

Internet Engineering Task Force (IETF)
Request for Comments: 8047
Category: Standards Track
ISSN: 2070-1721

T. Henderson, Ed.
University of Washington
C. Vogt
Independent
J. Arkko
Ericsson
February 2017

Host Multihoming with the Host Identity Protocol

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 is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8047>.

Copyright Notice

Copyright (c) 2017 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 and Scope	3
2. Terminology and Conventions	4
3. Protocol Model	4
4. Protocol Overview	4
4.1. Background	5
4.2. Usage Scenarios	6
4.2.1. Multiple Addresses	6
4.2.2. Multiple Security Associations	6
4.2.3. Host Multihoming for Fault Tolerance	7
4.2.4. Host Multihoming for Load Balancing	9
4.2.5. Site Multihoming	10
4.2.6. Dual-Host Multihoming	10
4.2.7. Combined Mobility and Multihoming	11
4.2.8. Initiating the Protocol in R1, I2, or R2	11
4.2.9. Using LOCATOR_SETs across Addressing Realms	13
4.3. Interaction with Security Associations	13
5. Processing Rules	14
5.1. Sending LOCATOR_SETs	14
5.2. Handling Received LOCATOR_SETs	16
5.3. Verifying Address Reachability	18
5.4. Changing the Preferred Locator	18
6. Security Considerations	19
7. References	21
7.1. Normative References	21
7.2. Informative References	21
Acknowledgments	22
Authors' Addresses	22

1. Introduction and Scope

The Host Identity Protocol (HIP) [RFC7401] 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.

One consequence of such a decoupling is that new solutions to network-layer mobility and host multihoming are possible. Basic host mobility is defined in [RFC8046] 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. [RFC8046] specifies the format of the HIP parameter (LOCATOR_SET parameter) used to convey IP addressing information between peers, the procedures for sending and processing this parameter to enable basic host mobility, and procedures for an address verification mechanism. 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 host multihoming in multihomed 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.

While HIP can potentially be used with transports other than the ESP transport format [RFC7402], this document largely assumes the use of ESP and leaves other transport formats for further study.

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.

This specification relies on implementing Sections 4 ("LOCATOR_SET Parameter Format") and 5 ("Processing Rules") of [RFC8046] as a starting point for this implementation.

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].

The following terms used in this document are defined in [RFC8046]: LOCATOR_SET, Locator, locator, Address, preferred locator, and Credit-Based Authorization.

3. Protocol Model

The protocol model for HIP support of host multihoming extends the model for host mobility described in Section 3 of [RFC8046]. This section only highlights the differences.

In host multihoming, a host has multiple locators simultaneously rather than sequentially, as in the case of mobility. By using the LOCATOR_SET parameter defined in [RFC8046], a host can inform its peers of additional (multiple) locators at which it can be reached. When multiple locators are available and announced to the peer, a host can designate a particular locator as a "preferred" locator, meaning that the host prefers that its peer send packets to the designated address before trying an alternative address. Although this document defines a basic mechanism for multihoming, it does not define all possible policies and procedures, such as which locators to choose when more than one is available, the operation of simultaneous mobility and multihoming, source address selection policies (beyond those specified in [RFC6724]), and the implications of multihoming on transport protocols.

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 transport [RFC7402], 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 [RFC7401], the use of the ESP transport format [RFC7402], and the HIP mobility specification [RFC8046]. 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 Security Parameter Index (SPI) acts as an index to the right host-to-host context.

4.1. Background

The multihoming scenarios can be explained in contrast to the non-multihoming case described in the base protocol specification [RFC7401]. We review the pertinent details here. In the base specification, when used with the ESP transport format, the HIP base exchange will set up a single SA in each direction. The IP addresses associated with the SAs are the same as those used to convey the HIP packets. For data traffic, a security policy database (SPD) and security association database (SAD) will likely exist, following the IPsec architecture. One distinction between HIP and IPsec, however, is that the host IDs, and not the IP addresses, are conceptually used as selectors in the SPD. In the outbound direction, as a result of SPD processing, when an outbound SA is selected, the correct IP destination address for the peer must also be assigned. Therefore, outbound SAs are conceptually associated with the peer IP address that must be used as the destination IP address below the HIP layer. In the inbound direction, the IP addresses may be used as selectors in the SAD to look up the SA, but they are not strictly required; the ESP SPI may be used alone. To summarize, in the non-multihoming case, there is only one source IP address, one destination IP address, one inbound SA, and one outbound SA.

