

XCON BOF
Internet-Draft
Expires: December 21, 2003

R. Even
Polycom
O. Levin
RADVISION
N. Ismail
Cisco Systems, Inc.
June 22, 2003

**Conferencing media policy requirements
draft-even-xcon-media-policy-requirements- 00.txt**

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/lid-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 December 21, 2003.

Copyright Notice

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

Abstract

This document defines the requirements for Media Policy, i.e. a set of rules associated with the media distribution of the conference. This document presents the requirements for the media manipulations that can be done using these rules by conference participants or third parties using any kind of media/conference policy control protocol. This document does not address the interface between the focus and the media policy.

Table of Contents

1.	Introduction	3
2.	Rational	3
3.	Terminology	3
4.	High Level Architecture	3
5.	Media Policy Data Model	5
5.1	General	5
6.	Media Policy Requirements	6
6.1	General Media Policy Requirements	6
6.2	Video specific requirements	7
7.	Security Considerations	7
	References	8
	Authors' Addresses	8
	Intellectual Property and Copyright Statements	9

1. Introduction

The Conferencing Framework [2] presents an overall framework and defines the terminology for SIP [1] tightly coupled conferencing. The conferencing framework architecture includes the media policy. This is a set of rules that describes the media distribution of a conference. This document presents the requirements for the media policy data model and for the manipulations on these rules by conference participants or third parties using any kind of media/conference policy control protocol. This document does not address the interface between the focus and the media policy and between the focus and the media mixer.

2. Rational

The media policy enables a conference participant or an application server to define and manipulate the content of the media streams going to the conference participants. This will enable applications like sidebars, announcement to specific participants, call centers and panel conferences.

3. Terminology

The draft relies on the terminology defined in the conferencing framework document[2].

4. High Level Architecture

The basic conferencing architecture used in this document is defined in the Conferencing architecture framework [2]. This document focuses on the media policy component and the requirements to manipulate the media policy by authorized entities.

An authorized entity can manipulate the media policy using a supplied application. Examples for such applications include a web application, an interactive voice response application, an interactive Instant Messaging (IM) base application, or an application that uses the media policy control protocol.

The Conference policy control protocol (CPCP) provides a standard way for an automated authorized entity to manipulate the media policy. The requirements and definition of the CPCP protocol are out of scope of this document.

The media policy is a set of rules that describes the media mixing or switching required for each participant in the conference. This includes the set of sources to be mixed or switched and the rules for their mixing or switching. The focus uses the media policy to

determine the proper configuration of the mixers. Authorized entities will be notified of changes to the media policy by subscribing to the conference event package. The information about the current contributing sources to the mixed streams can be learned by the information in the RTP header or by the conference event package [4]. The data structures that include the contributing sources of the current streams is in the focus of the mixer and is not in the scope of the work.

The initial state of the media policy data structure is defined at the conference creation time. It can be either provisioned or created by using a conference policy control protocol or/and other protocols being used to create the conference.

Typically, a focus has access to the media policy and is responsible for translating the media policy data into the actions towards the physical entities ("mixers").

Figure 1 describes an instance of media policy of a conference. The figure shows a single mixer and a single type of stream for ease of drawing but the model does not have such a restriction.




Figure 1: Media Policy in a Conference

5. Media Policy Data Model

5.1 General

The fundamental conferencing functionality is being able to combine (i.e. "to mix") in a media specific manner participants' streams that belong to a logical sub-function within a conference (such as participant's video, left audio stream, right audio stream, video streaming presentation, slide presentation) and are of the same media type (such as video, audio, etc.). In the case of using centralized-mixing the resultant stream(s) will be sent back to the participants. In the case of end-point mixing, the original streams, needed to produce the mixed media, will be distributed to the participants that will perform the actual mixing.

Typically, the maximum number of different mixers in a conference is preconfigured as part of the media conference policy. Mixers MAY be dynamically created and destroyed during the conference lifetime. This document will not describe the data model itself.

6. Media Policy Requirements

All the requirements are based on having a privilege mechanism that authorizes users to access and manipulate the media policy data. The requirements are for the manipulations that can be done on the media distribution by the mixer using the CPCP protocol. The protocol itself is not in the scope of this document. Authorization is also part of CPCP. The requirements are not for the mixer itself.

6.1 General Media Policy Requirements

REQ-GP1: An authorized participant MUST be able to specify its own unique topology.

REQ-GP2: It MUST be possible for a group of users to receive the same mix. This mix may be a conference common mix.

REQ-GP3: It MUST be possible, using the protocol, to dynamically modify the number of contributing streams associated with a mixer.

REQ-GP4: It MUST be possible, using the protocol, to define the mixing function for each participant in the conference.

REQ-GP6: It SHOULD be possible to send a participant multiple streams from one mixer. This requirement is to enable support for end-point mixing.

REQ-GP7: It SHOULD be possible to define relationships between different mixers. The relationships can be time synchronized such as specifying that the audio mixer and video mixer is a pair to establish lip-synchs.

REQ-GP8: It SHOULD be possible to define the number of different topologies and the number of streams in each of them that will be mixed in a mixer. For example the conference will support only one video topology that will go to all the participants, the video topology will support 2x2 display, or each participant will be able to receive his own audio topology that will include up to 4 contributing sources.

REQ-GP9: It SHOULD be possible to have more than one stream from the same type (video or audio) coming from the same user and to mix them separately. For example one video stream will be the video camera

showing the speaker while the other may be a presentation or a second camera showing the whole room.

6.2 Video specific requirements

Video is a bit different than audio when mixing is concerned. In multipoint video the common mixing modes are:

Video switching where one of the contributing sources is sent to all participants, the video source may be forced by the media policy control protocol or may be dynamic by using for example a voice activated video switching mode where the participants will see the loudest speaker.

"Continuous presence" or tiled windows display where the topology is composing one video stream that has a layout defining the shape and position of viewing windows that will be displayed to the participants. The layout includes N viewing windows so that in each of the windows there is one contributing stream. Even though the viewing windows can be of any shape we will address in this work only rectangular windows of any size. The windows may overlap.

The section defines the specific requirements for media policy and media policy control to enable "Continuous presence"

REQ-V1: It should be possible to define rectangular overlapping windows in a video mix.

REQ-V2: It should be possible to map a stream to a window based on some mode like having one window display the loudest speaker or the floor holder while for the remaining windows fixed input streams are used.

REQ-V3: It should be possible for authorized participants to change the layout of the video topology.

REQ-V4: It should be possible for authorized participants to define the mapping of a stream to a window.

7. Security Considerations

The media policy control protocol may enable unauthorized users to manipulate the media mixing of conferences, this may enable them to listen to conference or eject unsolicited media streams. The protocol should provide authentication of the users. The media policy data may include information about the sources and targets of mixer, if this information will be transferred in the protocol in the clear that may cause a security risk. The protocol should allow for

encryption of the media policy transferred in the media policy control protocol.

References

- [1] 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.
- [2] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol", [draft-rosenberg-sipping-conferencing-framework-01](#) (work in progress), February 2003.
- [3] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [4] Rosenberg, J. and H. Schulzrinne, "A session initiation protocol (SIP) event package for conference state", , June 2002.

Authors' Addresses

Roni Even
Polycom
94 Derech Em Hamoshavot
Petach Tikva 49130
Israel

EMail: roni.even@polycom.co.il

Orit Levin
RADVISION
266 Harristown Road
Glen Rock
NJ USA

EMail: orit@radvision.com

Nermeen Ismail
Cisco Systems, Inc.
170 West Tasman Drive
San Jose 95134
CA USA

EMail: nismail@cisco.com

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.