

SIPPING Working Group
Internet-Draft
Expires: August 31, 2003

O. Levin
RADVISION
R. Even
Polycom
March 2, 2003

High Level Requirements for Tightly Coupled SIP Conferencing
draft-levin-sipping-conferencing-requirements-03

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 31, 2003.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This document examines a wide range of conferencing requirements for tightly coupled SIP conferences. Separate documents will map the requirements to existing protocol primitives, define new protocol extensions, and introduce new protocols as needed. Together, these documents will provide a guide for building interoperable SIP conferencing applications.

Table of Contents

1.	Scope	3
2.	An Overview	3
3.	High Level Requirements	4
3.1	Discovery Phase	4
3.2	Conference Creation	5
3.3	Conference Termination	5
3.4	Participants Manipulations	5
3.4.1	Participation of a Conference-unaware User Agent	5
3.4.2	Dial-Out Scenarios	6
3.4.3	Dial-In Scenarios	6
3.4.4	Third Party Invitation to a Conference	6
3.4.5	Participants Removal	7
3.4.6	Participants Privacy	7
3.5	Conference State Information	7
3.5.1	Description	7
3.5.2	Dissemination of Changes	8
3.5.3	On-demand Information Dissemination	8
3.6	Focus Role Migration	9
3.7	Side-bar Conferences	9
3.8	Cascading of Conferences	10
3.9	SIMPLE and SIP Conferencing Coordination	10
4.	Security Considerations	10
5.	Contributors	10
	Normative References	10
	Informative References	11
	Authors' Addresses	11
	Intellectual Property and Copyright Statements	12

1. Scope

This document examines a wide range of conferencing requirements for tightly coupled SIP ([RFC 3261](#) [2]) conferencing.

The requirements are grouped by subjects in various areas allowing solutions to progress in parallel.

Separate documents will map the requirements to existing protocol primitives, define new protocol extensions, and introduce new protocols as needed.

Together, these documents will provide a guide for building interoperable SIP conferencing applications.

2. An Overview

A SIP conference is an association of SIP user agents (i.e. conference participants) with a central point (i.e. a conference focus) where the focus has a direct peer-wise relationship with each participant by means of a SIP dialog. Each dialog can belong to a different SIP session.

The focus is a SIP user agent which has abilities to host SIP conferences including their creation, maintenance, and manipulation using SIP call control means and potentially other non-SIP means.

In this tightly coupled model, the SIP conference graph is always a star. The conference focus maintains the correlation among conference's dialogs internally.

The conference focus can be implemented either by a participant or by a separate application server.

In the first case, a focus is typically capable of hosting a simple ad-hoc conference only. We envision that such basic conference can be established using SIP call control primitives only.

A dedicated conference server, in addition to the basic features, offers richer functionality including simultaneous conferences, large scalable conferences, reserved conferences, and managed conferences. A conferencing server can support any subset of the advanced conferencing functions presented in this document.

The media graph of a SIP conference can be centralized, de-centralized, or any combination of both and potentially differ per media type. In centralized case, the media sessions are established between the focus and each one of the participants. In de-centralized

(i.e. distributed) case, the media graph is a (multicast or multi-unicast) mesh among the participants. Consequently, the media processing (e.g. mixing) can be performed either by the focus alone or by the participants.

Conference participants and third parties can have different roles and privileges in a certain conference. For example, conferencing policy can state that the rights to disconnect from and to invite to a conference are limited to the conference chair only.

Throughout the document, by conference policies we mean a set of parameters and rules (e.g. maximum number of participants, needs chair-person supervision or not, password protected or not, duration, a way of media mixing, etc.) that are defined at the onset of a conference. Typically, conference policies would be specified by a conference creator and need special privileges to be manipulated.

Throughout the document, by a conference state we mean a set of information describing the conference in progress. This includes participants information (such as dialog identifiers), media sessions in progress, the current loudest speaker, the current chair, etc.

3. High Level Requirements

In addition to the requirements presented in this document, supplementary requirements for conferencing policy, media mixing and other manipulations, floor control, privileges control, etc. will be discussed in separate documents.

3.1 Discovery Phase

Some of the requirements presented in this section can be met either by configuration means or by using proprietary conventions. Nevertheless, we feel that standard means for implementing these functions by automata MUST be defined.

REQ -1: Discovery of a location of an arbitrary SIP conferencing server(s).

Editor s Note: No solution currently exists.

REQ -2: Given a SIP AOR of a certain entity, resolution whether the SIP entity has focus capabilities.

Editor s Note: No solution currently exists.

