Internet Engineering Task Force
Internet Draft
Document: draft-shankar-sip-forked-media-00.txt
Expires: Aug 2001

Mahesh Shankar
Cisco Systems
Feb 2001

SIP Forked Media

Status of this Memo

1. Abstract

This document provides guidelines and examples for initiating "forked" media sessions using multiple media description headers in the session description of a SIP message. The presence of multiple media and address description headers in the SDP implies that multiple media streams are opened for the session in question based on the media description headers.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [2].

3. The use of media description headers in forking media streams

In some application scenarios, a user contacts an application, places a new call in the context of the application, and returns to the application after the new call is finished. Some examples of such scenarios are: Calling card systems, Voicemail or Messaging systems which allows outgoing calls, and Voice Browsers or Voice Portals which allow outgoing calls.

Many of these applications employ speech recognition technology to listen for "hotwords" or other audio input from the Originating User, while he or she converses with the Target. For brevity, we will call these application servers "Voice Portals". In addition to receiving media, these applications may also send media in the form of announcements, tones, etc. which should be mixed with media sent by the Target before the media is rendered for the Originating User.

Note that while both the Voice Portal and the Target communicate with the Originating User simultaneously, this media relationship is not a conference. The Voice Portal does not receive any media sent by the Target, nor does the Target receive any media sent by the Voice Portal.

The Originating UA may choose to send or "fork" its own media to both peers, and locally mix the two forks of the received session locally. This implementation feature of the Originating UA is therefore called "media forking". The author notes that this media relationship happens to be similar to the mixing relationship required for multicast conferences; as a result this forking arrangement is sometimes alternatively referred to as multi-unicast.

The Session Description Protocol [3] defines the set of messages that would enable transfer of media through media channels. As such, each session has global attributes that would dictate the attributes of the global session, and apply to the session in general. The SDP message also defines attributes that would apply to a media session in particular.

Appendix B of RFC2543 (SIP) [4] implies that multiple instances of media description headers in the SDP message should result in either the opening of multiple media channels to the same location, or it should establish a forked media session.

Also, in the case of an end point initiating the forking media session, it should be capable of modifying the SDP message to add itself to the SDP message and moving the global connection information ("c= ") line, to be media specific.

4. Usage Examples

The following diagram illustrates the setup:

```
         ----       +--------+   SIP    +---------+
        /    \      |        |<-------->|         |
       |Origin+-----+   GW   |   RTP    | Voyager |
       |  UA  |     |        |<========>|         |
        \    /      +--------+          +---------+
         ----            ^^                   ^
                         ||                   |
                         ||                   | SIP
                         ||                   |
                         ||                   v
                         ||             +---------+
                         ||      RTP    |         |
                          =============>|  Target |
                                        |         |
                                        +---------+
```

When A places a call to Voice Portal, the call goes through gateway GW, and the ensuing call-flows illustrate this.

```
        Origin                      Portal                       Target
         |                             |                             |
         |          INVITE       A1    |                             |
         |---------------------------->|                             |
         |                             |                             |
         |       180 Ringing     A2    |                             |
         |<----------------------------|                             |
         |                             |                             |
         |          200 OK       A3    |                             |
         |<----------------------------|                             |
         |                             |                             |
         |            ACK        A4    |                             |
```

```
|---------------------------->|                             |
|                             |                             |
|         2-way media   A5    |                             |
|<===========================>|                             |
|                             |                             |
```


The Voice Portal initiates the forking media session.

```
|                             |   INVITE without SDP     B1 |
|                             |---------------------------->|
|                             |                             |
|                             |                             |
|                             |          200 OK          B2 |
|                             |<----------------------------|
```

The Voice Portal modifies the SDP message, and sends it as part of the INVITE message to the GW.

```
|                             |                             |
|                             |                             |
|  INVITE with media headers  |                             |
| from both Voice Portal and  |                             |
| Target                   C1 |                             |
|<----------------------------+                             |
|                             |                             |
| 200 OK with multiple media  |                             |
| headers for Voice Portal and|                             |
| Target                   C2 |                             |
|---------------------------->|                             |
|  ACK                     C3 |                             |
|<----------------------------+                             |
```

The Voice Portal modifies the SDP message to remove the media header that pertains to itÆs media channel, and sends the modified SDP message as part of the ACK for the initial INVITE without SDP.

```
|                             |                             |
|                             | ACK with single media header|
|                             | meant for Target.        B3 |
|                             +---------------------------->|
```

```
      |                              |                              |
      |   2-way media                |                              |
      |<==========================================================>|
      |                              |                              |
      |                              |                              |
      |                              |                              |
      |                              |                              |
      |                              |                              |
      |                              |                              |
```

The pertinent messages that are involved in the call-flow described above are shown below. Headers, SDP lines, tags, and other elements not directly applicable to the discusion are ommitted for brevity.

Messages A1-A5 and B1-B3 are standard [SIP] messages that are involved in any call setup scenario. Following are the relevant parts of the messages.

