Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 26, 2008

D. Wing
Cisco Systems
H. Kaplan
Acme Packet

February 23, 2008

SIP Identity using Media Path
draft-wing-sip-identity-media-02

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on August 26, 2008.

Abstract

This document defines a new SIP identity mechanism which operates through SBCs and B2BUAs. This new identity mechanism creates a signature over certain SIP headers and certain SDP lines. When the SIP body contains SDP, both the SIP signaling path and the media path are used to perform the identity function; when the SIP body contains non-SDP body parts, they are signed in their entirety.

Table of Contents

1. Introduction

SIP Identity [RFC4474] provides cryptographic identity for SIP requests. It provides this protection by signing certain SIP header fields (Contact, Date, Call-ID, CSeq, To, and From) and the SIP message body. The SIP message body typically contains the SDP. However, as discussed in [I-D.wing-sip-identity-analysis], RFC4474 does not work well if intermediate domains have B2BUAs or SBCs. As of this writing, most service providers utilize SBCs at network ingress and at network egress.

The mechanism described in this document provides cryptographic assurance of the endpoint's identity, and works through most B2BUAs and through most SBCs.

The mechanism described in this document signs only certain SDP attributes, and not all the same SIP headers. The remote endpoint is expected to validate the signature over the SIP headers and specified SDP attributes, to initiate a proof of possession test over the media path, which proves the session has been established with the "From:" party in the SIP header. Mechanisms to perform this proof of possession are shown using DTLS and using a small extension to ICE [I-D.ietf-mmusic-ice]. This mechanism is also extensible, in order to be usable by future mechanisms which need signed SDP attributes

Readers of this document are expected to be familiar with RFC4474, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", which defines the Identity and Identity-Info header fields. A future version of this document will have less reliance on RFC4474.

2. Notational 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 [RFC2119].

3. Operation

The operation of SIP-Identity-Media is similar to RFC4474 and uses authentication service proxies much like RFC4474. The basic steps are:

- A new header, Identity-Media, is created containing the names of certain SDP attributes from SDP bodyparts, and containing a hash of non-SDP bodyparts.

- Several SIP headers and the Identity-Media header are all signed (as detailed in Section 3.1), and the result is placed in Identity-Media-Signature.

- The receiving domain validates the signature, and if the request is an invitation to establish a media channel, performs a proof of identity validation using DTLS, TLS, ICE, HIP, or ZRTP over the media path.

The following figure shows how the Authentication Service and the media validation is performed. The figure assumes the endpoints themselves perform the media validation.

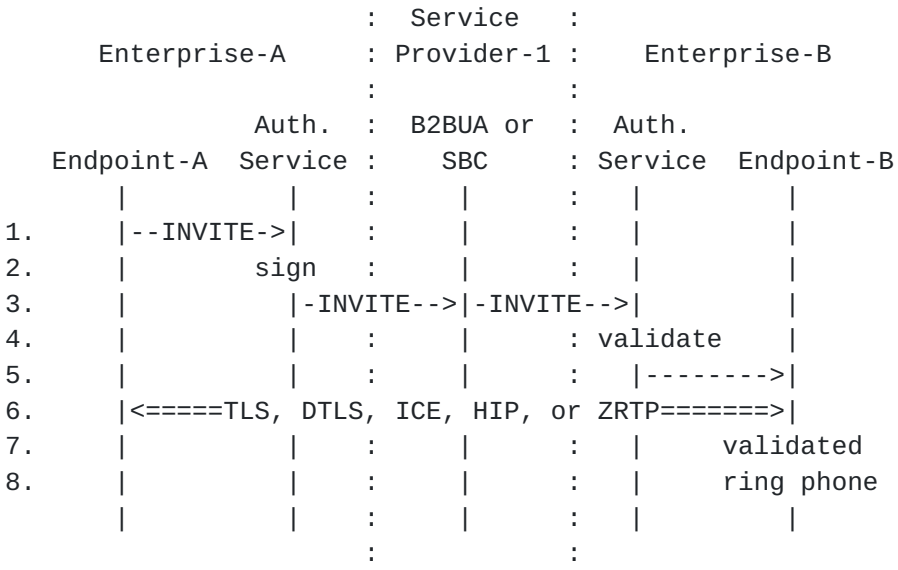

Figure 1: Message Flow

Step 1: Originating endpoint prepares to send an INVITE and chooses the identity-challenge technique it supports, and indicates that in the SDP it generates. Described in this document are identity challenges for TLS, DTLS, ICE, HIP, and ZRTP. It then sends the INVITE to its local SIP proxy.

