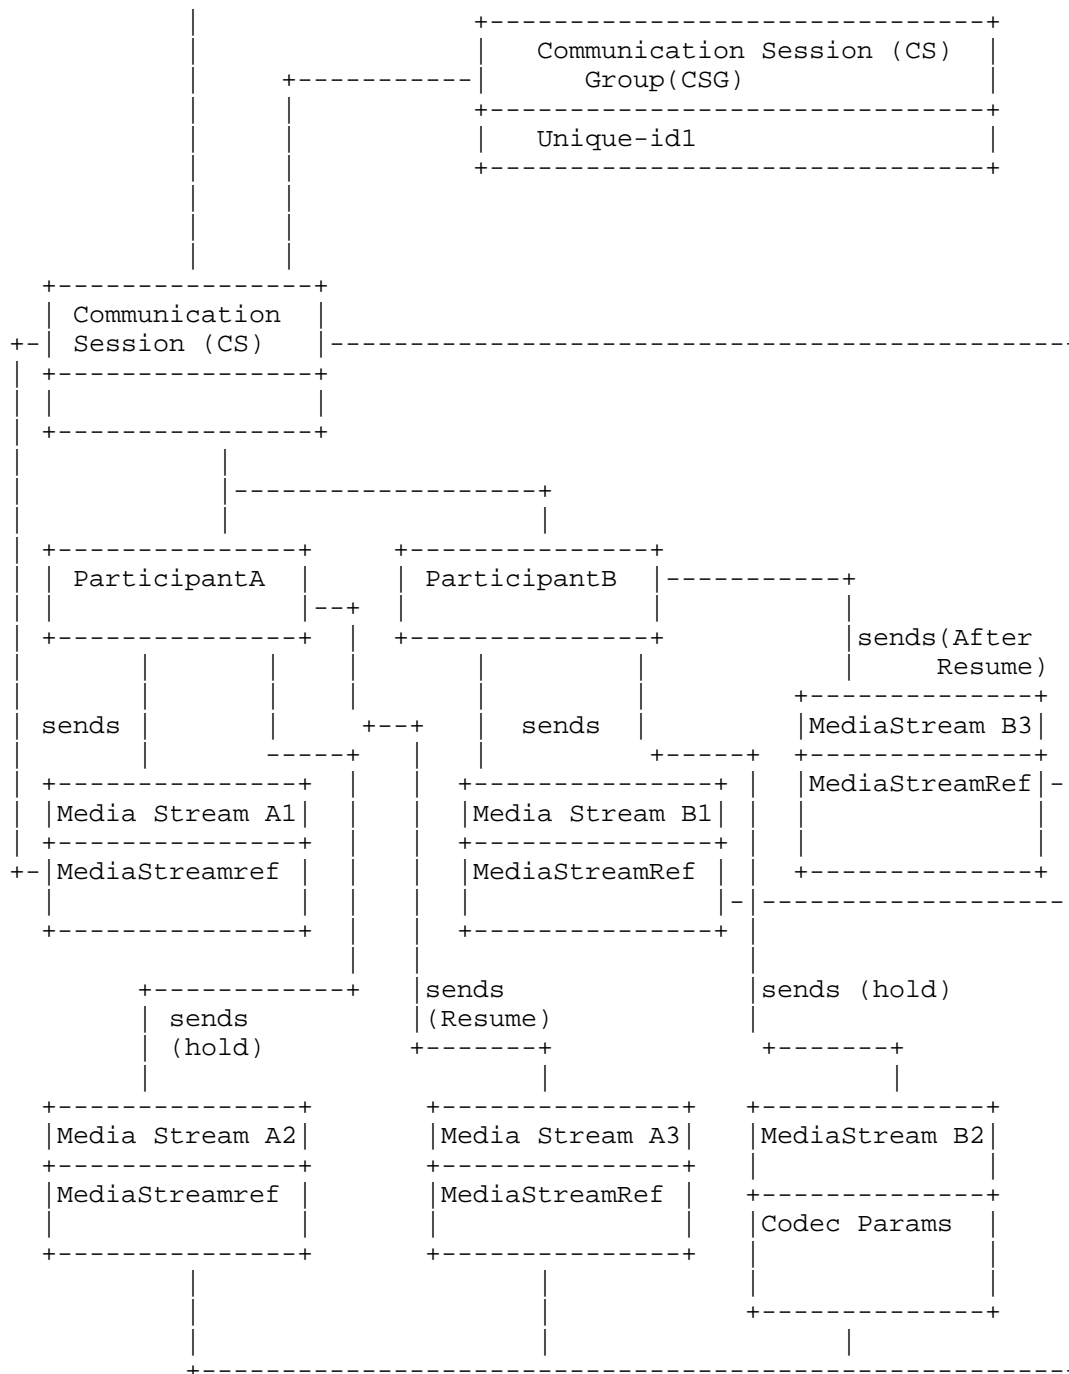
Basic call between two Participants A and B. In this use case each participant sends one Media Stream. For the sake of simplicity "receives" lines are not shown in this instance diagram. Media Streams sent by each participant is received all other participants of that CS.



12.2. Use case 2: Hold/Resume


Basic call between two Participants A and B and with Participant A or B doing a Hold/Resume. In this use case each participant sends one Media Stream. After Hold/Resume the properties of Media can change. For the sake of simplicity "receives" lines are not shown in this instance diagram. Media Streams sent by each participant is received all other participants of that CS.






12.3. Use case 3: Basic call with Transfer

Basic call between two Participants A and B and with Participant A transfer(consult transfer) to Participant C. In this use case each participant sends one Media Stream. After transfer the properties of Participant A Media can change. For the sake of simplicity "receives" lines are not shown in this instance diagram. Media Streams sent by each participant is received all other participants of that CS.





12.4. Conference Use Cases


Depending on who act as SRC and the information that an SRC has there can be several ways to model conference use cases. This section has instance diagrams for the following cases:

- o A CS where one of the participant (which is also SRC) is a user in a conference
- o A CS where one of the participant is focus (which is also SRC)
- o A CS where one of the participant is user and the SRC is a different entity like B2BUA
- o A CS where one of the participant is focus and the SRC is a different entity like B2BUA


NOTE: There MAY be other ways to model the same use cases depending on what information the SRC has.

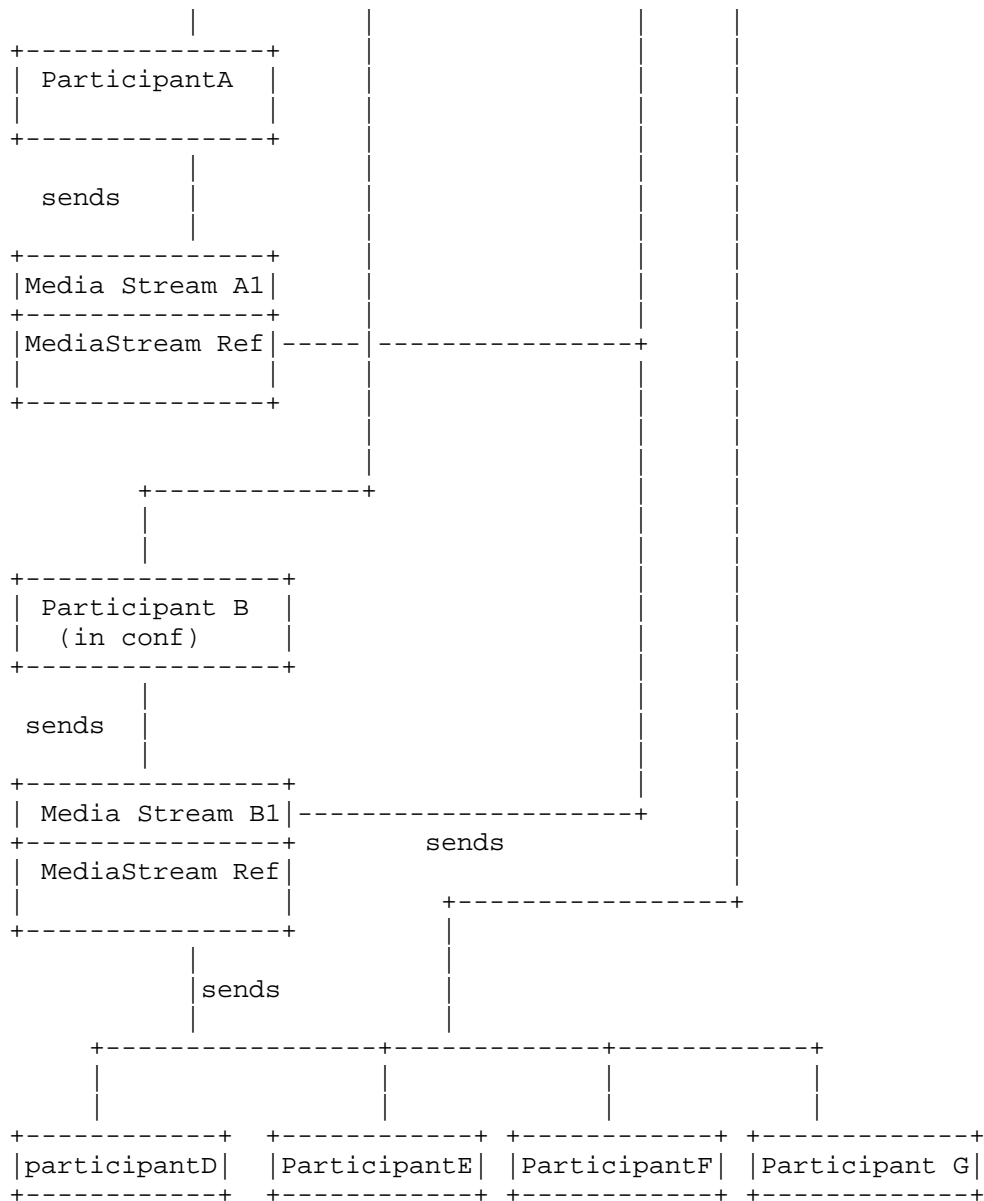
12.4.1. Case 1:

This is the usecase where there is a CS with one of the participant (who is also SRC) as a user in a conference. For the sake of simplicity the receive lines for each of the participant is not shown.



Instance Diagram:



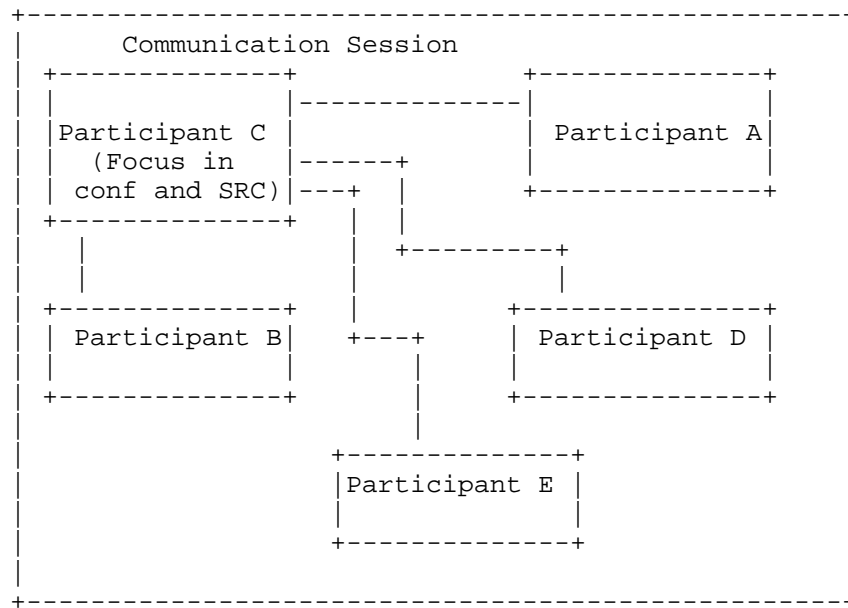


In this example we have two participants A and B who are part of a Communication Session(CS). One of the participants B is part of a conference and also acts as SRC. There can be two cases here. B can be a participant of the conference or B can be a focus. In this


instance diagram Participant B is a user in a conference. The SRC (Participant B) subscribes to conference event package to get the details of other participants. Participant B(SRC) sends the same through the metadata to SRS. In this instance diagram the Media Stream(mixed stream) sent from Participant B has media streams contributed by conference participants (D,E,F and G). For the sake of simplicity the "receives" line is not shown here. In this example the media stream sent by each participant(A or B) of CS is received by all other participant(A or B).


12.4.2. Case 2:

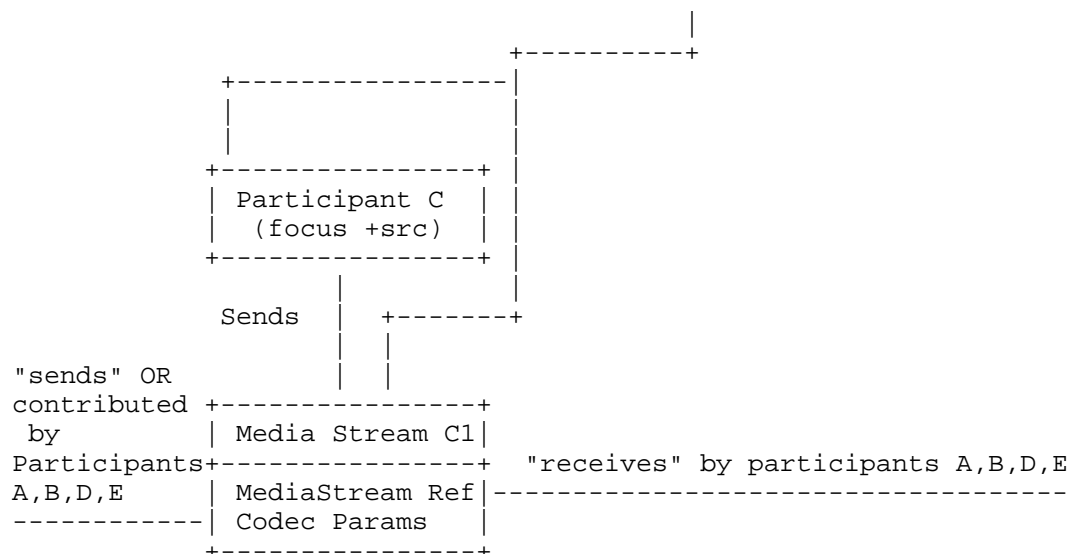
This is the usecase where there is a CS where one of the participant is focus (which is also SRC).