B1 INVITE without SDP Voice Portal -> Target

```
INVITE sip:Target@voice.com SIP/2.0
Via: SIP/2.0/UDP voice.com:5060
From: VoicePortal <sip:VoicePortal@voice.com>
To: Target <sip:Target@voice.com>
...
```

B2 200 OK Target -> Voice Portal

```
SIP/2.0 200 OK
...
From: VoicePortal <sip:VoicePortal@voice.com>
To: Target <sip:Target@voice.com>
...
Content-Type: application/sdp
Content-Length: ....

...
```

```
c=IN IP4 10.0.0.3
...
m=audio 16980 rtp/avt 0 101
a=rtpmap:0 PCMU/8000
...
```

C1 INVITE Voice Portal -> GW

```
INVITE sip:GW@voice.com SIP/2.0
Via: SIP/2.0/UDP voice.com:5060
From: VoicePortal <sip:VoicePortal@voice.com>
To: GW <sip:GW@voice.com>
...
Require: forked-media
...

m=audio 16980 rtp/avt 0 101
c=IN IP4 10.0.0.3
a=rtpmap:0 PCMU/8000
...
m=audio 19030 rtp/avt 0 101
c=IN IP4 10.0.0.2
a=...
```


C2 200 OK GW -> Voice Portal

```
SIP/2.0 200 OK
...
From: VoicePortal <sip:VoicePortal@voice.com>
To: GW <sip:GW@voice.com>]
...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
```

B3 ACK Voice Portal -> Target

```
ACK sip:Target@voice.com SIP/2.0
```

```
...
From: VoicePortal <sip:VoicePortal@voice.com>
To: Target <sip:Target@voice.com>
...
Content-Type: application/sdp
Content-Length: ....

m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
```

5. GW responses and use of Require:

For the purpose of establishing a successful forked-media session, it may be necessary to ensure that GW supports forked media sessions. Hence, Voyager may send a "Require:" header as part of its INVITE message to GW. The Require header for forked media may be defined as illustrated above, and simply as:

```
Require: forked-media
```

In cases where end-points or gateways do not support forked media, an appropriate response for an unsupported "Require: " header (420 Bad Extension) may be sent as:

```
SIP/2.0 420 Bad Extension
Unsupported: forked-media
```


With respect to the response of GW, which responds with multiple media description headers for the same port, it is worth noting that this would be indication of the ability of GW to support incoming RTP streams from two different sources on a single port (possibly through RTP mixing). Following this, the media description headers in the SDP message would act as an indication of the capability of the sender.

6. Example involving more than two forked media streams

For a more complete and complex forked media scenario, following is an illustration involving more than two forked media streams, with two Voice Portals and a Target.

```
          ----        +--------+   SIP    +---------+
         /    \       |        |<-------->|         |
        |Origin+-----+   GW    |   RTP    | Portal 1|
        |  UA  |      |        |<========>|         |
         \    /       +--------+          +---------+
          ----           ^^  ^^                ^
                         ||  ||                |
                         ||  ||                | SIP
                         ||  ||                |
                         ||  ||                |
                         ||  ||                v
                         ||  ||    RTP     +---------+
                         ||  ============ >|         |
                         ||                | Portal 2|
                         ||               >|         |
                         ||                +---------+
                         ||                    ^
                         ||                    |
                         ||                    |
                         ||                    v
                         ||                 +---------+
                         ||        RTP      |         |
                         ===================>|  Target |
                                            |         |
                                            +---------+
```


In the scenario shown above, the call-flow would like:

```
 GW            Portal1          Portal2          Target
 |                |                |                |
 | INVITE      A1 |                |                |
 |--------------->|                |                |
 |                |                |                |
 | 180 Ringing A2 |                |                |
```

```
    |<---------------|                |                |
    |                |                |                |
    | 200 OK      A3 |                |                |
    |<---------------|                |                |
    |                |                |                |
    | ACK          A4 |                |                |
    |--------------->|                |                |
    |                |                |                |
    | 2-way media A5 |                |                |
    |<==============>|                |                |
```

   Portal1 initiates a forked media session with Portal2.

```
    |                |  INVITE without|                |
    |                |  SDP        B1 |                |
    |                |--------------->|                |
    |                |                |                |
    |                |   200 OK    B2 |                |
    |                |<---------------|                |
    |                |                |                |
    |                |                |                |
    |                |                |                |
    |                |                |                |
    |INVITE with     |                |                |
    |multiple media  |                |                |
    |headers       C1|                |                |
    |<---------------|                |                |
    |                |                |                |
    |200 OK with   C2|                |                |
    |multiple media  |                |                |
    |--------------->|                |                |
    |  ACK         C3|                |                |
    |<---------------|                |                |
    |                | ACK with single|                |
    |                | media header B3|                |
    |                |--------------->|                |
    |                |                |                |
    | 2-way media    |                |                |
    |<===============================>|                |
    |                |                |                |
```