The HIP readdressing protocol [RFC8046] is an asymmetric protocol in which a mobile or multihomed host informs a peer host about changes of IP addresses on affected SPIs. IP address and ESP SPI information is carried in Locator fields in a HIP parameter called a LOCATOR_SET. The HIP mobility specification [RFC8046] describes how the LOCATOR_SET is carried in a HIP UPDATE packet.

To summarize the mobility elements of procedure, as background for multihoming, the basic idea of host mobility is to communicate a local IP address change to the peer when active HIP-maintained SAs are in use. To do so, the IP address must be conveyed, any association between the IP address and an inbound SA (via the SPI index) may be conveyed, and protection against flooding attacks must be ensured. The association of an IP address with an SPI is performed by a Locator Type of "1", which is a concatenation of an ESP SPI with an IP address.

An address verification method is specified in [RFC8046]. It is expected that addresses learned in multihoming scenarios also are subject to the same verification rules. At times, the scenarios 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 [RFC8046], an UNVERIFIED address may be used.

With this background, we next describe an additional protocol to facilitate scenarios in which one or both hosts have multiple IP addresses available. Increasingly, this is the common case with network-connected hosts on the Internet.

4.2. Usage Scenarios

4.2.1. Multiple Addresses

Hosts may have multiple IP addresses within different address families (IPv4 and IPv6) and scopes available to support HIP messaging and HIP-enabled SAs. The multiple addresses may be on a single network interface or multiple network interfaces. It is outside of the scope of this document to specify how a host decides which of possibly multiple addresses may be used to support a HIP association. Some IP addresses may be held back from usage due to privacy, security, or cost considerations.

When multiple IP addresses are shared with a peer, the procedures described in the HIP mobility specification [RFC8046] allow for a host to set a preferred locator ("P") bit, requesting that one of the multiple addresses be preferred for control- or data-plane traffic. It is also permitted to leave the preferred bit unset for all addresses, allowing the peer to make address selection decisions.

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_SET parameter.

To support mobility, as described in the HIP mobility specification [RFC8046], the LOCATOR_SET may be sent in a HIP UPDATE packet. To support multihoming, the LOCATOR_SET may also be sent in R1, I2, or R2 packets defined in the HIP protocol specification [RFC7401]. The reason to consider sending LOCATOR_SET parameters in base exchange packets is to convey all usable addresses for fault-tolerance or load-balancing considerations.

4.2.2. Multiple Security Associations

When multiple addresses are available between peer hosts, a question that arises is whether to use one or multiple SAs. The intent of this specification is to support different use cases but to leave the policy decision to the hosts.

When one host has n addresses and the other host has m addresses, it is possible to set up as many as $(n * m)$ SAs in each direction. In such a case, every combination of source and destination IP addresses would have a unique SA, and the possibility of the reordering of datagrams on each SA will be lessened (ESP SAs may have an anti-replay window [RFC4303] sensitive to reordering). However, the downside to creating a mesh of SAs is the signaling overhead required (for exchanging UPDATE messages conveying ESP_INFO parameters) and the state maintenance required in the SPD/SAD.

For load balancing, when multiple paths are to be used in parallel, it may make sense to create different SAs for different paths. In this use case, while a full mesh of $2 * (n * m)$ SAs may not be required, it may be beneficial to create one SA pair per load-balanced path to avoid anti-replay window issues.

For fault tolerance, it is more likely that a single SA and multiple IP addresses associated with that SA can be used, and the alternative addresses can be used only upon failure detection of the addresses in use. Techniques for path failure detection are outside the scope of this specification. An implementation may use ICMP interactions, reachability checks, or other means to detect the failure of a locator.

In summary, whether and how a host decides to leverage additional addresses in a load-balancing or fault-tolerant manner is outside the scope of the specification (although the academic literature on multipath TCP schedulers may provide guidance on how to design such a policy). However, in general, this document recommends that for fault tolerance, it is likely sufficient to use a single SA pair for all addresses, and for load balancing, to support a different SA pair for all active paths being balanced across.

4.2.3. Host Multihoming for Fault Tolerance

A (mobile or stationary) host may have more than one interface or global address. The host may choose to notify the peer host of the additional interface or address by using the LOCATOR_SET parameter. The LOCATOR_SET parameter may be included in an I2, R1, or R2 packet, or it may be conveyed, after the base exchange completes in an UPDATE packet.

