HIPRG H. Tschofenig
Internet-Draft Siemens
Expires: September 7, 2006 J. Ott
Helsinki University of Technology
H. Schulzrinne
Columbia U.
T. Henderson
The Boeing Company
G. Camarillo
Ericsson
March 6, 2006

Interaction between SIP and HIP
draft-tschofenig-hiprg-host-identities-03.txt

Status of this Memo

Copyright Notice

Abstract

This document investigates the interworking between the Session

Initiation Protocol (SIP) and the Host Identity Protocol (HIP) and the benefits that may arise from their combined operation.

The aspect of exchanging Host Identities (or Host Identity Tags) in SIP/SDP for later usage with the Host Identity Protocol Protocol (HIP) is described in more detail as an example of this interworking.

Table of Contents

1. Introduction

SIP [1] enables a pair of user agents to establish and maintain sessions. The communication typically involves SIP proxies before prior to communication between the end points taking place. As part of the initial exchange, a number of parameters are exchanged. Certain of these parameters are relevant to security. Examples of such parameters are keying material and other cryptographic information that is used in order to establish a security association for the protection of subsequent data traffic.

HIP (see [2] and [3]) propose an architecture with a cryptographic namespace and a layer between the network and the transport layer. This layer is used in order to shield applications from the impact of multi-homing, readdressing and mobility. A protocol, called the Host Identity Protocol, is used in order to establish state at the two end hosts. This state includes the establishment of IPsec SAs.

Several areas may benefit from the aforementioned interworking. These include the following.

Mobility:

Mobility support can be provided at different layers in the protocol stack. SIP can offer terminal mobility, as described in [4]. Prior to a call, mobility is handled by re-registration with the home registrar. For mid-call mobility, the moving node sends a re-INVITE directly to the correspondent host, or via the SIP proxies if, during the initial call setup, the proxy had inserted a Record-Route header. Session mobility in SIP is implemented through the usage of the SIP REFER method [9]. A discussion of session mobility with SIP is, for example, provided in [10]. The basic SIP security mechanisms are used in order to protect the signalling exchanges that refer to the above-mentioned terminal and session mobility.

The basic SIP mobility has two main limitations. Firstly, it is unable to move TCP sessions to new IP addresses. This could be accomplished by TCP extensions, such as TCP-Migrate or M-TCP or by the usage of SCTP (where possible). The second limitation is the low speed of handoffs.

One can shield the movement of the end hosts against each other though the usage of HIP. HIP itself, however, does not offer micromobility solutions or mechanisms to deal with the double-movement problem. Extensions have been defined, such as the HIP rendezvous concept [11] or Hi3 [12] that, among other things, deal with initial reachability and provide additional mobility

mechanisms. A later version of this document will also investigate the interworking of SIP with these HIP extensions. In summary, current HIP mobility mechanisms do not offer substantial additional features or security over what SIP provides. This applies especially to the typical scenario where reliable transport protocols are not used in SIP user data flows.

Middlebox Traversal:

The work on traversing Network Address Translators with SIP and media traffic has focused on MIDCOM and the Interactive Connectivity Establishment (ICE) methodology. ICE relies on other protocols, such as STUN [13] and TURN [14] in order to create a NAT binding.

HIP might be better suited for the traversal of HIP-aware NATs, since, in this setting, the NATs can inspect the HIP signaling exchange and create the necessary bindings. This approach is similar to the one proposed by the NSIS working group where a path-coupled signaling protocol is used to interact with these middleboxes to create NAT bindings (and firewall pin-holes). The NATFW-NSLP [15] is a protocol proposal that utilizes the NSIS protocol suite. The travesal of HIP unaware NATs is detailed in [16] and a discussion about NAT and firewall traversal of HIP-aware devices is given in [17].

