

[draft-camarillo-sipping-transc-framework-00.txt](#)

August 28, 2003

Expires: February, 2004

Framework for Transcoding with the Session Initiation Protocol

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>

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

Abstract

This document defines a framework for transcoding with SIP. This framework includes how to discover the need of transcoding services in a session and how to invoke those transcoding services. Two models for transcoding services invocation are discussed; the conference bridge model and the third party call control model. Both models meet the requirements for SIP regarding transcoding services invocation to support deaf, hard of hearing and speech-impaired individuals.

Table of Contents

<u>1</u>	Introduction	<u>3</u>
<u>2</u>	Discovery of the Need for Transcoding Services	<u>3</u>
<u>3</u>	Transcoding Services Invocation	<u>4</u>
<u>3.1</u>	Third Party Call Control Transcoding Model	<u>5</u>
<u>3.2</u>	Conference Bridge Transcoding Model	<u>5</u>
<u>4</u>	Security Considerations	<u>8</u>
<u>5</u>	Contributors	<u>8</u>
<u>6</u>	Authors' Addresses	<u>8</u>
<u>7</u>	Bibliography	<u>8</u>

1 Introduction

Two user agents involved in a SIP [[1](#)] dialog may find it impossible to establish a media session due to a variety of incompatibilities. Assuming that both user agents understand the same session description format (e.g., SDP), incompatibilities can be found at the user agent level and at the user level. At the user agent level, both terminals may not support any common codec or may not support common media types (e.g., a text-only terminal and an audio-only terminal). At the user level, a deaf person will not be able to understand what it is said over an audio stream.

In order to make communications possible in the presence of incompatibilities, user agents need to introduce intermediaries that provide transcoding services to a session. From the SIP point of view, the introduction of a transcoder is done in the same way to resolve both user level and user agent level incompatibilities. Therefore, the invocation mechanisms described in this document are generally applicable to any type of incompatibility related to how the information that needs to be communicated is encoded.

Furthermore, although this framework focuses on transcoding, the mechanisms described are applicable to media manipulation in general. It would be possible to use them, for example, to invoke a server that simply increased the volume of an audio stream.

This document does not describe media server discovery. That is an orthogonal problem that one can address using user agent provisioning or other methods.

The remainder of this document is organized as follows. [Section 2](#) deals with the discovery of the need of transcoding services for a particular session.

[Section 3.2](#) introduces the conference bridge transcoding invocation model, and [Section 3.1](#) introduces the third party call control model. Both models meet the requirements regarding transcoding services invocation in [RFC3351](#) [[2](#)] to support deaf, hard of hearing and speech-impaired individuals.

2 Discovery of the Need for Transcoding Services

According to the one-party consent model defined in [RFC 3238](#) [[3](#)], services that involve media manipulation invocation are best invoked by one of the end-points involved in the communication, as opposed to being invoked by an intermediary in the network. Following this principle, one of the end-points should be the one detecting that

transcoding is needed for a particular session.

In order to decide whether or not transcoding is needed, a user agent needs to know the capabilities of the remote user agent. A user agent acting as an offerer typically obtains this knowledge by downloading a presence document that includes media capabilities (e.g., Bob is available on a terminal that only supports audio) or by getting an SDP description of media capabilities as defined in [RFC 3264](#) [4]. Presence documents are typically received in a NOTIFY request as a result of a subscription. SDP media capabilities descriptions are typically received in a 200 (OK) response to an OPTIONS request or in a 488 (Not Acceptable Here) response to an INVITE.

It is recommended that an offerer does not invoke transcoding services before making sure that the answerer does not support the capabilities needed for the session. Making wrong assumptions about the answerer's capabilities can lead to situations where two transcoders are introduced (one by the offerer and one by the answerer) in a session that would not need any transcoding services at all.

An example of the situation above is a call between two GSM phones (without using transcoding-free operation). Both phones use a GSM codec, but the speech is converted from GSM to PCM by the originating MSC and from PCM back to GSM by the terminating MSC.

