Network Working Group
Internet-Draft
Expires: September 6, 2006

C. Jennings
D. Wing
Cisco Systems

March 5, 2006

Session Initiation Protocol (SIP) Offer/Answer with Multipart Alternative
draft-jennings-sipping-multipart-02

Status of this Memo

Copyright Notice

Abstract

SIP needs a mechanism for general backwards compatibility for moving from SDP to SDPng or moving from non end-to-end encrypted SDP to end-to-end encrypted SDP. This document specifies how a SIP offer uses multipart/alternative, and how an answer indicates which part was selected.

Table of Contents

1. Introduction and Overview

SIP (RFC 3261 [5]) uses an offer/answer negotiation mechanism described in [6]. This system carries offers in formats such as SDP [10] and various signed and encrypted versions of SDP. However, the current offer/answer scheme does not allow a backwards compatibility mode in which a SIP User Agent can make both old and new types of offers and allow the other User Agent to select the type of offer that it supports. This specification extends SIP to allow for these backwards compatible offer/answer schemes.

The mechanism for doing this is based on multipart alternative MIME types[2]. The User Agent making the offer uses a multipart alternative and includes a unique Content-ID for each of the body parts. The User Agent receiving the offer selects one of the parts in the offer and sends an answer based on that part. When the User Agent sends the answer, it indicates the Content-ID of the selected offer part.

The indication is done by using a new header field, Content-Answering-CID, defined in this document. This new header is similar to In-Reply-To (RFC822 [9]) but instead of indicating the Message-ID that elicited the email reply, Content-Answering-CID indicates the Content-ID of the body part that was interpreted and that generated the SIP answer.

This approach can allow a single offer to contain both SDP and SDPng[12][13]. It can also allow migration from offers that are not S/MIME protected (as described in RFC 3261 [5]) to ones that are, and allow SRTP [14] keying material to be passed in the S/MIME protected SDP using a mechanism such as sdescriptions [14]. As with all offers, the offerer's local policy and local capabilities would determine if an offer would, in fact, contain multiple alternatives.

2. Conventions

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 [3].

3. Mechanisms

The mechanisms described below apply only to SIP Offer/Answer [6] exchanges. As currently defined, SIP Offer/Answer exchanges have a Content-Disposition of "session" (either implied or explicit -- see section 20.11 of [5]) or "early-session" [8]. Unless specified

otherwise, other extensions to SIP Offer/Answer are allowed to interwork with the mechanisms described below.

3.1. Sending Offers

A User Agent that can support multiple types of offers SHOULD construct a multipart/alternative body with a body for each type it supports. Each body MUST include a Content-ID header that MUST be unique within this SIP dialog. It is RECOMMENDED that the Content-ID be generated by a combination of a random string and the User Agent's host name or IP address, in order to make them globally unique.

It is critical that multipart/alternative offers follow the semantics of multipart/alternative, notably the following text from RFC2046 [2]:

> "THE SEMANTICS OF CONTENT-ID IN MULTIPART/ALTERNATIVE: Each part of a "multipart/alternative" entity represents the same data, but the mappings between the two are not necessarily without information loss. For example, information is lost when translating ODA to PostScript or plain text. It is recommended that each part should have a different Content-ID value in the case where the information content of the two parts is not identical."

Because support for MIME multipart isn't mandated by SIP, User Agents should expect to receive error responses that indicate multipart/alternative wasn't supported ("415 Unsupported Media Type"). In such events, User Agents MAY retry the request using an offer that consists of only an SDP body (that is, without the multipart/alternative described in this document). See also Section 7.

Note: SIP forking may cause a problem with the above.

Per section 2.9 of [4] the top-level Content-Disposition header applies to all parts of a multipart. The body parts of a multipart/alternative MUST NOT have their own Content-Disposition, as this severely complicates the selection of the appropriate part by the receiver.

3.2. Receiving Offers and Sending Answers

If a User Agent receives multiple SDP offers in an multipart/alternative body, it MUST interpret these as it would a normal multipart alternative, as defined in RFC 2046 [2], which describes the User Agent starting from the last part, attempting to interpret it, and working backwards to the next-to-last part, and so on, until it can interpret a part. Interpreting a part requires being able to

successfully decrypt the part (if encrypted) and being able to understand the Content-Type.

If the multipart/alternative body doesn't contain a Content-Disposition header, a Content-Disposition of "handling=required" MUST be assumed, as with other MIME types that lack a Content-Disposition header (section 3.3 of [7]). If this header includes a "handling" parameter, or if "handling=required" was assumed, the "handling" parameter applies only to the part that the user agent chose to interpret. Specifically, the "handling" parameter does not apply to the parts that the user agent could not interpret or chose not to interpret, but rather to the ability to "handle" the multipart/alternative part as a whole.

