

The Session Initiation Protocol (SIP)
Conference Bridge Transcoding Model

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document describes how to invoke transcoding services using the conference bridge model. This way of invocation meets the requirements for SIP regarding transcoding services invocation to support deaf, hard of hearing, and speech-impaired individuals.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Caller's Invocation	3
3.1. Procedures at the User Agent	3
3.2. Procedures at the Transcoder	3
3.3. Example	4
3.4. Unsuccessful Session Establishment	6
4. Callee's Invocation	7
5. Security Considerations	7
6. Contributors	8
7. References	8
7.1. Normative References	8
7.2. Informative References	9

1. Introduction

RFC 5369 [RFC5369] describes how two SIP [RFC3261] UAs (User Agents) can discover incompatibilities that prevent them from establishing a session (e.g., lack of support for a common codec or for a common media type). When such incompatibilities are found, the UAs need to invoke transcoding services to successfully establish the session. The transcoding framework introduces two models to invoke transcoding services: the 3pcc (third-party call control) model [RFC4117] and the conference bridge model. This document specifies the conference bridge model.

In the conference bridge model for transcoding invocation, a transcoding server that provides a particular transcoding service (e.g., speech-to-text) behaves as a B2BUA (Back-to-Back User Agent) between both UAs and is identified by a URI. As shown in Figure 1, both UAs, A and B, exchange signalling and media with the transcoder T. The UAs do not exchange any traffic (signalling or media) directly between them.

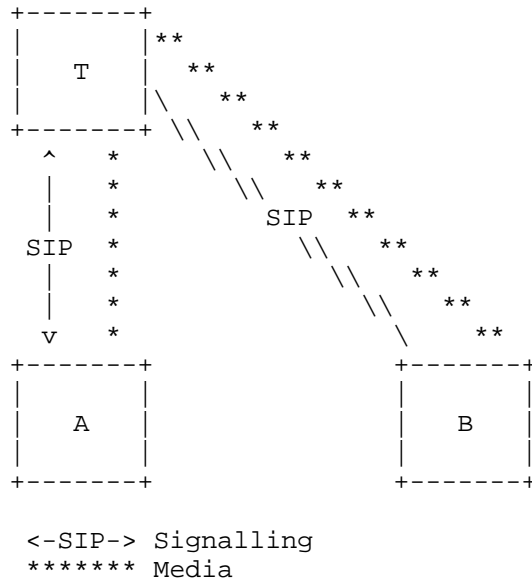


Figure 1: Conference bridge model

Sections 3 and 4 specify how the caller A or the callee B, respectively, can use the conference bridge model to invoke transcoding services from T.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119], and indicate requirement levels for compliant implementations.

3. Caller's Invocation

User agent A needs to perform two operations to invoke transcoding services from T for a session between user agent A and user agent B. User agent A needs to establish a session with T and provide T with user agent B's URI so that T can generate an INVITE towards user agent B.

3.1. Procedures at the User Agent

User agent A uses the procedures for RFC 5366 [RFC5366] to provide T with B's URI using the same INVITE that establishes the session between A and T. That is, user agent A adds to the INVITE a body part whose disposition type is recipient-list [RFC5363]. This body part consists of a URI-list that contains a single URI: user agent B's URI.

Note that, as described in the transcoding framework [RFC5369], the transcoding model described in this document is modeled as a two-party conference server. Consequently, this document focuses on two-party sessions that need transcoding. Multi-party sessions can be established using INVITE requests with multiple URIs in their bodies, as specified in [RFC5366].

3.2. Procedures at the Transcoder

On receiving an INVITE with a URI-list body, the transcoder follows the procedures in [RFC5366] to generate an INVITE request towards the URI contained in the URI-list body. Note that the transcoder acts as a B2BUA, not as a proxy.

Additionally, the transcoder MUST generate the From header field of the outgoing INVITE request using the same value as the From header field included in the incoming INVITE request, subject to the privacy requirements (see [RFC3323] and [RFC3325]) expressed in the incoming INVITE request. Note that this does not apply to the "tag" parameter.