   Portal2 now initiates a forked media session with Target.

```
    |                |                |  INVITE without|
```

```
    |                |                |  SDP         D1|
    |                |                |--------------->|
    |                |                |                |
    |                |                |   200 OK     D2|
    |                |                |<---------------|
    |                |                |                |
    |                |INVITE with     |                |
    |                |multiple media  |                |
    |                |headers       E1|                |
    |                |<---------------|                |
    |                |                |                |
    |                |                |                |
    |                |                |                |
    |INVITE with     |                |                |
    |multiple media  |                |                |
    |headers       F1|                |                |
    |<---------------|                |                |
    |                |                |                |
    |200 OK with   F2|                |                |
    |multiple media  |                |                |
    |--------------->|                |                |
    |  ACK         F3| 200 OK with  E2|                |
    |<---------------| multiple media |                |
    |                |--------------->|                |
    |  ACK           |              E3|                |
    |<---------------|----------------| ACK with single|
    |                |                | media header D3|
    |                |                |--------------->|
    |                |                |                |
    |                |                |                |
    | 2-way media    |                |                |
    |<================================================>|
    |                |                |                |
    |                |                |                |
    |                |                |                |
```

Messages A1-B3 have been covered in the previous example. Relevant parts of the rest of the messages are given below:

D1 INVITE without SDP Voice Portal2 -> Target

```
INVITE sip:Target@voice.com SIP/2.0
Via: SIP/2.0/UDP voice.com:5060
From: VoicePortal2 <sip:VoicePortal2@voice.com>
To: Target <sip:Target@voice.com>
...
```

D2 200 OK Target -> Voice Portal

```
SIP/2.0 200 OK
...
From: VoicePortal2 <sip:VoicePortal2@voice.com>
To: Target <sip:Target@voice.com>
...
Content-Type: application/sdp
Content-Length: ....

...
c=IN IP4 10.0.0.4
...
m=audio 19000 rtp/avt 0 101
a=rtpmap:0 PCMU/8000
...
```

E1 INVITE Voice Portal2 -> Voice Portal1

```
INVITE sip:VoicePortal1@voice.com SIP/2.0
Via: SIP/2.0/UDP voice.com:5060
From: VoicePortal2 <sip:VoicePortal2@voice.com>
To: VoicePortal1 <sip:VoicePortal1@voice.com>
...
Require: forked-media
...

m=audio 19000 rtp/avt 0 101
c=IN IP4 10.0.0.4
a=rtpmap:0 PCMU/8000
...
m=audio 19030 rtp/avt 0 101
c=IN IP4 10.0.0.3
a=...
```

F1 INVITE Voice Portal1 -> GW

```
INVITE sip:GW@voice.com SIP/2.0
Via: SIP/2.0/UDP voice.com:5060
From: VoicePortal1 <sip:VoicePortal1@voice.com>
To: GW <sip:GW@voice.com>
...
Require: forked-media
...
```

```
m=audio 19000 rtp/avt 0 101
c=IN IP4 10.0.0.4
a=rtpmap:0 PCMU/8000
```



```
...
m=audio 19030 rtp/avt 0 101
c=IN IP4 10.0.0.3
a=...
...
m=audio 16980 rtp/avt 0 101
c=IN IP4 10.0.0.2
a=...
```

F2 200 OK GW -> Voice Portal

```
SIP/2.0 200 OK
...
From: VoicePortal1 <sip:VoicePortal1@voice.com>
To: GW <sip:GW@voice.com>]
...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
```

E2 200 OK Voice Portal1 -> Voice Portal2

```
SIP/2.0 200 OK
...
From: VoicePortal2 <sip:VoicePortal2@voice.com>
To: VoicePortal1 <sip:VoicePortal1@voice.com>
...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
```

D3 ACK Voice Portal2 -> Target

```
ACK sip:Target@voice.com SIP/2.0
...
From: VoicePortal2 <sip:VoicePortal2@voice.com>
To: Target <sip:Target@voice.com>
...
Content-Type: application/sdp
Content-Length: ....
```



```
m=audio 17000 rtp/avt 0 101
c=IN IP4 10.0.0.1
a= ...
```

7. Security Considerations

This draft demonstrates examples of SIP and 3rd Party Call Control [5] call flows. The security implications of these documents are explained in the respective documents.

8. References

[1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

[2] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Request for Comments (Best Current Practice) 2119, Internet Engineering Task Force, Mar. 1997.

[3] M. Handley and V. Jacobson, "SDP: session description protocol," Request for Comments 2327, Internet Engineering Task Force, April 1998

[4] M. Handley, E. Schooler, and H. Schulzrinne, "SIP: Session Initiation Protocol", RFC2543, Internet Engineering Task Force, Nov 1998.

[5] J. Rosenberg, J. Peterson, H. Schulzrinne, "Third Party Call Control in SIP", Internet Draft <draft-rosenberg-sip-3pcc-01.txt>, IETF; Nov. 2000. Work in progress

10. Acknowledgments

Thanks to Rohan Mahy for original ideas regarding this draft.

11. Author's Addresses

Mahesh Shankar
Cisco Systems Inc.
170 W Tasman Drive MS: SJC21/3/3
San Jose, CA 95134
Phone: +1 408 527 6966
Email: mshankar@cisco.com