Instance Diagram:








In this example we have two participants A and B who are part of a Communication Session(CS). One of the participants (C) is focus of a conference and also acts as SRC. The SRC (Participant C) being the Focus of the conference has access to the details of other participants. SRC (Participant C) sends the same through the metadata to SRS. In this instance diagram the Media Stream(mixed stream) sent by C has media streams contributed by conference participants (A, B, D and E). Participants A, B,D and E sends Media Streams A1, B1, D1 and E1 respectively. The media stream sent by Participant C(Focus) is received by all other participants of CS. For the sake of simplicity the "receives" line is not shown linked to all other participants.

NOTE: SRC (Participant C) can send mixed stream or seperate streams to SRS

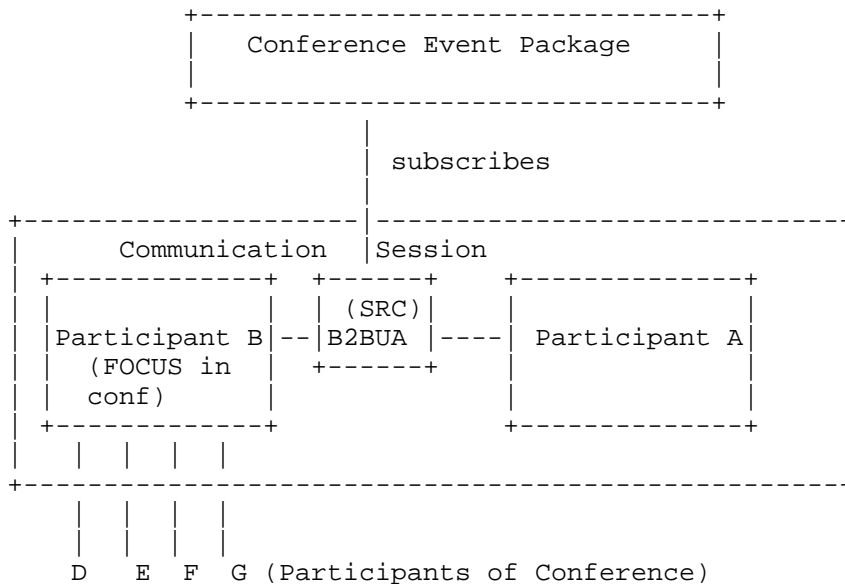
12.4.3. Case 3:

A CS where one of the participant is user and the SRC is a different entity like B2BUA. In this case the SRC may not know that one of the user is part of conference. Hence the instance diagram will not have information about the conference participants.



12.4.4. Case 4:

A CS where one of the participant is focus and the SRC is a different entity like B2BUA. In this case the participant which is focus sends "isfocus" in SIP message to SRC. The SRC subscribe to conference event package on seeing this "isfocus". SRC learns the details of other participants of conference from the conference package and send the same in metadata to SRS. The instance diagram for this use case is same as Case 1.



13. Appendix B: Metadata XML schema Instances

Note: This Appendix has to be moved to callflow document after the discussion in the mailing alias

This section describes the metadata model XML instances for different use cases of SIPREC. For the sake of simplicity the complete SIP messages are NOT shown here.

13.1. Use case 1: Basic Call

Basic call between two Participants A(Alice) and B(Bob) who are part of one session. In this use case each participant sends two Media Streams. Media Streams sent by each participant is received all other participants of that CS in this use-case. Below is the initial snapshot sent by SRC that has complete metadata. For the sake of completeness even snippets of SDP is shown. For the sake of simplicity these use-cases assume the RS stream is unmixed.

```
Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>complete</dataMode>
    <group group_id="7+OTCyoxTmqmqyA/lweDAg==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
      <!-- Standardized extension -->
```

```
<call-center xmlns='urn:ietf:params:xml:ns:callcenter'>
  <supervisor>sip:alice@cisco.com</supervisor>
</call-center>
<mydata xmlns='http://example.com/my'>
  <structure>FOO!</structure>
  <whatever>bar</whatever>
</mydata>
</group>
<session session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <group-ref>7+OTCyoxTmqmqyA/lweDAg==
</group-ref>
  <start-time>2010-12-16T23:41:07Z</start-time>
  <!-- Standardized extension -->
  <structure>FOO!</structure>
  <whatever>bar</whatever>
</session>
<participant
  participant_id="srfBEImCRp2QB23b7Mpk0w=="
  <nameID aor=sip:alice@atlanta.com>
    <name xml:lang="it">Alice</name>
  </nameID>
  <!-- Standardized extension -->
  <structure>FOO!</structure>
  <whatever>bar</whatever>
</participant>
<participantsessionassoc
  participant_id="srfBEImCRp2QB23b7Mpk0w=="
  session_id="hVpd7YQgRW2nD22h7q60JQ=="
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="srfBEImCRp2QB23b7Mpk0w=="
  <send>i1Pz3to5hGk8fuXl+PbwCw==</send>
  <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
  <recv>8zc6e0lYTlWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<participant
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  <nameID aor=sip:bob@biloxi.com>
    <name xml:lang="it">Bob</name>
  </nameID>
  <!-- Standardized extension -->
  <structure>FOO!</structure>
  <whatever>bar</whatever>
</participant>
<participantsessionassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
```



```

        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <associate-time>2010-12-16T23:41:07Z</associate-time>
    </participantsessionassoc>
    <participantstreamassoc
        participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
        <send>8zc6e0lYTlWIINA6GR+3ag==</send>
        <send>EiXGlc+4TruqqoDaNE76ag==</send>
        <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
        <recv>ilPz3to5hGk8fuXl+PbwCw==</recv>
    </participantstreamassoc>
    <stream id="UAAMm5GRQKSCMVvLyl4rFw=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <label>96</label>
    </stream>
    <stream id="ilPz3to5hGk8fuXl+PbwCw=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <label>97</label>
    </stream>
    <stream id="8zc6e0lYTlWIINA6GR+3ag=="
        session_id="zSfPoSvdSDCmU3A3TRDxAw==">
        <label>98</label>
    </stream>
    <stream id="EiXGlc+4TruqqoDaNE76ag=="
        session_id="zSfPoSvdSDCmU3A3TRDxAw==">
        <label>99</label>
    </stream>
</recording>

```

13.2. Use case 2: Hold/resume

Basic call between two Participants A and B. This is the continuation of above use-case. One of the participants(say A) goes on hold and then resumes as part of the same session. The metadata snapshot looks as below

During hold

```
Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=inactive
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=inactive
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....

<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <participantstreamassoc
      participant_id="srfBElmCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
      <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
    </participantstreamassoc>
    <participantstreamassoc
      participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <send>8zc6e0lYtLWIINA6GR+3ag==</send>
      <send>EiXGlc+4TruqqoDaNE76ag==</send>
    </participantstreamassoc>
  </recording>
```

During resume

The snapshot will look pretty much same as Use-case 1.

13.3. Use case 3: Basic Call with transfer

Basic call between two Participants A and B is connected as in Use-case 1. Transfer is initiated by one of the participants or by other entity(3PCC case). SRC sends a snapshot of the participant changes to SRS. In this instance participant A(Alice) drops out during the transfer and Participant C(Paul) joins the session. There can be two cases here, same session continues after transfer or a new session (e.g. REFER based transfer) is created

Transfer with same session retained - (.e.g. RE-INVITE based transfer). Participant A drops out and C is added to the same session. No change to session/group element. C will be new stream element which maps to RS SDP using the same labels in this instance.

```
Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <participantstreamassoc
      participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
      <send>8zc6e0lYtLWIINA6GR+3ag==</send>
      <send>EiXGlc+4TruqqoDaNE76ag==</send>
      <recv>60JAJm9UTvik0Ltlih/Gzw==</recv>
      <recv>AcR5FUd3Edi8cACQJy/3JQ==</recv>
    </participantstreamassoc>
```

```

<participant
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ==">
    <nameID aor=sip:paul@example.com>
      <name xml:lang="it">Paul</name>
    </nameID>
    <!-- Standardized extension -->
    <structure>FOO!</structure>
    <whatever>bar</whatever>
  </participant>
<participantsessionassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
    <associate-time>2010-12-16T23:41:07Z</associate-time>
  </participantsession>
<participantstreamassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ==">
    <send>60JAJm9UTvik0Ltlih/Gzw==</send>
    <send>AcR5FUd3Edi8cACQJy/3JQ==</send>
    <recv>8zc6e0lYTlWIINA6GR+3ag==</recv>
    <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
  </participantstreamassoc>
<stream stream_id="60JAJm9UTvik0Ltlih/Gzw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>96</label>
</stream>
<stream stream_id="AcR5FUd3Edi8cACQJy/3JQ=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>97</label>
</stream>
<stream stream_id="8zc6e0lYTlWIINA6GR+3ag=="
  session_id="zSfPoSvdSDCmU3A3TRDxAw==">
  <label>98</label>
</stream>
<stream stream_id="EiXGlc+4TruqqoDaNE76ag=="
  session_id="zSfPoSvdSDCmU3A3TRDxAw==">
  <label>99</label>
</stream>
</recording>

```

Transfer with new session - (.e.g. REFER based transfer). In this case new session is part of same grouping (done by SRC).

SRC may send an optional snapshot indicating stop for the old session.

```

    <?xml version="1.0" encoding="UTF-8"?>
    <recording xmlns='urn:ietf:params:xml:ns:recording'>
      <dataMode>Partial</dataMode>
      <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <group-ref>7+OTCyoxTmqmqyA/lweDAg==
        </group-ref>
        <stop-time>2010-12-16T23:41:07Z</stop-time>
        <!-- Standardized extension -->
        <structure>FOO!</structure>
        <whatever>bar</whatever>
      </session>
      <participantsessionassoc
        participant_id="srfBEImCRp2QB23b7Mpk0w=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
      </participantsessionassoc>
      <participantsessionassoc
        participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
      </participantsessionassoc>
    </recording>

```

SRC sends a snapshot to indicate the participant change and new session information after transfer. In this example the same RS is used.

Content-Type: application/SDP

```

...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000

```

```

a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <session session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
      <group-ref>7+OTCyoxTmqmqyA/lweDAg==</group-ref>
      <start-time>2010-12-16T23:41:07Z</start-time>
      <!-- Standardized extension -->
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </session>
    <participant
      participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
      <nameID aor=sip:Alice@atlanta.com/>
      <!-- Standardized extension -->
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </participant>
    <participantsessionassoc
      participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
      session_id="bfLZ+NTFEeCNxQTuRyQBmw=="
      <associate-time>2010-12-16T23:32:03Z</associate-time>
    </participantsessionassoc>
    <participantstreamassoc
      participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
      <send>8zc6e0lYtLWIINA6GR+3ag==</send>
      <send>EiXGlc+4TruqqoDaNE76ag==</send>
      <recv>60JAJm9UTvik0Ltlih/Gzw==</recv>
      <recv>AcR5FUd3Edi8cACQJy/3JQ==</recv>
    </participantstreamassoc>
    <participant
      participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
      <nameID aor=sip:Bob@biloxi.com/>
      <!-- Standardized extension -->
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </participant>
    <participantsessionassoc
      participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
      session_id="bfLZ+NTFEeCNxQTuRyQBmw=="
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </participantsessionassoc>
    <participantstreamassoc
      participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
      <send>60JAJm9UTvik0Ltlih/Gzw==</send>

```

```
<send>AcR5FUd3Edi8cACQJy/3JQ==</send>
<recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
<recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<stream stream_id="60JAJm9UTvik0Ltlih/Gzw=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>96</label>
</stream>
<stream stream_id="AcR5FUd3Edi8cACQJy/3JQ=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>97</label>
</stream>
<stream stream_id="8zc6e0lYtLWIINA6GR+3ag=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>98</label>
</stream>
<stream stream_id="EiXGlc+4TruqqoDaNE76ag=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>99</label>
</stream>
</recording>
```

13.4. Use Case 4: Call disconnect

This example shows a snapshot of metadata sent by an SRC at CS disconnect where the participants of CS are Alice and Bob

```
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>Partial</dataMode>
    <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <group-ref>7+OTCyoxTmqmqyA/lweDAg==</group-ref>
      <stop-time>2010-12-16T23:41:07Z</stop-time>
      <!-- Standardized extension -->
      <structure>FOO!</structure>
      <whatever>bar</whatever>
    </session>
    <participant
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      <nameID aor=sip:Alice@atlanta.com/>
      <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
    </participant>

    <participant
      participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
      <nameID aor=sip:Bob@biloxxy.com/>
      <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
    </participant>
  </recording>
```

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2141] Moats, R., "URN Syntax", RFC 2141, May 1997.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [RFC3339] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session

Description Protocol", RFC 4566, July 2006.

- [RFC4574] Levin, O. and G. Camarillo, "The Session Description Protocol (SDP) Label Attribute", RFC 4574, August 2006.
- [RFC4796] Hautakorpi, J. and G. Camarillo, "The Session Description Protocol (SDP) Content Attribute", RFC 4796, February 2007.
- [RFC3840] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", RFC 3840, August 2004.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.

14.2. Informative References

- [RFC6341] Rehor, K., Portman, L., Hutton, A., and R. Jain, "Use Cases and Requirements for SIP-Based Media Recording (SIPREC)", RFC 6341, August 2011.
- [I-D.ietf-siprec-architecture]
Hutton, A., Portman, L., Jain, R., and K. Rehor, "An Architecture for Media Recording using the Session Initiation Protocol", draft-ietf-siprec-architecture-06 (work in progress), September 2012.
- [RFC2648] Moats, R., "A URN Namespace for IETF Documents", RFC 2648, August 1999.
- [RFC3326] Schulzrinne, H., Oran, D., and G. Camarillo, "The Reason Header Field for the Session Initiation Protocol (SIP)", RFC 3326, December 2002.
- [RFC3325] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", RFC 3325, November 2002.

Authors' Addresses

Ram Mohan Ravindranath
Cisco Systems, Inc.
Cessna Business Park,
Kadabeesanahalli Village, Varthur Hobli,
Sarjapur-Marathahalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: rmohanr@cisco.com

Parthasarathi Ravindran
Nokia Siemens Networks
Bangalore, Karnataka
India

Email: partha@parthasarathi.co.in

Paul Kyzivat
Huawei
Hudson, MA
USA

Email: pkyzivat@alum.mit.edu

SIPREC
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2013

L. Portman
NICE Systems
H. Lum, Ed.
Genesys
C. Eckel
Cisco
A. Johnston
Avaya
A. Hutton
Siemens Enterprise
Communications
October 22, 2012

Session Recording Protocol
draft-ietf-siprec-protocol-08

Abstract

This document specifies the use of the Session Initiation Protocol (SIP), the Session Description Protocol (SDP), and the Real Time Protocol (RTP) for delivering real-time media and metadata from a Communication Session (CS) to a recording device. The Session Recording Protocol specifies the use of SIP, SDP, and RTP to establish a Recording Session (RS) between the Session Recording Client (SRC), which is on the path of the CS, and a Session Recording Server (SRS) at the recording device.

