

Session Initiation Protocol (SIP)
Authenticated Identity Body (AIB) Format

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.

Copyright Notice

Copyright (C) The Internet Society (2004).

Abstract

RFC 3261 introduces the concept of adding an S/MIME body to a Session Initiation Protocol (SIP) request or response in order to provide reference integrity over its headers. This document provides a more specific mechanism to derive integrity and authentication properties from an 'authenticated identity body', a digitally-signed SIP message, or message fragment. A standard format for such bodies (known as Authenticated Identity Bodies, or AIBs) is given in this document. Some considerations for the processing of AIBs by recipients of SIP messages with such bodies are also given.

Table of Contents

1.	Introduction	2
1.1.	Requirements Notation.	3
2.	AIB Format	4
3.	Example of a Request with AIB	5
4.	AIBs for Identifying Third-Parties	6
5.	Identity in non-INVITE Requests	7
6.	Identity in Responses	7
7.	Receiving an AIB	7
8.	Encryption of Identity	8
9.	Example of Encryption	8
10.	Security Considerations	9
11.	IANA Considerations	11
12.	References	11
12.1.	Normative References	11
12.2.	Informative References	11
13.	Acknowledgements	11
14.	Author's Address	12
15.	Full Copyright Statement	13

1. Introduction

Section 23.4 of RFC 3261 [1] describes an integrity mechanism that relies on signing tunneled 'message/sip' MIME bodies within SIP requests. The purpose of this mechanism is to replicate the headers of a SIP request within a body carried in that request in order to provide a digital signature over these headers. The signature on this body also provides authentication.

The core requirement that motivates the tunneled 'message/sip' mechanism is the problem of providing a cryptographically verifiable identity within a SIP request. The baseline SIP protocol allows a user agent to express the identity of its user in any of a number of headers. The primary place for identity information asserted by the sender of a request is the From header. The From header field contains a URI (like 'sip:alice@example.com') and an optional display-name (like "Alice") that identifies the originator of the request. A user may have many identities that are used in different contexts.

Typically, this URI is an address-of-record that can be de-referenced in order to contact the originator of the request; specifically, it is usually the same address-of-record under which a user registers their devices in order to receive incoming requests. This address-of-record is assigned and maintained by the administrator of the SIP service in the domain identified by the host portion of the address-of-record. However, the From field of a request can usually be set

arbitrarily by the user of a SIP user agent; the From header of a message provides no internal assurance that the originating user can legitimately claim the given identity. Nevertheless, many SIP user agents will obligingly display the contents of the From field as the identity of the originator of a received request (as a sort of caller identification function), much as email implementations display the From field as the sender's identity.

In order to provide the recipient of a SIP message with greater assurance of the identity of the sender, a cryptographic signature can be provided over the headers of the SIP request, which allows the signer to assert a verifiable identity. Unfortunately, a signature over the From header alone is insufficient because it could be cut-and-pasted into a replay or forwarding attack, and more headers are therefore needed to correlate a signature with a request. RFC 3261 therefore recommends copying all of the headers from the request into a signed MIME body; however, SIP messages can be large, and many of the headers in a SIP message would not be relevant in determining the identity of the sender or assuring reference integrity with the request, and moreover some headers may change in transit for perfectly valid reasons. Thus, this large tunneled 'message/sip' body will almost necessarily be at variance with the headers in a request when it is received by the UAS, and the burden is on the UAS to determine which header changes were legitimate, and which were security violations. It is therefore desirable to find a happy medium - to provide a way of signing just enough headers that the identity of the sender can be ascertained and correlated with the request. 'message/sipfrag' [4] provides a way for a subset of SIP headers to be included in a MIME body; the Authenticated Identity Body (AIB) format described in Section 2 is based on 'message/sipfrag'.

For reasons of end-to-end privacy, it may also be desirable to encrypt AIBs; procedures for this encryption are given in Section 8.

This document proposes that the AIB format should be used instead of the existing tunneled 'message/sip' mechanism described in RFC 3261, section 23.4, in order to provide the identity of the caller; if integrity over other, unrelated headers is required, then the 'message/sip' mechanism should be used.

1.1. Requirements Notation

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 BCP 14, RFC 2119 [2].

2. AIB Format

As a way of sharing authenticated identity among parties in the network, a special type of MIME body format, the Authenticated Identity Body (AIB) format, is defined in this section. AIBs allow a party in a SIP transaction to cryptographically sign the headers that assert the identity of the originator of a message, and provide some other headers necessary for reference integrity.

