

DISPATCH
Internet-Draft
Intended status: Informational
Expires: July 4, 2011

K. Rehor, Ed.
Cisco Systems
L. Portman, Ed.
NICE Systems
A. Hutton
Siemens Enterprise Communications
R. Jain
IPC Systems
December 31, 2010

Requirements for SIP-based Media Recording (SIPREC)
draft-ietf-siprec-req-06

Abstract

Session recording is a critical requirement in many business communications environments such as call centers and financial trading floors. In some of these environments, all calls must be recorded for regulatory and compliance reasons. In others, calls may be recorded for quality control or business analytics.

Recording is typically done by sending a copy of the session media to the recording devices. This document specifies requirements for extensions to SIP that will manage delivery of RTP media from an endpoint that originates media (or that has access to it) to a recording device. This is being referred to as SIP-based Media Recording.

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 July 4, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the

Internet-Draft

Requirements for SIPREC

December 2010

document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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

Internet-Draft

Requirements for SIPREC

December 2010

Table of Contents

1.	Requirements notation	4
2.	Introduction	4
3.	Definitions	5
4.	Use Cases	7
5.	Requirements	11
6.	Privacy Considerations	13
7.	Security Considerations	14
8.	IANA Considerations	15
9.	Acknowledgements	15
10.	Normative References	15
	Authors' Addresses	16

1. Requirements notation

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](#)] and indicate requirement levels for compliant mechanisms.

2. Introduction

Session recording is a critical operational requirement in many businesses, especially where voice is used as a medium for commerce and customer support. A prime example where voice is used for trade is the financial industry. The call recording requirements in this industry are quite stringent. The recorded calls are used for dispute resolution and compliance. Other businesses such as customer support call centers typically employ call recording for quality control or business analytics, with different requirements.

Depending on the country and its regulatory requirements, financial trading floors typically must record all calls. In contrast, call centers typically only record a subset of the calls, and calls must not fail regardless of the availability of the recording device.

Respecting the privacy rights and wishes of users engaged in a call is of paramount importance. In many jurisdictions participants have a right to know that the session is being recorded or might be recorded, and have a right to opt out, either by terminating the call or by demanding that the call not be recorded. Therefore this

document contains requirements for being able to notify users that a call is being recorded and for users to be able to request that a call not be recorded. Use cases where users participating in a call are not informed that the call is or might be recorded are outside the scope of this document. In particular, lawful intercept is outside the scope of this document.

Furthermore, the scale and cost burdens vary widely, in all markets, where the different needs for solution capabilities such as media injection, transcoding, and security-related needs do not conform well to a one-size-fits-all model. If a standardized solution supports all of the requirements from every recording market, but doing so would be expensive for markets with lesser needs, then proprietary solutions for those markets will continue to propagate. Care must be taken, therefore, to make a standards-based solution support optionality and flexibility.

This document specifies requirements for using SIP [[RFC3261](#)] between a Session Recording Client and a Session Recording Server to control

the recording of media that has been transmitted in the context of a Communication Session. A Communication Session is the "call" between participants. The Session Recording Client is the source of the recorded media. The Session Recording Server is the sink of recorded media. It should be noted that the requirements for the protocol between a Session Recording Server and Session Recording Client have very similar requirements (such as codec and transport negotiation, encryption key interchange, firewall traversal) as compared to regular SIP media sessions. The choice of SIP for session recording provides reuse of an existing protocol.

The recorded sessions can be any RTP media sessions including voice, DTMF (as defined by [[RFC4733](#)]), video, and text (as defined by [[RFC4103](#)]).

An archived session recording is typically comprised of the Communication Session media content and the Communication Session Metadata. The Communication Session Metadata allows recording archives to be searched and filtered at a later time and allows a session to be played back in a meaningful way, e.g., with correct synchronization between the media. The Communication Session Metadata needs to be conveyed from the Session Recording Client to

the Session Recording Server.

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, is outside the scope of this document.

3. 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 a logical function that typically archives media for extended durations of time and provides interfaces for search and retrieval of the archived media. An SRS is typically implemented as a multi-port device that is capable of receiving media from several sources simultaneously. An SRS is typically also 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 itself or another entity working on its behalf (such as a SIP Application Server) may act as the source of the recording metadata.

