

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: July 14, 2012

D. Zhang
Huawei
M. Komu
January 11, 2012

An Extension of HIP Base Exchange to Support Identity Privacy
draft-zhang-hip-privacy-protection-04

Abstract

In this document, an extension of HIP Base Exchange (BEX) is proposed protect the identity privacy of HIP hosts. Apart from describing the protocol and packet formats, the applicability and the security strength of the proposed approach are analyzed. This work is based on BLIND [YLI04].

Requirements Language

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 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on July 14, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Overview of the Protocol	4
4. Protocol Description	4
4.1. Base Exchange Extensions	5
4.1.1. Blind Initiator and Responder	5
4.1.2. Blind Initiator	6
4.1.3. Blind Responder	8
4.2. UPDATE Extensions	8
4.3. CLOSE	9
5. Packet Formats	9
5.1. Control Header Flags	9
6. Applicability Consideration	9
6.1. Opportunistic Base Exchange	9
6.2. Comparison of Disposable Identities vs. Blind	10
6.3. HIP-based Middleboxes	10
6.4. Immediate Carriage and Conveyance of Upper-layer Protocol	11
7. Security Considerations	11
8. Contributors	11
9. Acknowledgements	11
10. References	11
10.1. Normative References	11
10.2. Informative References	12
Authors' Addresses	12

1. Introduction

The Host Identity Protocol (HIP) [RFC5201] was proposed as a complete solution to address multiple critical issues (e.g., mobility, multi-homing, and security) which the current Internet infrastructure suffers from. In order to achieve this objective, HIP separates the semantics of host identifier from IP addresses by intercepting an "ID" layer in the middle of the network layer and the transport layer. Compared to other ID/Locator separating solutions (e.g., LISP, GSE, ILNP, etc.), HIP is security-inherent. Each HIP host has a public key pair; the public key is used as the Host Identifier (HI) transported over the Internet while the private key is maintained locally.

Additionally, a HIP host also needs to generate a 128-bits long Host Identity Tag (HIT) by hashing its HI. HITs are transported in the common parts of HIP headers and regarded by upper layer protocols (e.g., TCP) as ordinary IPv6 addresses. Before two HIP hosts communicate with each other, they use the HIP Base Exchange protocol (HIP BEX) to verify each other's identity and create shared keying material for subsequent communications. Normally, the HIT and HI of a host are much steadier than its locator (i.e., IP address). Therefore, the changes in the location of a host will not be detected by upper layer protocols.

In the current version of HIP BEX, the identities (i.e., HITs and HIs) of communicating partners are transported in plaintexts. This caused an identity privacy issue. In many scenarios, a user may want to keep its identity confidential to other unrelated entities. However, by eavesdropping HIP BEXs, it is possible for a third party to identify a HIP host even when the host is attached to different locations in the network. As a consequence, it is easier for an attacker to combine the host's activities to reason additional useful information.

The identity privacy issue mentioned here is closely related with the location privacy issues. As illustrated in [RFC4882], the movement of a mobile host can be detected if any constant information related with the host is detected by a third party. Such information can be a Security Parameter Index (SPI) in an IPsec [RFC4301] header, an Interface Identifier (IID) [RFC2462] in an IPv6 address that remains unchanged across networks, the home address of a host supporting mobile IP, the MAC address of a mobile host, an identifier of a mobile host adopted in a upper layer protocol, and so on. Therefore, in order to protect the privacy of a mobile HIP host, a comprehensive solution which cover multiple layers must be provided, and the identity privacy is one of the most important issues which need to be considered in such a solution. In the current HIP BEX, HIs and HITs

are the only permanent information transported in plaintext in different HIP BEXs. Although SPIs are also transported in plaintext, the valid period of a SPI is no longer than the associated IPsec SA. Therefore, without tracing the permanent identity of a host, SPIs mean much less for attackers.

Instead of attempting to address the overall privacy problem, the solution proposed in this document only addresses the identity privacy issue, that is, the solution provides protection against the attempts to track HIP hosts by inspecting the HITs and HIs transported in HIP BEXs.

2. Terminology

BEX: Base Exchange

HIP: Host Identity Protocol

HI: Host Identifier

HIT: Host Identity Tag

3. Overview of the Protocol

The proposed solution is an extension of BLIND, a framework for protecting the identity privacy of hosts that are identified with their public keys. In our solution, if an initiator of a HIP base exchange intends to protect its identity privacy, it will not transport its HI and HIT in plaintexts over the network. Instead, it generates a scramble HIT for itself. The scramble HIT is called a blinded HIT in this document. If the initiator intends to protect the identity privacy of its communicating partner, it also needs to generate a blinded HIT for the partner as well. A blinded HIT is generated by hashing the concatenation of a nonce and the associated HIT.