Step 2: Originating endpoint's authentication service creates a new header, Identity-Media, containing certain attribute names from the SDP (e.g., "a=fingerprint", "a=ice-pub-key"). The authentication service then creates a signature over certain SIP headers (e.g., From, To) and this new Identity-Media header. The resulting signature is inserted into the new Identity-Media-Signature header. An Identity-Info header is added, pointing to this domain's certificate. The INVITE, with these additional headers, is forwarded to the next administrative domain.

[NOTE: alternatively, we could allow the UAC to create the Identity- Media header with the attributes it wants signed, then have the auth server sign them and insert the signature header - this would be more flexible]

Step 3: The next administrative domain has an SBC (or B2BUA). The SBC modifies or rewrites certain SDP fields. Most typically an SBC will modify the "m" and "c" lines. These modifications do not break the signature, so long as the SBC doesn't remove the headers Identity-Media, Identity-Media-Signature, or Identity-Info, and do not remove or alter the signed attributes from the SDP.

Step 4: The terminating endpoint's authentication service receives the INVITE. It validates that the signature contained in the Identity-Media-Signature header, and validates that the signing certificate is owned by the originating domain from step 2. This validation is done by using the certificate pointed to in the Identity-Info header, which MUST match the domain in the From: address.

Step 5: If the validation was successful, the terminating endpoint's authentication service forwards the INVITE to the endpoint.

Step 6: The terminating endpoint chooses a compatible identity-challenge technique from the INVITE (TLS, DTLS, ICE, HIP, or ZRTP), and performs that challenge. Described in this document are identity challenges for TLS, DTLS, ICE, HIP, and ZRTP.

Step 7: All of the identity challenges (TLS, DTLS, ICE, HIP, and ZRTP) cause the exchange of either a certificate or a public key in the media path. The terminating endpoint compares the certificate or public key with the fingerprint in the (signed) Identity-Media header (originally created in step 2). If they match, the terminating endpoint completes the identity challenge exchange. After completion, the originating endpoint has proven (to the terminating endpoint) that the originating endpoint knows the private key associated with the certificate (or public key) signed in step 2. The terminating endpoint has now validated the identity of the originating endpoint.

Step 8: The terminating endpoint can reliably and honestly indicate calling party information ("caller-id") and ring the phone.

3.1. Identity Media Signature

In RFC4474, a signature is formed over some SIP headers and over the entire body (which most typically contains SDP). In this specification, some SIP headers are signed but only specific SDP attributes that provide cryptographic identity are signed (e.g., "a=fingerprint" and its value). The specific SDP attributes that are signed depends on which cryptographic identity technique(s) is used; see section Section 4.

The SIP headers that are signed, for the signature placed into the Identity-Media-Signature header are:

- The AoR of the UA sending the message, or addr-spec of the From header field (referred to occasionally here as the 'identity field').

- The addr-spec component of the To header field, which is the AoR to which the request is being sent.

- The SIP method.

- [NOTE: Contact, CSeq and Call-Id not included]

- The Date header field, with exactly one space each for each SP and the weekday and month items case set as shown in the BNF in RFC3261. RFC3261 specifies that the BNF for weekday and month is a choice amongst a set of tokens. The RFC2234 rules for the BNF specify that tokens are case sensitive. However, when used to construct the canonical string defined here, the first letter of each week and month MUST be capitalized, and the remaining two letters must be lowercase. This matches the capitalization provided in the definition of each token. All requests that use the Identity-Media mechanism MUST contain a Date header.

- The Identity-Media header field value.

The hash is formed of these elements:

```
digest-string = addr-spec "|" addr-spec "|"
                Method "|" SIP-date "|"
                attrib-bodyhash-list
```

The first addr-spec MUST be taken from the From header field value, the second addr-spec MUST be taken from the To header field value.

The Identity-Info header points to where the authentication service's certificate can be retrieved from.