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

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

Figure 1 pictorially represents the relationship between a Recording Session and Communication Session.

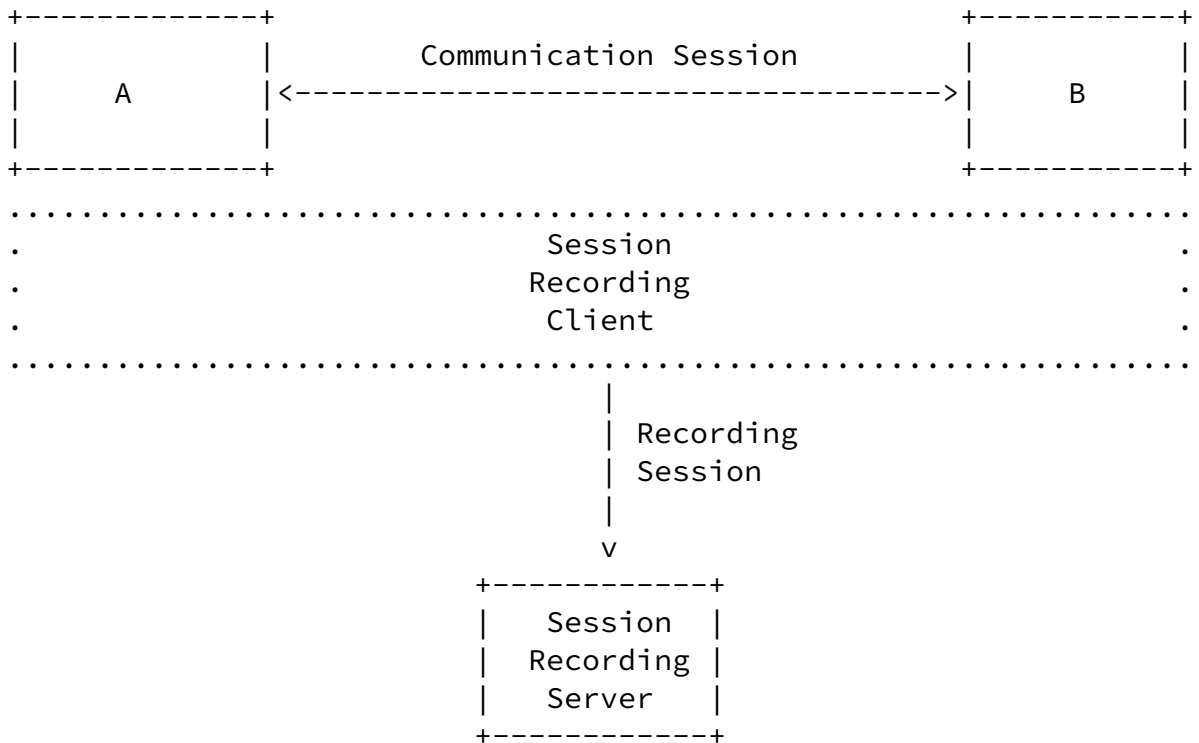


Figure 1

Metadata: Information that describes recorded media and the CS to which they relate.

Pause and Resume during a Communication Session: Pause: The action of temporarily discontinuing the transmission and collection of RS media
 Resume: The action of recommencing the transmission and collection of RS media

4. Use Cases

Use Case 1: Full-time Recording: One (or more, in the case of redundant recording) Recording Session for each Communication Session.

For example, the diagram below shows the lifecycle of Communication Sessions (CS) and the relationship to the Recording Sessions (RS)

```

CS |--- CS 1 ---|      |--- CS 2 ---|      |--- CS 3 ---|
RS |--- RS 1 ---|      |--- RS 2 ---|      |--- RS 3 ---|

```

Figure 2

Record every CS for specific extension/person.

The need to record all calls is typically due to business process purposes (such as transaction confirmation or dispute resolution) or to ensure compliance with governmental regulations. Applications include enterprise, contact center, and financial trading floors.

Also commonly known as Total Recording.

Use Case 2: Selective Recording: Start a Recording Session when a Communication Session to be recorded is established.