When more than one locator is provided to the peer host, the host MAY indicate which locator is preferred (the locator on which the host prefers to receive traffic). By default, the address that a host uses in the base exchange is its preferred locator (for the address

family and address scope in use during the base exchange) until indicated otherwise. It may be the case that the host does not express any preferred locators.

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. The host should try to use the peer's preferred locator unless policy or other circumstances prevent such usage. A preferred local locator may be overridden if source address selection rules on the destination address (peer's preferred locator) suggest the use of a different source address.

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 suggested 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 also be rekeyed at that time. Section 4.3 discusses the interaction between addresses and security associations in more detail.

Consider the case of two hosts, one single-homed and one multihomed. The multihomed host may decide to inform the single-homed host about its other address(es). It may choose to do so as follows.

If the multihomed host wishes to convey the additional address(es) for fault tolerance, it should include all of its addresses in Locator fields, indicating the Traffic Type, Locator Type, and whether the locator is a preferred locator. If it wishes to bind any particular address to an existing SPI, it may do so by using a Locator Type of "1" as specified in the HIP mobility specification [RFC8046]. It does not need to rekey the existing SA or request additional SAs at this time.

Figure 1 illustrates this scenario. Note that the conventions for message parameter notations in figures (use of parentheses and brackets) is defined in Section 2.2 of [RFC7401].




Figure 1: Basic Multihoming Scenario

In this scenario, the peer host associates the multiple addresses with the SA pair between it and the multihomed host. It may also undergo address verification procedures to transition the addresses to ACTIVE state. For inbound data traffic, it may choose to use the addresses along with the SPI as selectors. For outbound data traffic, it must choose among the available addresses of the multihomed host, considering the state of address verification [RFC8046] of each address, and also considering available information about whether an address is in a working state.

4.2.4. Host Multihoming for Load Balancing

A multihomed host may decide to set up new SA pairs corresponding to new addresses, for the purpose of load balancing. The decision to load balance and the mechanism for splitting load across multiple SAs is out of scope of this document. The scenario can be supported by sending the LOCATOR_SET parameter with one or more ESP_INFO parameters to initiate new ESP SAs. To do this, the multihomed host sends a LOCATOR_SET 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_SET 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_SET 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 a corresponding 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 2 illustrates this scenario.

Multihomed Host	Peer Host
	UPDATE(ESP_INFO, LOCATOR_SET, SEQ, [DIFFIE_HELLMAN])
	----->
	UPDATE(ESP_INFO, SEQ, ACK, [DIFFIE_HELLMAN,] ECHO_REQUEST)
	<-----
	UPDATE(ACK, ECHO_RESPONSE)
	----->

Figure 2: Host Multihoming for Load Balancing

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 LOCATOR_SETs

that establish new security associations on an interface with multiple addresses, a host uses the destination address of the UPDATE containing the LOCATOR_SET as the local address to which the LOCATOR_SET 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.

4.2.5. 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 Sections 4.2.3 and 4.2.4. Note that a single interface may have addresses corresponding to site multihoming while the host itself may also have multiple network 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 present additional site multihoming extensions to HIP; such extensions are for further study.

4.2.6. Dual-Host Multihoming

Consider the case in which both hosts are multihomed and would like to notify the peer of an additional address after the base exchange completes. It may be the case that both hosts choose to simply announce the second address in a LOCATOR_SET parameter using an UPDATE message exchange. It may also be the case that one or both hosts decide to ask for new SA pairs to be created using the newly announced address. In the case that both hosts request this, the result will be a full mesh of SAs as depicted in Figure 3. In such a scenario, 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_SET (containing the address addr1b) to host2, using destination address addr2a, and a new ESP_INFO, 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; the 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 3: Dual-Multihoming Case in which Each Host Uses LOCATOR_SET to Add a Second Address

4.2.7. Combined Mobility and Multihoming

Mobile hosts may 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 others may experience mobility (undergo IP address change).

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

4.2.8. Initiating the Protocol in R1, I2, or R2

A Responder host MAY include a LOCATOR_SET 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_SET parameters with a new preferred locator, the Initiator SHOULD directly set the new preferred locator to status ACTIVE without performing address verification first, and it 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. It is also possible for the host to send the LOCATOR_SET without any preferred bits set, in which case the exchange will continue as normal and the newly learned addresses will be in an UNVERIFIED state at the initiator.