3.2. Authentication Service

The authentication service examines the SIP message body to build the Identity-Media header. For each message body found, in the order found:

- if the body part is application/sdp, the authentication service retrieves only the cryptographic attributes from the SDP (as described in Section 4), and appends that information to the Identity-Media header.

- otherwise, for all other body parts, the body part is hashed using SHA-1, and the first 96 bytes are appended to the Identity-Media header using "BPH=".

For example, A SIP request with three bodyparts: text/plain, application/sdp, and image/jpg, the Identity-Media attribute would contain a bodypart hash of the text/plain part, certain SDP attribute lines from the application/sdp bodypart (a=fingerprint in this example), and a bodypart hash of the image/jpg bodypart:

```
Identity-Media: BPH="e32je3j23cjek3dz","a=fingerprint",
  BPH="8fj289r3i892381c"
```

This Identity-Media header, along with the headers and portions of headers described in Section 3.1 are all signed by the authentication service. The resulting signature is placed on the new Identity-Media-Signature header.

3.3. Validation

The validation service can be performed by the remote endpoint itself or by a device acting on behalf of the endpoint. The validation service first checks the signature in the Identity-Media-Signature field. If this is valid, the endpoint (or its validation service operating on its behalf) then initiates a DTLS, TLS, ICE, HIP, or ZRTP identity proof (Section 4). This causes the originating endpoint to prove possession of its private key that corresponds to the certificate (or public key) that was signed by the remote domain's authentication service.

4. Proof of Identity Techniques

Five techniques are described below, TLS, DTLS, ICE, HIP, and ZRTP. Each provides a means to cryptographically prove the identity signed by the authentication service in SIP is the same as the identity on the media path.

Each of these techniques work similarly -- each technique causes unique information to appear in the SDP -- a certificate fingerprint (DTLS, TLS), public key (ICE), or hash (ZRTP). The authentication service creates a new Identity-Media header and places into that header those SDP attribute names associated with that technique. The authentication service then creates a signature over specific SIP headers (see Section 3.1), and places that signature into the new Identity-Media-Signature header. The SIP request is then sent outside of the originating domain.

The receiving domain validates the Identity-Media-Signature. If successful, the SIP request is forwarded to the end system. The end system initiates a TLS, DTLS, ICE, HIP, or ZRTP session and validates that the (signed) certificate fingerprint presented in the SIP signaling matches the certificate presented in the TLS, DTLS, ICE, HIP, or ZRTP exchange. If they match, and the TLS, DTLS, ICE, HIP, or ZRTP exchange completes successfully, the local endpoint has validated the identity of the remote endpoint.

Note: Due to SIP forking, the calling party may receive many identity challenges, each incurring a public key operation to prove identity. Mechanisms to deal with this are for future study.

> Discussion point: It is anticipated that, during the course of standardization, a subset of these five techniques will be chosen as mandatory to implement for the purpose of establishing identity.

4.1. TLS

TLS uses the "fingerprint" attribute to provide a hash of the certificate in the SDP. The fingerprint attribute is defined by [RFC4572] for TLS.

4.2. DTLS

DTLS uses the same "fingerprint" attribute originally described for TLS. The syntax is described in [I-D.ietf-sip-dtls-srtp-framework].

4.2.1. SRTP after DTLS optional

[[Discussion Point: Is there interest in having identity without SRTP??]]

DTLS is only necessary to prove identity with DTLS; SRTP [RFC3711] does not need to be used afterwards. Obviously, using SRTP provides significant benefits over continuing to use RTP, because an attacker can inject bogus RTP after a successful validation of identity which is quite undesirable. The SDP for doing RTP after a DTLS exchange might be signaled in SDP by using "RTP/AVP" rather than "RTP/SAVP" (lines folded for readability):

```
v=0
o=- 25678 753849 IN IP4 192.0.2.1
s=
c=IN IP4 192.0.2.1
t=0 0
m=audio 3456 RTP/AVP 0 18
a=fingerprint:SHA-1
  4A:AD:B9:B1:3F:82:18:3B:54:02:12:DF:3E:5D:49:6B:19:E5:7C:AB
```

Of course, it would be desirable to more clearly indicate this somehow in SDP. The example above collides with non-standard, but deployed, "best-effort" media encryption mechanisms. SDP Capability Negotiation [I-D.ietf-mmusic-sdp-capability-negotiation] might be a useful consideration for this functionality.