In this example, Communication Sessions 1 and 3 are recorded but CS 2 is not.

```

CS |--- CS 1 ---|      |--- CS 2 ---|      |--- CS 3 ---|
RS |--- RS 1 ---|      |--- RS 2 ---|

```

Figure 3

Use Case 3: Dynamic Recording: Start/Stop a Recording Session during a Communication Session.

The Recording Session starts during a Communication Session, either manually via a user-controlled mechanism (e.g. button on user's phone) or automatically via an application (e.g. a Contact Center customer service application) or business event. A Recording Session either ends during the Communication Session, or when the Communication Session ends.




Figure 4

Also known as Mid-session or Mid-call Recording.

Use Case 4: Persistent Recording: A single Recording Session captures one or more Communication Sessions, in sequence (Fig. 6) or in parallel (Fig. 7).




Figure 5

A Recording Session records continuously without interruption. Silent periods must be reproduced upon playback (e.g. by recording the silent period, by not recording the silent periods but marking them as metadata for a player to utilize, etc.) Applications include financial trading desks and emergency (first-responder) service bureaus. The length of a Persistent Recording Sessions is independent from the length of the actual Communication Sessions. Persistent Recording Sessions avoid issues such as media clipping that can occur due to delays in Recording Session establishment.

The connection and attributes of media in the Recording Session are not dynamically signaled for each Communication Session before it can be recorded; however, codec re-negotiation is possible. CS details and CS metadata will still be signaled, and can be correlated to the recorded media. There will still need to be a means of correlating the recorded media connection/packets to the Communication Session.

In some cases, more than one concurrent Communication Session (on a single end-user apparatus, e.g. trading floor turret) is mixed into one Recording Session:

```

          |----- CS 1 -----|
            |----- CS 2 -----|
          |----- CS 3 -----|

RS  |----- Recording Session -----|

```

Figure 6

Use Case 5: Real-time Recording Controls.

For an active Recording Session, privacy or security reasons may demand not capturing a specific portion of a conversation. An example is for PCI (payment card industry) compliance where credit card info must be protected. One solution is to not record a caller speaking their credit card information.

An example of a real-time controls is Pause/Resume.

Use Case 6: IVR / Voice Portal Recording.

Self-service Interactive Voice Response (DTMF or ASR) applications may need to be recorded for application performance tuning or to meet compliance requirements.

Metadata about an IVR session recording must include session information and may include application context information (e.g. VoiceXML session variables, dialog names, etc.)

Use Case 7: Enterprise Mobility Recording.

Many agents and enterprise workers are not located on company premises.

Examples:

- o Home-based agents or enterprise workers.
- o Mobile phones of knowledge workers when they conduct work related (and legally required recording) calls. i.e. insurance agents, brokers, physicians.

Use Case 8: Geographically distributed or centralized recording.

Global banks with multiple branches up to thousands of small sites.

- o Only phones and network infrastructure in branches, no recording services.

- o Internal calls inside or between branches must be recorded.

- o Centralized recording system in data centers together with telephony infrastructure (e.g. PBX).

Use Case 9: Record complex call scenarios.

Record a call that is associated with another call.

Example:

- o Customer in conversation with Agent

- o Agent puts customer on hold in order to consult with a Supervisor.

- o Agent in conversation with Supervisor.

- o Agent disconnects from Supervisor, reconnects with Customer.

- o The Supervisor call must be associated with the original customer call.

Use case 10: High availability and continuous recording.

Specific deployment scenarios present different requirements for system availability, error handling, etc. including:

- o An SRS must always be available at call setup time.

- o No loss of media recording, including during failure of an SRS.

- o The Communication Session must be terminated (or suitable notification) in the event of a recording failure.

Use Case 11: Record multi-channel, multi-media session.

Some applications require the recording of more than one media stream, possibly of different types. Media is synchronized, either at storage or at playback.

Speech analytics technologies (e.g. word spotting, emotion detection, speaker identification) may require speaker-separated recordings for optimum performance.

Multi-modal Contact Centers may include audio, video, IM or other

interaction modalities.

In trading floors environments, in order to save resources, it may be preferable to mix multiple concurrent calls (Communication Sessions) on different handsets/speakers on the same turret into single recording session.