When the User Agent constructs the answer, it MUST include a Content-Answering-CID header field as defined in Section 4 with the same value as the Content-ID of the offer that was selected. If the answer itself is a multipart MIME message[8], the Content-Answering-CID header field MUST be in the same MIME part of the answer. To reduce complexity, only one answer is allowed even if the offer contained multiple alternatives; that is, the answer MUST NOT be a multipart/alternative. If the answer is being rejected, the User Agent SHOULD indicate its capabilities (section 21.4.26 of [5]).

3.3. Receiving Answers

When the User Agent receives an answer, it MUST look at the Content-Answering-CID header field value to find which answer has been used. If the answer is a multipart/alternative, the User Agent MUST reject the answer as malformed. It then proceeds with normal offer/answer processing.

When a UA offers a multipart session, and the user might reasonably be expected to behave differently depending on the part that was answered, the UA SHOULD inform the user of the part that was answered. For example, a multipart/alternative offer might contain one part with SDP for only audio, and another encrypted part with SDP for audio and video. If the second part, containing the audio and video stream, is answered, it is reasonable to illuminate an LED on the video camera.

4. Syntax

This specification defines a new MIME header called "Content-Answering-CID". This updates RFC 2045 [1] with:

```
answering-cid := "Content-Answering-CID" ":" msg-id
```

and adds "[answering-cid CRLF]" to the Identity headers in RFC 2045.

The Content-Answering-CID header is used in answers and has a msg-id value that is the same as the Content-ID value of the offer to which this answer is related

5. Example SIP SRTP Call

In this example, large parts of the message are omitted to highlight what is relevant to this specification. The lines in the example that are prefixed by $ represent encrypted blocks of data.

In this example, Alice calls Bob and offers both an RTP and an SRTP session. The SDP for the SRTP session contains the SRTP keying material, and the SDP is encrypted with S/MIME. It is assumed that Alice has Bob's public key.

Alice sends an INVITE to Bob that offers two alternative SDP bodies: the first part contains SDP for an RTP audio stream and the second encrypted part contains SDP for an SRTP audio and SRTP video stream. Per multipart/alternative semantics, the encrypted version is preferred because it is the last part. Both parts contain unique Content-ID headers. The top-most part indicates the disposition is "session", which applies to all of the parts within that top-most part.

The "$" indicates encrypted data. The a=crypto line is shown wrapped because of document formatting restrictions; it is actually one long line.

```
      INVITE sip:bob@biloxi.example.com SIP/2.0
    ...
      Content-Type: multipart/alternative; boundary=yradnuob
      Content-Disposition: session

      --yradnuob
      Content-ID: <83rqjqef3.218.1@10.1.1.1>
      Content-Type: application/sdp

      v=0
      o=alice 2890844526 2890844526 IN IP4 192.168.47.11
      s=-
      c=IN IP4 192.168.47.11
      t=0 0
      m=audio 51400 RTP/AVP 0
      a=rtpmap:0 PCMU/8000

      --yradnuob
      Content-ID: <83rqjqef3.218.2@10.1.1.1>
      Content-Type: application/pkcs7-mime

    $ Content-Type: application/sdp
    $
    $ v=0
    $ o=alice 2890844526 2890844526 IN IP4 192.168.47.11
    $ s=-
    $ c=IN IP4 192.168.47.11
    $ t=0 0
    $ m=video 51372 RTP/SAVP 31
    $ a=crypto:1 AES_CM_128_HMAC_SHA1_80
    $   inline:d0RmdmcmVCspeEc3QGZiNWpVLFJhQX1cfHAwJSoj|2^20|1:32
    $ m=audio 49170 RTP/SAVP 0
    $ a=crypto:1 AES_CM_128_HMAC_SHA1_80
    $   inline:NzB4d1BINUAvLEw6UzF3WSJ+PSdFcGdUJShpX1Zj|2^20|1:32

      --yradnuob--
```

Assuming that Bob's User Agent supports encryption and had Alice's public key, Bob would be able to decode and interpret Alice's second alternative body part. Bob would use this body part to construct an answer. Bob's answer includes the Content-Reply-to-CID header which indicates which alternative body part was chosen.

```
  200 OK
...
  Content-Answering-CID: <83rqjqef3.218.2@10.1.1.1>
  Content-Type: application/pkcs7-mime
  Content-Disposition: session

$ Content-Type: application/sdp
$
$ v=0
$ o=bob 2890844526 2890844526 IN IP4 192.168.47.11
$ s=-
$ c=IN IP4 192.168.51.1
$ t=0 0
$ m=video 27350 RTP/SAVP 31
$ a=crypto:1 AES_CM_128_HMAC_SHA1_80
$   inline:xiNb96JefmqJ8JneiqjqXqizje334+jkeiq298fA|2^20|1:32
$ m=audio 27352 RTP/SAVP 0
$ a=crypto:1 AES_CM_128_HMAC_SHA1_80
$   inline:eu23cfnze++jejekqnQQjefiuwfj938ejefQQfec|2^20|1:32
```