4.3. ICE

ICE doesn't have inherent support for public/private keys. If public keys were sent with other ICE attributes, there can be a real risk of an ICE connectivity check exceeding the MTU. ICE lacks a mechanism to fragment such large messages. It is also bandwidth inefficient to send multiple ICE connectivity checks containing public keys, either as retransmissions or with multiple candidates. Thus, for ICE, the public key is sent in SDP and the public key's fingerprint is exchanged on the media path -- opposite of TLS, DTLS, HIP, and ZRTP.

4.3.1. ICE Public Key SDP Attribute

The offerer includes its public key, which it will use for the subsequent PK-CHALLENGE and PK-RESPONSE, in its SDP. The syntax is a BASE64-encoded version of the endpoint's public key.

The new attribute is called "ice-pub-key", which may appear on the session level, media level, or both.

4.3.2. New STUN attributes

Two new STUN attributes are defined to carry the plaintext challenge and the encrypted response.

4.3.2.1. PK-CHALLENGE

This is sent in a STUN Binding Request, and contains two fields: the fingerprint of the public key exchanged in the SDP, and the public key challange. The fingerprint is included so that the remote peer can choose the correct key, in the event it used different public keys. The public key challange field are the bits to be encrypted by the remote peer's private key. Up to 256 bits can be included in the challenge.

The PK-CHALLENGE MUST be the same for each candidate address that is being tested for connectivity. If this requirement is not followed, the peer will incur a public key operation for every ICE connectivity check, which is not reasonable or necessary.

When the remote peer receives a STUN Binding Request containing this attribute, the contents of the PK-CHALLENGE are encrypted using the private key associated with the public key's fingerprint, and the result is sent in the PK-RESPONSE attribute of the Binding Response.

4.3.2.2. PK-RESPONSE

This is sent in a STUN Binding Response from the offerer to the answerer, and contains the encrypted result of the PK-CHALLENGE.

4.4. HIP

In [I-D.tschofenig-hiprg-host-identities], a new attribute "key-mgmt:host-identity-tag" is defined which contains the hash of the public key used in the subsequent HIP exchange. This can be utilized and signed exactly like the "fingerprint" attribute for TLS or DTLS.

4.5. ZRTP

In [I-D.zimmermann-avt-zrtp], a new attribute "zrtp-hello-hash" is defined which contains a hashed value of the ZRTP Hello packet. The entire ZRTP exchange is protected as described in Section 10 of [I-D.zimmermann-avt-zrtp]. After the ZRTP exchange has completed, the remote party's identity is proven to match the identity signed via SIP-Identity-Media.

5. ABNF

The following figure shows the syntax of the new SIP header fields using ABNF [RFC5234]

```
identity-media        = "Identity-Media" HCOLON
                        attrib-bodyhash-list
attrib-bodyhash-list  = attrib-bodyhash *(COMMA attrib-bodyhash)
attrib-bodyhash       = quoted-attrib | quoted-bodyparthash
quoted-attribute      = DQUOTE attribute DQUOTE  ; SDP "a=" line
quoted-bodyhash       = "BPH" EQUAL DQUOTE bodyparthash DQUOTE
bodyparthash          = 32HEXDIG

identity-media-sig    = "Identity-Media-Signature" HCOLON
                        signature
signature             = DQUOT 32HEXDIG DQUOT

Identity-Info = "Identity-Info" HCOLON ident-info
                *( SEMI ident-info-params )
ident-info = LAQUOT absoluteURI RAQUOT
ident-info-params = ident-info-alg / ident-info-extension
ident-info-alg = "alg" EQUAL token
ident-info-extension = generic-param
```

Figure 2: ABNF for new SIP headers

The following figure shows the syntax of the new SDP attribute containing the ICE public key. This is used only by endpoints implementing the ICE proof of identity technique (Section 4.3).

```
ice-pub-key        = token    ; BASE64 encoded public key
```

Figure 3: ABNF for new SDP attribute

6. Security Considerations

[[some of RFC4474's security considerations also apply.]]

6.1. Device Disclosure

Although the mechanism described in this paper allows SBCs to be used with a cryptographic identity scheme, it does expose the identity of the user's certificate. If a unique certificate is installed on each user's device, the remote party will be able to discern which device is terminating the call. This problem is more pronounced when SIP retargeting occurs in conjunction with Connected Identity [RFC4916].

