

Mediactrl
Internet-Draft
Intended status: Informational
Expires: August 27, 2008

M. Dolly
AT&T Labs
R. Even
Polycom
February 24, 2008

Media Server Control Protocol Requirements
draft-ietf-mediactrl-requirements-04.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 August 27, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document addresses the communication between an application server and media server. The current work in IETF working groups shows these logical entities but does not address the physical decomposition and the protocol between the entities.

This document presents the requirements for a media server control protocol (MCP) that enables an application server to use a media

Internet-Draft

MCP Requirements

February 2008

server. It will address the aspects of announcements, Interactive Voice Response and conferencing media services.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Requirements	4
3.1.	Media Control Requirements	4
3.2.	Media mixing Requirements	6
3.3.	IVR Requirements	7
3.4.	Operational Requirements	7
4.	IANA consideration	7
5.	Security Considerations	7
6.	Acknowledgment	8
7.	Informative References	8
	Authors' Addresses	8
	Intellectual Property and Copyright Statements	10

Internet-Draft

MCP Requirements

February 2008

1. Introduction

The IETF conferencing framework in [RFC4353](#) [[CARCH](#)] presents an architecture that is built of several functional entities. [RFC4353](#) [[CARCH](#)] does not specify the protocols between the functional entities since it is considered out of scope.

Based on [RFC4353](#) [[CARCH](#)] the document defines the requirements for a protocol that will enable one functional entity, known as an Application Server (AS), that includes the conference/media policy server, the notification server and the focus, all defined in [RFC 4353](#) [[CARCH](#)], to interact with one or more functional entities, called Media Server (MS), that serves as mixer or media server.

The Media server can also be used for announcements and Interactive Voice Response (IVR) functions.

Application Servers host one or more instances of a communications application. Media servers provide real time media processing functions. An example of the decomposition of a media server and an application server is described in the media control framework document [[mediactrl-fw](#)].

This document presents the requirements for a media server control protocol (MCP) that enables an application server to control a media server. It will address the aspects of announcements, IVR and conferencing media services.

The requirements are for the protocol and do not address the AS or MS functionality discussed in the media control framework.

Since the media server is a centralized component, the charter of the working group states that this work will not investigate distributed media processing algorithms or control protocols.

[2.](#) Terminology

The Media Server work uses, when appropriate, and expands on the terminology introduced in the conferencing framework[CARCH] and Centralized Conferencing (XCON) conferencing framework[xcon-framework]. The following additional terms are defined:

Application Server (AS) - A functional entity that hosts one or more instances of a communications application. The application server may include the conference policy server, the focus and the conference notification server as defined in [[CARCH](#)]. It may include

also communication applications that use IVR or announcements services.

Media Server (MS) - The media server includes the mixer as defined in [[CARCH](#)]. The media server plays announcements, it processes media streams for functions like DTMF detection and transcoding. The media server may also record media streams for supporting IVR functions like announcing participants

Media Resource Broker (MRB) - A logical entity that is responsible for both collection of appropriate published Media Server (MS) information and supplying of appropriate MS information to consuming entities. The MRB is an optional entity and will be discussed in a separate document.

Notification - A notification is used when there is a need to report event related information from the MS to the AS.

Request - A request is sent from the controlling entity, such as an Application Server, to another resource, such as a Media Server, asking that a particular type of operation be executed.

Response - A response is used to signal information such as an acknowledgement or error code in reply to a previously issued request.

[3.](#) Requirements

3.1. Media Control Requirements

The following are the media control requirements:

REQ-MCP-01 - The MS Control Protocol shall enable one or more Application Servers to request media services from one or more Media Servers.

REQ-MCP-02 The MS Control Protocol shall use a reliable transport protocol.

REQ-MCP-03 - The applications supported by the protocol shall include Conferencing and Interactive Voice Response media services.

Note: Though the protocol enables these services, the functionality is invoked through other mechanisms.

REQ-MCP-04 - Media types supported in the context of the applications shall include audio, tones, text and video. Tones media include in band audio or [RFC 4733](#) payload.

REQ-MCP-05- The MS control protocol should allow, but must not require, a media resource broker (MRB) or intermediate proxy to exist with the Application Server and Media Server.

REQ-MCP-06 - On the MS control channel, there shall be requests to the MS, responses from the MS and notifications to the AS.

REQ-MCP-07 - SIP/SDP shall be used to establish and modify media connections to a Media Server.

REQ-MCP-08 - It should be possible to support a single conference spanning multiple Media Servers.

Note: It is probable that spanning multiple MS can be accomplished by the AS and does not require anything in the protocol for the scenarios we have in mind. However, the concern is that if this requirement is treated too lightly, one may end up with a protocol that precludes its support.

REQ-MCP-09 - It must be possible to split call legs individually or in groups away from a main conference on a given Media Server, without performing re-establishment of the call legs to the MS (e.g., for purposes such as performing IVR with a single call leg or creating sub-conferences, not for creating entirely new conferences).

REQ-MCP-10 - The MS control protocol should be extendable, facilitating forward and backward compatibility.

