

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 16, 2015

T. Henderson, Ed.
University of Washington
C. Vogt
J. Arkko
Ericsson Research NomadicLab
November 12, 2014

Host Multihoming with the Host Identity Protocol
draft-ietf-hip-multihoming-04

Abstract

This document defines host multihoming extensions to the Host Identity Protocol (HIP), by leveraging protocol components defined for host mobility.

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 May 16, 2015.

Copyright Notice

Copyright (c) 2014 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction and Scope	2
2.	Terminology and Conventions	4
3.	Protocol Model	4
3.1.	Operating Environment	5
3.2.	Multihoming Overview	7
4.	Protocol Overview	7
4.1.	Host Multihoming	8
4.2.	Site Multihoming	9
4.3.	Dual host multihoming	10
4.4.	Combined Mobility and Multihoming	10
4.5.	Initiating the Protocol in R1 or I2	11
5.	Other Considerations	12
5.1.	Address Verification	12
5.2.	Preferred Locator	12
5.3.	Interaction with Security Associations	13
6.	Processing Rules	15
6.1.	Sending LOCATORS	15
6.2.	Handling Received LOCATORS	17
6.3.	Verifying Address Reachability	19
6.4.	Changing the Preferred Locator	19
7.	Security Considerations	20
8.	IANA Considerations	20
9.	Authors and Acknowledgments	20
10.	References	20
10.1.	Normative references	20
10.2.	Informative references	21
Appendix A.	Document Revision History	22

[1.](#) Introduction and Scope

The Host Identity Protocol [[I-D.ietf-hip-rfc4423-bis](#)] (HIP) supports an architecture that decouples the transport layer (TCP, UDP, etc.) from the internetworking layer (IPv4 and IPv6) by using public/private key pairs, instead of IP addresses, as host identities. When

a host uses HIP, the overlying protocol sublayers (e.g., transport layer sockets and Encapsulating Security Payload (ESP) Security Associations (SAs)) are instead bound to representations of these host identities, and the IP addresses are only used for packet forwarding. However, each host must also know at least one IP address at which its peers are reachable. Initially, these IP addresses are the ones used during the HIP base exchange [[I-D.ietf-hip-rfc5201-bis](#)].

One consequence of such a decoupling is that new solutions to network-layer mobility and host multihoming are possible. Host mobility is defined in [[I-D.ietf-hip-rfc5206-bis](#)] and covers the case in which a host has a single address and changes its network point-of-attachment while desiring to preserve the HIP-enabled security association. Host multihoming is somewhat of a dual case to host mobility, in that a host may simultaneously have more than one network point-of-attachment. There are potentially many variations of host multihoming possible. The scope of this document encompasses messaging and elements of procedure for some basic host multihoming scenarios of interest.

Another variation of multihoming that has been heavily studied is site multihoming. Solutions for site multihoming in IPv6 networks have been specified by the IETF shim6 working group. The shim6 protocol [[RFC5533](#)] bears many architectural similarities to HIP but there are differences in the security model and in the protocol. Future versions of this draft will summarize the differences more completely.

While HIP can potentially be used with transports other than the ESP transport format [[I-D.ietf-hip-rfc5202-bis](#)], this document largely assumes the use of ESP and leaves other transport formats for further study.

There are a number of situations where the simple end-to-end readdressing functionality defined herein is not sufficient. These include the initial reachability of a multihomed host, location privacy, simultaneous mobility of both hosts, and some modes of NAT traversal. In these situations, there is a need for some helper functionality in the network, such as a HIP rendezvous server [[I-D.ietf-hip-rfc5204-bis](#)]. Such functionality is out of the scope of this document. Finally, making underlying IP multihoming transparent to the transport layer has implications on the proper response of transport congestion control, path MTU selection, and Quality of Service (QoS). Transport-layer mobility triggers, and the proper transport response to a HIP multihoming address change, are outside the scope of this document.

2. Terminology and 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 [RFC 2119](#) [[RFC2119](#)].

Terminology is copied from [[I-D.ietf-hip-rfc5206-bis](#)].

LOCATOR. The name of a HIP parameter containing zero or more Locator fields. This parameter's name is distinguished from the Locator fields embedded within it by the use of all capital letters.

Locator. A name that controls how the packet is routed through the network and demultiplexed by the end host. It may include a concatenation of traditional network addresses such as an IPv6 address and end-to-end identifiers such as an ESP SPI. It may also include transport port numbers or IPv6 Flow Labels as demultiplexing context, or it may simply be a network address.

Address. A name that denotes a point-of-attachment to the network. The two most common examples are an IPv4 address and an IPv6 address. The set of possible addresses is a subset of the set of possible locators.

Preferred locator. A locator on which a host prefers to receive data. With respect to a given peer, a host always has one active Preferred locator, unless there are no active locators. By default, the locators used in the HIP base exchange are the Preferred locators.

Credit Based Authorization. A host must verify a mobile or multihomed peer's reachability at a new locator. Credit-Based Authorization authorizes the peer to receive a certain amount of data at the new locator before the result of such verification is known.

3. Protocol Model

This section is an overview; more detailed specification follows this section.

