HIP Working Group P. Nikander
Internet-Draft G. Camarillo
Intended status: Experimental J. Melen
Expires: January 31, 2009 Ericsson
July 30, 2008

HIP (Host Identity Protocol) Immediate Carriage and Conveyance of Upper-layer Protocol Signalling (HICCUPS)
draft-nikander-hip-hiccups-00.txt

Status of this Memo

Copyright Notice

Abstract

This memo defines how one can use HIP packet formats, and optionally the HIP base exchange, to securely convey arbitrary signalling messages over the Internet or various overlay networks.

Table of Contents

1. Introduction

There has recently been discussion at the IETF on how to design and route new signalling protocols. Typical to these discussions are that the requirements for supporting mobility, multi-homing, security, NAT traversal, or overlay routing go beyond of what is currently possible with plain IP, UDP, or TCP.

In this memo we briefly outline how the Host Identity Protocol (HIP) can be used, either in parts or as a whole, to convey signalling messages when the above mentioned properties are of paramount value.

2. Background

The HIP protocol defines a number of messages and parameters [RFC5201]. The parameters are encoded as TLVs, as shown in Section 2.1.2. Furthermore, the HIP header carries a Next Header field, allowing other arbitrary packets to be carried within HIP packets.

2.1. Message formats

2.1.1. HIP fixed header

The HIP packet format consists of a fixed header, followed by a variable number of parameters. The parameter format is described in Section 2.1.2, below.

The fixed header is defined in Section 5.1 of [RFC5201], and copied below.

```
      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     | Next Header   | Header Length |0| Packet Type |  VER. | RES.|1|
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |          Checksum             |           Controls            |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                Sender's Host Identity Tag (HIT)               |
     |                                                               |
     |                                                               |
     |                                                               |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |               Receiver's Host Identity Tag (HIT)              |
     |                                                               |
     |                                                               |
     |                                                               |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                                                               |
     /                        HIP Parameters                         /
     /                                                               /
     |                                                               |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

2.1.2. HIP parameter format

The HIP parameter format is defined in Section 5.2.1 of [RFC5201], and copied below.

```
      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |             Type            |C|             Length            |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                                                               |
     /                          Contents                             /
     /                                               +-+-+-+-+-+-+-+-+
     |                                               |    Padding    |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