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 25, 2013.

Copyright Notice

Copyright (c) 2012 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	4
2. Terminology	4
3. Definitions	4
4. Scope	4
5. Overview of operations	5
5.1. Delivering recorded media	5
5.2. Delivering recording metadata	7
5.3. Receiving recording indications and providing recording preferences	8
6. SIP Handling	9
6.1. Procedures at the SRC	9
6.1.1. Initiating a Recording Session	10
6.1.2. SIP extensions for recording indication and preference	10
6.2. Procedures at the SRS	11
6.3. Procedures for Recording-aware User Agents	11
7. SDP Handling	12
7.1. Procedures at the SRC	12
7.1.1. SDP handling in RS	12
7.1.1.1. Handling media stream updates	13
7.1.2. Recording indication in CS	14
7.1.3. Recording preference in CS	15
7.2. Procedures at the SRS	15
7.3. Procedures for Recording-aware User Agents	17
7.3.1. Recording indication	17
7.3.2. Recording preference	18
8. RTP Handling	19
8.1. RTP Mechanisms	19
8.1.1. RTCP	19
8.1.2. RTP Profile	19
8.1.3. SSRC	20
8.1.4. CSRC	20
8.1.5. SDP	21

8.1.5.1. CNAME	21
8.1.6. Keepalive	21
8.1.7. RTCP Feedback Messages	21
8.1.7.1. Full Intra Request	22
8.1.7.2. Picture Loss Indicator	22
8.1.7.3. Temporary Maximum Media Stream Bit Rate Request	22
8.1.8. Symmetric RTP/RTCP for Sending and Receiving	23
8.2. Roles	23
8.2.1. SRC acting as an RTP Translator	24
8.2.1.1. Forwarding Translator	25
8.2.1.2. Transcoding Translator	25
8.2.2. SRC acting as an RTP Mixer	26
8.2.3. SRC acting as an RTP Endpoint	26
8.3. RTP Session Usage by SRC	27
8.3.1. SRC Using Multiple m-lines	27
8.3.2. SRC Using SSRC Multiplexing	28
8.3.3. SRC Using Mixing	29
9. Metadata	30
9.1. Procedures at the SRC	30
9.2. Procedures at the SRS	32
9.2.1. Formal Syntax	34
10. Persistent Recording	34
11. IANA Considerations	34
11.1. Registration of Option Tags	34
11.1.1. siprec Option Tag	35
11.1.2. record-aware Option Tag	35
11.2. Registration of media feature tags	35
11.2.1. src feature tag	35
11.2.2. srs feature tag	36
11.3. New Content-Disposition Parameter Registrations	36
11.4. Media Type Registration	36
11.4.1. Registration of MIME Type application/rs-metadata	36
11.4.2. Registration of MIME Type application/rs-metadata-request	37
11.5. SDP Attributes	37
11.5.1. 'record' SDP Attribute	37
11.5.2. 'recordpref' SDP Attribute	37
12. Security Considerations	38
12.1. Authentication and Authorization	38
12.2. RTP handling	39
12.3. Metadata	39
12.4. Storage and playback	40
13. Acknowledgements	40
14. References	40
14.1. Normative References	40
14.2. Informative References	41
Authors' Addresses	42

1. Introduction

This document specifies the mechanism to record a Communication Session (CS) by delivering real-time media and metadata from the CS to a recording device. In accordance to the architecture [I-D.ietf-siprec-architecture], the Session Recording Protocol specifies the use of SIP, SDP, and RTP to establish a Recording Session (RS) between the Session Recording Client (SRC), which is on the path of the CS, and a Session Recording Server (SRS) at the recording device.

SIP is also used to deliver metadata to the recording device, as specified in [I-D.ietf-siprec-metadata]. Metadata is information that describes recorded media and the CS to which they relate.

The Session Recording Protocol intends to satisfy the SIP-based Media Recording requirements listed in [RFC6341].

In addition to the Session Recording Protocol, this document specifies extensions for user agents that are participants in a CS to receive recording indications and to provide preferences for recording.

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 [RFC2119].

3. Definitions

This document refers to the core definitions provided in the architecture document [I-D.ietf-siprec-architecture].

The RTP Handling section uses the definitions provided in "RTP: A Transport Protocol for Real-Time Application" [RFC3550].

4. Scope

The scope of the Session Recording Protocol includes the establishment of the recording sessions and the reporting of the metadata. The scope also includes extensions supported by User Agents participating in the CS such as indication of recording. The user agents need not be recording-aware in order to participate in a CS being recorded.

The following items, which are not an exhaustive list, do not represent the protocol itself and are considered out of the scope of the Session Recording Protocol:

- o Delivering recorded media in real-time as the CS media
- o Specifications of criteria to select a specific CS to be recorded or triggers to record a certain CS in the future
- o Recording policies that determine whether the CS should be recorded and whether parts of the CS are to be recorded
- o Retention policies that determine how long a recording is stored
- o Searching and accessing the recorded media and metadata
- o Policies governing how CS users are made aware of recording
- o Delivering additional recording session metadata through non-SIP mechanism

5. Overview of operations

This section is informative and provides a description of recording operations.

Section 6 describes SIP the handling in a recording session between a SRC and a SRS, and the procedures for recording-aware user agents participating in a CS. Section 7 describes the SDP in a recording session, and the procedures for recording indications and recording preferences. Section 8 describes the RTP handling in a recording session. Section 9 describes the mechanism to deliver recording metadata from the SRC to the SRS.

As mentioned in the architecture document [I-D.ietf-siprec-architecture], there are a number of types of call flows based on the location of the Session Recording Client. The following sample call flows provide a quick overview of the operations between the SRC and the SRS.

5.1. Delivering recorded media

When a SIP Back-to-back User Agent (B2BUA) with SRC functionality routes a call from UA(A) to UA(B), the SRC has access to the media path between the user agents. When the SRC is aware that it should be recording the conversation, the SRC can cause the B2BUA to bridge the media between UA(A) and UA(B). The SRC then establishes the

Recording Session with the SRS and sends replicated media towards the SRS.

An endpoint may also have SRC functionality, where the endpoint itself establishes the Recording Session to the SRS. Since the endpoint has access to the media in the Communication Session, the endpoint can send replicated media towards the SRS.

The following is a sample call flow that shows the SRC establishing a recording session towards the SRS. The call flow is essentially identical when the SRC is a B2BUA or as the endpoint itself. Note that the SRC can choose when to establish the Recording Session independent of the Communication Session, even though the following call flow suggests that the SRC is establishing the Recording Session (message #5) after the Communication Session is established.




Figure 1: Basic recording call flow

The above call flow can also apply to the case of a centralized

conference with a mixer. For clarity, ACKs to INVITES and 200 OKs to BYEs are not shown. The conference focus can provide the SRC functionality since the conference focus has access to all the media from each conference participant. When a recording is requested, the SRC delivers the metadata and the media streams to the SRS. Since the conference focus has access to a mixer, the SRC may choose to mix the media streams from all participants as a single mixed media stream towards the SRS.

An SRC can use a single recording session to record multiple communication sessions. Every time the SRC wants to record a new call, the SRC updates the recording session with a new SDP offer to add new recorded streams to the recording session, and correspondingly also update the metadata for the new call.

An SRS can also establish a recording session to an SRC, although it is beyond the scope of this document to define how an SRS would specify which calls to record.

5.2. Delivering recording metadata

The SRC is responsible for the delivery of metadata to the SRS. The SRC may provide an initial metadata snapshot about recorded media streams in the initial INVITE content in the recording session. Subsequent metadata updates can be represented as a stream of events in UPDATE or reINVITE requests sent by the SRC. These metadata updates are normally incremental updates to the initial metadata snapshot to optimize on the size of updates, however, the SRC may also decide to send a new metadata snapshot anytime.

Metadata is transported in the body of INVITE or UPDATE messages. Certain metadata, such as the attributes of the recorded media stream are located in the SDP of the recording session.

The SRS has the ability to send a request to the SRC to request for a new metadata snapshot update from the SRC. This can happen when the SRS fails to understand the current stream of incremental updates for whatever reason, for example, when SRS loses the current state due to internal failure. The SRS may optionally attach a reason along with the snapshot request. This request allows both SRC and SRS to synchronize the states with a new metadata snapshot so that further metadata incremental updates will be based on the latest metadata snapshot. Similar to the metadata content, the metadata snapshot request is transported as content in UPDATE or INVITE sent by the SRS in the recording session.




Figure 2: Delivering metadata via SIP UPDATE

5.3. Receiving recording indications and providing recording preferences

The SRC is responsible to provide recording indications to the participants in the CS. A recording-aware UA supports receiving recording indications via the SDP attribute `a=record`, and it can specify a recording preference in the CS by including the SDP attribute `a=recordpref`. The recording attribute is a declaration by the SRC in the CS to indicate whether recording is taking place. The recording preference attribute is a declaration by the recording-aware UA in the CS to indicate the recording preference.

To illustrate how the attributes are used, if a UA (A) is initiating

a call to UA (B) and UA (A) is also an SRC that is performing the recording, then UA (A) provides the recording indication in the SDP offer with a=record:on. Since UA (A) is the SRC, UA (A) receives the recording indication from the SRC directly. When UA (B) receives the SDP offer, UA (B) will see that recording is happening on the other endpoint of this session. Since UA (B) is not an SRC and does not provide any recording preference, the SDP answer does not contain a=record nor a=recordpref.

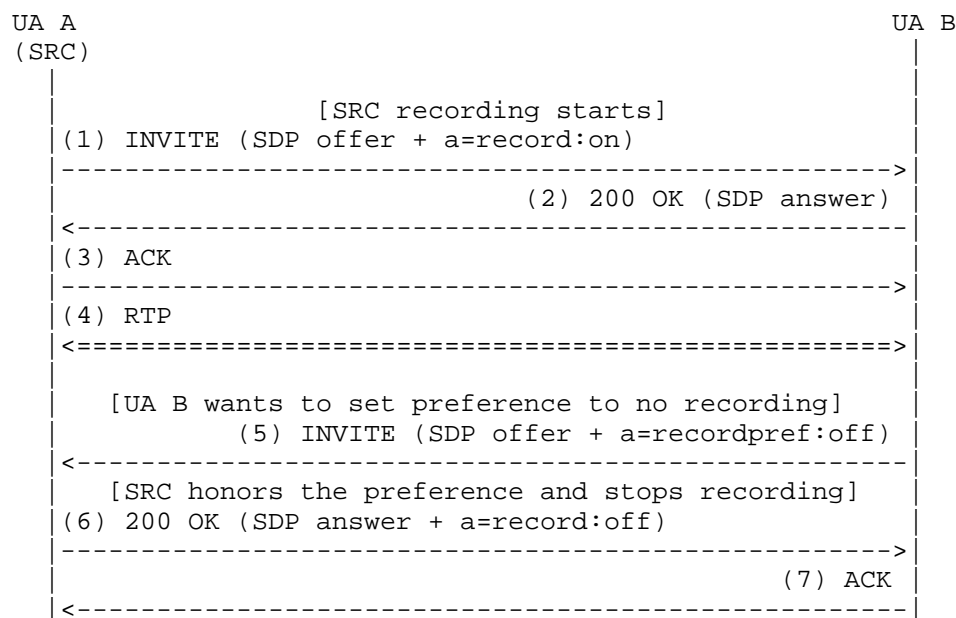


Figure 3: Recording indication and recording preference

After the call is established and recording is in progress, UA (B) later decides to change the recording preference to no recording and sends a reINVITE with the a=recordpref attribute. It is up to the SRC to honor the preference, and in this case SRC decides to stop the recording and updates the recording indication in the SDP answer.

6. SIP Handling

6.1. Procedures at the SRC

6.1.1. Initiating a Recording Session

A recording session is a SIP session with specific extensions applied, and these extensions are listed in the procedures for SRC and SRS below. When an SRC or an SRS receives a SIP session that is not a recording session, it is up to the SRC or the SRS to determine what to do with the SIP session.

The SRC can initiate a recording session by sending a SIP INVITE request to the SRS. The SRC and the SRS are identified in the From and To headers, respectively.

The SRC MUST include the '+sip.src' feature tag in the Contact URI, defined in this specification as an extension to [RFC3840], for all recording sessions. An SRS uses the presence of the '+sip.src' feature tag in dialog creating and modifying requests and responses to confirm that the dialog being created is for the purpose of a Recording Session. In addition, when an SRC sends a REGISTER request to a registrar, the SRC MUST include the '+sip.src' feature tag to indicate that it is a SRC.

Since SIP Caller Preferences extensions are optional to implement for routing proxies, there is no guarantee that a recording session will be routed to an SRC or SRS. A new options tag is introduced: "siprec". As per [RFC3261], only an SRC or an SRS can accept this option tag in a recording session. An SRC MUST include the "siprec" option tag in the Require header when initiating a Recording Session so that UA's which do not support the session recording protocol extensions will simply reject the INVITE request with a 420 Bad Extension.

When an SRC receives a new INVITE, the SRC MUST only consider the SIP session as a recording session when both the '+sip.srs' feature tag and 'siprec' option tag are included in the INVITE request.

6.1.2. SIP extensions for recording indication and preference

For the communication session, the SRC MUST provide recording indication to all participants in the CS. A participant UA in a CS can indicate that it is recording-aware by providing the "record-aware" option tag, and the SRC MUST provide recording indications in the new SDP a=record attribute described in the SDP Handling section. In the absence of the "record-aware" option tag, meaning that the participant UA is not recording-aware, an SRC MUST provide recording indications through other means such as playing a tone inband, if the SRC is required to do so (e.g. based on policies).

An SRC in the CS may also indicate itself as a session recording

client by including the '+sip.src' feature tag. A recording-aware participant can learn that a SRC is in the CS, and can set the recording preference for the CS with the new SDP a=recordpref attribute described in the SDP Handling section below.

6.2. Procedures at the SRS

When an SRS receives a new INVITE, the SRS MUST only consider the SIP session as a recording session when both the '+sip.src' feature tag and 'siprec' option tag are included in the INVITE request.

The SRS can initiate a recording session by sending a SIP INVITE request to the SRC. The SRS and the SRC are identified in the From and To headers, respectively.