The overall protocol model is the same as in Section 3 of [[I-D.ietf-hip-rfc5206-bis](#)]; this section only highlights the differences.

3.1. Operating Environment

The Host Identity Protocol (HIP) [[I-D.ietf-hip-rfc5201-bis](#)] is a key establishment and parameter negotiation protocol. Its primary applications are for authenticating host messages based on host identities, and establishing security associations (SAs) for the ESP transport format [[I-D.ietf-hip-rfc5202-bis](#)] and possibly other protocols in the future.

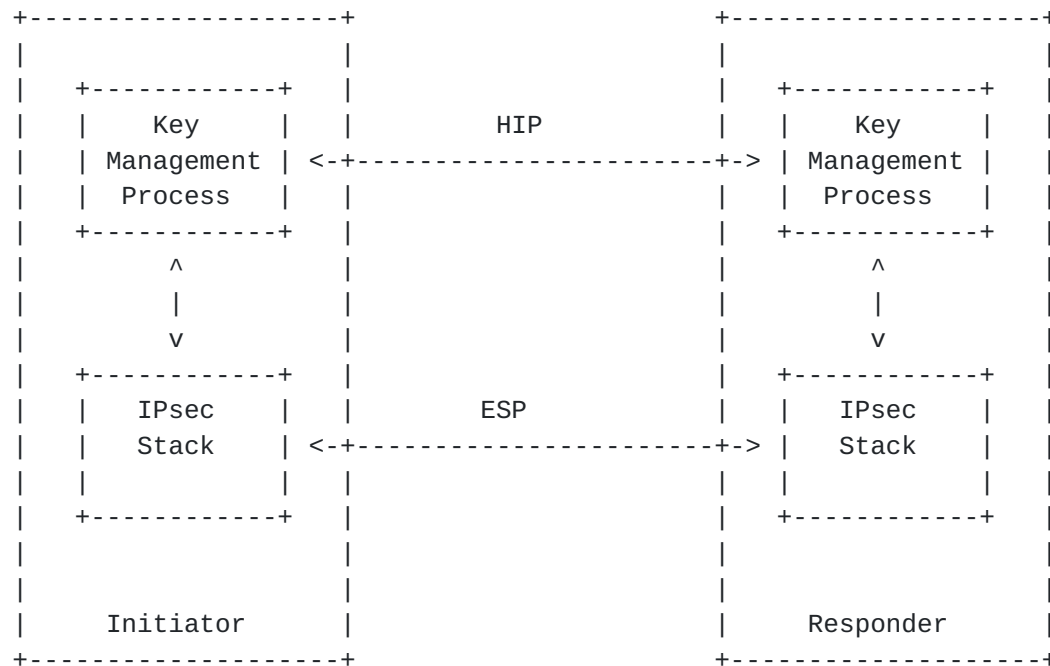


Figure 1: HIP Deployment Model

The general deployment model for HIP is shown above, assuming operation in an end-to-end fashion. This document specifies extensions to the HIP protocol to enable end-host mobility and basic multihoming. In summary, these extensions to the HIP base protocol enable the signaling of new addressing information to the peer in HIP messages. The messages are authenticated via a signature or keyed hash message authentication code (HMAC) based on its Host Identity.




Figure 2: Architecture for HIP Multihoming (MH)

Figure 2 depicts a layered architectural view of a HIP-enabled stack using the ESP transport format. In HIP, upper-layer protocols (including TCP and ESP in this figure) are bound to Host Identity Tags (HITs) and not IP addresses. The HIP sublayer is responsible for maintaining the binding between HITs and IP addresses. The SPI is used to associate an incoming packet with the right HITs. The block labeled "MH" is introduced below.

Consider the case when a host is multihomed (has more than one globally routable address) and has multiple addresses available at the HIP layer as alternative locators for fault tolerance. Examples include the use of (possibly multiple) IPv4 and IPv6 addresses on the same interface, or the use of multiple interfaces attached to different service providers. Such host multihoming generally necessitates that a separate ESP SA is maintained for each interface in order to prevent packets that arrive over different paths from falling outside of the ESP anti-replay window [[RFC4303](#)]. Multihoming thus makes it possible that the bindings shown on the right side of Figure 2 are one to many (in the outbound direction, one HIT pair to multiple SPIs, and possibly then to multiple IP addresses). However, only one SPI and address pair can be used for any given packet, so the job of the "MH" block depicted above is to dynamically manipulate these bindings. Beyond locally managing such multiple bindings, the peer-to-peer HIP signaling protocol needs to be flexible enough to define the desired mappings between HITs, SPIs, and addresses, and needs to ensure that UPDATE messages are sent along the right network paths so that any HIP-aware middleboxes can observe the SPIs. This document does not specify the "MH" block, nor does it specify detailed elements of procedure for how to handle various multihoming (perhaps combined with mobility) scenarios. The "MH" block may apply

to more general problems outside of HIP. However, this document does describe a basic multihoming case (one host adds one address to its initial address and notifies the peer) and leave more complicated scenarios for experimentation and future documents.