REQ -3: Given a global identifier of a particular conference,

locating the conference focus.

REQ -4: Given a global identifier of a particular conference, obtaining the conference properties.

REQ -5: Given a global identifier of a particular conference, obtaining the conference state information.

3.2 Conference Creation

Given a focus location, a means **MUST** be defined for an interested entity (including a user agent) to implement the procedures below:

REQ -1: Creation of an ad-hoc conference identifier and the conference with default properties.

REQ -2: Creation of an ad-hoc conference identifier and the conference with particular properties.

REQ -3: Creation of a reserved conference identifier for a conference with default properties.

REQ -4: Creation of a reserved conference identifier for a conference with particular properties.

3.3 Conference Termination

REQ -1: Given a conference identifier, a means **MUST** be defined for a user agent to disconnect all participants from the conference and terminate the conference including the release of the associated resources.

REQ -2: A means **MAY** be defined for requesting a focus to revert a two-party conference to a basic SIP point-to-point session including the release of the associated conferencing resources.

3.4 Participants Manipulations

Some of the requirements presented in this section can be met by human intervention, configuration means, or by using proprietary conventions. Nevertheless, we feel that standard means for implementing these functions by automata **MUST** be defined.

3.4.1 Participation of a Conference-unaware User Agent

REQ -1: Focus **MUST** be able to invite and disconnect an [RFC 3261](#) compliant only SIP user agent to and from a SIP conference.

REQ -2: [RFC 3261](#) compliant only SIP user agent MUST be able to dial-in a particular SIP conference. In this case, only the human knows that he/she is connected to the conference.

3.4.2 Dial-Out Scenarios

REQ -1: A means MUST be defined for a focus to invite another user agent to one of the focus conferences. This procedure MUST result in establishing of a single SIP dialog between the two.

REQ -2: Given an existent SIP dialog between two user agents, where at least one with focus capabilities, a means MUST be defined for the conference focus to invite the other user agent to one of the focus conferences without additional SIP dialog establishment.

REQ -3: An invitation to a user agent to join a conference MUST include a standard indication that it is "a conference" and the conference identifier.

3.4.3 Dial-In Scenarios

REQ -1: A means MUST be defined for a user agent to create an ad-hoc conference with default properties (as per "Conference Creation" REQ -1 above) and to become its participant using a single SIP dialog.

REQ -2: Given a reserved conference identifier, a means MUST be defined for a user agent to activate the conference and to become its participant using a single SIP dialog.

REQ -3: Given a conference identifier of an active conference, a means MUST be defined for a user agent to dial-in the conference and to become its participant using a single SIP dialog between the two.

REQ -4: Given an identifier of one of the dialogs of a particular active conference, a means MUST be defined for a user agent to dial-in the conference and to become its participant.

3.4.4 Third Party Invitation to a Conference

REQ -1: Given a conference identifier, a means MUST be defined for a user agent to invite another user agent to this conference.

REQ -2: Given an identifier of one of the dialogs of a particular active conference, a means MUST be defined for a user agent to invite another user agent to this conference.

REQ -3: Given a conference identifier, a means SHOULD be defined for a user agent to invite a list of user agents to this conference (a

so-called "mass invitation").

3.4.5 Participants Removal

REQ -1: A means **MUST** be defined for a conference focus to remove a conference participant from the conference.

REQ -2: Given a conference identifier, a means **MUST** be defined for a user agent to remove a participant from the conference.

REQ -3: Given an identifier of one of the dialogs of a particular active conference, a means **MUST** be defined for a user agent to remove a participant from the conference.

REQ -4: Given a conference identifier, a means **MUST** be defined for a user agent to remove all the participants from the conference.

REQ -5: Given a conference identifier and a sub-list of participants, a means **MAY** be defined for a user agent to remove the specified participants from the conference (a so-called "mass ejection").

3.4.6 Participants Privacy

A conference focus **SHOULD** support the procedures described in this section. A conference participant **MAY** support the procedures described in this section.

REQ -1: A conference participant joins the conference "anonymously", i.e. his/her presence can be announced but without disclosing his/her identity.

REQ -2: A conference participant requests a focus for anonymous participation in the conference.

REQ -3: A conference participant joins a conference in a "hidden mode", i.e. his/her both presence and identity are not to be disclosed to other participants.

REQ -4: A conference participant requests a focus for participation in the conference in a hidden mode.

3.5 Conference State Information

3.5.1 Description

By a conference state we mean a virtual database describing the conference in progress. This includes different conference aspects - participants information (such as dialog identifiers and state),

media sessions in progress (such as current stream contributing sources and encoding schemes), the current loudest speaker, the current chair, etc. Conference state is the latest conference snapshot triggered by changes in participants state, conference policy changes, etc.