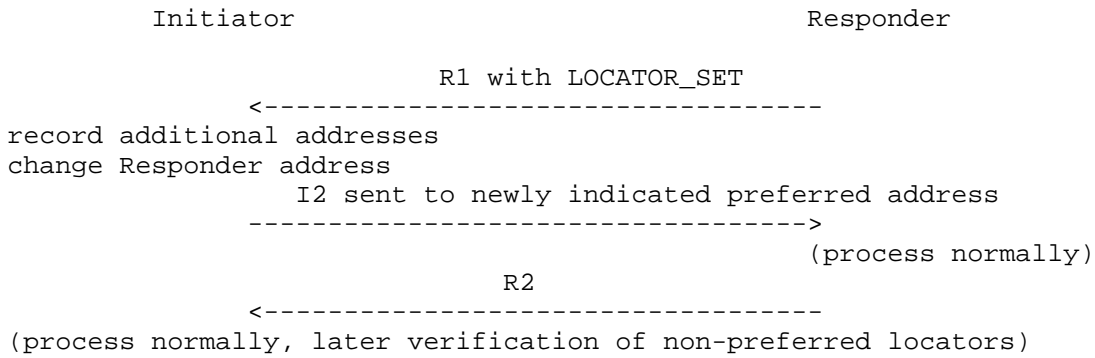


Figure 4: LOCATOR_SET Inclusion in R1

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

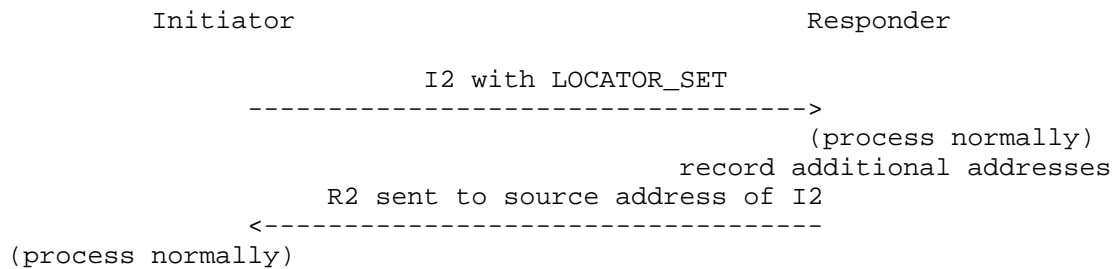


Figure 5: LOCATOR_SET Inclusion in I2

The I1 and I2 may be arriving from different source addresses if the LOCATOR_SET 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 [RFC7401]. As a solution, the Responder's puzzle indexing mechanism must be flexible enough to accommodate the situation when R1 includes a LOCATOR_SET parameter.

Finally, the R2 may be used to carry the LOCATOR_SET parameter. In this case, the LOCATOR_SET is covered by the HIP_MAC_2 and HIP_SIGNATURE. Including LOCATOR_SET in R2 as opposed to R1 may have some advantages when a host prefers not to divulge additional locators until after the I2 is successfully processed.

When the LOCATOR_SET parameter is sent in an UPDATE packet, the receiver will respond with an UPDATE acknowledgment. When the LOCATOR_SET parameter is sent in an R1, I2, or R2 packet, the base exchange retransmission mechanism will confirm its successful delivery.

4.2.9. Using LOCATOR_SETs across Addressing Realms

It is possible for HIP associations to use these mechanisms to migrate their HIP associations and security associations from addresses in the IPv4 addressing realm to IPv6, or vice versa. It may be possible for a state to arise in which both hosts are only using locators in different addressing realms, but in such a case, some type of mechanism for interworking between the different realms must be employed; such techniques are outside the scope of the present text.

4.3. Interaction with Security Associations

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, or to perform load balancing.

A basic property of HIP SAs is that the inbound IP address is not used to look up the incoming SA. 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 extended or disabled.

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.

5. Processing Rules

Basic processing rules for the LOCATOR_SET parameter are specified in [RFC8046]. This document focuses on multihoming-specific rules.

5.1. Sending LOCATOR_SETs

The decision of when to send a LOCATOR_SET, and which addresses to include, is a local policy issue. [RFC8046] recommends that a host "send a LOCATOR_SET whenever it recognizes a change of its IP addresses in use on an active HIP association and [when it] assumes that the change is going to last at least for a few seconds." It is possible to delay the exposure of additional locators to the peer, and to send data from previously unannounced locators, as might arise in certain mobility or multihoming situations.