3.2. Multihoming Overview

In host multihoming, a host has multiple locators simultaneously rather than sequentially, as in the case of mobility. By using the LOCATOR parameter defined in [[I-D.ietf-hip-rfc5206-bis](#)], a host can inform its peers of additional (multiple) locators at which it can be reached, and can declare a particular locator as a "preferred" locator. Although this document defines a basic mechanism for multihoming, it does not define detailed policies and procedures, such as which locators to choose when more than one pair is available, the operation of simultaneous mobility and multihoming, source address selection policies (beyond those specified in [[RFC3484](#)]), and the implications of multihoming on transport protocols and ESP anti-replay windows.

4. Protocol Overview

In this section, we briefly introduce a number of usage scenarios for HIP multihoming. These scenarios assume that HIP is being used with the ESP transform [[I-D.ietf-hip-rfc5202-bis](#)], although other scenarios may be defined in the future. To understand these usage scenarios, the reader should be at least minimally familiar with the HIP protocol specification [[I-D.ietf-hip-rfc5201-bis](#)]. However, for the (relatively) uninitiated reader, it is most important to keep in mind that in HIP the actual payload traffic is protected with ESP, and that the ESP SPI acts as an index to the right host-to-host context.

The scenarios below assume that the two hosts have completed a single HIP base exchange with each other. Both of the hosts therefore have one incoming and one outgoing SA. Further, each SA uses the same pair of IP addresses, which are the ones used in the base exchange.

The readdressing protocol is an asymmetric protocol where a mobile or multihomed host informs a peer host about changes of IP addresses on affected SPIs. The readdressing exchange is designed to be piggybacked on existing HIP exchanges. The majority of the packets on which the LOCATOR parameters are expected to be carried are UPDATE packets. However, some implementations may want to experiment with sending LOCATOR parameters also on other packets, such as R1, I2, and NOTIFY.

The scenarios below at times describe addresses as being in either an ACTIVE, VERIFIED, or DEPRECATED state. From the perspective of a host, newly-learned addresses of the peer must be verified before put into active service, and addresses removed by the peer are put into a deprecated state. Under limited conditions described in [\[I-D.ietf-hip-rfc5206-bis\]](#), an UNVERIFIED address may be used.

Hosts that use link-local addresses as source addresses in their HIP handshakes may not be reachable by a mobile peer. Such hosts SHOULD provide a globally routable address either in the initial handshake or via the LOCATOR parameter.

4.1. Host Multihoming

A (mobile or stationary) host may sometimes have more than one interface or global address. The host may notify the peer host of the additional interface or address by using the LOCATOR parameter. To avoid problems with the ESP anti-replay window, a host SHOULD use a different SA for each interface or address used to receive packets from the peer host when multiple locator pairs are being used simultaneously rather than sequentially.

When more than one locator is provided to the peer host, the host SHOULD indicate which locator is preferred (the locator on which the host prefers to receive traffic). By default, the addresses used in the base exchange are preferred until indicated otherwise.

In the multihoming case, the sender may also have multiple valid locators from which to source traffic. In practice, a HIP association in a multihoming configuration may have both a preferred peer locator and a preferred local locator, although rules for source address selection should ultimately govern the selection of the source locator based on the destination locator.

Although the protocol may allow for configurations in which there is an asymmetric number of SAs between the hosts (e.g., one host has two interfaces and two inbound SAs, while the peer has one interface and one inbound SA), it is RECOMMENDED that inbound and outbound SAs be created pairwise between hosts. When an ESP_INFO arrives to rekey a particular outbound SA, the corresponding inbound SA should be also rekeyed at that time. Although asymmetric SA configurations might be experimented with, their usage may constrain interoperability at this time. However, it is recommended that implementations attempt to support peers that prefer to use non-paired SAs. It is expected that this section and behavior will be modified in future revisions of this protocol, once the issue and its implications are better understood.

Consider the case between two hosts, one single-homed and one multihomed. The multihomed host may decide to inform the single-homed host about its other address. It is RECOMMENDED that the multihomed host set up a new SA pair for use on this new address. To do this, the multihomed host sends a LOCATOR with an ESP_INFO, indicating the request for a new SA by setting the OLD SPI value to zero, and the NEW SPI value to the newly created incoming SPI. A Locator Type of "1" is used to associate the new address with the new SPI. The LOCATOR parameter also contains a second Type "1" locator, that of the original address and SPI. To simplify parameter processing and avoid explicit protocol extensions to remove locators, each LOCATOR parameter MUST list all locators in use on a connection (a complete listing of inbound locators and SPIs for the host). The multihomed host waits for an ESP_INFO (new outbound SA) from the peer and an ACK of its own UPDATE. As in the mobility case, the peer host must perform an address verification before actively using the new address. Figure 3 illustrates this scenario.