Use Case 12: Real-time media processing.

Recorder must support real-time media processing, such as speech analytics.

Recording and real-time analytics of trading floor interactions (including video and instant messaging). Real time analytics is required for automatic intervention (stopping interaction or alert) if for example, trader is not following regulations.

Speaker separation is required in order to reliably detect who is saying specific phrases.

5. Requirements

The following are requirements for SIP-based Media Recording:

- o REQ-001 The mechanism MUST provide a means for "using the SIP protocol for" establishing, maintaining and terminating Recording Sessions between a Session Recording Client and a Session Recording Server.
- o REQ-002 The mechanism MUST support the ability to record all CSs in their entirety.

- o REQ-003 The mechanism MUST support the ability to record selected CSs in their entirety, according to policy.
- o REQ-004 The mechanism MUST support the ability to record selected parts of selected CSs.
- o REQ-005 The mechanism MUST support the ability to record a CS without an intentional loss of media (for example, clipping media at the beginning of the CS) and without impacting the quality or timing of the CS (for example, delaying the start of the CS while preparation for recording takes place). See Use Case 4 in [Section 5](#).
- o REQ-006 The mechanism MUST support the recording of IVR sessions.
- o REQ-007 The mechanism MUST support the recording of RTP media types

voice, DTMF (as defined by [\[RFC4733\]](#)), video, and text (as defined by [\[RFC4103\]](#)).

- o REQ-008 The mechanism MUST support the ability for an SRC to deliver mixed audio streams from multiple Communication Sessions to an SRS.

Note: A mixed audio stream is where several Communication Sessions are carried in a single Recording Session. A mixed media stream is typically produced by a mixer function. The RS MAY be informed about the composition of the mixed streams through session metadata.

- o REQ-009: The mechanism MUST support the ability for an SRC to deliver mixed audio streams from different parties of a given Communication Session to an SRS.
- o REQ-010 The mechanism MUST support the ability to deliver to the SRS multiple media streams for a given CS.
- o REQ-011 The mechanism MUST support the ability to pause and resume the transmission and collection of RS media.
- o REQ-012 The mechanism MUST provide the SRS with metadata describing CSs that are being recorded, including the media being used and the identities of parties involved.

- o REQ-013 The mechanism MUST provide the SRS with the means to correlate RS media with CS participant media described in metadata.
- o REQ-014 Metadata format must be agnostic of the transport protocol.
- o REQ-015: The mechanism MUST support a means to cancel and discard the recording and associated metadata for a CS.
- o REQ-016: The mechanism MUST support a means to cancel and discard the recording but not the associated metadata for a CS.
- o REQ-017 The mechanism MUST support a means for an authorized participant involved in a CS to request, prior to the start of recording, that the CS not be recorded
- o REQ-018 The mechanism MUST provide a means of indicating to the participants involved in a CS that their session is being recorded.

Examples include: inject tones into the CS from the SRC, play a message at the beginning of a session, a visual indicator on a display, etc.

- o REQ-019 The mechanism MUST provide a way for metadata to be conveyed to the SRS incrementally during the CS.
- o REQ-020 The mechanism MUST NOT prevent high availability deployments.
- o REQ-021 The mechanism MUST support means to relate Recording Session(s) with Communication Session(s).
- o REQ-022 The mechanism MUST provide the SRS the starting wall clock time for each RS media stream corresponding to the CS participant media.
- o REQ-023 The mechanism MUST provide the SRS the wall clock time when the Recording Session is paused and resumed.
- o REQ-024 The mechanism MUST support functionality such that if the CS is encrypted, the RS may use the same or different encryption

keys.

o REQ-025 The mechanism MUST provide means for an SRS to authenticate the SRC on RS initiation.

o REQ-026 The mechanism MUST provide means for an SRC to authenticate the SRS on RS initiation.

o REQ-027 The mechanism MUST ensure that the integrity of the metadata sent from SRC to SRS is an accurate representation of the original CS metadata.

o REQ-028 The mechanism MUST ensure that the integrity of the media sent from SRC to SRS is an accurate representation of the original CS media.

o REQ-029 The mechanism MUST ensure the confidentiality of the Metadata sent from SRC to SRS.

o REQ-030 The mechanism MUST provide a means to support RS confidentiality.