Denial of Service Prevention:

SIP follows the offer/answer model. The offerer generates an offer that contains, among other things, the IP address the answerer is expected to send its media to. The offer/answer model can be used in order to perform denial of service attacks against third parties. The offerer generates an offer with the IP address of the victim and the answerer, on reception of such offer, starts sending media to the victim. If the session consists of media sent over a connection-oriented transport protocol such as TCP, the victim is unlikely to respond to the connection establishment request (e.g. the initial SYN in TCP) and the connection is not established. However, if the session consists of media sent over RTP and UDP, the sender just floods the victim with RTP packets. The ICMP "not reachable" messages generated by the victim may or may not stop the attack. Firewalls in the path may discard these packets or the RTP library of the sender may ignore them. The use of HIP would prevent this type of attack because the victim would not accept the incoming HIP message. Of course, in order to further address this type of attack no user agent in the network

should accept session descriptions that only provide IP addresses in order to send RTP data. Sessions that did not use HIP would need to use a different method to deal with this attack. An example of such a method consists of using ICE (Interactive Connectivity Establishment) [18] as a return routability test before starting to send media. Other approaches as part of the HIP overlay infrastructure in combination with HIP Registration servers might also provide the same effect or even a higher degree of security.

SRTP and HIP:

The Host Identity Protocol is able to establish IPsec security associations, as described in [19]. In the case of SIP for voice communication, end-to-end media protection using SRTP may be applied. HIP supports a variety of cryptographic protection mechanisms for the data traffic, including IPsec, SRTP. The establishment of the necessary parameters and the keying material for enabling SRTP communication is outlined in a separate document [20].

Exchanging Host Identities with SIP:

HIP can benefit from existing SIP infrastructure because it enables the distribution of Host Identities or Host Identity Tags, as described in Section 3.

2. Terminology

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 [5].

3. Exchanging Host Identities with SIP

3.1. Concept

In order to provide security between two HIP end hosts beyond opportunistic encryption it is necessary to securely retrieve the Host Identities. A number of mechanisms can be used including directories (such as DNS) or more advanced concepts for example based

on Distributed Hash Tables typically used in peer-to-peer networks.

This document suggests to exchange the Host Identities (or Host Identity Tags) as part of the initial SIP exchange inside the SDP payload. As such, the Host Identities can also be bound to the user identities - a concept not used in HIP.

Figure 1 illustrates the main idea:

```
                   +-----------+            +-----------+
            HI/HIT |SIP        |   HI/HIT   |SIP        | HI/HIT
           +------>|Proxy      |<---------->|Proxy      |<------+
           |       |Server X   |   TLS      |Server Y   |       |
           |       +-----------+ (+auth.id.)+-----------+       |
           |                                                    |
           | TLS or                                      TLS or |
           | SIP Digest                              SIP Digest |
           |                                        (+auth.id.) |
           |                                                    |
           v                                                    v
       +-----------+          SIP and HIP                +-----------+
       |SIP        | <---------------------------------> |SIP        |
       |User Agent |          RTP                        |User Agent |
       |Alice      | <=================================> |Bob        |
       +-----------+                                     +-----------+

       Legend:
       <--->: Signaling Traffic
       <===>: Data Traffic
```

Figure 1: SIP Trapezoid

The initial SIP signaling messages between Alice and Bob often take place via the proxy servers. This exchange may be protected with TLS (between SIP proxies but also between SIP UAs and SIP proxies) or with SIP digest authentication between SIP UAs and the outbound proxy. Further SIP security mechanisms should be used in combination with this proposal. The security consideration section, see Section 4, provides a discussion about the possible approaches to secure the Host Identity Tag and to relate it ongoing session.

This allows two hosts to securely exchange keys even if there are

only domain-level public and private keys, as well as secure associations within a domain, thus avoiding the need for a global user-level PKI.

This initial message exchange is used to exchange Host Identities between the end points within the SDP payload.

Subsequently, when both user agents Alice and Bob communicate directly with each other they are able to reuse the Host Identity for the HIP message exchange.

If the SIP communication does not involve third parties (i.e., SIP proxies) and is therefore executed directly between the two SIP UAs then it is not useful to exchange Host Identities in the SDP payloads since the HIP exchange already took place before the first SIP message can be exchanged between the two peers. Still HIP might provide some advantages for the end-to-end communication, such as providing security at the lower layer and mobility and multi-homing support.

The security of this approach relies on two properties:

The signaling messages and the data traffic traverse a different path. Hence, an adversary needs to be located where it is able to see both, the signaling and the the data traffic.