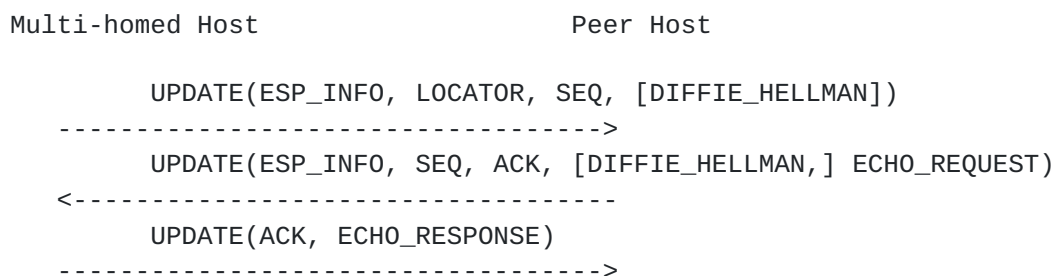


Figure 3: Basic Multihoming Scenario

In multihoming scenarios, it is important that hosts receiving UPDATES associate them correctly with the destination address used in the packet carrying the UPDATE. When processing inbound LOCATORS that establish new security associations on an interface with multiple addresses, a host uses the destination address of the UPDATE containing the LOCATOR as the local address to which the LOCATOR plus ESP_INFO is targeted. This is because hosts may send UPDATES with the same (locator) IP address to different peer addresses -- this has the effect of creating multiple inbound SAs implicitly affiliated with different peer source addresses.

4.2. Site Multihoming

A host may have an interface that has multiple globally routable IP addresses. Such a situation may be a result of the site having multiple upper Internet Service Providers, or just because the site provides all hosts with both IPv4 and IPv6 addresses. The host should stay reachable at all or any subset of the currently available global routable addresses, independent of how they are provided.

This case is handled the same as if there were different IP addresses, described above in [Section 4.1](#). Note that a single interface may experience site multihoming while the host itself may have multiple interfaces.

Note that a host may be multihomed and mobile simultaneously, and that a multihomed host may want to protect the location of some of its interfaces while revealing the real IP address of some others.

This document does not presently specify additional site multihoming extensions to HIP; further alignment with the IETF shim6 working group may be considered in the future.

[4.3.](#) Dual host multihoming

Consider the case in which both hosts would like to add an additional address after the base exchange completes. In Figure 4, consider that host1, which used address addr1a in the base exchange to set up SPI1a and SPI2a, wants to add address addr1b. It would send an UPDATE with LOCATOR (containing the address addr1b) to host2, using destination address addr2a, and a new set of SPIs would be added between hosts 1 and 2 (call them SPI1b and SPI2b -- not shown in the figure). Next, consider host2 deciding to add addr2b to the relationship. Host2 must select one of host1's addresses towards which to initiate an UPDATE. It may choose to initiate an UPDATE to addr1a, addr1b, or both. If it chooses to send to both, then a full mesh (four SA pairs) of SAs would exist between the two hosts. This is the most general case; it often may be the case that hosts primarily establish new SAs only with the peer's Preferred locator. The readdressing protocol is flexible enough to accommodate this choice.

```

      -<- SPI1a --                -- SPI2a ->-
host1 <                > addr1a <---> addr2a <                > host2
      ->- SPI2a --                -- SPI1a -<-

                        addr1b <---> addr2a  (second SA pair)
                        addr1a <---> addr2b  (third SA pair)
                        addr1b <---> addr2b  (fourth SA pair)

```

Figure 4: Dual Multihoming Case in Which Each Host Uses LOCATOR to Add a Second Address

[4.4.](#) Combined Mobility and Multihoming

It looks likely that in the future, many mobile hosts will be simultaneously mobile and multihomed, i.e., have multiple mobile interfaces. Furthermore, if the interfaces use different access

technologies, it is fairly likely that one of the interfaces may appear stable (retain its current IP address) while some other(s) may experience mobility (undergo IP address change).

The use of LOCATOR plus ESP_INFO should be flexible enough to handle most such scenarios, although more complicated scenarios have not been studied so far.

4.5. Initiating the Protocol in R1 or I2

A Responder host MAY include a LOCATOR parameter in the R1 packet that it sends to the Initiator. This parameter MUST be protected by the R1 signature. If the R1 packet contains LOCATOR parameters with a new Preferred locator, the Initiator SHOULD directly set the new Preferred locator to status ACTIVE without performing address verification first, and MUST send the I2 packet to the new Preferred locator. The I1 destination address and the new Preferred locator may be identical. All new non-preferred locators must still undergo address verification once the base exchange completes.