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. If hosts are deployed in an operational environment in which HIP-aware NATs and firewalls (that may perform parameter inspection) exist, and different such devices may exist on different paths, hosts may take that knowledge into consideration about how addresses are grouped, and may send the same LOCATOR_SET in separate UPDATES on the different paths. However, more detailed guidelines about how to operate in the presence of such HIP-aware NATs and firewalls are a topic for further study. 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. The grouping policy is outside of the scope of this document.

Locators corresponding to tunnel interfaces (e.g., IPsec tunnel interfaces or Mobile IP home addresses) or other virtual interfaces MAY be announced in a LOCATOR_SET, but implementations SHOULD avoid announcing such locators as preferred locators if more direct paths may be obtained by instead preferring locators from non-tunneling interfaces if such locators provide a more direct path to the HIP peer.

[RFC8046] specifies that hosts MUST NOT announce broadcast or multicast addresses in LOCATOR_SETs. 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_SET parameter that serves as a complete representation of the addresses and associated 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 MAY choose to simply announce this address to the peer, for fault tolerance. To do this, the multihomed host creates a LOCATOR_SET parameter including the existing address and SPI as a Type "1" Locator, and the new address as a Type "0" Locator. The host sends this in an UPDATE message with the SEQ parameter, which is acknowledged by the peer.
2. The host MAY set up a new SA pair between this new address and an 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_SET 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 LOCATOR_SETs is unsupported. Note that the inclusion of LOCATOR_SET in an R1 packet requires the use of Type "0" Locators since no SAs are set up at that point.

5.2. Handling Received LOCATOR_SETs

A host SHOULD be prepared to receive a LOCATOR_SET parameter in the following HIP packets: R1, I2, R2, and UPDATE.

This document describes sending both ESP_INFO and LOCATOR_SET parameters in an UPDATE. The ESP_INFO parameter is included when there is a need to rekey or key a new SPI and can otherwise be included for the possible benefit of HIP-aware middleboxes. The LOCATOR_SET 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_SET 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_SET 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 5.1.

The processing of ESP_INFO and LOCATOR_SET parameters is intended to be modular and support future generalization to the inclusion of multiple ESP_INFO and/or multiple LOCATOR_SET parameters. A host SHOULD first process the ESP_INFO before the LOCATOR_SET, 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_SET 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_SET 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_SET parameter are processed. For each locator listed in the LOCATOR_SET 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.

For each Type "1" address listed in the LOCATOR_SET 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. If there exist remaining addresses corresponding to the SPI that were NOT listed in the LOCATOR_SET parameter, the host sets the status of such addresses to DEPRECATED.

For each Type "0" address listed in the LOCATOR_SET parameter, if the status of the address is DEPRECATED, or the address was not previously known, the status is changed to UNVERIFIED. The host MAY choose to associate this address with one or more SAs. The association with different SAs is a local policy decision, unless the peer has indicated that the address is preferred, in which case the address should be put into use on an SA that is prioritized in the security policy database.

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

Once the host has processed the locators, if the LOCATOR_SET 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 5.4.

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_SET parameter.

5.3. Verifying Address Reachability

Address verification is defined in [RFC8046].

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 [RFC8046]. Once address verification succeeds, the status of the new preferred locator changes to ACTIVE.

5.4. 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_SET parameter that has the preferred 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.

6. Security Considerations

This document extends the scope of host mobility solutions defined in [RFC8046] to also include host multihoming, and as a result, many of the same security considerations for mobility also pertain to multihoming. In particular, [RFC8046] describes how HIP host mobility is resistant to different types of impersonation attacks and denial-of-service (DoS) attacks.

The security considerations for this document are similar to those of [RFC8046] because the strong authentication capabilities for mobility also carry over to end-host multihoming. [RFC4218] provides a threat analysis for IPv6 multihoming, and the remainder of this section first describes how HIP host multihoming addresses those previously described threats, and then it discusses some additional security considerations.

The high-level threats discussed in [RFC4218] involve redirection attacks for the purposes of packet recording, data manipulation, and availability. There are a few types of attackers to consider: on-path attackers, off-path attackers, and malicious hosts.