An AIB is a MIME body of type 'message/sipfrag' - for more information on constructing sipfrags, including examples, see [4]. This MIME body MUST have a Content-Disposition [3] disposition-type of 'aib', a new value defined in this document specifically for authenticated identity bodies. The Content-Disposition header SHOULD also contain a 'handling' parameter indicating that this MIME body is optional (i.e., if this mechanism is not supported by the user agent server, it can still attempt to process the request).

AIBs using the 'message/sipfrag' MIME type MUST contain the following headers when providing identity for an INVITE request: From, Date, Call-ID, and Contact; they SHOULD also contain the To and CSeq header. The security properties of these headers, and circumstances in which they should be used, are described in Section 10. AIBs MAY contain any other headers that help to uniquely identify the transaction or provide related reference integrity. An example of the AIB format for an INVITE is:

```
Content-Type: message/sipfrag
Content-Disposition: aib; handling=optional
```

```
From: Alice <sip:alice@example.com>
To: Bob <sip:bob@example.net>
Contact: <sip:alice@pc33.example.com>
Date: Thu, 21 Feb 2002 13:02:03 GMT
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
```

Unsigned AIBs MUST be treated by any recipients according to the rules set out in Section 7 for AIBs that do not validate. After the AIB has been signed, it SHOULD be added to existing MIME bodies in the request (such as SDP), if necessary by transitioning the outermost MIME body to a 'multipart/mixed' format.

3. Example of a Request with AIB

The following shows a full SIP INVITE request with an AIB:

```
INVITE sip:bob@example.net SIP/2.0
Via: SIP/2.0/UDP pc33.example.com;branch=z9hG4bKnashds8
To: Bob <sip:bob@example.net>
From: Alice <sip:alice@example.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Max-Forwards: 70
Date: Thu, 21 Feb 2002 13:02:03 GMT
Contact: <sip:alice@pc33.example.com>
Content-Type: multipart/mixed; boundary=unique-boundary-1
```

--unique-boundary-1

```
Content-Type: application/sdp
Content-Length: 147
```

```
v=0
o=UserA 2890844526 2890844526 IN IP4 example.com
s=Session SDP
c=IN IP4 pc33.example.com
t=0 0
m=audio 49172 RTP/AVP 0
a=rtpmap:0 PCMU/8000
```

--unique-boundary-1

```
Content-Type: multipart/signed;
  protocol="application/pkcs7-signature";
  micalg=sha1; boundary=boundary42
Content-Length: 608
```

--boundary42

```
Content-Type: message/sipfrag
Content-Disposition: aib; handling=optional
```

```
From: Alice <sip:alice@example.com>
To: Bob <sip:bob@example.net>
Contact: <sip:alice@pc33.example.com>
Date: Thu, 21 Feb 2002 13:02:03 GMT
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
```

--boundary42

```
Content-Type: application/pkcs7-signature; name=smime.p7s
Content-Transfer-Encoding: base64
```

Content-Disposition: attachment; filename=smime.p7s;
 handling=required

ghyHhHUujhJh77n8HHGTrfvbnj756tbB9HG4VQpfyF467GhIGfHfYT6
 4VQpfyF467GhIGfHfYT6jH77n8HHGghyHhHUujhJh756tbB9HGTrfvbnj
 n8HHGTrfvhJh776tbB9HG4VQbnj7567GhIGfHfYT6ghyHhHUujpFyF4
 7GhIGfHfYT64VQbnj756

--boundary42--

--unique-boundary-1--

4. AIBs for Identifying Third-Parties

There are special-case uses of the INVITE method in which some SIP messages are exchanged with a third party before an INVITE is sent, and in which the identity of the third party needs to be carried in the subsequent INVITE. The details of addressing identity in such contexts are outside the scope of this document. At a high level, it is possible that identity information for a third party might be carried in a supplemental AIB. The presence of a supplemental AIB within a message would not preclude the appearance of a 'regular' AIB as specified in this document.

Example cases in which supplemental AIBs might appear include:

The use of the REFER [5] method, for example, has a requirement for the recipient of an INVITE to ascertain the identity of the referrer who caused the INVITE to be sent.

Third-party call control (3PCC [6]) has an even more complicated identity problem. A central controller INVITES one party, gathers identity information (and session context) from that party, and then uses this information to INVITE another party. Ideally, the controller would also have a way to share a cryptographic identity signature given by the first party INVITED by the controller to the second party invited by the controller.

In both of these cases, the Call-ID and CSeq of the original request (3PCC INVITE or REFER) would not correspond with that of the request in by the subsequent INVITE, nor would the To or From. In both the REFER case and the 3PCC case, the Call-ID and CSeq cannot be used to guarantee reference integrity, and it is therefore much harder to correlate an AIB to a subsequent INVITE request.