The session description the transcoder includes in the outgoing INVITE request depends on the type of transcoding service that particular transcoder provides. For example, a transcoder resolving audio codec incompatibilities would generate a session description listing the audio codecs the transcoder supports.

When the transcoder receives a final response for the outgoing INVITE requests, it generates a new final response for the incoming INVITE request. This new final response SHOULD have the same status code as the one received in the response for the outgoing INVITE request.

If a transcoder receives an INVITE request with a URI-list with more than one URI, it SHOULD return a 488 (Max 1 URI allowed in URI-list) response.

3.3. Example

Figure 2 shows the message flow for the caller's invocation of a transcoder T. The caller A sends an INVITE (1) to the transcoder (T) to establish the session A-T. Following the procedures in [RFC5366], the caller A adds a body part whose disposition type is recipient-list [RFC5363].

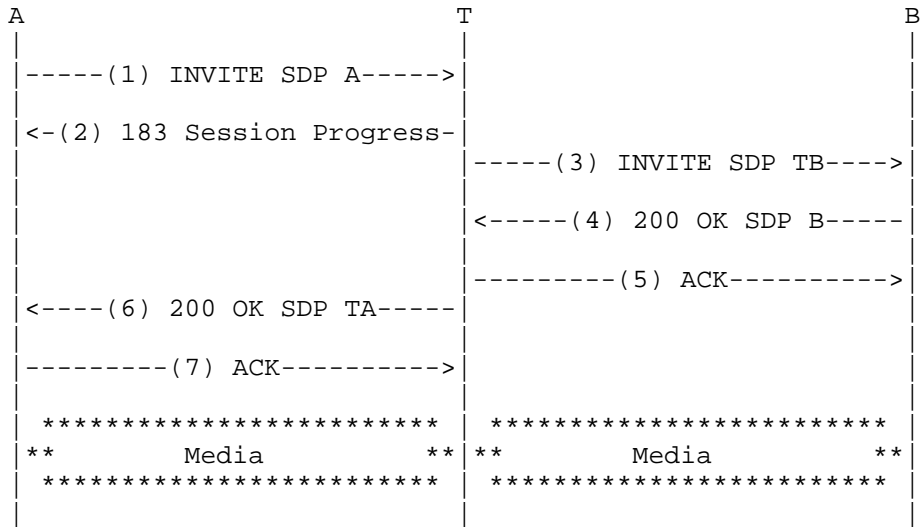


Figure 2: Successful invocation of a transcoder by the caller

The following example shows an INVITE with two body parts: an SDP [RFC4566] session description and a URI-list.

```

INVITE sip:transcoder@example.com SIP/2.0
Via: SIP/2.0/TCP client.chicago.example.com
    ;branch=z9hG4bKhjhs8ass83
Max-Forwards: 70
To: Transcoder <sip:transcoder@example.org>
From: A <sip:A@chicago.example.com>;tag=32331
Call-ID: d432fa84b4c76e66710
CSeq: 1 INVITE
Contact: <sip:A@client.chicago.example.com>
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER,
    SUBSCRIBE, NOTIFY
Allow-Events: dialog
Accept: application/sdp, message/sipfrag
Require: recipient-list-invite
Content-Type: multipart/mixed;boundary="boundary1"
Content-Length: 556

--boundary1
Content-Type: application/sdp

v=0
o=example 2890844526 2890842807 IN IP4 chicago.example.com
s=-
c=IN IP4 192.0.2.1
t=0 0
m=audio 50000 RTP/AVP 0
a=rtpmap:0 PCMU/8000

--boundary1
Content-Type: application/resource-lists+xml
Content-Disposition: recipient-list

<?xml version="1.0" encoding="UTF-8"?>
<resource-lists xmlns="urn:ietf:params:xml:ns:resource-lists"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <list>
    <entry uri="sip:B@example.org" />
  </list>
</resource-lists>
--boundary1--

```

On receiving the INVITE, the transcoder generates a new INVITE towards the callee. The transcoder acts as a B2BUA, not as a proxy. Therefore, this new INVITE (3) belongs to a different transaction than the INVITE (1) received by the transcoder.