REQ -1: Conference state virtual database MUST have a modular definition, i.e. it MUST be possible to access different conference aspects independently.

REQ -2: It MUST be possible to aggregate information relating to different conference aspects in a single report.

REQ -3: A mechanism for extensible definition and registration of conference state elements MUST be present.

REQ -4: A default minimal conference state MUST be defined. It SHOULD include a list of current conference participants only.

3.5.2 Dissemination of Changes

REQ -1: A means MUST be defined for reporting the conference state changes to interested parties (including conference participants) in a timely manner.

REQ -2: A means MUST be defined for a SIP user agent to express its interest in selected state changes only.

REQ -3: A means MUST be defined for a SIP user agent to express the minimum interval between receiving state change reports.

REQ -4: It MUST be possible to aggregate recent changes in a single reporting event.

REQ -5: A default minimal conference state changes MUST be defined. They SHOULD include events about participants joining and leaving the conference only.

3.5.3 On-demand Information Dissemination

REQ -1: A means MUST be defined to disseminate any conference state information to interested parties (including SIP user agents) on-demand.

REQ -2: A means MUST be defined for an interested party (including SIP user agents) to request conference state information of a particular conference defined by the conference identifier.

REQ -3: A means MUST be defined for an interested party (including SIP user agents) to specify the subset of the conference state information, it wants and capable to receive.

3.6 Focus Role Migration

Editor s Note: We should decide whether the requirements below can be met by using SIP or non-SIP means.

REQ -1: A procedure for delegating a focus role by the current focus to another participant MUST be defined.

REQ -2: A procedure for requesting a conference focus to transfer its role to another participant MUST be defined.

REQ -3: A procedure for on-demand unconditional transfer of the focus role to a different participant MUST be defined.

REQ -4: A detection procedure for a focus failure condition MUST be defined.

3.7 Side-bar Conferences

A standard means MUST be defined in order to implement the operations defined in this section below.

REQ -1: A user agent (not a conference participant) joins a side-bar within the conference by SIP means.

REQ -2: A user agent (not a conference participant) is invited to a side-bar within the conference by SIP means.

REQ -3: A conference participant creates a side-bar conference with one or more participants in a conference by SIP means.

REQ -4: A conference participant joins a side-bar within the conference by SIP means.

REQ -5: A conference participant is invited to a side-bar within the conference by SIP means.

REQ -6: A conference-unaware user agent (a participant or not) creates and participates in side-bar conferences. It MAY be achieved by non-SIP means.

REQ -7: A conference participant creates side-bar conferences within the conference without establishing any additional SIP dialogs with the focus. It MAY be achieved by non-SIP means.

REQ -8: A conference participant joins any number of side-bars within the conference without establishing any additional SIP dialogs with the focus. It MAY be achieved by non-SIP means.

REQ -9: A conference participant is invited to any number of side-bars within the conference without establishing any additional SIP dialogs with the focus. It MAY be achieved by non-SIP means.

3.8 Cascading of Conferences

"Cascading of Conferences" is a term that has different meanings in different contexts. Some examples are listed below:

- Peer-to-peer chaining of signaling. (Many ways exist to build the media graph in this case.)
- Conferences have hierarchal signaling relations. (Many ways exists to build the media graph in this case.)
- "Cascading" is used to distribute the media "mixing" only. The distribution of signaling is not required.

As it can be seen from the examples, each will define a different set of requirements.

Editor s Note: We need to discuss which of the architectures require our attention as a part of the SIP conferencing force.

3.9 SIMPLE and SIP Conferencing Coordination

REQ -1: SIMPLE-based Presence and Instant Messaging architecture SHOULD fit into the general SIP Conferencing architecture.

REQ -2: A scenario where a multimedia SIP conference and a multiparty IM conversation take place among the same group of participants MUST be addressed.

REQ -3: A scenario where a side-bar or/and a sub-IM-conference is being held as a part of SIP conference MUST be addressed.

4. Security Considerations

Editor s Note: Will be provided in the next version of the document.

5. Contributors

This work is based on the discussions among the members of the SIP Conferencing design team.

Normative References

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

- [2] 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.

Informative References

Authors' Addresses

Orit Levin
RADVISION
266 Harristown Road
Glen Rock, NJ 75024

EMail: orit@radvision.com

Roni Even
Polycom
94 Derech Em Hamoshavot
Petach Tikva, Israel

EMail: roni.even@polycom.co.il

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF
MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the
Internet Society.