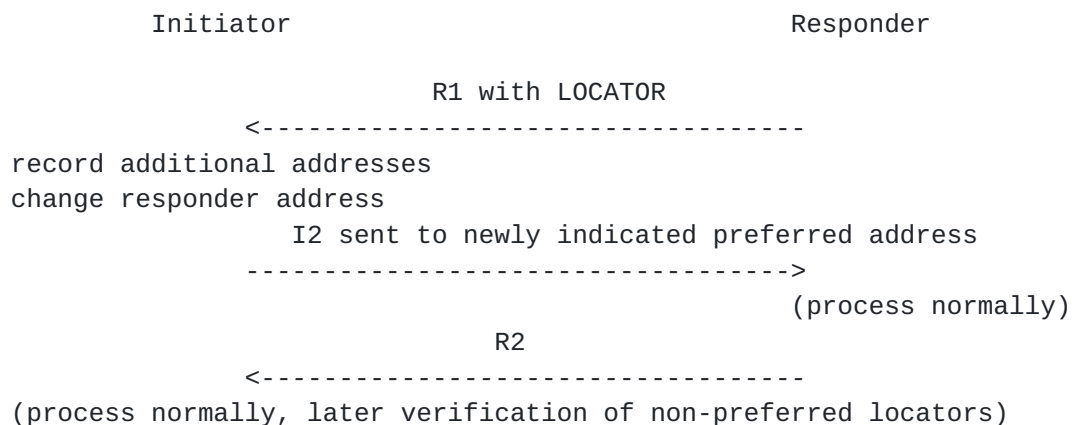


Figure 5: LOCATOR Inclusion in R1

An Initiator MAY include one or more LOCATOR parameters in the I2 packet, independent of whether or not there was a LOCATOR parameter in the R1. These parameters MUST be protected by the I2 signature. Even if the I2 packet contains LOCATOR parameters, the Responder MUST still send the R2 packet to the source address of the I2. The new Preferred locator SHOULD be identical to the I2 source address. If the I2 packet contains LOCATOR parameters, all new locators must undergo address verification as usual, and the ESP traffic that subsequently follows should use the Preferred locator.

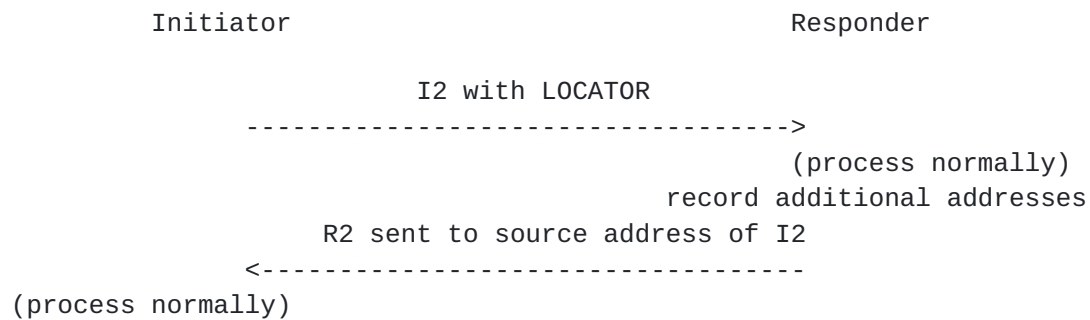


Figure 6: LOCATOR Inclusion in I2

The I1 and I2 may be arriving from different source addresses if the LOCATOR parameter is present in R1. In this case, implementations simultaneously using multiple pre-created R1s, indexed by Initiator IP addresses, may inadvertently fail the puzzle solution of I2 packets due to a perceived puzzle mismatch. See, for instance, the example in [Appendix A](#) of [\[I-D.ietf-hip-rfc5201-bis\]](#). As a solution, the Responder's puzzle indexing mechanism must be flexible enough to accommodate the situation when R1 includes a LOCATOR parameter.

5. Other Considerations

5.1. Address Verification

An address verification method is specified in [\[I-D.ietf-hip-rfc5206-bis\]](#). It is expected that addresses learned in multihoming scenarios also are subject to the same verification rules.

5.2. Preferred Locator

When a host has multiple locators, the peer host must decide which to use for outbound packets. It may be that a host would prefer to receive data on a particular inbound interface. HIP allows a particular locator to be designated as a Preferred locator and communicated to the peer.

In general, when multiple locators are used for a session, there is the question of using multiple locators for failover only or for load-balancing. Due to the implications of load-balancing on the transport layer that still need to be worked out, this document assumes that multiple locators are used primarily for failover. An implementation may use ICMP interactions, reachability checks, or other means to detect the failure of a locator.

5.3. Interaction with Security Associations

This document uses the HIP LOCATOR protocol parameter, specified in [[I-D.ietf-hip-rfc5206-bis](#)]), that allows the hosts to exchange information about their locator(s) and any changes in their locator(s). The logical structure created with LOCATOR parameters has three levels: hosts, Security Associations (SAs) indexed by Security Parameter Indices (SPIs), and addresses.

The relation between these levels for an association constructed as defined in the base specification [[I-D.ietf-hip-rfc5201-bis](#)] and ESP transform [[I-D.ietf-hip-rfc5202-bis](#)] is illustrated in Figure 7.

```

      -<- SPI1a --                               -- SPI2a ->-
host1 <                                     > addr1a <---> addr2a <                                     > host2
      ->- SPI2a --                               -- SPI1a -<-

```

Figure 7: Relation between Hosts, SPIs, and Addresses (Base Specification)

In Figure 7, host1 and host2 negotiate two unidirectional SAs, and each host selects the SPI value for its inbound SA. The addresses addr1a and addr2a are the source addresses that the hosts use in the base HIP exchange. These are the "preferred" (and only) addresses conveyed to the peer for use on each SA. That is, although packets sent to any of the hosts' interfaces may be accepted on the inbound SA, the peer host in general knows of only the single destination address learned in the base exchange (e.g., for host1, it sends a packet on SPI2a to addr2a to reach host2), unless other mechanisms exist to learn of new addresses.