If this isn't desired, there are two solutions:

- All devices under the control of the user will need to have the same certificate (and associated private key) installed on them.

- The device needs to manufacture a new self-signed certificate (or public key) for each call, and populate the appropriate SDP attributes with that certificate (or public key). This is possible because the identity service described in this paper does not require the same certificate or public key to be used on every call.

6.2. Modification of SDP

One issue with only signing specific SDP attributes is that a man in the middle can modify the un-signed SDP for nefarious purposes, beyond simply changing m=/c= lines. In particular, an attacker could set the c= connection line used for DTLS-SRTP fingerprint to 0.0.0.0 and the m= media line to port 0, essentially disabling that offered media session. The attacker could also add a set of c=/m= lines for non-SRTP media, and thus make a non-SRTP offer with a perfectly valid identity signature. Or an attacker could insert SDP capability negotiation attributes to create a best-effort type SRTP offer, with SRTP (rather than RTP) being the lowest preference.

This draft prevents such downgrade attacks by requiring the called UA use DTLS-SRTP, HIP, ICE, or TLS on the media path to establish identity. Thus, an attacker performing the attacks described above will not successfully fool the called UA because the (intended) victim will use DTLS-SRTP (or HIP, ICE, or TLS) on the media path, and the attacker does not possess the private key of the legitimate caller.

7. Operational Differences from RFC4474

RFC4474 imposes one public key operation for the authentication service and one for validation. If Connected Identity [RFC4916] is used, only one additional public key operation is necessary for the header signature validation; the expense of the DTLS, TLS, or ICE public key operation has already been incurred by both parties and is not repeated.

RFC4474 includes the Contact URI in the signed headers. That is not required by this mechanism because it adds no security property, and will fail validation when crossing SBCs and B2BUA's. It is of dubious security value because Via/Record-Route can be inserted for response interception regardless, and some requests don't contain a

Contact anyway (e.g., MESSAGE). It does not provide any replay/ copy-paste protection either, for the same reasons.

RFC4474 includes the CSeq in the signed headers. That is not required by this mechanism because it adds little security, and will fail validation when crossing SBCs and B2BUA's in some cases. It is of little security value because it provides no protection from cut-paste attack for different targets, and although it would prevent replay attack within the same session, since the media key-related SDP portions are signed anyway, replaying the request will not do anything useful.

RFC4474 includes the Call-Id in the signed headers. That is not required by this mechanism because it adds little security, and will fail validation when crossing SBCs and B2BUA's in some cases. It is of little security value because it provides no protection from cut-paste attack for different targets, and although it would prevent replay attack for the same target, since the media key-related SDP portions are signed anyway, replaying the request will not do anything useful.

The mechanism described in this document has the following advantages over RFC4474:

- Only the edge network needs to create signatures on SIP requests -- not every intervening SBC,
- The original cryptographically-provable identity is preserved across any number of SBCs, B2BUA's, etc.
- SBCs, B2BUA's, and other "middle-boxes" in intermediate domains do not need to be upgraded or changed in order for the originating and terminating domains to use this new mechanism.

8. Limitations

For the identity procedure described in this document to function, every device -- including Session Border Controllers -- on the path MUST permit DTLS, TLS, ICE, HIP, or ZRTP messages to be exchanged in the media path. Further, those devices MUST NOT interfere with the signed SDP attributes or the new SIP headers necessary for Identity Media to operate.

For the technique described in this document to function, all on-path SIP elements -- SBCs, B2BUAs, and SIP proxies -- MUST NOT interfere with the signed headers. The identity mechanism described in this document is not harmed if on-path SIP elements alter the SDP (e.g.,

by deleting non-signed attributes, connection addresses, etc.).

9. Examples

9.1. DTLS

This example shows how two a=fingerprint lines in SDP would populate the Identity-Media SIP header field. The following is an example of an INVITE created by the endpoint.