The SRS MUST include the '+sip.srs' feature tag in the Contact URI, as per [RFC3840], for all recording sessions. An SRC uses the presence of this feature tag in dialog creating and modifying requests and responses to confirm that the dialog being created is for the purpose of a Recording Session (REQ-30). In addition, when an SRS sends a REGISTER request to a registrar, the SRS MUST include the '+sip.srs' feature tag to indicate that it is a SRS.

An SRS MUST include the "siprec" option tag in the Require header as per [RFC3261] when initiating a Recording Session so that UA's which do not support the session recording protocol extensions will simply reject the INVITE request with a 420 Bad Extension.

6.3. Procedures for Recording-aware User Agents

A recording-aware user agent is a participant in the CS that supports the SIP and SDP extensions for receiving recording indication and for requesting recording preferences for the call. A recording-aware UA MUST indicate that it can accept reporting of recording indication provided by the SRC with a new option tag "record-aware" when initiating or establishing a CS, meaning including the "record-aware" tag in the Supported header in the initial INVITE request or response.

A recording-aware UA MUST be prepared to provide recording indication to the end user through an appropriate user interface an indication whether recording is on, off, or paused for each medium. Some user agents that are automations (e.g. IVR, media server, PSTN gateway) may not have a user interface to render recording indication. When such user agent indicates recording awareness, the UA SHOULD render recording indication through other means, such as passing an inband tone on the PSTN gateway, putting the recording indication in a log file, or raising an application event in a VoiceXML dialog. These

user agents MAY also choose not to indicate recording awareness, thereby relying on whatever mechanism an SRC chooses to indicate recording, such as playing a tone inband.

7. SDP Handling

7.1. Procedures at the SRC

The SRC and SRS follows the SDP offer/answer model in [RFC3264]. The procedures for SRC and SRS describe the conventions used in a recording session.

7.1.1. SDP handling in RS

Since the SRC does not expect to receive media from the SRS, the SRC typically sets each media stream of the SDP offer to only send media, by qualifying them with the a=sendonly attribute, according to the procedures in [RFC3264].

The SRC sends recorded streams of participants to the SRS, and the SRC MUST provide a label attribute (a=label), as per [RFC4574], on each media stream in order to identify the recorded stream with the rest of the metadata. The a=label attribute identifies each recorded media stream, and the label name is mapped to the Media Stream Reference in the metadata as per [I-D.ietf-siprec-metadata]. The scope of the a=label attribute only applies to the SDP and Metadata conveyed in the bodies of the SIP request or response that the label appeared in. Note that a recorded stream is distinct from a CS stream; the metadata provides a list of participants that contributes to each recorded stream.

The following is an example SDP offer from SRC with both audio and video recorded streams. Note that the following example contains unfolded lines longer than 72 characters. These are captured between <allOneLine> tags.

```
v=0
o=SRC 2890844526 2890844526 IN IP4 198.51.100.1
s=-
c=IN IP4 198.51.100.1
t=0 0
m=audio 12240 RTP/AVP 0 4 8
a=sendonly
a=label:1
m=video 22456 RTP/AVP 98
a=rtpmap:98 H264/90000
<allOneLine>
a=fmtp:98 profile-level-id=42A01E;
      sprop-parameter-sets=Z0IACpZTBmI,aMljiA==
</allOneLine>
a=sendonly
a=label:2
m=audio 12242 RTP/AVP 0 4 8
a=sendonly
a=label:3
m=video 22458 RTP/AVP 98
a=rtpmap:98 H264/90000
<allOneLine>
a=fmtp:98 profile-level-id=42A01E;
      sprop-parameter-sets=Z0IACpZTBmI,aMljiA==
</allOneLine>
a=sendonly
a=label:4
```

Figure 4: Sample SDP offer from SRC with audio and video streams

7.1.1.1. Handling media stream updates

Over the lifetime of a recording session, the SRC can add and remove recorded streams from the recording session for various reasons. For example, when a CS stream is added or removed from the CS, or when a CS is created or terminated if a recording session handles multiple CSes. To remove a recorded stream from the recording session, the SRC sends a new SDP offer where the port of the media stream to be removed is set to zero, according to the procedures in [RFC3264]. To add a recorded stream to the recording session, the SRC sends a new SDP offer by adding a new media stream description or by reusing an old media stream which had been previously disabled, according to the procedures in [RFC3264].

The SRC can temporarily discontinue streaming and collection of recorded media from the SRC to the SRS for reason such as masking the

recording. In this case, the SRC sends a new SDP offer and sets the media stream to inactive (a=inactive) for each recorded stream to be paused, as per the procedures in [RFC3264]. To resume streaming and collection of recorded media, the SRC sends a new SDP offer and sets the media streams with a=sendonly attribute. Note that when a CS stream is muted/unmuted, this information is conveyed in the metadata by the SRC. The SRC SHOULD NOT modify the media stream with a=inactive for mute since this operation is reserved for pausing the RS media.

7.1.2. Recording indication in CS

While there are existing mechanisms for providing an indication that a CS is being recorded, these mechanisms are usually delivered on the CS media streams such as playing an in-band tone or an announcement to the participants. A new 'record' SDP attribute is introduced to allow the SRC to indicate recording state to a recording-aware UA in CS.

The 'record' SDP attribute appears at the media level or session level in either SDP offer or answer. When the attribute is applied at the session level, the indication applies to all media streams in the SDP. When the attribute is applied at the media level, the indication applies to the media stream only, and that overrides the indication if also set at the session level. Whenever the recording indication needs to change, such as termination of recording, then the SRC MUST initiate a reINVITE or UPDATE to update the SDP a=record attribute.

The following is the ABNF of the 'record' attribute:

```
attribute /= record-attr
; attribute defined in RFC 4566
record-attr = "record:" indication
indication = "on" / "off" / "paused"
on Recording is in progress.
off No recording is in progress.
paused Recording is in progress but media is paused.
```

7.1.3. Recording preference in CS

When the SRC receives the a=recordpref SDP in an SDP offer or answer, the SRC chooses to honor the preference to record based on local policy at the SRC. Whether or not the SRC honors the recording preference, the SRC MUST update the a=record attribute to indicate the current state of the recording (on/off/paused).

7.2. Procedures at the SRS

Typically the SRS only receives RTP streams from the SRC; therefore, the SDP offer/answer from the SRS normally sets each media stream to receive media, by setting them with the a=recvonly attribute, according to the procedures of [RFC3264]. When the SRS is not ready to receive a recorded stream, the SRS sets the media stream as inactive in the SDP offer or answer by setting it with a=inactive attribute, according to the procedures of [RFC3264]. When the SRS is ready to receive recorded streams, the SRS sends a new SDP offer and sets the media streams with a=recvonly attribute.

The following is an example of SDP answer from SRS for the SDP offer from the above sample. Note that the following example contain unfolded lines longer than 72 characters. These are captured between <allOneLine> tags.

```

v=0
o=SRS 0 0 IN IP4 198.51.100.20
s=-
c=IN IP4 198.51.100.20
t=0 0
m=audio 10000 RTP/AVP 0
a=recvonly
a=label:1
m=video 10002 RTP/AVP 98
a=rtpmap:98 H264/90000
<allOneLine>
a=fmtp:98 profile-level-id=42A01E;
      sprop-parameter-sets=Z0IACpZTBmI,aMljiA==
</allOneLine>
a=recvonly
a=label:2
m=audio 10004 RTP/AVP 0
a=recvonly
a=label:3
m=video 10006 RTP/AVP 98
a=rtpmap:98 H264/90000
<allOneLine>
a=fmtp:98 profile-level-id=42A01E;
      sprop-parameter-sets=Z0IACpZTBmI,aMljiA==
</allOneLine>
a=recvonly
a=label:4

```

Figure 5: Sample SDP answer from SRS with audio and video streams

Over the lifetime of a recording session, the SRS can remove recorded streams from the recording session for various reasons. To remove a recorded stream from the recording session, the SRS sends a new SDP offer where the port of the media stream to be removed is set to zero, according to the procedures in [RFC3264].

The SRS SHOULD NOT add recorded streams in the recording session when SRS sends a new SDP offer. Similarly, when the SRS starts a recording session, the SRS SHOULD initiate the INVITE without an SDP offer to let the SRC generate the SDP offer with recorded streams.

The following sequence diagram shows an example where the SRS is initially not ready to receive recorded streams, and later updates the recording session when the SRS is ready to record.




Figure 6: SRS responding to offer with a=inactive

7.3. Procedures for Recording-aware User Agents

7.3.1. Recording indication

When a recording-aware UA receives an SDP offer or answer that includes the a=record attribute, the UA MUST provide the recording indication to the end user whether the recording is on, off, or paused for each medium based on the most recently received a=record SDP attribute for that medium.

If a call is traversed through one or more SIP B2BUA, and it happens that there are more than one SRC in the call path, the recording indication attribute does not provide any hint as to which SRC is performing the recording, meaning the endpoint only knows that the call is being recorded. This attribute is also not used as an indication to negotiate which SRC in the call path will perform recording and is not used as a request to start/stop recording if there are multiple SRCs in the call path.

7.3.2. Recording preference

A participant in a CS MAY set the recording preference in the CS to be recorded or not recorded at session establishment or during the session. A new 'recordpref' SDP attribute is introduced, and the participant in CS may set this recording preference attribute in any SDP offer/answer at session establishment time or during the session. The SRC is not required to honor the recording preference from a participant based on local policies at the SRC, and the participant can learn the recording indication through the a=record SDP attribute as described in the above section.

The SDP a=recordpref attribute can appear at the media level or session level and can appear in an SDP offer or answer. When the attribute is applied at the session level, the recording preference applies to all media stream in the SDP. When the attribute is applied at the media level, the recording preference applies to the media stream only, and that overrides the recording preference if also set at the session level. The user agent can change the recording preference by changing the a=recordpref attribute in subsequent SDP offer or answer. The absence of the a=recordpref attribute in the SDP indicates that the UA has no recording preference.

The following is the ABNF of the recordpref attribute:

```
attribute /= recordpref-attr  
  
; attribute defined in RFC 4566  
  
recordpref-attr = "a=recordpref:" pref  
  
pref = "on" / "off" / "pause" / "nopreference"
```

on Sets the preference to record if it has not already been started. If the recording is currently paused, the preference is to resume recording.

off Sets the preference for no recording. If recording has already been started, then the preference is to stop the recording.

pause If the recording is currently in progress, sets the preference to pause the recording.

nopreference To indicate that the UA has no preference on recording.

8. RTP Handling

This section provides recommendations and guidelines for RTP and RTCP in the context of SIPREC. In order to communicate most effectively, the Session Recording Client (SRC), the Session Recording Server (SRS), and any Recording aware User Agents (UAs) SHOULD utilize the mechanisms provided by RTP in a well-defined and predicable manner. It is the goal of this document to make the reader aware of these mechanisms and provide recommendations and guidelines.

8.1. RTP Mechanisms

This section briefly describes important RTP/RTCP constructs and mechanisms that are particularly useful within the content of SIPREC.

8.1.1. RTCP

The RTP data transport is augmented by a control protocol (RTCP) to allow monitoring of the data delivery. RTCP, as defined in [RFC3550], is based on the periodic transmission of control packets to all participants in the RTP session, using the same distribution mechanism as the data packets. Support for RTCP is REQUIRED, per [RFC3550], and it provides, among other things, the following important functionality in relation to SIPREC:

1) Feedback on the quality of the data distribution

This feedback from the receivers may be used to diagnose faults in the distribution. As such, RTCP is a well-defined and efficient mechanism for the SRS to inform the SRC, and for the SRC to inform Recording aware UAs, of issues that arise with respect to the reception of media that is to be recorded.

2) Carries a persistent transport-level identifier for an RTP source called the canonical name or CNAME

The SSRC identifier may change if a conflict is discovered or a program is restarted; in which case receivers can use the CNAME to keep track of each participant. Receivers may also use the CNAME to associate multiple data streams from a given participant in a set of related RTP sessions, for example to synchronize audio and video. Synchronization of media streams is also facilitated by the NTP and RTP timestamps included in RTCP packets by data senders.

8.1.2. RTP Profile

The RECOMMENDED RTP profiles for the SRC, SRS, and Recording aware UAs are "Extended Secure RTP Profile for Real-time Transport Control

Protocol (RTCP)-Based Feedback (RTP/SAVPF)", [RFC5124] when using encrypted RTP streams, and "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", [RFC4585] when using non encrypted media streams. However, as this is not a requirement, some implementations may use "The Secure Real-time Transport Protocol (SRTP)", [RFC3711] and "RTP Profile for Audio and Video Conferences with Minimal Control", AVP [RFC3551]. Therefore, it is RECOMMENDED that the SRC, SRS, and Recording aware UAs not rely entirely on SAVPF or AVPF for core functionality that may be at least partially achievable using SAVP and AVP.

AVPF and SAVPF provide an improved RTCP timer model that allows more flexible transmission of RTCP packets in response to events, rather than strictly according to bandwidth. AVPF based codec control messages provide efficient mechanisms for an SRC, SRS, and Recording aware UAs to handle events such as scene changes, error recovery, and dynamic bandwidth adjustments. These messages are discussed in more detail later in this document.

SAVP and SAVPF provide media encryption, integrity protection, replay protection, and a limited form of source authentication. They do not contain or require a specific keying mechanism.

8.1.3. SSRC

The synchronization source (SSRC), as defined in [RFC3550] is carried in the RTP header and in various fields of RTCP packets. It is a random 32-bit number that is required to be globally unique within an RTP session. It is crucial that the number be chosen with care in order that participants on the same network or starting at the same time are not likely to choose the same number. Guidelines regarding SSRC value selection and conflict resolution are provided in [RFC3550].

The SSRC may also be used to separate different sources of media within a single RTP session. For this reason as well as for conflict resolution, it is important that the SRC, SRS, and Recording aware UAs handle changes in SSRC values and properly identify the reason of the change. The CNAME values carried in RTCP facilitate this identification.

8.1.4. CSRC

The contributing source (CSRC), as defined in [RFC3550], identifies the source of a stream of RTP packets that has contributed to the combined stream produced by an RTP mixer. The mixer inserts a list of the SSRC identifiers of the sources that contributed to the generation of a particular packet into the RTP header of that packet.

This list is called the CSRC list. It is RECOMMENDED that a SRC or Recording aware UA, when acting a mixer, sets the CSRC list accordingly, and that the SRC and SRS interpret the CSRC list appropriately when received.