In general, the bindings that exist in an implementation corresponding to this document can be depicted as shown in Figure 8. In this figure, a host can have multiple inbound SPIs (and, not shown, multiple outbound SPIs) associated with another host. Furthermore, each SPI may have multiple addresses associated with it. These addresses that are bound to an SPI are not used to lookup the incoming SA. Rather, the addresses are those that are provided to the peer host, as hints for which addresses to use to reach the host on that SPI. The LOCATOR parameter is used to change the set of addresses that a peer associates with a particular SPI.



Figure 8: Relation between Hosts, SPIs, and Addresses (General Case)

A host may establish any number of security associations (or SPIs) with a peer. The main purpose of having multiple SPIs with a peer is to group the addresses into collections that are likely to experience fate sharing. For example, if the host needs to change its addresses on SPI2, it is likely that both address21 and address22 will simultaneously become obsolete. In a typical case, such SPIs may correspond with physical interfaces; see below. Note, however, that especially in the case of site multihoming, one of the addresses may become unreachable while the other one still works. In the typical case, however, this does not require the host to inform its peers about the situation, since even the non-working address still logically exists.

A basic property of HIP SAs is that the inbound IP address is not used to lookup the incoming SA. Therefore, in Figure 8, it may seem unnecessary for address31, for example, to be associated only with SPI3 -- in practice, a packet may arrive to SPI1 via destination address address31 as well. However, the use of different source and destination addresses typically leads to different paths, with different latencies in the network, and if packets were to arrive via an arbitrary destination IP address (or path) for a given SPI, the reordering due to different latencies may cause some packets to fall outside of the ESP anti-replay window. For this reason, HIP provides a mechanism to affiliate destination addresses with inbound SPIs, when there is a concern that anti-replay windows might be violated. In this sense, we can say that a given inbound SPI has an "affinity" for certain inbound IP addresses, and this affinity is communicated to the peer host. Each physical interface SHOULD have a separate SA, unless the ESP anti-replay window is loose.

Moreover, even when the destination addresses used for a particular SPI are held constant, the use of different source interfaces may also cause packets to fall outside of the ESP anti-replay window, since the path traversed is often affected by the source address or

interface used. A host has no way to influence the source interface on which a peer sends its packets on a given SPI. A host SHOULD consistently use the same source interface and address when sending to a particular destination IP address and SPI. For this reason, a host may find it useful to change its SPI or at least reset its ESP anti-replay window when the peer host readdresses.

An address may appear on more than one SPI. This creates no ambiguity since the receiver will ignore the IP addresses during SA lookup anyway. However, this document does not specify such cases.

When the LOCATOR parameter is sent in an UPDATE packet, then the receiver will respond with an UPDATE acknowledgment. When the LOCATOR parameter is sent in an R1 or I2 packet, the base exchange retransmission mechanism will confirm its successful delivery. LOCATORS may experimentally be used in NOTIFY packets; in this case, the recipient MUST consider the LOCATOR as informational and not immediately change the current preferred address, but can test the additional locators when the need arises. The use of the LOCATOR in a NOTIFY message may not be compatible with middleboxes.

6. Processing Rules

Processing rules are specified in [[I-D.ietf-hip-rfc5206-bis](#)]. Future versions of this document will specify multihoming-specific processing rules here.

6.1. Sending LOCATORS

The decision of when to send LOCATORS is basically a local policy issue. However, it is RECOMMENDED that a host send a LOCATOR whenever it recognizes a change of its IP addresses in use on an active HIP association, and assumes that the change is going to last at least for a few seconds. Rapidly sending LOCATORS that force the peer to change the preferred address SHOULD be avoided.

When a host decides to inform its peers about changes in its IP addresses, it has to decide how to group the various addresses with SPIs. The grouping should consider also whether middlebox interaction requires sending the same LOCATOR in separate UPDATES on different paths. Since each SPI is associated with a different Security Association, the grouping policy may also be based on ESP anti-replay protection considerations. In the typical case, simply basing the grouping on actual kernel level physical and logical interfaces may be the best policy. Grouping policy is outside of the scope of this document.

Note that the purpose of announcing IP addresses in a LOCATOR is to provide connectivity between the communicating hosts. In most cases, tunnels or virtual interfaces such as IPsec tunnel interfaces or Mobile IP home addresses provide sub-optimal connectivity.

Furthermore, it should be possible to replace most tunnels with HIP based "non-tunneling", therefore making most virtual interfaces fairly unnecessary in the future. Therefore, virtual interfaces SHOULD NOT be announced in general. On the other hand, there are clearly situations where tunnels are used for diagnostic and/or testing purposes. In such and other similar cases announcing the IP addresses of virtual interfaces may be appropriate.

Hosts MUST NOT announce broadcast or multicast addresses in LOCATORS. Link-local addresses MAY be announced to peers that are known to be neighbors on the same link, such as when the IP destination address of a peer is also link-local. The announcement of link-local addresses in this case is a policy decision; link-local addresses used as Preferred locators will create reachability problems when the host moves to another link. In any case, link-local addresses MUST NOT be announced to a peer unless that peer is known to be on the same link.