     Type         Type code for the parameter
       C          Critical bit, part of the Type.
     Length       Length of the parameter, in bytes.
     Contents     Parameter specific, defined by Type.
     Padding      Padding, 0-7 bytes, added if needed.
```

2.2. HIP base exchange, updates, and state removal

The HIP base exchange is a four message half-stateless authentication and key exchange protocol that creates shared, mutually authenticated keying material at the communicating parties. These keying materials, together with associated public keys and IP addresses, form a HIP Security Association (SA). The details of the protocol are defined in the HIP base exchange specification [RFC5201].

In addition to creating the HIP SA, the base exchange message may carry additional parameters that are used to create additional state. For example, the HIP ESP specification [RFC5202] defines how HIP can be used to create end-to-end, host-to-host IPsec ESP Security Associations, used to carry data packets. However, it is important to understand that the HIP base exchange is by no means bound to IPsec; using IPsec ESP to carry data traffic forms just a baseline and ensures interoperability between initial HIP implementations.

Once there is a HIP SA between two HIP-enabled hosts, they can exchange further HIP control messages. Typically, UPDATE messages are used. The contents of UPDATE messages is completely open; for example, the HIP mobility and multi-homing specification [RFC5206] defines how to use UPDATE messages to change the set of IP addresses associated with a HIP SA.

In addition to the base exchange and updates, the HIP base protocol specification also defines how one can remove a HIP SA once it is no longer needed.

2.3. Basic ways to extend HIP

As a protocol, HIP has been defined as a secure, extensible protocol that can be used for any kinds of host-to-host negotiations. Since HIP packets can carry additional payloads, it can also be used to carry upper layer, application specific signalling negotiations. However, as the HIP packets by default always carry a digital signature in order to facilitate third party packet authentication, they are somewhat expensive to produce and therefore typically not suitable for bulk data traffic.

The protocol contains the following basic extension mechanisms:

o The currently defined Host Identity value in HIP is a sole public key. However, as explained in the architecture specification [RFC4423], in theory the Host Identity can also consist of some other data. In practise, the public key can be extended with additional identifying data or alternative identifiers.

- To facilitate the HIP protocol machinery, each HIP packet carries an 8-bit packet type. Currently only a few of these packet types are used. Consequently, for extensions that require more states at the HIP base exchange and protocol level, the best way to extend is to define new packet types.
- The fixed header carries a 16-bit Controls field, which can be used to introduce new base level features that are orthogonal to the protocol state machine.
- Each HIP packet can carry zero or more parameters. Each parameter type is identified with a 16-bit Type value. As only few of these are defined, perhaps the generally best way to extend HIP is to define new parameter types and define what kind of HIP packets may be used to exchange them. As a part of this, many of the existing parameter values can be used to help defining new extensions, see below.
- The Next Header field in the fixed header allows a HIP packet to carry arbitary data; for example, simple SIP messages may be exchange over HIP in this way.
- The HIP registration extension [RFC5203] defines a generic protocol that can be used to announce availability of HIP based services and to register as a user to such a service. The extension has itself been designed to be extensible, allowing it to be used for announcing and using different services.
- The SEQ and ACK parameters allow several request-reply pairs to be reliably and parallelly exchanged over a single HIP SA.
- The SIG and HMAC parameters allow HIP-based message exchanges to be authenticated.
- The ENCRYPTED parameter allows any HIP parameters to be optionally encrypted.
- The CERT parameter allows HIP peers to exchange certificates.

2.4. Present limitations to extendability

As HIP extensions are a relatively unexplored area, there may still be features in the HIP protocol that make extensions harder than necessary. The author is presently aware of the following limitations:

- HIP itself does not support fragmentation but relies on underlying IP-layer fragmentation. This may lead to reliability problems in the case where a message cannot be easily split over multiple HIP messages.
- HIP currently requires always that the four-message base exchange is executed at the first encounter of hosts that have not communicated before. This may add an additional round trip time to protocols based on a single message exchange. However, the four-message exchange is essential to preserve the half-stateless, DoS protection nature of the base exchange; see Section 4.

- HIP currently requires that all messages (but I1) are digitally signed. This adds to the packet size and the processing capacity needed to send packets. However, in applications where security is not paramount, it is possible to use very short keys, thereby reducing resource consumption.

2.5. Mobility, multi-homing, and NAT traversal

The HIP mobility and multi-homing specification [RFC5206] defines how one can move the end-points of an existing HIP association from one IP address to another (due to e.g. host mobility) or to associate multiple IP addresses with an end-point (e.g. to help with multi-homing or NAT traversal).

2.6. Routing HIP packets

Each HIP packet carries two identifiers: the Host Identity Tag (HIT) of the sender and that of the receiver. The HITs are 128-bit long entities, consisting of a fixed prefix as defined in [RFC4843], and a 100-bits long hash of an upper-layer Host Identifier value.

In the base Internet, HIP packets are routed as any IP traffic, based on the IP addresses in the IP header preceeding the HIP header.

When more flexible routing constructions are needed, such as for overlay networks, it is possible to create and maintain forwarding state based on the HITs. For one particular example of how this can be done, one can consider the Host Identity Indirection Infrastructure (Hi3) proposal [paper-hi3], which basically combines HIP with the Internet Indirection Infrastructure (i3) [paper-i3]. Another example if the HIP BONE framework [I-D.camarillo-hip-bone].

3. Using HIP to carry signalling protocol messages

Above we have briefly described the basic facilities provided by HIP and succintly explained various options to expand it. In this section we discuss, in general terms, how one can use the HIP extension capabilities to use HIP, either in whole or in parts, to facilitate signalling message exchange.

We start with a few brief examples, and then continue to some more generic observations, and finally outline potential benefits and drawbacks that may stem from using HIP to carry signalling messages.

3.1. Examples

The SHIM6 protocol [I-D.ietf-shim6-proto] uses the same packet format and parameter formats as HIP does. The protocols have been carefully designed to be compatible, so that it should be very easy to adopt features from one protocol to another. Furthermore, most early SHIM6 implementations are based on existing open-source HIP implementations, basically borrowing the underlying implementation architecture.

The Lightweight HIP [I-D.heer-hip-lhip] proposal specifies a new security model for HIP, using hash chains instead of public keys. Other than that, the proposal preserves HIP semantics and packet formats, and is fully compatible with HIP, thereby providing a different way of securing HIP-based mobility, multi-homing, NAT-traversal, registration, etc.