The signaling traffic is often protected.

3.2. SDP Extension

This document proposes to enhance the SDP [6] 'k' or 'a' parameter.

The 'k' parameter has the following structure:

k=<method>:<encryption key>

This document defines two new method fields:

k=host-identity:<HIP Host Identity>

k=host-identity-tag:<hash of the public key>

Alternatively, the 'a' parameter could be used like [7] proposes. An example for MIKEY [21] is given in the reference, which could be reused for HIP. As defined in [22], the 'a' parameter has the following structure:

a=<attribute>:<value>

Similar to the MIKEY example in [7], this document defines two new method fields:

a=key-mgmt:host-identity <HIP Host Identity>

a=key-mgmt:host-identity-tag <hash of the public key>

Both, the Host Identity and the Host Identity Tag are defined in [3]. The Host Identity contains the public key and a number of cryptographic parameters (such as used algorithms and Diffie-Hellmann public parameters). The Host Identity is base64 encoded.

FOR DISCUSSION:

The usage of the k parameter as defined in [8] is deprecated. [6] is more appropriate but like 'k=', they come with the caveat that they require a secured e2e signaling path (or SDP is S/MIME protected). One alternative is the usage of MIKEY for the exchange as defined in [7].

Furthermore, and probably more important, it is important to said what the Host Identity is supposed to be used with. They may help avoiding re-INVITEs when underlying IP addresses change to update the 'Contact:' address as well as the addresses in the 'c=' lines for the various media.

However, multiple devices may take part in the different media sessions (your laptop doing video in parallel to your hardware IP phone). To support these cases, it may be necessary to exchange _several_ HI(T)s within SDP and denote what they shall be used for. Such a mapping could naturally be achieved for each media stream (even using 'k=' attributes); at simple 'a=' attributes (or the mechanisms from [6]/ [7] would be preferred.

SDP only deals with media streams and does not have a notion of user or main device in the background. Hence, the SIP HI(T) may need to go into SIP signaling (rather than be carried in SDP).

Logically, this appears to belong to the 'Contact:' header which may be conveyed protected in an S/MIME body (signed and

encrypted).

3.3. Example

This example contains the full details of the example session setup taken from Section 4 of [1]. The message flow is shown in Figure 1 of [1] and resembles the architecture shown in Figure 1. Note that these flows show the minimum required set of header fields; some other header fields such as Allow and Supported would normally be present.

In our example Alice uses the following Host Identity Tag (7214148E0433AFE2FA2D48003D31172E) and Bob uses (44A5C522D7EDEDF962E55A0677DB1346) as the HIT. These HITs correspond to the following Host Identities (for convenience we reuse the XML representation format used by the Boeing implementation).

```
------
Alice:
------
```


```
<host_identity alg="DSA" alg_id="3" length="128"
 anon="no" incoming="yes">

<name>sip:alice@atlanta.com</name>

<P>D757262C4584C44C211F18BD96E5F061C4F0A423F7FE6B6B85B34CEF72CE14
A0D3A222FE08CECE65BE6C265854889DC1EDBD13EC8B274DA9F75BA26CCB98772
3602787E92BA84421F22C3C89CB9B06FD60FE01941DDD77FE6B12893DA76EEBC1
D128D97F0678D772B5341C8506F358214B16A2FAC4B368950387811C7DA33</P>

<Q>C773218C737EC8EE993B4F2DED30F48EDACE915F</Q>

<G>82269009E14EC474BAF2932E69D3B1F18517AD9594184CCDFCEAE96EC4D5EF
9313384B47093C52B20CD35D02492B3959EC6499625BC4FA5082E22C5B374E16D
D00132CE71020217091AC717B612391C76C1FB2E88317C1BD8171D41ECB83E210
C03CC9B32E81056C21621C73D6DAAC028F4B1585DA7F42519718CC9B09EEF</G>

<PUB>A4666AED5F5E753773DC961EDD0412A03F1F8D7CEC70A057076062804B86
619D3DA4E7610EBBDB05F44C5784622D1B86600DFCC1431BC4451D4FD31329354
```