Once the host has decided on the groups and assignment of addresses to the SPIs, it creates a LOCATOR parameter that serves as a complete representation of the addresses and affiliated SPIs intended for active use. We now describe a few cases introduced in [Section 4](#). We assume that the Traffic Type for each locator is set to "0" (other values for Traffic Type may be specified in documents that separate the HIP control plane from data plane traffic). Other mobility and multihoming cases are possible but are left for further experimentation.

1. Host multihoming (addition of an address). We only describe the simple case of adding an additional address to a (previously) single-homed, non-mobile host. The host SHOULD set up a new SA pair between this new address and the preferred address of the peer host. To do this, the multihomed host creates a new inbound SA and creates a new SPI. For the outgoing UPDATE message, it inserts an ESP_INFO parameter with an OLD SPI field of "0", a NEW SPI field corresponding to the new SPI, and a KEYMAT Index as selected by local policy. The host adds to the UPDATE message a LOCATOR with two Type "1" Locators: the original address and SPI active on the association, and the new address and new SPI being added (with the SPI matching the NEW SPI contained in the ESP_INFO). The Preferred bit SHOULD be set depending on the policy to tell the peer host which of the two locators is preferred. The UPDATE also contains a SEQ parameter and optionally a DIFFIE_HELLMAN parameter, and follows rekeying

procedures with respect to this new address. The UPDATE message SHOULD be sent to the peer's Preferred address with a source address corresponding to the new locator.

The sending of multiple LOCATORs, locators with Locator Type "0", and multiple ESP_INFO parameters is for further study. Note that the inclusion of LOCATOR in an R1 packet requires the use of Type "0" locators since no SAs are set up at that point.

6.2. Handling Received LOCATORs

A host SHOULD be prepared to receive a LOCATOR parameter in the following HIP packets: R1, I2, UPDATE, and NOTIFY.

This document describes sending both ESP_INFO and LOCATOR parameters in an UPDATE. The ESP_INFO parameter is included when there is a need to rekey or key a new SPI, and is otherwise included for the possible benefit of HIP-aware middleboxes. The LOCATOR parameter contains a complete map of the locators that the host wishes to make or keep active for the HIP association.

In general, the processing of a LOCATOR depends upon the packet type in which it is included. Here, we describe only the case in which ESP_INFO is present and a single LOCATOR and ESP_INFO are sent in an UPDATE message; other cases are for further study. The steps below cover each of the cases described in [Section 6.1](#).

The processing of ESP_INFO and LOCATOR parameters is intended to be modular and support future generalization to the inclusion of multiple ESP_INFO and/or multiple LOCATOR parameters. A host SHOULD first process the ESP_INFO before the LOCATOR, since the ESP_INFO may contain a new SPI value mapped to an existing SPI, while a Type "1" locator will only contain a reference to the new SPI.

When a host receives a validated HIP UPDATE with a LOCATOR and ESP_INFO parameter, it processes the ESP_INFO as follows. The ESP_INFO parameter indicates whether an SA is being rekeyed, created, deprecated, or just identified for the benefit of middleboxes. The host examines the OLD SPI and NEW SPI values in the ESP_INFO parameter:

1. (no rekeying) If the OLD SPI is equal to the NEW SPI and both correspond to an existing SPI, the ESP_INFO is gratuitous (provided for middleboxes) and no rekeying is necessary.
2. (rekeying) If the OLD SPI indicates an existing SPI and the NEW SPI is a different non-zero value, the existing SA is being rekeyed and the host follows HIP ESP rekeying procedures by

creating a new outbound SA with an SPI corresponding to the NEW SPI, with no addresses bound to this SPI. Note that locators in the LOCATOR parameter will reference this new SPI instead of the old SPI.

3. (new SA) If the OLD SPI value is zero and the NEW SPI is a new non-zero value, then a new SA is being requested by the peer. This case is also treated like a rekeying event; the receiving host must create a new SA and respond with an UPDATE ACK.
4. (deprecating the SA) If the OLD SPI indicates an existing SPI and the NEW SPI is zero, the SA is being deprecated and all locators uniquely bound to the SPI are put into the DEPRECATED state.

If none of the above cases apply, a protocol error has occurred and the processing of the UPDATE is stopped.

Next, the locators in the LOCATOR parameter are processed. For each locator listed in the LOCATOR parameter, check that the address therein is a legal unicast or anycast address. That is, the address MUST NOT be a broadcast or multicast address. Note that some implementations MAY accept addresses that indicate the local host, since it may be allowed that the host runs HIP with itself.

The below assumes that all locators are of Type "1" with a Traffic Type of "0"; other cases are for further study.

For each Type "1" address listed in the LOCATOR parameter, the host checks whether the address is already bound to the SPI indicated. If the address is already bound, its lifetime is updated. If the status of the address is DEPRECATED, the status is changed to UNVERIFIED. If the address is not already bound, the address is added, and its status is set to UNVERIFIED. Mark all addresses corresponding to the SPI that were NOT listed in the LOCATOR parameter as DEPRECATED.

As a result, at the end of processing, the addresses listed in the LOCATOR parameter have either a state of UNVERIFIED or ACTIVE, and any old addresses on the old SA not listed in the LOCATOR parameter have a state of DEPRECATED.