Thus, in these cases some other headers might be used to provide reference integrity between the headers in a supplemental AIB with the headers of a 3PCC or REFER-generated INVITE, but this usage is

outside of the scope of this document. In order for AIBs to be used in these third-party contexts, further specification work is required to determine which additional headers, if any, need to be included in an AIB in a specific third-party case, and how to differentiate the primary AIB in a message from a third-party AIB.

5. Identity in non-INVITE Requests

The requirements for populating an AIB in requests within a dialog generally parallel those of the INVITE: From, Call-ID, Date, and Contact header fields are REQUIRED.

Some non-INVITE requests, however, may have different identity requirements. New SIP methods or extensions that leverage AIB security MUST identify any special identity requirements in the Security Considerations of their specification.

6. Identity in Responses

Many of the practices described in the preceding sections can be applied to responses as well as requests. Note that a new set of headers must be generated to populate the AIB in a response. The From header field of the AIB in the response to an INVITE MUST correspond to the address-of-record of the responder, NOT to the From header field received in the request. The To header field of the request MUST NOT be included. A new Date header field and Contact header field should be generated for the AIB in a response. The Call-ID and CSeq should, however, be copied from the request.

Generally, the To header field of the request will correspond to the address-of-record of the responder. In some architectures where re-targeting is used, however, this need not be the case. Some recipients of response AIBs may consider it a cause for security concern if the To header field of the request is not the same as the address-of-record in the From header field of the AIB in a response.

7. Receiving an AIB

When a user agent receives a request containing an AIB, it MUST verify the signature, including validating the certificate of the signer, and compare the identity of the signer (the subjectAltName) with, in the INVITE case, the domain portion of the URI in the From header field of the request (for non-INVITE requests, other headers MAY be subject to this comparison). The two should correspond exactly; if they do not, the user agent MUST report this condition to its user before proceeding. User agents MAY distinguish between plausibly minor variations (the difference between 'example.com' and 'sip.example.com') and major variations ('example.com' vs.

'example.org') when reporting these discrepancies in order to give the user some idea of how to handle this situation. Analysis and comparison of the Date, Call-ID, and Contact header fields as described in Section 10 MUST also be performed. Any discrepancies or violations MUST be reported to the user.

When the originating user agent of a request receives a response containing an AIB, it SHOULD compare the identity in the From header field of the AIB of the response with the original value of the To header field in the request. If these represent different identities, the user agent SHOULD render the identity in the AIB of the response to its user. Note that a discrepancy in these identity fields is not necessarily an indication of a security breach; normal re-targeting may simply have directed the request to a different final destination. Implementors therefore may consider it unnecessary to alert the user of a security violation in this case.

8. Encryption of Identity

Many SIP entities that support the use of S/MIME for signatures also support S/MIME encryption, as described in RFC 3261, Section 23.4.3.

While encryption of AIBs entails that only the holder of a specific key can decrypt the body, that single key could be distributed throughout a network of hosts that exist under common policies. The security of the AIB is therefore predicated on the secure distribution of the key. However, for some networks (in which there are federations of trusted hosts under a common policy), the widespread distribution of a decryption key could be appropriate. Some telephone networks, for example, might require this model.

When an AIB is encrypted, the AIB SHOULD be encrypted before it is signed. Implementations MUST still accept AIBs that have been signed and then encrypted.

9. Example of Encryption

The following is an example of an encrypted and signed AIB (without any of the preceding SIP headers). In a rendition of this body sent over the wire, the text wrapped in asterisks would be in ciphertext.

```
Content-Type: multipart/signed;  
  protocol="application/pkcs7-signature";  
  micalg=sha1; boundary=boundary42  
Content-Length: 568  
Content-Disposition: aib; handling=optional  
  
--boundary42
```



```
Content-Type: application/pkcs7-mime; smime-type=enveloped-data;
  name=smime.p7m
Content-Transfer-Encoding: base64
Content-Disposition: attachment; filename=smime.p7m
  handling=required
Content-Length: 231
```

```
*****
* Content-Type: message/sipfrag *
* Content-Disposition: aib; handling=optional *
* *
* From: sip:alice@example.com *
* Call-ID: a84b4c76e66710 *
* Contact: sip:alice@device21.example.com *
* Date: Thu, 21 Feb 2002 13:02:03 GMT *
*****
```

--boundary42

```
Content-Type: application/pkcs7-signature; name=smime.p7s
Content-Transfer-Encoding: base64
Content-Disposition: attachment; filename=smime.p7s;
  handling=required
```

```
ghyHhHUujhJhj77n8HHGTrfvbnj756tbB9HG4VQpfyF467GhIGfHfYT6
4VQpfyF467GhIGfHfYT6jh77n8HHGghyHhHUujhJhj756tbB9HGTrfvbnj
n8HHGTrfvhJhj776tbB9HG4VQbnj7567GhIGfHfYT6ghyHhHUujpF4
7GhIGfHfYT64VQbnj756
```

--boundary42--

10. Security Considerations

The purpose of an AIB is to provide an identity for the sender of a SIP message. This identity is held in the From header field of an AIB. While other headers are also included, they are provided solely to assist in detection of replays and cut-and-paste attacks leveraged to impersonate the caller. The contents of the From header field of a valid AIB are suitable for display as a "Caller ID" for the sender of the SIP message.