07A9B24718CB82BAE93A4CDD9CC4C8B9A41C000AB53D52A65E8383F54F5BF92A8
21EA776A207C6991EF23808C00DB820977D97CAC01CB96307274E2386001327
</PUB>

<HIT>7214148E0433AFE2FA2D48003D31172E</HIT>
</host_identity>

Bob:

<host_identity alg="DSA" alg_id="3" length="96"
 anon="no" incoming="yes">

<name>sip:bob@biloxi.com</name>

<P>F13ACC1693AFD04B9E1E8D2A9DEA6DE8DE4C276BE2BF15B6CFF6E269B0169
378CB0DDDE23D187827015DC67E6768193914B823BDF215D0DAD7A151E434F9E
128DAFB9DEFAE07874621E70D7ED2D34B80A95FA8312B9564E4D118FB525664C
77D</P>

<Q>C773218C737EC8EE993B4F2DED30F48EDACE915F</Q>

<G>241F32CF48F424B1A75D33B7AE6088E745D9E24E653AE2CAEBE67E4AA1C11
15BA0CC25055A63C139235A95B36EFBC2064AF304C0F8A431D151B2B5854DE61
5168B45B9EAEBF9A88354CA7876E52D169E14E502BEA0CBB98B55AD2AB61620F
498</G>

<PUB>E481C20D8FBAA84F9C7ED8B5598F60F5A7D03951CA4783841EB8ADDC63D
DE11A2F3555C5641F465160AB1E016756D826B0F8CE4FDE33BA17F6FFFA751DA
1389A10E5599802AB1EBE4FD943405819A74FD6F1C9EA2815EE6B651610DF107
5D19F</PUB>

<HIT>44A5C522D7EDEDF962E55A0677DB1346</HIT>

</host_identity>

F1 INVITE Alice -> atlanta.com proxy

```
INVITE sip:bob@biloxi.com SIP/2.0
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
Max-Forwards: 70
To: Bob <sip:bob@biloxi.com>
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:alice@pc33.atlanta.com>
Content-Type: application/sdp
Content-Length: ...

v=0
o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com
s=Session SDP
t=0 0
c=IN IP4 pc33.atlanta.com
m=audio 3456 RTP/AVP 0 1 3 99
a=rtpmap:0 PCMU/8000
k=host-identity-tag:7214148E0433AFE2FA2D48003D31172E
```

F2 100 Trying atlanta.com proxy -> Alice

```
SIP/2.0 100 Trying
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Content-Length: 0
```


F3 INVITE atlanta.com proxy -> biloxi.com proxy

```
INVITE sip:bob@biloxi.com SIP/2.0
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
  ;branch=z9hG4bK77ef4c2312983.1
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
Max-Forwards: 69
To: Bob <sip:bob@biloxi.com>
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:alice@pc33.atlanta.com>
Content-Type: application/sdp
Content-Length: ...

v=0
o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com
```

```
s=Session SDP
t=0 0
c=IN IP4 pc33.atlanta.com
m=audio 3456 RTP/AVP 0 1 3 99
a=rtpmap:0 PCMU/8000
k=host-identity-tag:7214148E0433AFE2FA2D48003D31172E
```

F4 100 Trying biloxi.com proxy -> atlanta.com proxy

```
SIP/2.0 100 Trying
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Content-Length: 0
```


F5 INVITE biloxi.com proxy -> Bob

```
INVITE sip:bob@192.0.2.4 SIP/2.0
Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
Max-Forwards: 68
To: Bob <sip:bob@biloxi.com>
```

```
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:alice@pc33.atlanta.com>
Content-Type: application/sdp
Content-Length: ...

v=0
o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com
s=Session SDP
t=0 0
c=IN IP4 pc33.atlanta.com
m=audio 3456 RTP/AVP 0 1 3 99
a=rtpmap:0 PCMU/8000
k=host-identity-tag:7214148E0433AFE2FA2D48003D31172E
```

F6 180 Ringing Bob -> biloxi.com proxy

```
SIP/2.0 180 Ringing
Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
 ;received=192.0.2.3
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
Contact: <sip:bob@192.0.2.4>
CSeq: 314159 INVITE
Content-Length: 0
```