Once the host has processed the locators, if the LOCATOR parameter contains a new Preferred locator, the host SHOULD initiate a change of the Preferred locator. This requires that the host first verifies reachability of the associated address, and only then changes the Preferred locator; see [Section 6.4](#).

If a host receives a locator with an unsupported Locator Type, and when such a locator is also declared to be the Preferred locator for

the peer, the host SHOULD send a NOTIFY error with a Notify Message Type of LOCATOR_TYPE_UNSUPPORTED, with the Notification Data field containing the locator(s) that the receiver failed to process. Otherwise, a host MAY send a NOTIFY error if a (non-preferred) locator with an unsupported Locator Type is received in a LOCATOR parameter.

6.3. Verifying Address Reachability

Address verification is defined in [[I-D.ietf-hip-rfc5206-bis](#)].

When address verification is in progress for a new Preferred locator, the host SHOULD select a different locator listed as ACTIVE, if one such locator is available, to continue communications until address verification completes. Alternatively, the host MAY use the new Preferred locator while in UNVERIFIED status to the extent Credit-Based Authorization permits. Credit-Based Authorization is explained in [[I-D.ietf-hip-rfc5206-bis](#)]. Once address verification succeeds, the status of the new Preferred locator changes to ACTIVE.

6.4. Changing the Preferred Locator

A host MAY want to change the Preferred outgoing locator for different reasons, e.g., because traffic information or ICMP error messages indicate that the currently used preferred address may have become unreachable. Another reason may be due to receiving a LOCATOR parameter that has the "P" bit set.

To change the Preferred locator, the host initiates the following procedure:

1. If the new Preferred locator has ACTIVE status, the Preferred locator is changed and the procedure succeeds.
2. If the new Preferred locator has UNVERIFIED status, the host starts to verify its reachability. The host SHOULD use a different locator listed as ACTIVE until address verification completes if one such locator is available. Alternatively, the host MAY use the new Preferred locator, even though in UNVERIFIED status, to the extent Credit-Based Authorization permits. Once address verification succeeds, the status of the new Preferred locator changes to ACTIVE and its use is no longer governed by Credit-Based Authorization.
3. If the peer host has not indicated a preference for any address, then the host picks one of the peer's ACTIVE addresses randomly or according to policy. This case may arise if, for example,

ICMP error messages that deprecate the Preferred locator arrive, but the peer has not yet indicated a new Preferred locator.

4. If the new Preferred locator has DEPRECATED status and there is at least one non-deprecated address, the host selects one of the non-deprecated addresses as a new Preferred locator and continues. If the selected address is UNVERIFIED, the address verification procedure described above will apply.

7. Security Considerations

Security considerations are addressed in [[I-D.ietf-hip-rfc5206-bis](#)].

8. IANA Considerations

None.

9. Authors and Acknowledgments

This document contains content that was originally included in [RFC5206](#). Pekka Nikander and Jari Arkko originated [RFC5206](#), and Christian Vogt and Thomas Henderson (editor) later joined as co-authors. Also in [RFC5206](#), Greg Perkins contributed the initial draft of the security section, and Petri Jokela was a co-author of the initial individual submission.

The authors thank Miika Komu, Mika Kousa, Jeff Ahrenholz, and Jan Melen for many improvements to the document.

10. References

10.1. Normative references

[[I-D.ietf-hip-rfc5201-bis](#)]

Moskowitz, R., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", [draft-ietf-hip-rfc5201-bis-14](#) (work in progress), October 2013.

[[I-D.ietf-hip-rfc5202-bis](#)]

Jokela, P., Moskowitz, R., and J. Melen, "Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP)", [draft-ietf-hip-rfc5202-bis-05](#) (work in progress), November 2013.

[[I-D.ietf-hip-rfc5206-bis](#)]

Henderson, T., Vogt, C., and J. Arkko, "Host Mobility with the Host Identity Protocol", [draft-ietf-hip-rfc5206-bis-06](#) (work in progress), July 2013.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.

[10.2](#). Informative references

- [I-D.ietf-hip-rfc4423-bis]
Moskowitz, R. and M. Komu, "Host Identity Protocol Architecture", [draft-ietf-hip-rfc4423-bis-08](#) (work in progress), April 2014.
- [I-D.ietf-hip-rfc5204-bis]
Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", [draft-ietf-hip-rfc5204-bis-04](#) (work in progress), June 2014.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", [RFC 5533](#), June 2009.

Appendix A. Document Revision History

To be removed upon publication

Revision	Comments
draft-00	Initial version with multihoming text imported from RFC5206 .
draft-01	Document refresh; no other changes.
draft-02	Document refresh; no other changes.
draft-03	Document refresh; no other changes.
draft-04	Document refresh; no other changes.

Authors' Addresses

Thomas R. Henderson (editor)
University of Washington
Campus Box 352500
Seattle, WA
USA

E-Mail: tomhend@u.washington.edu

Christian Vogt
Ericsson Research NomadicLab
Hirsalantie 11
JORVAS FIN-02420
FINLAND

EMail: christian.vogt@ericsson.com

Jari Arkko
Ericsson Research NomadicLab
JORVAS FIN-02420
FINLAND

Phone: +358 40 5079256

E-Mail: jari.arkko@ericsson.com