8.1.5. SDES

The Source Description (SDES), as defined in [RFC3550], contains an SSRC/CSRC identifier followed by a list of zero or more items, which carry information about the SSRC/CSRC. End systems send one SDES packet containing their own source identifier (the same as the SSRC in the fixed RTP header). A mixer sends one SDES packet containing a chunk for each contributing source from which it is receiving SDES information, or multiple complete SDES packets if there are more than 31 such sources.

8.1.5.1. CNAME

The Canonical End-Point Identifier (CNAME), as defined in [RFC3550], provides the binding from the SSRC identifier to an identifier for the source (sender or receiver) that remains constant. It is important the SRC and Recording aware UAs generate CNAMEs appropriately and that the SRC and SRS interpret and use them for this purpose. Guidelines for generating CNAME values are provided in "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMEs)" [RFC6222].

8.1.6. Keepalive

It is anticipated that media streams in SIPREC may exist in an inactive state for extended periods of times for any of a number of valid reasons. In order for the bindings and any pinholes in NATs/firewalls to remain active during such intervals, it is RECOMMENDED that the SRC, SRS, and Recording aware UAs follow the keep-alive procedure recommended in "Application Mechanism for Keeping Alive the NAT Mappings Associated to RTP/RTCP Control Protocol (RTCP) Flows" [RFC6263] for all RTP media streams.

8.1.7. RTCP Feedback Messages

"Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)" [RFC5104] specifies extensions to the messages defined in AVPF [RFC4585]. Support for and proper usage of these messages is important to SRC, SRS, and Recording aware UA implementations. Note that these messages are applicable only when using the AVFP or SAVPF RTP profiles

8.1.7.1. Full Intra Request

A Full Intra Request (FIR) Command, when received by the designated media sender, requires that the media sender sends a Decoder Refresh Point at the earliest opportunity. Using a decoder refresh point implies refraining from using any picture sent prior to that point as a reference for the encoding process of any subsequent picture sent in the stream.

Decoder refresh points, especially Intra or IDR pictures for H.264 video codecs, are in general several times larger in size than predicted pictures. Thus, in scenarios in which the available bit rate is small, the use of a decoder refresh point implies a delay that is significantly longer than the typical picture duration.

8.1.7.1.1. SIP INFO for FIR

"XML Schema for Media Control" [RFC5168] defines an Extensible Markup Language (XML) Schema for video fast update. Implementations are discouraged from using the method described except for backward compatibility purposes. Implementations SHOULD use FIR messages instead.

8.1.7.2. Picture Loss Indicator

Picture Loss Indication (PLI), as defined in [RFC4585], informs the encoder of the loss of an undefined amount of coded video data belonging to one or more pictures. Using the FIR command to recover from errors is explicitly disallowed, and instead the PLI message SHOULD be used. FIR SHOULD be used only in situations where not sending a decoder refresh point would render the video unusable for the users. Examples where sending FIR is appropriate include a multipoint conference when a new user joins the conference and no regular decoder refresh point interval is established, and a video switching MCU that changes streams.

8.1.7.3. Temporary Maximum Media Stream Bit Rate Request

A receiver, translator, or mixer uses the Temporary Maximum Media Stream Bit Rate Request (TMMBR) to request a sender to limit the maximum bit rate for a media stream to the provided value. Appropriate use of TMMBR facilitates rapid adaptation to changes in available bandwidth.

8.1.7.3.1. Renegotiation of SDP bandwidth attribute

If it is likely that the new value indicated by TMMBR will be valid for the remainder of the session, the TMMBR sender is expected to

perform a renegotiation of the session upper limit using the session signaling protocol. Therefore for SIPREC, implementations are RECOMMENDED to use TMMBR for temporary changes, and renegotiation of bandwidth via SDP offer/answer for more permanent changes.

8.1.8. Symmetric RTP/RTCP for Sending and Receiving

Within an SDP offer/answer exchange, RTP entities choose the RTP and RTCP transport addresses (i.e., IP addresses and port numbers) on which to receive packets. When sending packets, the RTP entities may use the same source port or a different source port as those signaled for receiving packets. When the transport address used to send and receive RTP is the same, it is termed "symmetric RTP" [RFC4961]. Likewise, when the transport address used to send and receive RTCP is the same, it is termed "symmetric RTCP" [RFC4961].

When sending RTP, it is REQUIRED to use symmetric RTP. When sending RTCP, it is REQUIRED to use symmetric RTCP. Although an SRS will not normally send RTP, it will send RTCP as well as receive RTP and RTCP. Likewise, although an SRC will not normally receive RTP from the SRS, it will receive RTCP as well as send RTP and RTCP.

Note: Symmetric RTP and symmetric RTCP are different from RTP/RTCP multiplexing [RFC5761].

8.2. Roles

An SRC has the task of gathering media from the various UAs in one or more Communication Sessions (CSs) and forwarding the information to the SRS within the context of a corresponding Recording Session (RS). There are numerous ways in which an SRC may do this is, including but not limited to, appearing as a UA within a CS, or as a B2BUA between UAs within a CS.

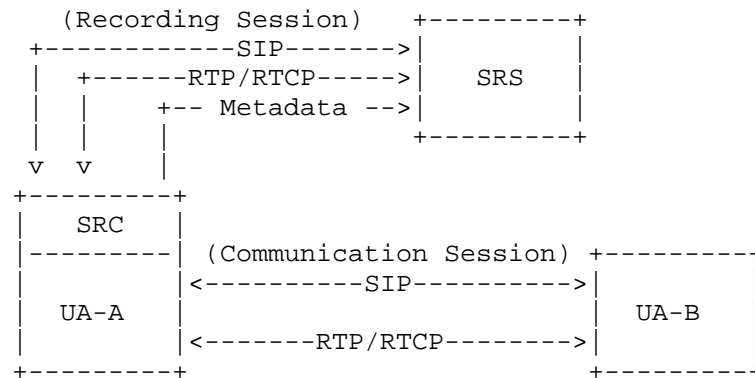


Figure 7: UA as SRC

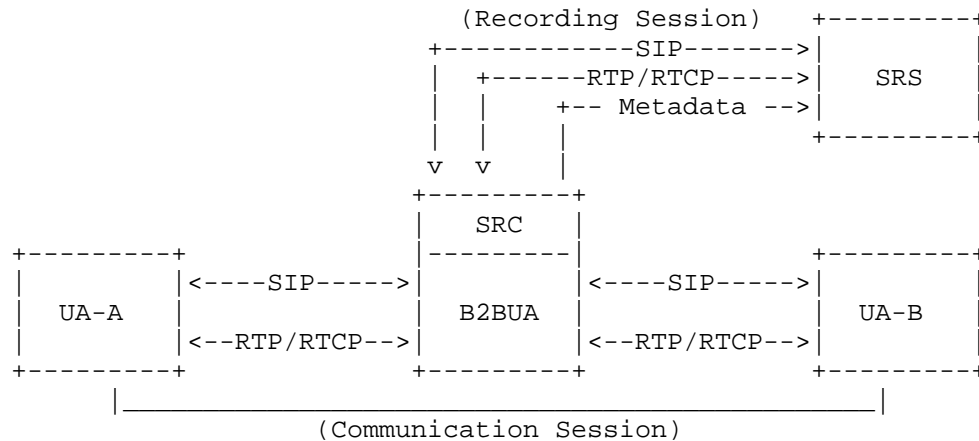


Figure 8: B2BUA as SRC

The following subsections define a set of roles an SRC may choose to play based on its position with respect to a UA within a CS, and an SRS within an RS. A CS and a corresponding RS are independent sessions; therefore, an SRC may play a different role within a CS than it does within the corresponding RS.

8.2.1. SRC acting as an RTP Translator

The SRC may act as a translator, as defined in [RFC3550]. A defining characteristic of a translator is that it forwards RTP packets with their SSRC identifier intact. There are two types of translators, one that simply forwards, and another that performs transcoding (e.g., from one codec to another) in addition to forwarding.

8.2.1.1. Forwarding Translator

When acting as a forwarding translator, RTP received as separate streams from different sources (e.g., from different UAs with different SSRCs) cannot be mixed by the SRC and MUST be sent separately to the SRS. All RTCP reports MUST be passed by the SRC between the UAs and the SRS, such that the UAs and SRS are able to detect any SSRC collisions.

RTCP Sender Reports generated by a UA sending a stream MUST be forwarded to the SRS. RTCP Receiver Reports generated by the SRS MUST be forwarded to the relevant UA.

UAs may receive multiple sets of RTCP Receiver Reports, one or more from other UAs participating in the CS, and one from the SRS participating in the RS. A Recording aware UA SHOULD be prepared to process the RTCP Receiver Reports from the SRS, whereas a recording unaware UA may discard such RTCP packets as not of relevance.

If SRTP is used on both the CS and the RS, decryption and/or re-encryption may occur. For example, if different keys are used, it will occur. If the same keys are used, it need not occur. Section 12 provides additional information on SRTP and keying mechanisms.

If packet loss occurs, either from the UA to the SRC or from the SRC to the SRS, the SRS SHOULD detect and attempt to recover from the loss. The SRC does not play a role in this other than forwarding the associated RTP and RTCP packets.

8.2.1.2. Transcoding Translator

When acting as a transcoding translator, an SRC MAY perform transcoding (e.g., from one codec to another), and this may result in a different rate of packets between what the SRC receives and what the SRC sends. As when acting as a forwarding translator, RTP received as separate streams from different sources (e.g., from different UAs with different SSRCs) cannot be mixed by the SRC and MUST be sent separately to the SRS. All RTCP reports MUST be passed by the SRC between the UAs and the SRS, such that the UAs and SRS are able to detect any SSRC collisions.

RTCP Sender Reports generated by a UA sending a stream MUST be forwarded to the SRS. RTCP Receiver Reports generated by the SRS MUST be forwarded to the relevant UA. The SRC may need to manipulate the RTCP Receiver Reports to take account of any transcoding that has taken place.

UAs may receive multiple sets of RTCP Receiver Reports, one or more from other UAs participating in the CS, and one from the SRS participating in the RS. A Recording aware UA SHOULD be prepared to process the RTCP Receiver Reports from the SRS, whereas a recording unaware UA may discard such RTCP packets as not of relevance.

If SRTP is used on both the CS and the RS, decryption and/or re-encryption may occur. For example, if different keys are used, it will occur. If the same keys are used, it need not occur. Section 12 provides additional information on SRTP and keying mechanisms.

If packet loss occurs, either from the UA to the SRC or from the SRC to the SRS, the SRS SHOULD detect and attempt to recover from the loss. The SRC does not play a role in this other than forwarding the associated RTP and RTCP packets.

8.2.2. SRC acting as an RTP Mixer

In the case of the SRC acting as a RTP mixer, as defined in [RFC3550], the SRC combines RTP streams from different UA and sends them towards the SRS using its own SSRC. The SSRCs from the contributing UA SHOULD be conveyed as CSRCs identifiers within this stream. The SRC may make timing adjustments among the received streams and generate its own timing on the stream sent to the SRS. Optionally an SRC acting as a mixer can perform transcoding, and can even cope with different codings received from different UAs. RTCP Sender Reports and Receiver Reports are not forwarded by an SRC acting as mixer, but there are requirements for forwarding RTCP Source Description (SDES) packets. The SRC generates its own RTCP Sender and Receiver reports toward the associated UAs and SRS.

The use of SRTP between the SRC and the SRS for the RS is independent of the use of SRTP between the UAs and SRC for the CS. Section 12 provides additional information on SRTP and keying mechanisms.

If packet loss occurs from the UA to the SRC, the SRC SHOULD detect and attempt to recover from the loss. If packet loss occurs from the SRC to the SRS, the SRS SHOULD detect and attempt to recover from the loss.

8.2.3. SRC acting as an RTP Endpoint

The case of the SRC acting as an RTP endpoint, as defined in [RFC3550], is similar to the mixer case, except that the RTP session between the SRC and the SRS is considered completely independent from the RTP session that is part of the CS. The SRC can, but need not, mix RTP streams from different participants prior to sending to the

SRS. RTCP between the SRC and the SRS is completely independent of RTCP on the CS.

The use of SRTP between the SRC and the SRS for the RS is independent of the use of SRTP between the UAs and SRC for the CS. Section 12 provides additional information on SRTP and keying mechanisms.

If packet loss occurs from the UA to the SRC, the SRC SHOULD detect and attempt to recover from the loss. If packet loss occurs from the SRC to the SRS, the SRS SHOULD detect and attempt to recover from the loss.

8.3. RTP Session Usage by SRC

There are multiple ways that an SRC may choose to deliver recorded media to an SRS. In some cases, it may use a single RTP session for all media within the RS, whereas in others it may use multiple RTP sessions. The following subsections provide examples of basic RTP session usage by the SRC, including a discussion of how the RTP constructs and mechanisms covered previously are used. An SRC may choose to use one or more of the RTP session usages within a single RS. The set of RTP session usages described is not meant to be exhaustive.

8.3.1. SRC Using Multiple m-lines

When using multiple m-lines, an SRC includes each m-line in an SDP offer to the SRS. The SDP answer from the SRS MUST include all m-lines, with any rejected m-lines indicated with a zero port, per [RFC3264]. Having received the answer, the SRC starts sending media to the SRS as indicated in the answer. Alternatively, if the SRC deems the level of support indicated in the answer to be unacceptable, it may initiate another SDP offer/answer exchange in which an alternative RTP session usage is negotiated.

In order to preserve the mapping of media to participant within the CSs in the RS, the SRC SHOULD map each unique CNAME within the CSs to a unique CNAME within the RS. Additionally, the SRC SHOULD map each unique combination of CNAME/SSRC within the CSs to a unique CNAME/SSRC within the RS. In doing so, the SRC may act as an RTP translator or as an RTP endpoint.

The following figure illustrates a case in which each UA represents a participant contributing two RTP sessions (e.g. one for audio and one for video), each with a single SSRC. The SRC acts as an RTP translator and delivers the media to the SRS using four RTP sessions, each with a single SSRC. The CNAME and SSRC values used by the UAs within their media streams are preserved in the media streams from

the SRC to the SRS.