F7 180 Ringing biloxi.com proxy -> atlanta.com proxy

```
SIP/2.0 180 Ringing
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
```

```
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
Contact: <sip:bob@192.0.2.4>
CSeq: 314159 INVITE
Content-Length: 0
```

F8 180 Ringing atlanta.com proxy -> Alice

```
SIP/2.0 180 Ringing
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
Contact: <sip:bob@192.0.2.4>
CSeq: 314159 INVITE
Content-Length: 0
```

```
F9 200 OK Bob -> biloxi.com proxy

SIP/2.0 200 OK
Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
 ;received=192.0.2.3
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:bob@192.0.2.4>
Content-Type: application/sdp
Content-Length: ...

v=0
o=bob 2890844527 2890844527 IN IP4 192.0.2.4
s=Session SDP
c=IN IP4 192.0.2.4
t=3034423619 0
m=audio 3456 RTP/AVP 0
a=rtpmap:0 PCMU/8000
k=host-identity-tag:44A5C522D7EDEDF962E55A0677DB1346
```

F10 200 OK biloxi.com proxy -> atlanta.com proxy

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP bigbox3.site3.atlanta.com
 ;branch=z9hG4bK77ef4c2312983.1
 ;received=192.0.2.2
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:bob@192.0.2.4>
Content-Type: application/sdp
Content-Length: ...

v=0
o=bob 2890844527 2890844527 IN IP4 192.0.2.4
s=Session SDP
c=IN IP4 192.0.2.4
t=3034423619 0
m=audio 3456 RTP/AVP 0
a=rtpmap:0 PCMU/8000
k=host-identity-tag:44A5C522D7EDEDF962E55A0677DB1346
```

F11 200 OK atlanta.com proxy -> Alice

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
 ;received=192.0.2.1
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 INVITE
Contact: <sip:bob@192.0.2.4>
Content-Type: application/sdp
Content-Length: ...

v=0
```

```
o=bob 2890844527 2890844527 IN IP4 192.0.2.4
s=Session SDP
c=IN IP4 192.0.2.4
t=3034423619 0
m=audio 3456 RTP/AVP 0
a=rtpmap:0 PCMU/8000
k=host-identity-tag:44A5C522D7EDEDF962E55A0677DB1346
```


F12 ACK Alice -> Bob

```
ACK sip:bob@192.0.2.4 SIP/2.0
Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds9
Max-Forwards: 70
To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
From: Alice <sip:alice@atlanta.com>;tag=1928301774
Call-ID: a84b4c76e66710
CSeq: 314159 ACK
Content-Length: 0
```

The media session between Alice and Bob is now established.

The exchanged HITs are now placed in the pool of known HITs at both end hosts. As such there is also a binding established between URI and HIT at this point.

Next a regular HIP base exchange between Alice and Bob is started. As part of the exchange the two end hosts inspect their known-HITs pool and find the previously exchanged parameters.

```
Alice -> Bob: I1: Trigger exchange

Alice <- Bob: R1: {Puzzle, D-H(R), HI(R), ESP Transform,
              HIP Transform }SIG

Alice -> Bob: I2: {Solution, LSI(I), SPI(I), D-H(I),
              ESP Transform, HIP Transform, {H(I)}SK }SIG

Alice <- Bob: R2: {LSI(R), SPI(R), HMAC}SIG
```

As a result of this exchange, two IPsec SAs (one for each direction)

is established. RTP media traffic can be exchanged between the two end hosts, Alice and Bob, protected by IPsec. If end host mobility takes place then a HIP readdressing exchange takes place which is not detected at the upper layer by UDP/RTP or SIP.

4. Security Considerations

The standard HIP strategy for authenticating the communicating parties is to give the Initiator and the Responder a Host Identity and to assure the authenticity of the Host Identity via external mechanisms, such as DNSSEC (if the Host Identities are stored in the DNS). The Initiator then verifies the Host Identity and checks its validity. The complexity of ensuring that the Host Identity has not been tampered with is pushed to DNS (and DNSSEC), as the only mechanism specified for ensuring that the public key is genuine. The infrastructure provided for SIP can provide a similar, but more deployment friendly, functionality when combined with already available SIP security mechanisms.