This document mandates the inclusion of the Contact, Date, Call-ID, and From header fields within an AIB, and recommends the inclusion of CSeq and To header fields, when 'message/sipfrag' is used to represent the identity of a request's sender. If these headers are omitted, some important security properties of AIB are lost. In general, the considerations related to the inclusion of various headers in an AIB are the same as those given in RFC 3261 for

including headers in tunneled 'message/sip' MIME bodies (see Section 23 in particular).

The From header field indicates the identity of the sender of the message; were this header to be excluded, the creator of the AIB essentially would not be asserting an identity at all. The Date and Contact headers provide reference integrity and replay protection, as described in RFC 3261, Section 23.4.2. Implementations of this specification MUST follow the rules for acceptance of the Date header field in tunneled 'message/sip' requests described in RFC 3261, Section 23.4.2; this ensures that outdated AIBs will not be replayed (the suggested interval is that the Date header must indicate a time within 3600 seconds of the receipt of a message). Implementations MUST also record Call-IDs received in AIBs, and MUST remember those Call-IDs for at least the duration of a single Date interval (i.e., 3600 seconds). Accordingly, if an AIB is replayed within the Date interval, receivers will recognize that it is invalid because of a Call-ID duplication; if an AIB is replayed after the Date interval, receivers will recognize that it is invalid because the Date is stale. The Contact header field is included to tie the AIB to a particular device instance that generated the request. Were an active attacker to intercept a request containing an AIB, and cut-and-paste the AIB into their own request (reusing the From, Contact, Date, and Call-ID fields that appear in the AIB), they would not be eligible to receive SIP requests from the called user agent, since those requests are routed to the URI identified in the Contact header field.

The To and CSeq header fields provide properties that are generally useful, but not for all possible applications of AIBs. If a new AIB is issued each time a new SIP transaction is initiated in a dialog, the CSeq header field provides a valuable property (replay protection for this particular transaction). If, however, one AIB is used for an entire dialog, subsequent transactions in the dialog would use the same AIB that appeared in the INVITE transaction. Using a single AIB for an entire dialog reduces the load on the generator of the AIB. The To header field usually designates the original URI that the caller intended to reach, and therefore it may vary from the Request-URI if re-targeting occurs at some point in the network. Accordingly, including the To header field in the AIB helps to identify cut-and-paste attacks in which an AIB sent to a particular destination is re-used to impersonate the sender to a different destination. However, the inclusion of the To header field probably would not make sense for many third-party AIB cases (as described in Section 4), nor is its inclusion necessary for responses.

11. IANA Considerations

This document defines a new MIME Content-Disposition disposition-type value of 'aib'. This value is reserved for MIME bodies that contain an authenticated identity, as described in section Section 2.

12. References

12.1. Normative References

- [1] 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.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Troost, R., Dorner, S., and K. Moore, "Communicating Presentation Information in Internet Messages: The Content-Disposition Header Field", RFC 2183, August 1997.
- [4] Sparks, R., "Internet Media Type message/sipfrag", RFC 3420, November 2002.

12.2. Informative References

- [5] Sparks, R., "The Session Initiation Protocol (SIP) Referred-By Mechanism", RFC 3892, September 2004.
- [6] Rosenberg, J., Peterson, J., Schulzrinne, H., and G. Camarillo, "Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)", BCP 85, RFC 3725, April 2004.

13. Acknowledgements

The author would like to thank Robert Sparks, Jonathan Rosenberg, Mary Watson, and Eric Rescorla for their comments. Rohan Mahy also provided some valuable guidance.

14. Author's Address

Jon Peterson
NeuStar, Inc.
1800 Sutter St
Suite 570
Concord, CA 94520
US

Phone: +1 925/363-8720
EMail: jon.peterson@neustar.biz
URI: <http://www.neustar.biz/>

15. Full Copyright Statement

Copyright (C) The Internet Society (2004).

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/S HE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY 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 IETF's procedures with respect to rights in IETF 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.

Acknowledgement

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