6. Example SIP SDPng Call

This shows an offer containing SDP and SDPng:

```
    INVITE sip:bob@biloxi.example.com SIP/2.0
  ...
    Content-Type: multipart/alternative; boundary=yradnuob
    Content-Disposition: session

    --yradnuob
    Content-ID: <98efj3.1@10.1.1.1>
    Content-Type: application/sdp

    v=0
    o=alice 2890844526 2890844526 IN IP4 192.168.47.11
    s=-
    c=IN IP4 192.168.47.11
    t=0 0
    m=audio 51400 RTP/AVP 0 33
    a=rtpmap:0 PCMU/8000
    a=rtpmap:3 GSM/8000

    --yradnuob
    Content-ID: <98efj3.2@10.1.1.1>
    Content-Type: application/sdpng

    <?xml version="1.0" encoding="UTF-8"?>
    <sdpng xmlns="http://www.iana.org/sdpng"
  ...
    </sdpng>

    --yradnuob--
```

Here is the answer which indicates via the Content-Answering-CID that the SDP body part was interpreted.

```
  200 OK
...
  Content-Answering-CID: <98efj3.1@10.1.1.1>
  Content-Type: application/sdp
  Content-Disposition: session

  v=0
  o=bob 990821536 1230844577 IN IP4 bob.example.com
  s=-
  c=IN IP4 192.168.1.2
  t=0 0
  m=audio 55111 RTP/AVP 0 33
  a=rtpmap:0 PCMU/8000
  a=rtpmap:3 GSM/8000
```

7. Security Considerations

In SIP, there is a risk of an active bid-down attack. The active attacker can modify an SRTP offer, or SRTP answer, in order to make the offerer believe the answerer cannot understand SRTP. Such an attack is possible with or without multipart/alternative offers described in this paper. In such an attack without a multipart/ alternative offer, the offerer might send a new RTP offer. In such an attack with a multipart/alternative offer (containing both an RTP offer and an encrypted offer), an attacker might guess the encrypted offer is an SRTP offer and might reasonably assume the offerer's policy allows an RTP session. To protect against such an attack, the offer can be protected (e.g., using the SIPS URI or using Identity[15]), and the answer can be similarly protected. The addition of multipart/alternative doesn't change this risk, or the requirement to appropriately protect such offers and answers, rather it only provides a hint about the offerer's policy which might allow an RTP session to be established. See also Section 3.3 for user interface guidance.

8. IANA Considerations

The MIME Content-Answering-CID header does not require any IANA actions.

9. Acknowledgments

Thanks for comments from Flemming Andreasen, Paul Kyzivat, and Mark Baugher.

10. References

10.1. Normative References

[1] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.

[2] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[4] Troost, R., Dorner, S., and K. Moore, "Communicating Presentation Information in Internet Messages: The Content-Disposition Header Field", RFC 2183, August 1997.

[5] 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.

[6] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

[7] Burger, E., "Critical Content Multi-purpose Internet Mail Extensions (MIME) Parameter", RFC 3459, January 2003.

[8] Camarillo, G., "The Early Session Disposition Type for the Session Initiation Protocol (SIP)", RFC 3959, December 2004.

10.2. Informational References

[9] Crocker, D., "Standard for the format of ARPA Internet text messages", STD 11, RFC 822, August 1982.

[10] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

[11] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-06 (work in progress), October 2005.

[12] Kutscher, D., Ott, J., and C. Bormann, "Session Description and Capability Negotiation", draft-ietf-mmusic-sdpng-08 (work in progress), February 2005.

[13] Ott, J. and C. Perkins, "SDPng Transition", draft-ietf-mmusic-sdpng-trans-04 (work in progress), May 2003.

[14] Andreasen, F., "Session Description Protocol Security Descriptions for Media Streams", draft-ietf-mmusic-sdescriptions-12 (work in progress), September 2005.

[15] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", draft-ietf-sip-identity-06 (work in progress), October 2005.

Authors' Addresses

Cullen Jennings
Cisco Systems
170 West Tasman Drive
MS: SJC-21/2
San Jose, CA 95134
USA

Phone: +1 408 902-3341
Email: fluffy@cisco.com

Dan Wing
Cisco Systems
170 West Tasman Drive
MS: SJC-21/2
San Jose, CA 95134
USA

Email: dwing@cisco.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.