REQ-MCP-11 - The MS control protocol shall include an authentication component to ensure that only an authorized AS can communicate with the MS and vice versa.

REQ-MCP-12 - The MS control protocol shall use some form of transport protection to ensure the confidentiality and integrity of the data between the AS and MS.

REQ-MCP-13 - Different Application Servers may have different privileges for using a MS. The protocol should prevent the AS for doing unauthorized operations on a MS.

REQ-MCP-14 - The MS control protocol requires mechanisms to protect the MS resources used by one AS from another AS since the solution need to support multiple AS controlling one MS.

REQ-MCP-15 - During session establishment, there shall be a capability to negotiate parameters that are associated with media streams. This requirement should enable also an AS managing conference to specify the media streams allowed in the conference.

REQ-MCP-16 - The AS shall be able to instruct the MS to perform streams operations like mute and gain control.

REQ-MCP-17 - The AS shall be able to instruct the MS to play a specific announcement.

REQ-MCP-18 - The AS shall be able to request the MS to create, delete, and manipulate a mixing, IVR or announcement session.

REQ-MCP-19 - The AS shall be able to instruct the MS to play announcements to a single user or to a conference mix.

REQ-MCP-20 - The MS control protocol should enable the AS to ask the MS for session summary report. The report may include resources usage and quality metrics.

REQ-MCP-21 - The MS shall be able to notify the AS of events received in the media stream if requested by AS. (Examples - STUN request, Flow Control, etc.)

[3.2.](#) Media mixing Requirements

REQ-MCP-22 - The AS shall be able to define a conference mix, MS may offer different mixing topologies. The conference mix may be defined on a conference or user level.

REQ-MCP-23 - The AS may be able to define a custom video layout built of rectangular sub windows.

REQ-MCP-24 - For video the AS shall be able to map a stream to a specific sub-window or to define to the MS how to decide which stream will go to each sub window.

REQ-MCP-25 - The MS shall be able to notify the AS who are the active sources of the media; for example who is the active speaker or who is being viewed in a conference. The speaker and the video source may be different, for example a person describing a video stream from a remote camera managed by a different user.

REQ-MCP-26 - The MS shall be able to inform the AS which layouts it supports.

REQ-MCP-27 - The MS control protocol should enable the AS to instruct the MS to record a specific conference mix.

[3.3.](#) IVR Requirements

REQ-MCP-28 - The AS shall be able to instruct the MS to perform one or more IVR script and receive the results. The script may be in a server or contained in the control message.

REQ-MCP-29 - The AS shall be able to manage the IVR session by sending requests to play announcements to the MS and receiving the response (e.g., DTMF). The IVR session flow in this case is handled by the AS by starting a next phase based on the response it receives from the MS on the current phase.

REQ-MCP-30 - The AS should be able to instruct the MS to record a short participant stream and play it back. This is not a recording requirement.

[3.4.](#) Operational Requirements

These requirements may be applicable to the MRB but can be used by an AS if it has one to one connection to the MS.

REQ-MCP-31 - The MS control protocol must allow the AS to audit the MS state, during an active session.

REQ-MCP-32 - The MS shall be able to inform the AS about its status during an active session.

[4.](#) IANA consideration

There are no IANA considerations.

[5.](#) Security Considerations

This document discusses high-level requirements for MCP. The MCP has some specific security requirements, which will be summarized here at a very high level.

All of the operations and functions described in this document need to be authorized by a MS or a AS. It is expected that MS resources will be governed by a set of authorization rules defined as part of

the AS / MS policy. In order for the policy to be implemented, the

MS needs to be able to authenticate requests. Normal SIP mechanisms including Digest authentication and certificates can be used as specified in [RFC3261](#)[\[RFC3261\]](#) These MCP security requirements will be discussed in detail in the framework and protocol documents.

[6.](#) Acknowledgment

This draft represents the work from two previous personal drafts, [draft-dolly-xcon-mediactrlframe-02](#) and [draft-even-media-server-req-02](#). The authors would like to acknowledge the work of Gary Munson from AT &T Labs and James Rafferty from Cantata who helped with drafting [draft-dolly-xcon-mediactrlframe-02](#) on which this work is based.

[7.](#) Informative References

- [CARCH] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", [RFC 4353](#), February 2006.
- [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.
- [mediactrl-fw] Melanchuk, T., "An Architectural Framework for Media Server Control", [draft-ietf-mediactrl-architecture-02](#) (work in progress), February 2008.
- [xcon-framework] Barnes, M., "A Framework for Centralized Conferencing", [draft-ietf-xcon-framework-10](#) (work in progress), November 2007.

Authors' Addresses

Martin Dolly
AT&T Labs
200 Laurel Avenue
Middletown, NJ 07748
USA

Phone:
Email: mdolly@att.com
URI:

Roni Even
Polycom
94 Derech Em Hamoshavot
Petach Tikva 49130
Israel

Email: roni.even@polycom.co.il

Internet-Draft

MCP Requirements

February 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

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

ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

Dolly & Even

Expires August 27, 2008

[Page 10]