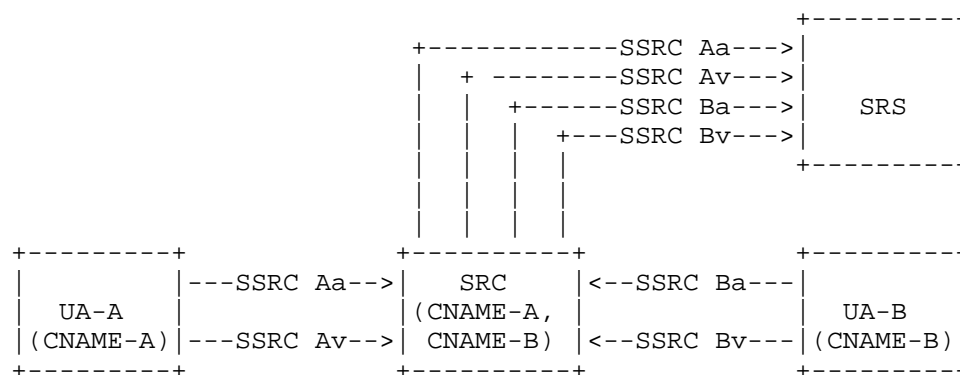


Figure 9: SRC Using Multiple m-lines

8.3.2. SRC Using SSRC Multiplexing

When using SSRC multiplexing, an SRC multiplexes RTP packets of the same media type from multiple RTP sessions into a single RTP session with multiple SSRC values. The SRC includes one m-line for each RTP session in an SDP offer to the SRS. The SDP answer from the SRS MUST include all m-lines, with any rejected m-lines indicated with the zero port, per [RFC3264]. Having received the answer, the SRC starts sending media to the SRS as indicated in the answer.

In order to preserve the mapping of media to participant within the CSs in the RS, the SRC SHOULD map each unique combination of CNAME/SSRC within the CSs to a unique SSRC within the RS. The CNAMEs used in the CSs are not preserved within the RS. The SRS relies on the SIPREC metadata to determine the participants included within each multiplexed stream. The SRC MUST avoid SSRC collisions, rewriting SSRCs if necessary. In doing to, the SRC acts as an RTP endpoint.

In the event the SRS does not support SSRC multiplexing, the SRC becomes aware of this when it receives RTCP receiver reports from the SRS indicating the absence of any packets for one or more of the multiplexed SSRC values. If the SRC deems the level of support indicated in the RTCP receiver report to be unacceptable, it may initiate another SDP offer/answer exchange in which an alternative RTP session usage is negotiated.

The following figure illustrates a case in which each UA represents a participant contributing two RTP sessions (e.g. one for audio and another for video), each with a single SSRC. The SRC delivers the

media to the SRS using two RTP sessions, multiplexing one stream with the same media type from each participant into a single RTP session containing two SSRCs. The SRC uses its own CNAME and SSRC values, but it preserves the mapping of unique CNAME/SSRC used by the UAs within their media streams in the media streams from the SRC to the SRS.

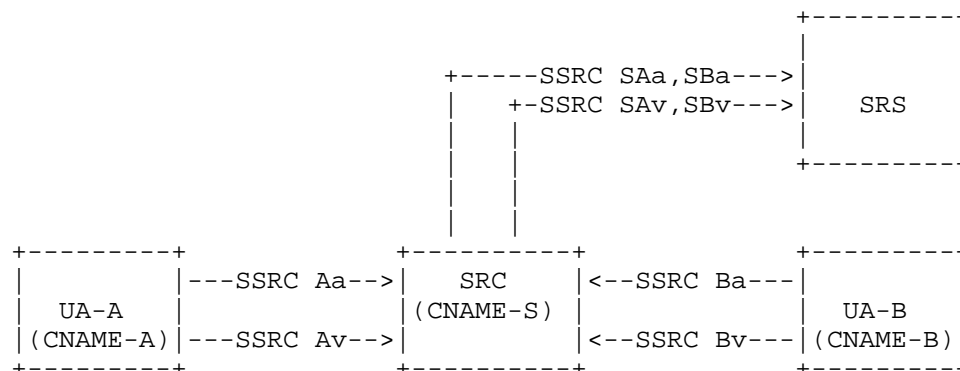


Figure 10: SRC Using SSRC Multiplexing

8.3.3. SRC Using Mixing

When using mixing, the SRC combines RTP streams from different participants and sends them towards the SRS using its own SSRC. The SSRCs from the contributing participants SHOULD be conveyed as CSRCs identifiers. The SRC includes one m-line for each RTP session in an SDP offer to the SRS. The SDP answer from the SRS MUST include all m-lines, with any rejected m-lines indicated with the zero port, per [RFC3264]. Having received the answer, the SRC starts sending media to the SRS as indicated in the answer.

In order to preserve the mapping of media to participant within the CSs in the RS, the SRC SHOULD map each unique CNAME within the CSs to a unique CNAME within the RS. Additionally, the SRC SHOULD map each unique combination of CNAME/SSRC within the CSs to a unique CNAME/SSRC within the RS. The SRC MUST avoid SSRC collisions, rewriting SSRCs if necessary when used as CSRCs in the RS. In doing to, the SRC acts as an RTP mixer.

In the event the SRS does not support this usage of CSRC values, it relies entirely on the SIPREC metadata to determine the participants included within each mixed stream.

The following figure illustrates a case in which each UA represents a

participant contributing two RTP sessions (e.g. one for audio and one for video), each with a single SSRC. The SRC acts as an RTP mixer and delivers the media to the SRS using two RTP sessions, mixing media from each participant into a single RTP session containing a single SSRC and two CSRCs.

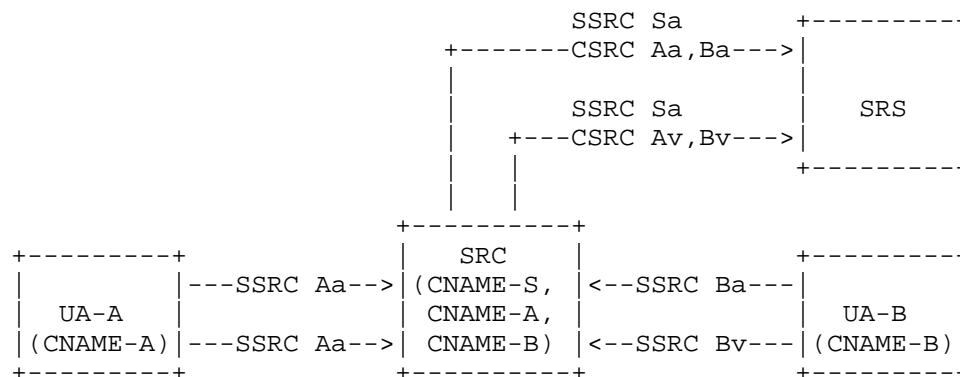


Figure 11: SRC Using Mixing

9. Metadata

9.1. Procedures at the SRC

The SRC MUST deliver metadata to the SRS in a recording session; the timing of which SRC sends the metadata depends on when the metadata becomes available. Metadata SHOULD be provided by the SRC in the initial INVITE request when establishing the recording session, and subsequent metadata updates can be provided by the SRC in reINVITE and UPDATE requests ([RFC3311]) and responses in the recording session. There are cases that metadata is not available in the initial INVITE request sent by the SRC, for example, when a recording session is established in the absence of a communication session, and the SRC would update the recording session with metadata whenever metadata becomes available.

Certain metadata attributes are contained in the SDP, and others are contained in a new content type "application/rs-metadata". The format of the metadata is described as part of the mechanism in [I-D.ietf-siprec-metadata]. A new "disposition-type" of Content-Disposition is defined for the purpose of carrying metadata and the value is "recording-session". The "recording-session" value indicates that the "application/rs-metadata" content contains metadata to be handled by the SRS, and the disposition can be carried

in either INVITE or UPDATE requests or responses sent by the SRC.

Metadata sent by the SRC can be categorized as either a full metadata snapshot or partial update. A full metadata snapshot describes all the recorded streams and all metadata associated with the recording session. When the SRC sends a full metadata snapshot, the SRC MUST send an INVITE or an UPDATE request ([RFC3311]) with an SDP offer and the "recording-session" disposition. A partial update represents an incremental update since the last metadata update sent by the SRC. A partial update sent by the SRC can be an INVITE request or response with an SDP offer, or an INVITE/UPDATE request or response containing a "recording-session" disposition, or an INVITE request containing both an SDP offer and the "recording-session" disposition.

The following is an example of a full metadata snapshot sent by the SRC in the initial INVITE request:

```
INVITE sip:recorder@example.com SIP/2.0
Via: SIP/2.0/TCP src.example.com;branch=z9hG4bKdf6b622b648d9
From: <sip:2000@example.com>;tag=35e195d2-947d-4585-946f-098392474
To: <sip:recorder@example.com>
Call-ID: d253c800-b0dlea39-4a7dd-3f0e20a
CSeq: 101 INVITE
Max-Forwards: 70
Require: siprec
Accept: application/sdp, application/rs-metadata,
       application/rs-metadata-request
Contact: <sip:2000@src.example.com>;+sip.src
Content-Type: multipart/mixed;boundary=foobar
Content-Length: [length]

--foobar
Content-Type: application/sdp

v=0
o=SRS 2890844526 2890844526 IN IP4 198.51.100.1
s=-
c=IN IP4 198.51.100.1
t=0 0
m=audio 12240 RTP/AVP 0 4 8
a=sendonly
a=label:1

--foobar
Content-Type: application/rs-metadata
Content-Disposition: recording-session

[metadata content]
```

Figure 12: Sample INVITE request for the recording session

9.2. Procedures at the SRS

The SRS receives metadata updates from the SRC in INVITE and UPDATE requests. Since the SRC can send partial updates based on the previous update, the SRS needs to keep track of the sequence of updates from the SRC.

In the case of an internal failure at the SRS, the SRS may fail to recognize a partial update from the SRC. The SRS may be able to recover from the internal failure by requesting for a full metadata snapshot from the SRC. Certain errors, such as syntax errors or semantic errors in the metadata information, are likely caused by an

error on the SRC side, and it is likely the same error will occur again even when a full metadata snapshot is requested. In order to avoid repeating the same error, the SRS can simply terminate the recording session when a syntax error or semantic error is detected in the metadata.

When the SRS explicitly requests for a full metadata snapshot, the SRS MUST send an UPDATE request without an SDP offer. A metadata snapshot request contains a content with the content disposition type "recording-session". Note that the SRS MAY generate an INVITE request without an SDP offer but this MUST NOT include a metadata snapshot request. The format of the content is "application/rs-metadata-request", and the body format is chosen to be a simple text-based format. The following shows an example:

```
UPDATE sip:2000@src.exmaple.com SIP/2.0
Via: SIP/2.0/UDP srs.example.com;branch=z9hG4bKdf6b622b648d9
To: <sip:2000@exmaple.com>;tag=35e195d2-947d-4585-946f-098392474
From: <sip:recorder@example.com>;tag=1234567890
Call-ID: d253c800-b0dlea39-4a7dd-3f0e20a
CSeq: 1 UPDATE
Max-Forwards: 70
Require: siprec
Contact: <sip:recorder@srs.example.com>;+sip.srs
Accept: application/sdp, application/rs-metadata
Content-Disposition: recording-session
Content-Type: application/rs-metadata-request
Content-Length: [length]

SRS internal error
```

Figure 13: Metadata Request

The SRS MAY include the reason why a metadata snapshot request is being made to the SRC in the reason line. This reason line is free form text, mainly designed for logging purposes on the SRC side. The processing of the content by the SRC is entirely optional since the content is for logging only, and the snapshot request itself is indicated by the use of the application/rs-metadata-request content type.

When the SRC receives the request for a metadata snapshot, the SRC MUST provide a full metadata snapshot in a separate INVITE or UPDATE transaction, along with an SDP offer. All subsequent metadata updates sent by the SRC MUST be based on the new metadata snapshot.

9.2.1. Formal Syntax

The formal syntax for the application/rs-metadata-request MIME is described below using the augmented Backus-Naur Form (BNF) as described in [RFC5234].

```
snapshot-request = srs-reason-line CRLF
```

```
srs-reason-line = [TEXT-UTF8-TRIM]
```

10. Persistent Recording

Persistent recording is a specific use case outlined in REQ-005 or Use Case 4 in [RFC6341], where a recording session can be established in the absence of a communication session. The SRC continuously records media in a recording session to the SRS even in the absence of a CS for all user agents that are part of persistent recording. By allocating recorded streams and continuously sending recorded media to the SRS, the SRC does not have to prepare new recorded streams with new SDP offer when a new communication session is created and also does not impact the timing of the CS. The SRC only needs to update the metadata when new communication sessions are created.

When there is no communication sessions running on the devices with persistent recording, there is no recorded media to stream from the SRC to the SRS. In certain environments where Network Address Translator (NAT) is used, typically a minimum of flow activity is required to maintain the NAT binding for each port opened. Agents that support Interactive Connectivity Establishment (ICE) solves this problem. For non-ICE agents, in order not to lose the NAT bindings for the RTP/RTCP ports opened for the recorded streams, the SRC and SRS SHOULD follow the recommendations provided in [RFC6263] to maintain the NAT bindings.

11. IANA Considerations

11.1. Registration of Option Tags

This specification registers two option tags. The required information for this registration, as specified in [RFC3261], is as follows.

11.1.1. siprec Option Tag

Name: siprec

Description: This option tag is for identifying the SIP session for the purpose of recording session only. This is typically not used in a Supported header. When present in a Require header in a request, it indicates that the UAS MUST be either a SRC or SRS capable of handling the contexts of a recording session.

11.1.2. record-aware Option Tag

Name: record-aware

Description: This option tag is to indicate the ability for the user agent to receive recording indicators in media level or session level SDP. When present in a Supported header, it indicates that the UA can receive recording indicators in media level or session level SDP.

11.2. Registration of media feature tags

This document registers two new media feature tags in the SIP tree per the process defined in [RFC2506] and [RFC3840]

11.2.1. src feature tag

Media feature tag name: sip.src

ASN.1 Identifier: 25

Summary of the media feature indicated by this tag: This feature tag indicates that the user agent is a Session Recording Client for the purpose for Recording Session.

Values appropriate for use with this feature tag: boolean

The feature tag is intended primarily for use in the following applications, protocols, services, or negotiation mechanisms: This feature tag is only useful for a Recording Session.

Examples of typical use: Routing the request to a Session Recording Server.

Security Considerations: Security considerations for this media feature tag are discussed in Section 11.1 of RFC 3840.

11.2.2. srs feature tag

Media feature tag name: sip.srs

ASN.1 Identifier: 26

Summary of the media feature indicated by this tag: This feature tag indicates that the user agent is a Session Recording Server for the purpose for Recording Session.

Values appropriate for use with this feature tag: boolean

The feature tag is intended primarily for use in the following applications, protocols, services, or negotiation mechanisms: This feature tag is only useful for a Recording Session.

Examples of typical use: Routing the request to a Session Recording Client.