4. Protocol Description

In order to benefit the discussion, assume there is a HIP host called Initiator which intends to communicate with a HIP host called Responder. The HITs of Initiator are referred to as HIT-Is, and the HITs of Responder are referred to as HIT-Rs. Additionally, the blinded HITs of Initiator and Responder are referred to as B-HIT-Is and B-HIT-Rs respectively. In the discussion of this section, it is assumed that Initiator has got a HIT-R through an out-of-band method

before initiating a BEX. Otherwise, it needs to communicate with Responder in an opportunistic mode. The related issues with the opportunistic mode are discussed in section 6.1. Additionally, in this work, Initiator and Responder do not need to know each other's HI beforehand. Such information will be transported in the BEX in an encrypted way.

4.1. Base Exchange Extensions

In order to distinguish the proposed approach from the ordinary BEX protocol, two control header bits, I and R are introduced. In the HIP packet, I indicates that the HI and HIT of the initiator of an HIP BEX transported in the HIP header are scrambled, and R indicates that the HI and HIT of the responder of an HIP BEX transported in the HIP header are scrambled.

4.1.1. Blind Initiator and Responder

In the scenarios where the identity privacy of both communicating partners needs to be protected, Initiator needs to generate a blind HIT (i.e., B-HIT-I) for itself and a blind HIT (i.e., B-HIT-R) for Responder before initiating a BEX.

In order to achieve this, Initiator first selects a random number nonce, N. Then, Initiator generates a B-HIT-I by SHA-1 hashing the concatenation of N and HIT-I, that is, $B-HIT-I = SHA-1(N, HIT-I)$. In the same way, Initiator generates a B-HIT-R for Responder. $B-HIT-R = SHA-1(N, HIT-R)$.

An extended BEX handshake is illustrated in Figure 1. In the I1 packet transported in the step 1, the blinded HITs of Initiator and Responder, and the nonce, N, are transported in plaintexts.

After receiving the I1 packet, Responder needs to calculate an SHA-1 hash value from the concatenation of N and HIT-R, and compare it with the B-HIT-R transported in the I1 packet. Note that if the Responder has multiple HITs, this process may need to be performed repeatedly. If there is no hash identical to the B-HIT-R, I1 packet will be discarded. If a hash value matches the B-HIT-R, Responder sends a pre-generated R1 packet of the associated HI to Initiator (see step 2 in Figure 1). In order to avoid Deny of Service attacks, Responder should not maintain any state information at this step. However, in practice, although it is not recommended, Responder can also select to maintain the mapping from the pseudonym and the associated HI in order to simplify the process of the I2 packet.

If Initiator has already got the HI of Responder, it can use the HI to assess the validity of the signature of the R1 packet. Otherwise,

Initiator cannot verify the signature of the R1 packet until it gains the HI from the R2 packet. No matter whether Initiator can assess the signature, Initiator generates a symmetric key, KDH, using the Diffie-Hellman algorithm and adopts the key to calculate the keying material with the HIT-R and B-HIT-I:

Key1=SHA1 (KDH, HIT-R, B-HIT-I, 1), ...

Keyn=SHA1 (KDH, HIT-R, B-HIT-I, n),

Keying material=Key1 XOR ... XOR Keyn.

The keying material is then used to generate a symmetric key to encrypt the HI of Initiator. The encrypted information is sent to Responder in an I2 packet (see step 3). Additionally, the I2 packet also contains the nonce which was transported in I1 and the pseudonym which was transported in R1. After receiving the I2 packet, Responder computes a key in the same way as the Initiator. Using the key, Responder decrypts the HIT and HI information of Initiator, and verifies the correctness of B-HIT-I. Moreover, Responder encrypts its HI using the obtained symmetric key and transports the encrypted HI to Initiator in a R2 packet (see step 4). After Initiator receives the R2 packets, it decrypts the Responder's HI. If Initiator does not know the HI of Responder, Initiator can assess its correctness against HIT-R and then verify the validity of the R1 packet.

Finally, the whole BEX handshake completes successfully. In the packets transported in the exchange, both R and S are set.