6. Privacy Considerations

Respecting the privacy rights and wishes of users engaged in a call is of paramount importance. In many jurisdictions participants have a right to know that the session is being recorded or might be recorded, and have a right to opt out, either by terminating the call or by demanding that the call not be recorded. Therefore this

document contains requirements for being able to notify users that a call is being recorded and for users to be able to request that a call not be recorded. Use cases where users participating in a call are not informed that the call is or might be recorded are outside the scope of this document. In particular, lawful intercept is outside the scope of this document.

Requirements for participant notification of recording varies widely by jurisdiction. In a given deployment, not all users will be authorized to stop the recording of a CS (although any user can terminate a CS). Typically users within the domain that is carrying

out the recording will be subject to policies of that domain concerning whether CSs are recorded. For example, in a call centre, agents will be subject to policies of the call centre and may or may not have the right to prevent the recording of a CS or part of a CS. Users calling into the call centre, on the other hand, will typically have to ask the agent not to record the CS. If the agent is unable to prevent recording, or if caller does not trust the agent, the only option generally is to terminate the CS.

Privacy considerations also extend to what happens to a recording once it has been created. Typical issues are who can access the recording (e.g., receive a copy of the recording, view the metadata, play back the media, etc.), for what purpose can the recording be used (e.g., for non-repudiation, for training purposes, for quality control purposes, etc.) and for how long the recording is to be retained before deletion. These are typically policies of the domain that makes the recording, rather than policies of individual users involved in a recorded CS, whether those users be in the same domain or in a different domain. Taking the call centre example again, agents might be made aware of call centre policy regarding retention and use of recordings as part of their employment contract, and callers from outside the call centre might be given some information about policy when notified that a CS will be recorded (e.g., through an announcement that says that calls may be recorded for quality purposes).

This document does not specify any requirements for a user engaged in a CS to be able to dictate policy for what happens to a recording, or for such information to be conveyed from an SRC to an SRS. It is assumed that the SRS has access to policy applicable to its environment and can ensure that recordings are stored and used in accordance with that policy.

7. Security Considerations

Session recording has substantial security implications, for the SIP

UA's being recorded, the SRC, and the SRS.

For the SIP UA's involved in the Communication Session, the requirements in this draft enable the UA to identify that a

Communication Session is being recorded and for the UA to request that a given Communication Session is not subject to recording.

Since humans don't typically look at or know about protocol signaling such as SIP, and indeed the SIP session might have originated through a PSTN Gateway without any ability to pass on in-signaling indications of recording, users can be notified of recording in the media itself through voice announcements, a visual indicator on the endpoint, or other means.

With regards to security implications of the protocol(s), clearly there is a need for authentication, authorization, eavesdropping protection, and non-repudiation for the solution. The SRC needs to know the SRS it is communicating with is legitimate, and vice-versa, even if they are in different domains. Both the signaling and media for the Recording Session need the ability to be authenticated and protected from eavesdropping and non-repudiation. Requirements are detailed in the requirements section.

8. IANA Considerations

This document has no IANA actions.

9. Acknowledgements

Thanks to Dan Wing, Alan Johnson, Vijay Gurbani, Cullen Jennings, Hadriel Kaplan, Henry Lum, Dave Smith, Martin Palmer, Alissa Cooper, Deepanshu Gautam, Paul Kyzivat, Parthasarathi R, and Ram Mohan R for their significant contributions and assistance with this document and Working Group, and to all the members of the DISPATCH WG and SIPREC WG mailing lists for providing valuable input to this work.

10. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

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

[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.

Authors' Addresses

Ken Rehor (editor)
Cisco Systems
170 West Tasman Dr.
Mail Stop SJC30/2/
San Jose, CA 95134
USA

Email: krehor@cisco.com

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

Email: leon.portman@nice.com

Andrew Hutton
Siemens Enterprise Communications

Email: andrew.hutton@siemens-enterprise.com
URI: <http://www.siemens-enterprise.com>

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

Email: rajnish.jain@ipc.com