Security Considerations: Security considerations for this media feature tag are discussed in Section 11.1 of RFC 3840.

11.3. New Content-Disposition Parameter Registrations

This document registers a new "disposition-type" value in Content-Disposition header: recording-session.

recording-session the body describes the metadata information about the recording session

11.4. Media Type Registration

11.4.1. Registration of MIME Type application/rs-metadata

This document registers the application/rs-metadata MIME media type in order to describe the recording session metadata. This media type is defined by the following information:

Media type name: application

Media subtype name: rs-metadata

Required parameters: none

Options parameters: none

11.4.2. Registration of MIME Type application/rs-metadata-request

This document registers the application/rs-metadata-request MIME media type in order to describe a recording session metadata snapshot request. This media type is defined by the following information:

Media type name: application

Media subtype name: rs-metadata-request

Required parameters: none

Options parameters: none

11.5. SDP Attributes

This document registers the following new SDP attributes.

11.5.1. 'record' SDP Attribute

Contact names: Leon Portman leon.portman@nice.com, Henry Lum
henry.lum@genesyslab.com

Attribute name: record

Long form attribute name: Recording Indication

Type of attribute: session or media level

Subject to charset: no

This attribute provides the recording indication for the session or media stream.

Allowed attribute values: on, off, paused

11.5.2. 'recordpref' SDP Attribute

Contact names: Leon Portman leon.portman@nice.com, Henry Lum
henry.lum@genesyslab.com

Attribute name: recordpref

Long form attribute name: Recording Preference

Type of attribute: session or media level

Subject to charset: no

This attribute provides the recording preference for the session or media stream.

Allowed attribute values: on, off, pause, nopreference

12. Security Considerations

The recording session is fundamentally a standard SIP dialog [RFC3261], therefore, the recording session can reuse any of the existing SIP security mechanism available for securing the session signaling, the recorded media as well as the metadata. The use cases and requirements document [RFC6341] outlines the general security considerations, and the following describe specific security recommendations.

The SRC and SRS MUST support SIP with TLS and MAY support SIPS with TLS as per [RFC5630]. The Recording Session SHOULD be at least as secure as the Communication Session, meaning using at least the same strength of cipher suite as the CS if the CS is secured. For example, if the CS uses SIPS for signalling and RTP/SAVP for media, then the RS does not downgrade the level of security in the RS to SIP or plain RTP since doing so will mean an automatic security downgrade for the CS. In deployments where the SRC and the SRS are in the same administrative domain and the same physical switch that prevents outside user access, some SRC may choose lower the level of security when establishing the recording session. While physically securing the SRC and SRS may prevent an outside attacker from accessing important call recordings, this still does not prevent an inside attacker from accessing the internal network to gain access to the call recordings.

12.1. Authentication and Authorization

The recording session reuses the SIP mechanism to challenge requests that are based on HTTP authentication. The mechanism relies on 401 and 407 SIP responses as well as other SIP header fields for carrying challenges and credentials.

At the transport level, the recording session uses TLS authentication to validate the authenticity of the SRC and SRS. The SRC and SRS MUST implement TLS mutual authentication for establishing the recording session, and whether the SRC/SRS chooses to use authentication is a deployment decision. In deployments where the SRC and the SRS are in the same administrative domain, the deployment may choose not to authenticate each other or only to have SRC authenticate the SRS as there is an inherent trust relation between the SRC and the SRS when they are hosted in the same administrative

domain. In deployments where the SRS can be hosted on a different administrative domain, then it is important to perform mutual authentication to ensure the authenticity of both the SRC and the SRS before transmitting any recorded media. The risk of not authenticating the SRS is that the recording may be sent to a compromised SRS and that sensitive call recording will be obtained by an attacker. On the other hand, the risk of not authenticating the SRC is that an SRS will accept calls from an unknown SRC and allow potential forgery of call recordings.

The SRS may have its own set of recording policies to authorize recording requests from the SRC. The use of recording policies is outside the scope of the Session Recording Protocol.

12.2. RTP handling

In many scenarios it will be critical that the media transported between the SRC and SRS to be protected. Media encryption is an important element in the overall SIPREC solution; therefore SRC and SRS MUST support RTP/SAVP [RFC3711] and RTP/SAVPF [RFC5124]. RTP/SAVP and RTP/SAVPF provide media encryption, integrity protection, replay protection, and a limited form of source authentication. They do not contain or require a specific keying mechanism.

When RTP/SAVP or RTP/SAVPF is used, RS can choose to use the same or different security keys than the ones used in the CS. Some SRCs are designed to simply replicate RTP packets from the CS media stream to the SRS, and the SRC will be reusing the same keys as the CS. In this case, the SRC MUST secure the SDP with SDP Security Descriptions (SDS) [RFC4568] in the RS with at least the same level of security as the CS. The risk of lowering the level of security in the RS for this case is that it will effectively become a downgrade attack on the CS since the same key is used for both CS and RS.

For SRCs that perform transcoding or mixing of media before sending to the SRS, the SRC MUST negotiate a different security key than the one being used in the CS, to ensure that the security in the CS is not compromised by the SRC when reusing the same security key.

12.3. Metadata

Metadata contains sensitive information such as the address of record of the participants and other extension data placed by the SRC. It is essential to protect the content of the metadata in the RS. Since metadata is a content type transmitted in SIP signalling, metadata SHOULD be protected at the transport level by SIPS/TLS.

12.4. Storage and playback

While storage and playback of the call recording is beyond the scope of this document, it is worthwhile to mention here that it is also important for the recording storage and playback to provide a level of security that is comparable to the communication session. It would defeat the purpose of securing both the communication session and the recording session mentioned in the previous sections if the recording can be easily played back with a simple unsecured HTTP interface without any form of authentication or authorization.

13. Acknowledgements

We want to thank John Elwell, Paul Kyzivat, Partharsarathi R, Ram Mohan R, Hadriel Kaplan, Adam Roach, Miguel Garcia, Thomas Stach, Muthu Perumal, Dan Wing, and Magnus Westerlund for their valuable comments and inputs to this document.

14. References

14.1. Normative References

- [I-D.ietf-siprec-metadata]
R, R., Ravindran, P., and P. Kyzivat, "Session Initiation Protocol (SIP) Recording Metadata",
draft-ietf-siprec-metadata-08 (work in progress),
October 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2506] Holtman, K., Mutz, A., and T. Hardie, "Media Feature Tag Registration Procedure", BCP 31, RFC 2506, March 1999.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [RFC3840] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", RFC 3840, August 2004.

- [RFC4574] Levin, O. and G. Camarillo, "The Session Description Protocol (SDP) Label Attribute", RFC 4574, August 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

14.2. Informative References

- [I-D.ietf-siprec-architecture]
Hutton, A., Portman, L., Jain, R., and K. Rehor, "An Architecture for Media Recording using the Session Initiation Protocol", draft-ietf-siprec-architecture-06 (work in progress), September 2012.
- [RFC3311] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.
- [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.
- [RFC4961] Wing, D., "Symmetric RTP / RTP Control Protocol (RTCP)", BCP 131, RFC 4961, July 2007.
- [RFC5104] Wenger, S., Chandra, U., Westerlund, M., and B. Burman, "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)", RFC 5104, February 2008.
- [RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, February 2008.

- [RFC5168] Levin, O., Even, R., and P. Hagendorf, "XML Schema for Media Control", RFC 5168, March 2008.
- [RFC5630] Audet, F., "The Use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", RFC 5630, October 2009.
- [RFC5761] Perkins, C. and M. Westerlund, "Multiplexing RTP Data and Control Packets on a Single Port", RFC 5761, April 2010.
- [RFC6222] Begen, A., Perkins, C., and D. Wing, "Guidelines for Choosing RTP Control Protocol (RTCP) Canonical Names (CNAMES)", RFC 6222, April 2011.
- [RFC6263] Marjou, X. and A. Sollaud, "Application Mechanism for Keeping Alive the NAT Mappings Associated with RTP / RTP Control Protocol (RTCP) Flows", RFC 6263, June 2011.
- [RFC6341] Rehor, K., Portman, L., Hutton, A., and R. Jain, "Use Cases and Requirements for SIP-Based Media Recording (SIPREC)", RFC 6341, August 2011.

Authors' Addresses

Leon Portman
NICE Systems
8 Hapnina
Ra'anana 43017
Israel

Email: leon.portman@nice.com

Henry Lum (editor)
Genesys
1380 Rodick Road, Suite 201
Markham, Ontario L3R4G5
Canada

Email: henry.lum@genesyslab.com

Charles Eckel
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States

Email: eckelcu@cisco.com

Alan Johnston
Avaya
St. Louis, MO 63124

Email: alan.b.johnston@gmail.com

Andrew Hutton
Siemens Enterprise Communications
Brickhill Street
Milton Keynes MK15 0DJ
United Kingdom

Email: andrew.hutton@siemens-enterprise.com

SIPREC
Internet-Draft
Intended status: Standards Track
Expires: January 17, 2013

Ram Mohan. Ravindranath
Cisco Systems, Inc.
Parthasarathi. Ravindran
Sonus Networks
Paul. Kyzivat
Huawei
July 16, 2012

Session Initiation Protocol (SIP) Recording Call Flows
draft-ram-siprec-callflows-01

Abstract

Session recording is a critical requirement in many communications environments such as call centers and financial trading. In some of these environments, all calls must be recorded for regulatory, compliance, and consumer protection reasons. Recording of a session is typically performed by sending a copy of a media stream to a recording device. This document lists call flows that has snapshot of metadata sent from SRC to SRS, the metadata format for which is described in [I-D.ietf-siprec-metadata] . This is purely an informational document that is written to support the model defined the metadata draft.

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 17, 2013.

Copyright Notice

Copyright (c) 2012 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. Overview	3
2. Terminology	3
3. SIP Recording Metadata Example	3
3.1. Complete SIP Recording Metadata Example	3
4. Metadata XML schema Instances	5
4.1. Sample Call flow	5
4.2. Example 1: Basic Call	6
4.3. Example 2: Hold/resume	9
4.4. Example 3: Basic Call with transfer	12
4.5. Example 4: Call disconnect	16
4.6. Example 5: Multiple CS into single RS with mixed stream	17
5. Security Considerations	19
6. IANA Considerations	19
7. Acknowledgement	19
8. References	19
8.1. Normative References	19
8.2. Informative References	19
Authors' Addresses	19

1. Overview

[I-D.ietf-siprec-metadata] document focuses on the Recording metadata which describes the communication session. The document lists a few examples and shows the snapshots of metadata sent from SRC to SRS. For the sake of simplicity the entire SIP [RFC3261] messages are not shown at various points, instead only a snip of the SIP/SDP message and the XML snapshot of metadata is shown. This document is informational and is not normative on any aspect of SIP.

2. Terminology

The terms using in this defined are defined in [I-D.ietf-siprec-metadata]. No new terms/definitions are introduced in this document.

3. SIP Recording Metadata Example

3.1. Complete SIP Recording Metadata Example

The following example provides all the tuples involved in Recording Metadata XML body. Subsequent sections will show snapshot of metadata sent for different scenarios.

```
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>complete</dataMode>
    <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </session>
    <participant>
      participant_id="srfBElmCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ=="
      <nameID aor=sip:bob@biloxi.com>
        <name xml:lang="it">Bob B</name>
      </nameID>
    </participant>
    <participantsessionassoc>
      participant_id="srfBElmCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ=="
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </participantsessionassoc>
    <participantstreamassoc>
      participant_id="srfBElmCRp2QB23b7Mpk0w=="
      <send>ilPz3to5hGk8fuXl+PbwCw==</send>
```

```

    <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
    <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
    <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<participant
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <nameID aor=sip:paul@biloxi.com>
    <name xml:lang="it">Paul</name>
  </nameID>
</participant>
<participantsessionassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
  <send>8zc6e0lYtLWIINA6GR+3ag==</send>
  <send>EiXGlc+4TruqqoDaNE76ag==</send>
  <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
  <recv>i1Pz3to5hGk8fuXl+PbwCw==</recv>
</participantstreamassoc>
<stream stream_id="UAAMm5GRQKSCMVvLyl4rFw=="
  session session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>96</label>
</stream>
<stream stream_id="i1Pz3to5hGk8fuXl+PbwCw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>97</label>
</stream>
<stream stream_id="8zc6e0lYtLWIINA6GR+3ag=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>98</label>
</stream>
<stream stream_id="EiXGlc+4TruqqoDaNE76ag=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <label>99</label>
</stream>
</recording>

```


SIP Recording Metadata Example XML body

4. Metadata XML schema Instances

This section describes the metadata model XML instances for different use cases of SIPREC. For the sake of simplicity the complete SIP messages are NOT shown here.

4.1. Sample Call flow


The following is a sample call flow that shows the SRC establishing a recording session towards the SRS. The call flow is essentially identical when the SRC is a B2BUA or as the endpoint itself. Note that the SRC can choose when to establish the Recording Session independent of the Communication Session, even though the following call flow suggests that the SRC is establishing the Recording Session (message #5) after the Communication Session is established.



The above call flow can also apply to the case of a centralized conference with a mixer. For clarity, ACKs to INVITES and 200 OKs to BYEs are not shown. The conference focus can provide the SRC functionality since the conference focus has access to all the media from each conference participant. When a recording is requested, the SRC delivers the metadata and the media streams to the SRS. Since the conference focus has access to a mixer, the SRC may choose to mix the media streams from all participants as a single mixed media stream towards the SRS.

An SRC can use a single recording session to record multiple communication sessions. Every time the SRC wants to record a new call, the SRC updates the recording session with a new SDP offer to add new recorded streams to the recording session, and correspondingly also update the metadata for the new call.