3.2. Observations

Based on the argumentation and examples above, our thinking can be summarised into the following observations:

- The HIP base protocol [RFC5201], with the basic extensions [RFC5206][RFC5203][RFC5204], offers a well-researched, experimental protocol that provides reasonable DoS resistence, public-key-based mutual authentication, host mobility and multi-homing with inherent route-optimisation and multi-home-agent support.
- The HIP mobility and multi-homing work across IPv4 and IPv6, thereby providing good IP version transition support for any protocols that utilise HIP.
- The HIP base protocol is suitable for low-volume, highly secure signalling-type traffic where interaction with middle boxes is important. The main reasons for these characteristics are that all packets contain a public-key signature, designed to be verifiable by middle boxes.
- The HIP base protocol is not suitable for high-volume data traffic. Instead, it is recommended that an extension is used to establish separate security associations for data protocols. Currently the only such extension is the ESP extension [RFC5202].
- The HIP base protocol has been designed to be extensible with different methods, as described in Section 2.3.
- The HIP packets are identified by the source and destination HITs, which are essentially hashes over some other identifiers (typically public keys). This allows the HIP packets to be routed on the bases of these identifiers, as long as the routing system supports routing on flat names.

- The logical location of HIP directly at the top of IPv4 and IPv6, together with its ability to simultaneously work on both, combined with mobility, multi-homing, and NAT-traversal functions, provides a good bases for universal connectivity in the current Internet, independent of the applications.
- Architecturally, any signalling protocol whose purpose is to control data traffic that flows over IPv4 and IPv6 can be converted to run on the top of HIP, while simutaneously either continuing the data traffic completely unmodified or converting it to run on the top of some security protocol, such as IPsec, SRTP, or perhaps even TLS. While doing the protocol conversion, the signalling protocol may benefit from the DoS resistance, security, mobility, multi-homing, IPv4/IPv6 interoperability, and NAT-traversal features of HIP.

3.3. Main benefits and drawbacks

In this section we list the main features of HIP that may be beneficial or harmful, depending on the point of view.

- Packet routability on flat, hash-based names.
- Strong authentication, visible to third parties.
- Creation of keying material, available for other protocols.
- Support for host mobility, over IPv4 and IPv6.
- Support for host multi-homing, over IPv4 and IPv6.
- Good compatibility with legacy IPv4 and IPv6 applications.
- Extensibility.

4. Security considerations

The HIP security model is based on the assumption that the peer hosts (or applications running on them) have secure access to their peer's public keys. How this access is established falls beyond the scope of HIP, and may be arranged, for example, opportunistically, using leap-of-faith, using extenal key distribution, or using third parties and certificates.

The HIT security model is based on the assumption that the used hash algorithm (currently SHA-1) is secure against second preimage attacks, thereby providing assurance that a given HIT corresponds to a given public key. In practical terms, this means that whenever an application has securely acquired a HIT and is using the HIT to name the peer, the underlying HIP machinery makes sure that all communications takes place with the entity denoted by the HIT, or does not place at all.

One two HIP hosts have access to their peer's public keys and know at least one currently reachable IP address of the peer or the peer's

rendezvous server, the HIP hosts can establish a HIP Security Association. The association formation is carried by the HIP base exchange protocol, based on the SIGMA family of cryptographic key exchange protocols. The protocol contains methods to mitigate some types of CPU and state exhausting denial-of-service attacks.

Currently, the HIP base exchange protocol is the simplest known protocol that provides the level of authentication, key formation, integrity protection, and DoS resistance that the protocol provides. Furthermore, there are strong reasons to believe that it is not possible to design significantly simpler protocols that accomplish the same characteristics.

The HIP mobility and multi-homing extension creates a secure, dynamic, one-to-many binding between the peer's host identity and the IP addresses through which the peer is currently reachable at. Security is based on public key signaturs, HMACs based on session keys, return routability, and credit based authorisation.

The native HIP NAT traversal proposal (SPINAT) provides a secure, stateful address translation facility between addressing domains. The legacy HIP NAT traversal proposal is vulnerable to same kind of session stealing attacks as plain NAT or STUN and TURN; however, since the signalling protocol itself is secure and since it is possible to use secure data transfer protocols (such as ESP), the only result of such session stealing is a short period of denial-of-service, until the HIP multi-homing facility manages to create new connectivity.

For some applications, the HIP security model can be replaced by the Lightweight HIP (LHIP) security model, which is based on opportunistic hash chains. For the security properties of this alternative, see the LHIP specification [I-D.heer-hip-lhip].

5. IANA considerations

This memo defines no IANA actions.

6. Acknowledgments

TBD.

7. Informative references

[RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol

(HIP) Architecture", RFC 4423, May 2006.

[RFC4843] Nikander, P., Laganier, J., and F. Dupont, "An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers (ORCHID)", RFC 4843, April 2007.

[RFC5201] Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol", RFC 5201, April 2008.

[RFC5202] Jokela, P., Moskowitz, R., and P. Nikander, "Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP)", RFC 5202, April 2008.

[RFC5206] Nikander, P., Henderson, T., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol", RFC 5206, April 2008.

[RFC5203] Laganier, J., Koponen, T., and L. Eggert, "Host Identity Protocol (HIP) Registration Extension", RFC 5203, April 2008.

[RFC5204] Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", RFC 5204, April 2008.

[I-D.ietf-shim6-proto]
Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", draft-ietf-shim6-proto-10 (work in progress), February 2008.

[I-D.heer-hip-lhip]
Heer, T., "LHIP Lightweight Authentication Extension for HIP", draft-heer-hip-lhip-00 (work in progress), March 2007.

[I-D.camarillo-hip-bone]
Camarillo, G., Nikander, P., and J. Hautakorpi, "HIP BONE: Host Identity Protocol (HIP) Based Overlay Networking Environment", draft-camarillo-hip-bone-00 (work in progress), December 2007.

[paper-hi3]
Nikander, P., Arkko, J., and B. Ohlman, "Host Identity Indirection Infrastructure (Hi3)", 2004.

[paper-i3]
Stoica, I., Adkins, D., Zhuang, S., Shenker, S., and S. Surana, "Internet Indirection Infrastructure (i3)", 2002.

Authors' Addresses

Pekka Nikander
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: Pekka.Nikander@ericsson.com

Gonzalo Camarillo
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: Gonzalo.Camarillo@ericsson.com

Jan Melen
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: Jan.Melen@ericsson.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).