The design described in this document is intended to leverage the authenticity of the signaling channel (while not requiring confidentiality). As long as each side of the connection can verify the integrity of the SDP INVITE then the HIP base exchange handshake cannot be hijacked via a man-in-the-middle attack. This integrity protection is easily provided by the caller to the callee via the SIP Identity [23] mechanism. However, it is less straightforward for the responder.

Ideally Alice would want to know that Bob's SDP had not been tampered with and who it was from so that Alice's User Agent could indicate to Alice that there was a secure phone call to Bob. This is known as the

SIP Response Identity problem and is still a topic of ongoing work in the SIP community. When a solution to the SIP Response Identity problem is finalized, it SHOULD be used here. In the meantime there are several approaches that can be used to mitigate this problem: Use UPDATE, Use SIPS, Use S/MIME, and do nothing. Each one is discussed here followed by the security implications of that approach.

4.1. UPDATE

In this approach, Bob sends an answer, then immediately follows up with an UPDATE that includes the Host Identity Tag and uses the SIP Identity mechanism to assert that the message is from Bob's. The downside of this approach is that it requires the extra round trip of the UPDATE. However, it is simple and secure even when the proxies are not trusted.

4.2. SIPS

In this approach, the signaling is protected by TLS from hop to hop. As long as all proxies are trusted, this provides integrity for the Host Identity Tag. It does not provide a strong assertion of who Alice is communicating with. However, as much as the target domain

can be trusted to correctly populate the From header field value, Alice can use that. The security issue with this approach is that if one of the Proxies wished to mount a man-in-the-middle attack, it could convince Alice that she was talking to Bob when really the media was flowing through a man in the middle media relay. However, this attack could not convince Bob that he was taking to Alice.

4.3. S/MIME

RFC 3261 [1] defines a S/MIME security mechanism for SIP that could be used to sign that the fingerprint was from Bob. This would be secure. However, so far there have been no deployments of S/MIME for SIP.

4.4. Single-sided Verification

In this approach, no integrity is provided for the fingerprint from Bob to Alice. In this approach, an attacker that was on the signaling path could tamper with the fingerprint and insert

themselves as a man-in-the-middle on the media. Alice would know that she had a secure call with someone but would not know if it was with Bob or a man-in-the-middle. Bob would know that an attack was happening. The fact that one side can detect this attack means that in most cases where Alice and Bob both wish the communications to be encrypted there is not a problem. Keep in mind that in any of the possible approaches Bob could always reveal the media that was received to anyone. We are making the assumption that Bob also wants secure communications. In this do nothing case, Bob knows the media has not been tampered with or intercepted by a third party and that it is from Alice. Alice knows that she is talking to someone and that whoever that is has probably checked that the media is not being intercepted or tampered with. This approach is certainly less than ideal but very usable for many situations. An alternative available to Alice and Bob is to use human speech to verified each others' identity then verify each others' Host Identity Tags also using human speech. Assuming that it is difficult to impersonate another's speech and seamlessly modify the audio contents of a call, this approach is relatively safe. On the other hand, SIP is not only used for voice communication.

Note that this proposal is closely aligned towards the usage of the 'k' parameter in SDP [8]. As a difference, an asymmetric key is exchanged unlike the proposals illustrated in Section 6 of [8]. Section 5.12 of [22] is relevant for this discussion.

Please note that this approach is in a certain sense a re-instantiation of the Purpose-Built-Key (PBK) idea (see [24]). With PBK a hash of a public key is sent from node A to node B. If there

was no adversary between A and B at that time to modify the transmitted hash value then subsequent communication interactions which use the public key are secure. This proposal reuses the same idea but focuses on the interworking between different protocols. In fact it would be possible to use the same approach to exchange the hash of an S/MIME certificate which can later be used in subsequent SIP signaling message exchanges.

5. IANA Considerations

[Editor's Note: A future version of this document will provide a discussion about IANA considerations.]

6. Open Issues

A list of open issues can be found at:
http://www.tschofenig.com:8080/sip-hip/

7. Contributors

We would like to thank Vesa Torvinen for his contributions to the initial version of this document.