When the transcoder receives a final response (4) from the callee, it generates a new final response (6) for INVITE (1). This new final response (6) has the same status code as the one received in the response from the callee (4).

3.4. Unsuccessful Session Establishment

Figure 3 shows a similar message flow as the one in Figure 3. Nevertheless, this time the callee generates a non-2xx final response (4). Consequently, the transcoder generates a non-2xx final response (6) towards the caller as well.

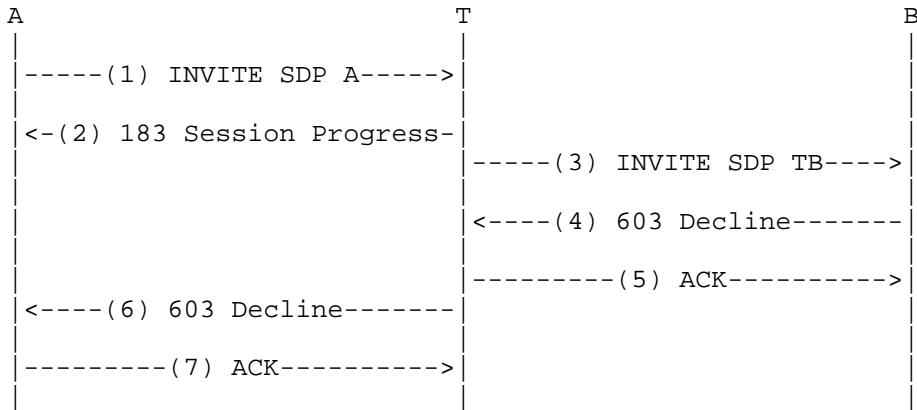


Figure 3: Unsuccessful session establishment

The ambiguity in this flow is that, if the provisional response (2) gets lost, the caller does not know whether the 603 (Decline) response means that the initial INVITE (1) was rejected by the transcoder or that the INVITE generated by the transcoder (4) was rejected by the callee. The use of the "History-Info" header field [RFC4244] between the transcoder and the caller resolves the previous ambiguity.

Note that this ambiguity problem could also have been resolved by having transcoders act as a pure conference bridge. The transcoder would respond with a 200 (OK) to the INVITE request from the caller, and it would generate an outgoing INVITE request towards the callee. The caller would get information about the result of the latter INVITE request by subscribing to the conference event package [RFC4575] at the transcoder. Although this flow would have resolved the ambiguity problem without requiring support for the "History-Info" header field, it is more complex, requires a higher number of messages, and introduces higher session setup delays. That is why it was not chosen to implement transcoding services.

4. Callee's Invocation

If a UA receives an INVITE with a session description that is not acceptable, it can redirect it to the transcoder by using a 302 (Moved Temporarily) response. The Contact header field of the 302 (Moved Temporarily) response contains the URI of the transcoder plus a "?body=" parameter. This parameter contains a recipient-list body with B's URI. Note that some escaping (e.g., for Carriage Returns and Line Feeds) is needed to encode a recipient-list body in such a parameter. Figure 4 shows the message flow for this scenario.