4.2. Example 1: Basic Call



Basic call between two Participants A(Ram) and B(Partha) who are part of one session. In this use case each participant sends two Media Streams. Media Streams sent by each participant is received all

other participants of that CS in this use-case. Below is the initial snapshot sent by SRC in the INVITE to SRS that has complete metadata. For the sake of simplicity snippets of SDP is shown. Here the RS stream is unmixed.

```
F1 INVITE SRC -----> SRS

Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>complete</dataMode>
    <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <start-time>2010-12-16T23:41:07Z</start-time>
    </session>
    <participant
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <nameID aor=sip:ram@example.com>
        <name xml:lang="it">RamMohan R</name>
      </nameID>
    </participant>
    <participantsessionassoc
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
```

```
</participantsessionassoc>
<participantstreamassoc
  participant_id="srfBElmCRp2QB23b7Mpk0w==">
  <send>ilPz3to5hGk8fuXl+PbwCw==</send>
  <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
  <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<participant
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ==">
  <nameID aor=sip:partha@example.com>
    <name xml:lang="it">Parthasarathi R</name>
  </nameID>
</participant>
<participantsessionassoc
  participant="zSfPoSvdSDCmU3A3TRDxAw=="
  session="hVpd7YQgRW2nD22h7q60JQ==">
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant="zSfPoSvdSDCmU3A3TRDxAw==">
  <send>8zc6e0lYtLWIINA6GR+3ag==</send>
  <send>EiXGlc+4TruqqoDaNE76ag==</send>
  <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
  <recv>ilPz3to5hGk8fuXl+PbwCw==</recv>
</participantstreamassoc>
<stream id="UAAMm5GRQKSCMVvLyl4rFw=="
  session="hVpd7YQgRW2nD22h7q60JQ==">
  <label>96</label>
</stream>
<stream id="ilPz3to5hGk8fuXl+PbwCw=="
  session="hVpd7YQgRW2nD22h7q60JQ==">
  <label>97</label>
</stream>
<stream id="8zc6e0lYtLWIINA6GR+3ag=="
  session="zSfPoSvdSDCmU3A3TRDxAw==">
  <label>98</label>
</stream>
<stream id="EiXGlc+4TruqqoDaNE76ag=="
  session="zSfPoSvdSDCmU3A3TRDxAw==">
  <label>99</label>
</stream>
</recording>
```

4.3. Example 2: Hold/resume

Basic call between two Participants A and B. This is the continuation of above use-case. One of the participants(say A) goes on hold and then resumes as part of the same session. The metadata snapshot looks as below

During hold

F4 RE-INVITE SRC----->SRS

Content-Type: application/SDP

...

m=audio 49170 RTP/AVP 0

a=rtpmap:0 PCMU/8000

a=label:96

a=inactive

...

m=video 49174 RTP/AVPF 96

a=rtpmap:96 H.264/90000

a=label:97

a=inactive

...

m=audio 51372 RTP/AVP 0

a=rtpmap:0 PCMU/8000

a=label:98

a=sendonly

...

m=video 49176 RTP/AVPF 96

a=rtpmap:96 H.264/90000

a=label:99

a=sendonly

....

<?xml version="1.0" encoding="UTF-8"?>

<recording xmlns='urn:ietf:params:xml:ns:recording'>

<dataMode>partial</dataMode>

<participantstreamassoc

participant="srfBEImCRp2QB23b7Mpk0w=="

session="hVpd7YQgRW2nD22h7q60JQ==">

<recv>8zc6e0lYtLWIINA6GR+3ag==</recv>

<recv>EiXGlc+4TruqgoDaNE76ag==</recv>

</participantstreamassoc>

<participantstreamassoc

participant="zSfPoSvdSDCmU3A3TRDxAw=="

session="hVpd7YQgRW2nD22h7q60JQ==">

<send>8zc6e0lYtLWIINA6GR+3ag==</send>

<send>EiXGlc+4TruqgoDaNE76ag==</send>

</participantstreamassoc>

</recording>

During resume

The snapshot will look pretty much same as above with just the SDP

dir change.

F5 RE-INVITE SRC----->SRS

```
Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....

<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <participantstreamassoc
      participant="srfBEImCRp2QB23b7Mpk0w=="
      session="hVpd7YQgRW2nD22h7q60JQ==">
      <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
      <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
    </participantstreamassoc>
    <participantstreamassoc
      participant="zSfPoSvdSDCmU3A3TRDxAw=="
      session="hVpd7YQgRW2nD22h7q60JQ==">
      <send>8zc6e0lYtLWIINA6GR+3ag==</send>
      <send>EiXGlc+4TruqqoDaNE76ag==</send>
    </participantstreamassoc>
  </recording>
```

4.4. Example 3: Basic Call with transfer

Basic call between two Participants A and B is connected as in Use-case 1. Transfer is initiated by one of the participants or by other entity(3PCC case). SRC sends a snapshot of the participant changes to SRS. In this instance participant A(Ram) drops out during the transfer and Participant C(Paul) joins the session. There can be two cases here, same session continues after transfer or a new session (e.g. REFER based transfer) is created

Transfer with same session retained - (.e.g. RE-INVITE based transfer). Participant A drops out and C is added to the same session. No change to session/group element. C will be new stream element which maps to RS SDP using the same labels in this instance.

```
Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>partial</dataMode>
    <participantstreamassoc
      participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
      <send>8zc6e0lYtLWIINA6GR+3ag==</send>
      <send>EiXGlc+4TruqgoDaNE76ag==</send>
      <recv>60JAJm9UTvik0Ltlih/Gzw==</recv>
      <recv>AcR5FUd3Edi8cACQJy/3JQ==</recv>
    </participantstreamassoc>
```

```

<participant
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  session_id="hVpd7YQgRW2nD22h7q60JQ=="
  <nameIDaor=sip:paul@example.com>
    <name xml:lang="it">Paul Kyzivat</name>
  </nameID>
  <send>60JAJm9UTvik0Ltlih/Gzw==</send>
  <send>AcR5FUd3Edi8cACQJy/3JQ==</send>
  <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participant>
<participantsessionassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  session_id="hVpd7YQgRW2nD22h7q60JQ=="
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsession>
<participantstreamassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  <send>60JAJm9UTvik0Ltlih/Gzw==</send>
  <send>AcR5FUd3Edi8cACQJy/3JQ==</send>
  <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<stream stream_id="60JAJm9UTvik0Ltlih/Gzw=="
  session_id="hVpd7YQgRW2nD22h7q60JQ=="
  <label>96</label>
</stream>
<stream stream_id="AcR5FUd3Edi8cACQJy/3JQ=="
  session_id="hVpd7YQgRW2nD22h7q60JQ=="
  <label>97</label>
</stream>
<stream stream_id="8zc6e0lYtLWIINA6GR+3ag=="
  session_id="zSfPoSvdSDCmU3A3TRDxAw=="
  <label>98</label>
</stream>
<stream stream_id="EiXGlc+4TruqqoDaNE76ag=="
  session_id="zSfPoSvdSDCmU3A3TRDxAw=="
  <label>99</label>
</stream>
</recording>

```

Transfer with new session - (.e.g. REFER based transfer). In this case new session is part of same grouping (done by SRC).

SRC may send an optional snapshot indicating stop for the old session.

```

    <?xml version="1.0" encoding="UTF-8"?>
    <recording xmlns='urn:ietf:params:xml:ns:recording'>
      <dataMode>Partial</dataMode>
      <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <stop-time>2010-12-16T23:41:07Z</stop-time>
      </session>
      <participantsessionassoc
        participant_id="srfBEImCRp2QB23b7Mpk0w=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
      </participantsessionassoc>
      <participantsessionassoc
        participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
        session_id="hVpd7YQgRW2nD22h7q60JQ==">
        <disasociate-time>2010-12-16T23:41:07Z</disasociate-time>
      </participantsessionassoc>
    </recording>

```

SRC sends a snapshot to indicate the participant change and new session information after transfer. In this example the same RS is used.

Content-Type: application/SDP

```

...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
m=video 49174 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:97
a=sendonly
...
m=audio 51372 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:98
a=sendonly
...
m=video 49176 RTP/AVPF 96
a=rtpmap:96 H.264/90000
a=label:99
a=sendonly
....
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>

```

```

    <dataMode>partial</dataMode>
<session session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <start-time>2010-12-16T23:41:07Z</start-time>
</session>
<participant
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <nameID aor=sip:partha@example.com/>
</participant>
<participantsessionassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <associate-time>2010-12-16T23:32:03Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="zSfPoSvdSDCmU3A3TRDxAw==">
  <send>8zc6e0lYtLWIINA6GR+3ag==</send>
  <send>EiXGlc+4TruqqoDaNE76ag==</send>
  <recv>60JAJm9UTvik0Ltlih/Gzw==</recv>
  <recv>AcR5FUd3Edi8cACQJy/3JQ==</recv>
</participantstreamassoc>
<participant
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <nameID aor=sip:paul@example.com/>
</participant>
<participantsessionassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <associate-time>2010-12-16T23:41:07Z</associate-time>
</participantsessionassoc>
<participantstreamassoc
  participant_id="AtnmlZRnOC6Pm5MApkrDzQ==">
  <send>60JAJm9UTvik0Ltlih/Gzw==</send>
  <send>AcR5FUd3Edi8cACQJy/3JQ==</send>
  <recv>8zc6e0lYtLWIINA6GR+3ag==</recv>
  <recv>EiXGlc+4TruqqoDaNE76ag==</recv>
</participantstreamassoc>
<stream stream_id="60JAJm9UTvik0Ltlih/Gzw=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>96</label>
</stream>
<stream stream_id="AcR5FUd3Edi8cACQJy/3JQ=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
  <label>97</label>
</stream>
<stream stream_id="8zc6e0lYtLWIINA6GR+3ag=="
  session_id="bfLZ+NTFEeCNxQTuRyQBmw==">

```


```

        <label>98</label>
    </stream>
    <stream stream_id="EiXGlc+4TruqgoDaNE76ag=="
        session_id="bfLZ+NTFEeCNxQTuRyQBmw==">
        <label>99</label>
    </stream>
</recording>

```

4.5. Example 4: Call disconnect

This example shows a snapshot of metadata sent by an SRC at CS disconnect where the participants of CS are Ram and Partha



F1 BYE SRC -----> SRS

```

<?xml version="1.0" encoding="UTF-8"?>
<recording xmlns='urn:ietf:params:xml:ns:recording'>
  <dataMode>Partial</dataMode>
  <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
    <stop-time>2010-12-16T23:41:07Z</stop-time>
  </session>
  <participant
    participant_id="srfBELmCRp2QB23b7Mpk0w=="
    session_id="hVpd7YQgRW2nD22h7q60JQ=="
    <nameID aor=sip:ram@example.com/>
    <disassociate-time>2010-12-16T23:41:07Z</disassociate-time>
  </participant>

  <participant
    participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
    session_id="hVpd7YQgRW2nD22h7q60JQ=="
    <nameID aor=sip:partha@example.com/>
    <disasociate-time>2010-12-16T23:41:07Z</disasociate-time>
  </participant>
</recording>

```

4.6. Example 5: Multiple CS into single RS with mixed stream

In trading turrent, the multiple audio stream in different communication sessions are mixed into single recording session.

```

F1 INVITE SRC -----> SRS

Content-Type: application/SDP
...
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMU/8000
a=label:96
a=sendonly
...
<?xml version="1.0" encoding="UTF-8"?>
  <recording xmlns='urn:ietf:params:xml:ns:recording'>
    <dataMode>complete</dataMode>
    <group group_id="7+OTCyoxTmqmqyA/1weDAg==">
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </group>
    <session session_id="hVpd7YQgRW2nD22h7q60JQ==">
      <group-ref>7+OTCyoxTmqmqyA/1weDAg==</group-ref>
      <start-time>2010-12-16T23:41:07Z</start-time>
    </session>
    <session session_id="zzlafnvvj1CHllaHF6mn8kkSS==">
      <group-ref>7+OTCyoxTmqmqyA/1weDAg==</group-ref>
      <start-time>2010-12-16T23:43:07Z</start-time>
    </session>
    <participant
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ=="
      <nameID aor=sip:ram@example.com>
        <name xml:lang="it">RamMohan R</name>
      </nameID>
    </participant>
    <participantsessionassoc
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      session_id="hVpd7YQgRW2nD22h7q60JQ=="
      <associate-time>2010-12-16T23:41:07Z</associate-time>
    </participantsessionassoc>
    <participantstreamassoc
      participant_id="srfBEImCRp2QB23b7Mpk0w=="
      <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
      <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
    </participantstreamassoc>
  </participant>

```

```
    participant_id="zSfPoSvdSDCmU3A3TRDxAw=="
    session_id="hVpd7YQgRW2nD22h7q60JQ==">
    <nameID aor=sip:partha@example.com>
      <name xml:lang="it">Parthasarathi R</name>
    </nameID>
  </participant>
  <participantsessionassoc
    participant="zSfPoSvdSDCmU3A3TRDxAw=="
    session="hVpd7YQgRW2nD22h7q60JQ==">
    <associate-time>2010-12-16T23:41:07Z</associate-time>
  </participantsessionassoc>
  <participantstreamassoc
    participant="zSfPoSvdSDCmU3A3TRDxAw=="
    <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
    <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
  </participantstreamassoc>
  <participant
    participant_id="EiXGlc+4TruqqoDaNE76ag=="
    session_id="zzlafnvj1CHllaHF6mn8kkSS==">
    <nameID aor=sip:paul@example.com>
      <name xml:lang="it">Paul kyzivat</name>
    </nameID>
  </participant>
  <participantsessionassoc
    participant="EiXGlc+4TruqqoDaNE76ag=="
    session="zzlafnvj1CHllaHF6mn8kkSS==">
    <associate-time>2010-12-16T23:43:07Z</associate-time>
  </participantsessionassoc>
  <participantstreamassoc
    participant="EiXGlc+4TruqqoDaNE76ag=="
    <send>UAAMm5GRQKSCMVvLyl4rFw==</send>
    <recv>UAAMm5GRQKSCMVvLyl4rFw==</recv>
  </participantstreamassoc>

  <stream id="UAAMm5GRQKSCMVvLyl4rFw=="
    session="hVpd7YQgRW2nD22h7q60JQ==">
    <label>96</label>
  </stream>
  <stream id="UAAMm5GRQKSCMVvLyl4rFw=="
    session="zzlafnvj1CHllaHF6mn8kkSS==">
    <label>96</label>
  </stream>
</recording>
```


5. Security Considerations

The security considerations for metadata are defined in [I-D.ietf-siprec-metadata]

6. IANA Considerations

This document has no IANA considerations

7. Acknowledgement

TBD

8. References

8.1. Normative References

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

8.2. Informative References

[I-D.ietf-siprec-metadata]
R, R., Ravindran, P., and P. Kyzivat, "Session Initiation Protocol (SIP) Recording Metadata", draft-ietf-siprec-metadata-07 (work in progress), July 2012.

Authors' Addresses

Ram Mohan Ravindranath
Cisco Systems, Inc.
Cessna Business Park,
Kadabeesanahalli Village, Varthur Hobli,
Sarjapur-Marathahalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: rmohanr@cisco.com

Parthasarathi Ravindran
Sonus Networks
Prestige Shantiniketan - Business Precinct
Whitefield Road
Bangalore, Karnataka 560066
India

Email: pravindran@sonusnet.com

Paul Kyzivat
Huawei

Email: pkyzivat@alum.mit.edu