(lines folded for readability)

```
INVITE sip:bob@biloxi.example.org SIP/2.0
Via: SIP/2.0/TLS pc33.atlanta.example.com;branch=z9hG4bKnashds8
To: Bob <sip:bob@biloxi.example.org>
From: Alice <sip:alice@atlanta.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.atlanta.example.com>
Content-Type: application/sdp
Content-Length: 147

v=0
o=- 6418913922105372816 2105372818 IN IP4 192.0.2.1
s=example2
c=IN IP4 192.0.2.1
t=0 0
m=audio 54113 RTP/SAVP 0
a=fingerprint:SHA-1
  4A:AD:B9:B1:3F:82:18:3B:54:02:12:DF:3E:5D:49:6B:19:E5:7C:AB
m=video 54115 RTP/SAVP 0
a=fingerprint:SHA-1
  4A:AD:B9:B1:3F:82:18:3B:54:02:12:DF:3E:5D:49:6B:19:E5:7C:AB
```

Figure 4: Example with DTLS

The SIP proxy performing the Media Identity authentication service would then insert the following three SIP headers into the message. The Identity-Media header contains all of the SDP attribute lines that are signed and the Identity-Media header contains the signature of all of the relevant SIP headers and of the Identity-Media header. Lines are folded for readability:

```
Identity-Info: <https://atlanta.example.com/atlanta.cer>
   ;alg=rsa-sha1
Identity-Media: "a=fingerprint","a=fingerprint"
Identity-Media-Signature:
 "ZYNBbHC00VMZr2kZt6VmCvPonWJMGvQTBDqghoWeLxJfzB2a1pxAr3VgrB0SsSAa
  ifsRdiOPoQZYOy2wrVghuhcsMbHWUSFxI6p6q5TOQXHMmz6uEo3svJsSH49thyGn
  FVcnyaZ++yRlBYYQTLqWzJ+KVhPKbfU/pryhVn9Yc6U="
```

Figure 5: SIP Headers Inserted by Authentication Service

9.2. ICE

With ICE, the public key is exchanged in the signaling path (in SDP) rather than in the media path (as is done with TLS, DTLS, HIP, and ZRTP).

This is the INVITE as it left the SIP user agent (lines folded for readability):

```
INVITE sip:bob@biloxi.example.org SIP/2.0
Via: SIP/2.0/TLS pc33.atlanta.example.com;branch=z9hG4bKnashds8
To: Bob <sip:bob@biloxi.example.org>
From: Alice <sip:alice@atlanta.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.atlanta.example.com>
Content-Type: application/sdp
Content-Length: 147

v=0
o=- 6418913922105372816 2105372818 IN IP4 192.0.2.1
s=example2
c=IN IP4 192.0.2.1
t=0 0
a=ice-pwd:asd88fgpdd777uzjYhagZg
a=ice-ufrag:8hhY
a=pub-key:ejfiwj289ceucuezeceEJFjefkcjeiquiefekureickejfeefe
  uirujejfecejejejkfeJJCEIUQQIEFJCQUCJCEQUURIE09dnjkeefjek
m=audio 54113 RTP/AVP 0
a=candidate:1 1 UDP 2130706431 192.0.2.1 54113 typ host
```

Figure 6: Example with ICE

The SIP proxy performing the Media Identity authentication service would then insert the following three SIP headers into the message. The Identity-Media header contains the ICE public key attribute and the Identity-Media header contains the signature of all of the relevant SIP headers and of the Identity-Media header (lines are folded for readability):

```
Identity-Info: <https://atlanta.example.com/atlanta.cer>
  ;alg=rsa-sha1
Identity-Media: "a=pub-key"
Identity-Media-Signature:
 "jjsRdiOPoQZYOy2wrVghuhcsMbHWUSFxI+p6q5TOQXHMmz6uEo3svJsSH49th8qc
  efQBbHC00VMZr2k+t6VmCvPonWJMGvQTBDqghoWeLxJfzB2a1pxAr3VgrB0Ssjcd
  VcunyaZucyRlBYYQTLqWzJ+KVhPKbfU/pryhVn9Jcqe="
```

Figure 7: Headers Inserted by Authentication Service

9.3. Request without SDP

This example shows how a SIP request without SDP is signed.