Note that transcoding services can be symmetric (e.g., speech-to-text plus text-to-speech) or asymmetric (e.g., a one-way speech-to-text transcoding for a hearing impaired user that can talk).

3 Transcoding Services Invocation

Once the need for transcoding for a particular session has been identified as described in [Section 2](#), one of the user agents needs to invoke transcoding services.

As stated previously, transcoder location is outside the scope of this document. Therefore, we assume that the user agent invoking transcoding services knows the URI of a server that provides them.

Invoking transcoding services from a server (T) for a session between two user agents (A and B) involves establishing two media sessions; one between A and T and another between T and B. How to invoke T's services (i.e., how to establish both A-T and T-B sessions) depends on how we model the transcoding service. We have considered two models for invoking a transcoding service. The first is to use third party call control [5], also referred to as 3pcc. The second is to

use a (dial-in and possibly dial-out) conference bridge that negotiates the appropriate media parameters on each individual leg (i.e., A-T and T-B).

[Section 3.1](#) analyzes the applicability of the third party call control model and [Section 3.2](#) analyzes the applicability of the conference bridge transcoding invocation model.

[3.1](#) Third Party Call Control Transcoding Model

In the 3pcc transcoding model, defined in ([draft-camarillo-sipping-transc-3pcc](#)), the user agent invoking the transcoding service has a signalling relationship with the transcoder and another signalling relationship with the remote user agent. There is no signalling relationship between the transcoder and the remote user agent, as shown in Figure 1.

This model is suitable for advanced end points that are able to perform third party call control. It allows end-points to invoke transcoding services on a stream basis. That is, the media streams that need transcoding are routed through the transcoder while the streams that do not need it are sent directly between the end points. This model also allows to invoke one transcoder for the sending direction and a different one for the receiving direction of the same stream.

Invoking a transcoder in the middle of an ongoing session is also quite simple. This is useful when session changes occur (e.g., an audio session is upgraded to an audio/video session) and the end-points cannot cope with the changes (e.g., they had common audio codecs but no common video codecs).

The privacy level that is achieved using 3pcc is high, since the transcoder does not see the signalling between both end-points. In this model, the transcoder only has access to the information that is strictly needed to perform its function.

[3.2](#) Conference Bridge Transcoding Model

In a centralized conference, there are a number of media streams between the conference server and each participant of a conference. For a given media type (e.g., audio) the conference server sends over each individual stream the media received over the rest of the streams, typically performing some mixing. If the capabilities of all the end-points participating in the conference are not the same, the conference server may have to send audio to different participants using different audio codecs.

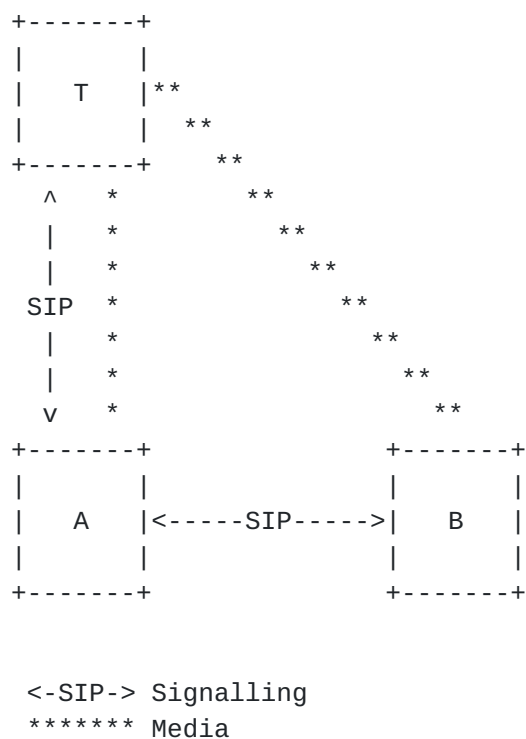


Figure 1: Third Party Call Control Model

Consequently, we can model a transcoding service as a two-party conference server that may change not only the codec in use, but also the format of the media (e.g., audio to text).