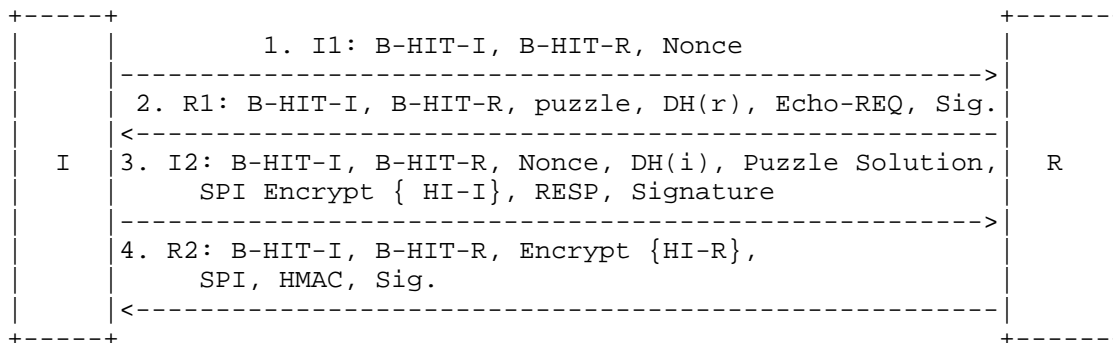


Figure 1. An Extended HIP Base Exchange

4.1.2. Blind Initiator

In the many circumstances, only the identity privacy of Initiator needs to be protected. For instance, a client may want to keep its identity untraceable to any third party while the server which the

client tries to communicate with intends to eliminate the overhead introduced by encrypting its HIT and HI. In this case, the client only needs to generate a blind HIT for itself before initiating a BEX. The process of generating B-HIT-I is as same as what is illustrated in section 4.1.1. Then Initiator sends an I1 packet to Responder (see step 1 in Figure 2). The packet consists of B-HIT-I, the actual HIT of Responder (HIT-R), and the nonce used in generating B-HIT-I. After receiving I1, Responder sends back Initiator a R1 packet (see step 2). The process of generating the R1 packet is identical to the ordinary BEX. Upon receiving R1, Initiator verifies the validity of R1 and calculates the keying material in the same way illustrated in section 4.1.1. In addition, Initiator encrypts its HI using a key derived from the keying material and transports the encrypted information in I2 (see step 3). Therefore, after receiving I2, Responder can compute a symmetric key to verify the correctness of B-HIT-I. The symmetric key is also used to calculate the HMAC of the R2 packet. Therefore, by verifying the HMAC of R2, Initiator can prove that Responder has shared a symmetric key with it. In this exchange, only the control flag I is set.

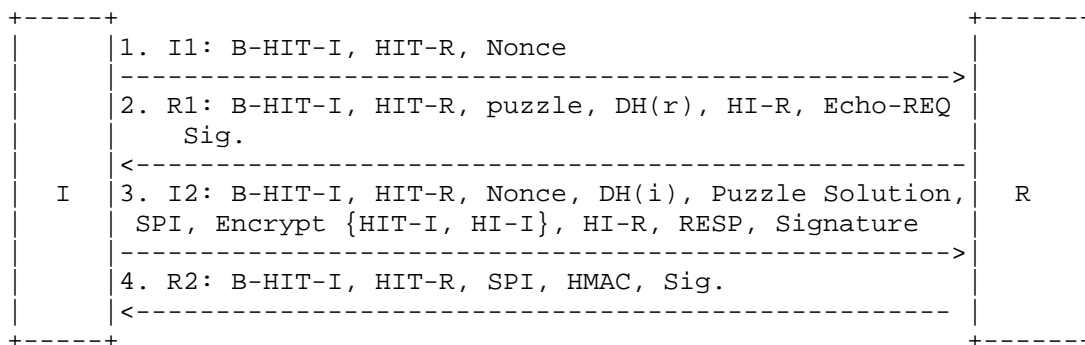


Figure 2. An Extended HIP Base Exchange

Typically, an Initiator can decide whether to protect its identity privacy according to its local policy. Before initiating a HIP exchange, the Initiator can also attempt to learn whether the responder intends to protect its identity privacy (e.g., from resolution systems or referrers). However, if there is no trustable method for an initiator to learn the privacy protecting policies of its communicating partner in advance, the initiator should carry out BEX in the way described in section 4.1.1. Otherwise, if the initiator uncovers the HIT or HI of Responder in I1, there is no opportunity left for the responder to decide whether to protect its identity privacy. When receiving such an I1 packet, the responder can select to 1) drop the packet directly if it does not support the extended HIP BEX, 2) carry out the handshake with the initiator in the way illustrated in the above section if it supports the extended

HIP BEX and prefers to protect its identity privacy, or 3) send a notify back if it supports the extended HIP BEX and does not intend to protect its identity privacy. In the third case, the notify packet is used to illustrate its identity privacy protecting policies. In the notify packet, the responder can proactively disclose its HIT. Therefore, after receiving the notify packet, the initiator can decide whether to restart a new HIP BEX.