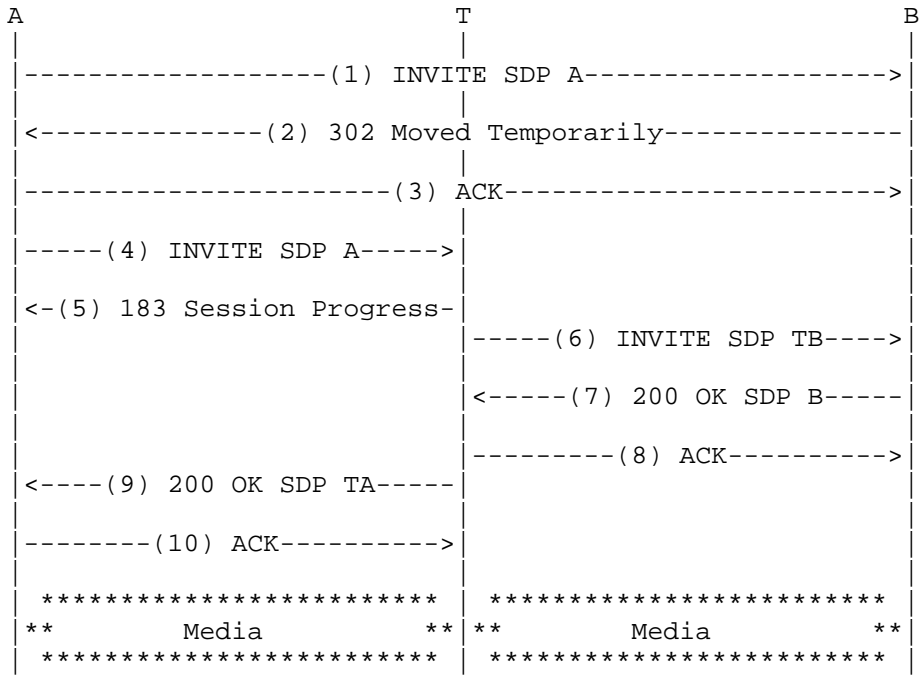


Figure 4: Callee's invocation of a transcoder

Note that the syntax resulting from encoding a body into a URI as described earlier is quite complex. It is actually simpler for callees to invoke transcoding services using the 3pcc transcoding model [RFC4117] instead.

5. Security Considerations

Transcoders implementing this specification behave as a URI-list service as described in [RFC5366]. Therefore, the security considerations for URI-list services discussed in [RFC5363] apply here as well.

In particular, the requirements related to list integrity and unsolicited requests are important for transcoding services. User agents SHOULD integrity protect URI-lists using mechanisms such as S/MIME [RFC3850] or TLS [RFC5246], which can also provide URI-list confidentiality if needed. Additionally, transcoders MUST authenticate and authorize users and MAY provide information about the identity of the original sender of the request in their outgoing requests by using the SIP identity mechanism [RFC4474].

The requirement in [RFC5363] to use opt-in lists (e.g., using RFC 5360 [RFC5360]) deserves special discussion. The type of URI-list service implemented by transcoders following this specification does not produce amplification (only one INVITE request is generated by the transcoder on receiving an INVITE request from a user agent) and does not involve a translation to a URI that may be otherwise unknown to the caller (the caller places the callee's URI in the body of its initial INVITE request). Additionally, the identity of the caller is present in the INVITE request generated by the transcoder. Therefore, there is no requirement for transcoders implementing this specification to use opt-in lists.

6. 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.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [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.
- [RFC3323] Peterson, J., "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, November 2002.

- [RFC3325] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", RFC 3325, November 2002.
- [RFC3850] Ramsdell, B., Ed., "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Certificate Handling", RFC 3850, July 2004.
- [RFC4117] Camarillo, G., Burger, E., Schulzrinne, H., and A. van Wijk, "Transcoding Services Invocation in the Session Initiation Protocol (SIP) Using Third Party Call Control (3pcc)", RFC 4117, June 2005.
- [RFC5369] Camarillo, G., "Framework for Transcoding with the Session Initiation Protocol", RFC 5369, October 2008.
- [RFC5363] Camarillo, G. and A.B. Roach, "Framework and Security Considerations for Session Initiation Protocol (SIP) URI-List Services", RFC 5363, October 2008.
- [RFC5366] Camarillo, G. and A. Johnston, "Conference Establishment Using Request-Contained Lists in the Session Initiation Protocol (SIP)", RFC 5366, October 2008.
- [RFC4244] Barnes, M., Ed., "An Extension to the Session Initiation Protocol (SIP) for Request History Information", RFC 4244, November 2005.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

7.2. Informative References

- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, Ed., "A Session Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, August 2006.
- [RFC5360] Rosenberg, J., "A Framework for Consent-Based Communications in the Session Initiation Protocol (SIP)", RFC 5360, October 2008.

Author's Address

Gonzalo Camarillo
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

E-Mail: Gonzalo.Camarillo@ericsson.com

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.