Message as sent by the UAC (lines folded for readability)

```
MESSAGE sip:user2@example.com SIP/2.0
Via: SIP/2.0/TCP user1pc.example.com;branch=z9hG4bK776sgdkse
Max-Forwards: 70
From: sip:user1@example.com;tag=49583
To: sip:user2@example.com
Call-ID: asd88asd77a@1.2.3.4
CSeq: 1 MESSAGE
Content-Type: text/plain
Content-Length: 18

Watson, come here.
```

Figure 8: Example with no SDP

The authentication service would add the following headers to the above message:

```
Identity-Info: <https://atlanta.example.com/atlanta.cer>
  ;alg=rsa-sha1
Identity-Media:
  BPH="MZr2k+t6VmCvPonWJMGvQTBDqghoWeLxJfzB2a1pxA"
Identity-Media-Signature:
 "diOPoQZYOy2wrVghuhcsMbHWUSFxI+p6q5TOQXHMmz6uEo3svJsSH49th8qcjjsR
  bHC00VMZr2k+t6efQBVmCvPonWJMGvQTBDqghoWeLxJfzB2a1pxAr3VgrB09JcVc
  unyaZucyRlBYYQTLqWzJ+KVhPKbfU/pryhVnqeSsjcd="
```

Figure 9: added headers

10. Acknowledgements

The mechanism described in this paper is derived from Jon Peterson and Cullen Jennings' [RFC4474], which was formerly a document of the SIP working group.

Thanks to Hans Persson for his suggestions which improved this document.

11. IANA Considerations

This document will add new IANA registrations for its new STUN

attributes.

[[This section will be completed in a later version of this document.]]

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

[I-D.ietf-sip-dtls-srtp-framework] Fischl, J., Tschofenig, H., and E. Rescorla, "Framework for Establishing an SRTP Security Context using DTLS", draft-ietf-sip-dtls-srtp-framework-00 (work in progress), November 2007.

[RFC4572] Lennox, J., "Connection-Oriented Media Transport over the Transport Layer Security (TLS) Protocol in the Session Description Protocol (SDP)", RFC 4572, July 2006.

[RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

[RFC4916] Elwell, J., "Connected Identity in the Session Initiation Protocol (SIP)", RFC 4916, June 2007.

[I-D.tschofenig-hiprg-host-identities] Tschofenig, H., "Interaction between SIP and HIP", draft-tschofenig-hiprg-host-identities-05 (work in progress), June 2007.

[I-D.zimmermann-avt-zrtp] Zimmermann, P., Johnston, A., and J. Callas, "ZRTP: Media Path Key Agreement for Secure RTP", draft-zimmermann-avt-zrtp-05 (work in progress), February 2008.

[I-D.ietf-mmusic-ice]
Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-19 (work in progress), October 2007.

12.2. Informational References

[I-D.ietf-mmusic-sdp-capability-negotiation]
Andreasen, F., "SDP Capability Negotiation", draft-ietf-mmusic-sdp-capability-negotiation-08 (work in progress), December 2007.

[I-D.wing-sip-identity-analysis]
Wing, D. and H. Kaplan, "An Analysis of SIP Identity with SIP Back-to-Back User Agents and Session Border Controllers", draft-wing-sip-identity-analysis-00 (work in progress), January 2008.

Appendix A. ToDo List

- Add Table-2 of RFC3261
- re-use RFC4474 response code for failures, or invent new ones?
- describe what occurs if both SIP-Identity-Media and SIP-Identity are both used?

Appendix B. Changes From Previous Versions

B.1. Changes from 00 to 01

- Removed "Contact" header from signature. SBCs need to change it.
- Removed "Call ID" header from signature. This header often contains an IP address, so many SBCs change it.
- Removed "CSeq" header from signature. This header is sometimes changed by SBCs and B2BUA's.
- include SDP attribute names in Identity-Media signature. This allows any attribute to be signed.
- Old "Identity-Fingerprints" header renamed to "Identity-Media", and only attribute names are listed in it, not attribute values.

- Old "Identity-Media" header renamed to "Identity-Media-Signature".
- Described how to sign SIP requests without an SDP body part, and with a mix of SDP and non-SDP bodyparts.

B.2. Changes from 01 to 02

- Describe how modification of SDP is prevented (section 7.2).
- Moved B2BUA and SBC analysis to separate document, [I-D.wing-sip-identity-analysis].
- Added ZRTP as another authentication technique.

Authors' Addresses

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: dwing@cisco.com

Hadriel Kaplan
Acme Packet
71 Third Ave.
Burlington, MA 01803
USA

Phone:
Fax:
Email: hkaplan@acmepacket.com
URI:

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.