4.1.3. Blind Responder

In the circumstance where an initiator which does not intend to protect its identity privacy attempts to contact another host which intends to protect its identity privacy, the initiator can only scramble the HIT of the responder and transport its own HIT in plaintext.

Figure 3 illustrates such a HIP BEX between Initiator and Responder. In the first step, Initiator sends HIT-I, B-HIT-R, and the nonce used to generate B-HIT-R in R1 to Responder. After receiving I1, Responder tries to find out the associated HI using the method indicated in section 4.1.1. If the HI is found, Responder then sends a pre-generated R1 packet of the associated HI back to Initiator (see step 2). In R1, B-HIT-R is encapsulated. After calculating the puzzle, Initiator sends an I2 packet back to Responder (see step 3). In step 4, responder sends its HI in the encapsulated part of R2 packet to Initiator.

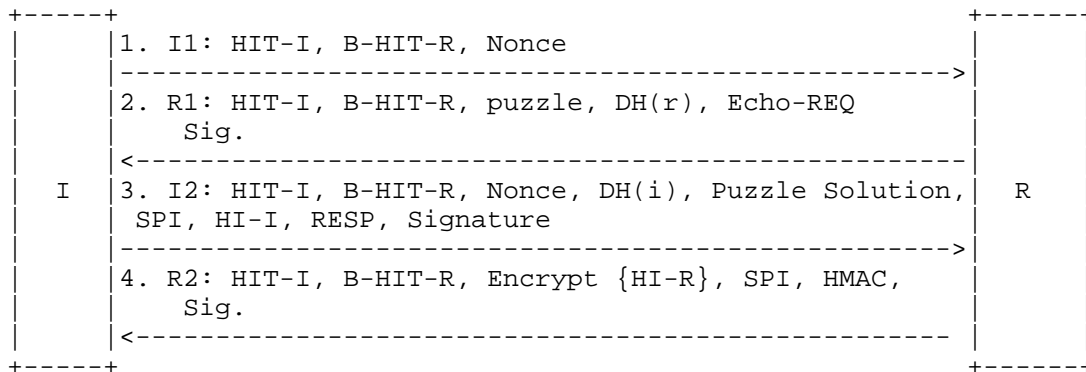


Figure 3. An Extended HIP Base Exchange

4.2. UPDATE Extensions

After two hosts achieve a HIP BEX, they may also need to change their SPIs or HITs for certain reasons. Such information can be transported in update packets. Because the confidentiality of SPIs and HITs needs to be protected, such information should be

transported in an encrypted way.

Additionally, according to different security requirement, the hosts changing its IP address also needs to select a new nonce to generate new scramble HIT(s). The nonce and scrambled are used in the HIP header. The associated flags are set as well. For instance, if the identity privacy of both communicating parties needs to be protected, a new pair of scramble HITs need to be generated. The new pair of scrambled HITs and the nonce are transported within the packet header. After receiving the packet, the receiver needs to first find out the associate private HITs and then locate the proper keys to verify the signature and decrypt the SPIs.


4.3. CLOSE

When an existing HIP association is no longer needed, it can be closed using closing mechanism defined in [RFC5201].

5. Packet Formats

5.1. Control Header Flags

In order to distinguish the packets in extended BEX from those in ordinary BEX, two control header flags are defined here:



I: if this bit is set to 1, the initiator's HIT transported in the packet is scrambled, and the HI in this packet is encrypted. Otherwise, the initiator's HIT and HI are transported within the common part of the packet header in an unencrypted way.

R: if this bit is set to 1, the responder's HIT transported in the packet is scrambled, and the HI in this packet is encrypted. Otherwise, the responder's HIT and HI are transported without encryption. Note that in opportunistic mode, this flag in an I1 packet only indicate the initiator support scrambled HITs.