8. Acknowledgments

The authors would like to thank Steffen Fries, Aarthi Nagarajan, Murugaraj Shanmugam, Franz Muenz, Jochen Grimminger and Joachim Kross for their feedback.

The content of the security consideration section is based on DTLS-SIP.

9. References

9.1. 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] Moskowitz, R. and P. Nikander, "Host Identity Protocol Architecture", draft-ietf-hip-arch-03 (work in progress), August 2005.

[3] Moskowitz, R., "Host Identity Protocol", draft-ietf-hip-base-05 (work in progress), March 2006.

[4] Schulzrinne, H. and E. Wedlund, "Application-Layer Mobility using SIP, ACM MC2R",   , July 2000.

[5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

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

[7] Arkko, J., "Key Management Extensions for Session Description Protocol (SDP) and Real Time Streaming Protocol (RTSP)",

draft-ietf-mmusic-kmgmt-ext-15 (work in progress), June 2005.

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

9.2. Informative References

[9] Sparks, R., "The Session Initiation Protocol (SIP) Refer Method", RFC 3515, April 2003.

[10] Shacham, R., "Session Initiation Protocol (SIP) Session Mobility", draft-shacham-sipping-session-mobility-01 (work in progress), July 2005.

[11] Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", draft-ietf-hip-rvs-04 (work in progress), October 2005.

[12] Nikander, P., "Host Identity Indirection Infrastructure (Hi3)", draft-nikander-hiprg-hi3-00 (work in progress), June 2004.

[13] Rosenberg, J., "Simple Traversal of UDP Through Network Address Translators (NAT) (STUN)", draft-ietf-behave-rfc3489bis-02 (work in progress), July 2005.

[14] Rosenberg, J., "Traversal Using Relay NAT (TURN)", draft-rosenberg-midcom-turn-08 (work in progress), September 2005.

[15] Stiemerling, M., "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", draft-ietf-nsis-nslp-natfw-09 (work in progress), February 2006.

[16] Stiemerling, M., "Middlebox Traversal Issues of Host Identity Protocol (HIP) Communication", draft-stiemerling-hip-nat-05 (work in progress), July 2005.

[17] Tschofenig, H. and M. Shanmugam, "Traversing HIP-aware NATs and Firewalls: Problem Statement and Requirements", draft-tschofenig-hiprg-hip-natfw-traversal-03 (work in progress), October 2005.

[18] 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.

[19] Jokela, P., "Using ESP transport format with HIP", draft-ietf-hip-esp-02 (work in progress), March 2006.

[20] Tschofenig, H., "Using SRTP transport format with HIP", draft-tschofenig-hiprg-hip-srtp-01 (work in progress), October 2005.

[21] Arkko, J., Carrara, E., Lindholm, F., Naslund, M., and K. Norrman, "MIKEY: Multimedia Internet KEYing", RFC 3830, August 2004.

[22] Handley, M., "SDP: Session Description Protocol", draft-ietf-mmusic-sdp-new-26 (work in progress), January 2006.

[23] 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.

[24] Bradner, S., Mankin, A., and J. Schiller, "A Framework for Purpose-Built Keys (PBK)", draft-bradner-pbk-frame-06 (work in progress), June 2003.

[25] Jennings, C. and J. Peterson, "Certificate Management Service for The Session Initiation Protocol (SIP)", draft-ietf-sipping-certs-02 (work in progress), July 2005.

Authors' Addresses

Hannes Tschofenig
Siemens
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

Email: Hannes.Tschofenig@siemens.com

Joerg Ott
Helsinki University of Technology
Otakaari 5A
Espoo FI-02150
Finland

Email: jo@netlab.hut.fi

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
USA

Phone: +1 212 939 7042
Email: schulzrinne@cs.columbia.edu
URI: http://www.cs.columbia.edu/~hgs

Thomas R. Henderson
The Boeing Company
P.O. Box 3707
Seattle, WA
USA

Email: thomas.r.henderson@boeing.com

Gonzalo Camarillo
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: Gonzalo.Camarillo@ericsson.com

Intellectual Property Statement

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.

Disclaimer of Validity

Copyright Statement

Acknowledgment

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