Using this model, T behaves as a B2BUA and the whole A-T-B session is established as described in ([draft-camarillo-sipping-transc-b2bua](#)). Figure 2 shows the signalling relationships between the end-points and the transcoder.

In the conferencing bridge model, the end-point invoking the transcoder is generally involved in less signalling exchanges than in the 3pcc model. This may be an important feature for end-point using

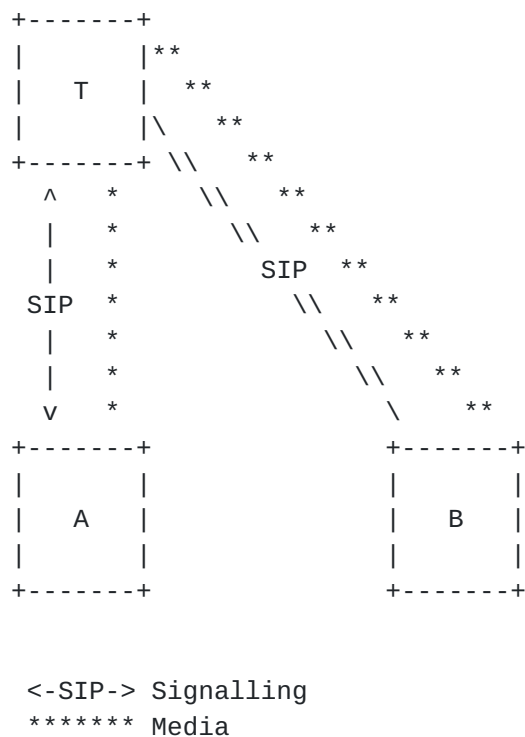


Figure 2: Conference Bridge Control Model

low bandwidth or high-delay access links (e.g., some wireless accesses).

However, this model is less flexible than the 3pcc model. It is not possible to use different transcoders for different streams or for different directions of a stream.

Invoking a transcoder in the middle of an ongoing session or changing from one transcoder to another requires the remote end-point to support the Replaces [6] extension. At present, not many user agents support it.

Simple end-points that cannot perform 3pcc and thus cannot use the 3pcc model, of course, need to use the conference bridge model.

4 Security Considerations

This document does not introduce any new security considerations.

5 Contributors

This document is the result of discussions amongst the conferencing design team. The members of this team include Eric Burger, Henning Schulzrinne and Arnoud van Wijk.

6 Authors' Addresses

Gonzalo Camarillo
Ericsson
Advanced Signalling Research Lab.
FIN-02420 Jorvas
Finland
electronic mail: Gonzalo.Camarillo@ericsson.com

7 Bibliography

- [1] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. R. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: session initiation protocol," [RFC 3261](#), Internet Engineering Task Force, June 2002.
- [2] N. Charlton, M. Gasson, G. Gybels, M. Spanner, and A. van Wijk, "User requirements for the session initiation protocol (SIP) in support of deaf, hard of hearing and speech-impaired individuals," [RFC 3351](#), Internet Engineering Task Force, Aug. 2002.
- [3] S. Floyd and L. Daigle, "IAB architectural and policy considerations for open pluggable edge services," [RFC 3238](#), Internet Engineering Task Force, Jan. 2002.
- [4] J. Rosenberg and H. Schulzrinne, "An offer/answer model with session description protocol (SDP)," [RFC 3264](#), Internet Engineering Task Force, June 2002.
- [5] J. Rosenberg, J. L. Peterson, H. Schulzrinne, and G. Camarillo, "Best current practices for third party call control in the session initiation protocol," internet draft, Internet Engineering Task Force, July 2003. Work in progress.
- [6] B. Biggs, R. W. Dean, and R. Mahy, "The session initiation protocol (SIP) Engineering Task Force, Aug. 2003. Work in progress.

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 assigns.

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.