6. Applicability Consideration

6.1. Opportunistic Base Exchange

When an Initiator does not know the HIT of its communicating partner beforehand, it can try to initiate the handshaking in the opportunistic mode.

If the initiator intends to protect its identity privacy, it can send its blind HIT and the associated nonce in an I1 packet to the responder. Both the I and R flag in the I1 packet are set. If the responder would like to protect its identity privacy, it can use the nonce to generate a blind HIT and send it back in a R1 packet. The operations of both communicating partner in the rest of the handshaking is identical to what is described in section 4.1.1. If the initiator does not want to protect its identity privacy, it can send its plain HIT in an I1 packet to the responder. The R flag in the I1 packet can be set to indicate that the initiator supports scrambled HITs. If the responder would like to protect its identity privacy, it can choose a nonce and use it to generate a blind HIT for itself. Both the blind HIT and the nonce are sent back in a R1 packet. The operations of both communicating partner in the rest of the handshaking is identical to what is described in section 4.1.3.

In the both cases above, if the responder supports scrambled HITs but does not want to protect its identity privacy, it can just send its plain HIT in the R1 packet and unset the R flag.

6.2. Comparison of Disposable Identities vs. Blind

During the design of the proposed approach, the solutions using ephemeral identities are also considered. A host can attempt to prevent itself from being tracked by using different HIs in different BEXs. However, some upper layer server applications may use HIs or HITs to identify hosts. When a host uses ephemeral HI, it may be difficult for the applications to find proper state to provide service correctly.

Additionally, it is difficult for a responder to use ephemeral HI to protect its identity, as the initiator normally need to know the responder's HIT to initiate a HIP BEX.

6.3. HIP-based Middleboxes

Currently, there is only a type of middlebox (RVS) specified in the HIP architecture. Our solution introduces additional overheads to a RVS but will not damage its functionality. When a RVS received an I1 packet containing a blind HIT of the responder, the RVS has to use the nonce to find out the associated HIT. Considering the large number of the HITs that a RVS may hold, an Initiator can select to disclose some bites of the plain HIT of the responder to reduce the overhead imposed on the RVS.

6.4. Immediate Carriage and Conveyance of Upper-layer Protocol

If a host receives a hiccups based packet, it must respond with an R1 as described in [I-D.nikander-hip-hiccups].

7. Security Considerations

Our solution should be a component of a multi-layer privacy protection solution. Although the confidentiality of consistent information transported in higher layer protocol can be protected by the key derived from HIP BEX, one still have to carefully avoid the consistent information transported below HIP layer to be disclosed to adversaries. For instance, an attacker may identify the FQDN of a host by querying the reverse DNS system with the IP address of the host.

In addition, our privacy extension may be incompatible with HIP based firewalls [RFC5207], relay servers [RFC5770], and other authentication service provided by middle boxes [I-D.heer-hip-middle-auth].

8. Contributors

This work is based on the research of Jukka Ylitalo.

9. Acknowledgements

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4882] Koodli, R., "IP Address Location Privacy and Mobile IPv6: Problem Statement", RFC 4882, May 2007.
- [RFC5201] Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson,

"Host Identity Protocol", RFC 5201, April 2008.

10.2. Informative References

- [I-D.heer-hip-middle-auth]
Hummen, R., Heer, T., and M. Komu, "End-Host Authentication for HIP Middleboxes", draft-heer-hip-middle-auth-04 (work in progress), October 2011.
- [I-D.nikander-hip-hiccups]
Nikander, P., Camarillo, G., and J. Melen, "HIP (Host Identity Protocol) Immediate Carriage and Conveyance of Upper-layer Protocol Signaling (HICCUPS)", draft-nikander-hip-hiccups-04 (work in progress), August 2009.
- [RFC5207] Stiemerling, M., Quittek, J., and L. Eggert, "NAT and Firewall Traversal Issues of Host Identity Protocol (HIP) Communication", RFC 5207, April 2008.
- [RFC5770] Komu, M., Henderson, T., Tschofenig, H., Melen, J., and A. Keranen, "Basic Host Identity Protocol (HIP) Extensions for Traversal of Network Address Translators", RFC 5770, April 2010.
- [YLI04] Ylitalo, J., "LIND: A Complete Identity Protection Framework for End-points", April 2004.

Authors' Addresses

Dacheng Zhang
Huawei

Email: zhangdacheng@huawei.com

Miika Komu
Laboratory of Software Technology
Department of Computer Science and Engineering, Aalto University
Finland

Email: miika@iki.fi