[RFC4218] also makes the comment that in identifier/locator split solutions such as HIP, application security mechanisms should be tied to the identifier, not the locator, and attacks on the identifier mechanism and on the mechanism binding locators to the identifier are of concern. This document does not consider the former issue (application-layer security bindings) to be within scope. The latter issue (locator bindings to identifier) is directly addressed by the cryptographic protections of the HIP protocol, in that locators associated to an identifier are listed in HIP packets that are signed using the identifier key.

Section 3.1 of [RFC4218] lists several classes of security configurations in use in the Internet. HIP maps to the fourth (strong identifier) and fifth ("leap-of-faith") categories, the

latter being associated with the optional opportunistic mode of HIP operation. The remainder of Section 3 describes existing security problems in the Internet and comments that the goal of a multihoming solution is not to solve them specifically but rather not to make any of them worse. HIP multihoming should not increase the severity of the identified risks. One concern for both HIP mobility and multihoming is the susceptibility of the mechanisms to misuse flooding-based redirections due to a malicious host. The mechanisms described in [RFC8046] for address verification are important in this regard.

Regarding the new types of threats introduced by multihoming (Section 4 of [RFC4218]), HIP multihoming should not introduce new concerns. Classic and premeditated redirection are prevented by the strong authentication in HIP messages. Third-party DoS attacks are prevented by the address verification mechanism. Replay attacks can be avoided via use of replay protection in ESP SAs. In addition, accepting packets from unknown locators is protected by either the strong authentication in the HIP control packets or by the ESP-based encryption in use for data packets.

The HIP mechanisms are designed to limit the ability to introduce DoS on the mechanisms themselves (Section 7 of [RFC4218]). Care is taken in the HIP base exchange to avoid creating state or performing much work before hosts can authenticate one another. A malicious host involved in HIP multihoming with another host might attempt to misuse the mechanisms for multihoming by, for instance, increasing the state required or inducing a resource limitation attack by sending too many candidate locators to the peer host. Therefore, implementations supporting the multihoming extensions should consider avoiding accepting large numbers of peer locators and rate limiting any UPDATE messages being exchanged.

The exposure of a host's IP addresses through HIP mobility and multihoming extensions may raise the following privacy concern. The administrator of a host may be trying to hide its location in some context through the use of a VPN or other virtual interfaces. Similar privacy issues also arise in other frameworks such as WebRTC and are not specific to HIP. Implementations SHOULD provide a mechanism to allow the host administrator to block the exposure of selected addresses or address ranges.

Finally, some implementations of VPN tunneling have experienced instances of 'leakage' of flows that were intended to have been protected by a security tunnel but are instead sent in the clear, perhaps because some of the addresses used fall outside of the range of addresses configured for the tunnel in the security policy or association database. Implementors are advised to take steps to

ensure that the usage of multiple addresses between hosts does not cause accidental leakage of some data session traffic outside of the ESP-protected envelope.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC7401] Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <<http://www.rfc-editor.org/info/rfc7401>>.
- [RFC7402] Jokela, P., Moskowitz, R., and J. Melen, "Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP)", RFC 7402, DOI 10.17487/RFC7402, April 2015, <<http://www.rfc-editor.org/info/rfc7402>>.
- [RFC8046] Henderson, T., Ed., Vogt, C., and J. Arkko, "Host Mobility with the Host Identity Protocol", RFC 8046, DOI 10.17487/RFC8046, February 2017, <<http://www.rfc-editor.org/info/rfc8046>>.

7.2. Informative References

- [RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions", RFC 4218, DOI 10.17487/RFC4218, October 2005, <<http://www.rfc-editor.org/info/rfc4218>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<http://www.rfc-editor.org/info/rfc4303>>.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, DOI 10.17487/RFC5533, June 2009, <<http://www.rfc-editor.org/info/rfc5533>>.

Acknowledgments

This document contains content that was originally included in RFC 5206. Pekka Nikander and Jari Arkko originated RFC 5206, and Christian Vogt and Thomas Henderson (editor) later joined as coauthors. Also in RFC 5206, Greg Perkins contributed the initial draft of the security section, and Petri Jokela was a coauthor of the initial individual submission.

The authors thank Miika Komu, Mika Kousa, Jeff Ahrenholz, and Jan Melen for many improvements to the document. Concepts from a paper on host multihoming across address families, by Samu Varjonen, Miika Komu, and Andrei Gurtov, contributed to this revised specification.

Authors' Addresses

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

Email: tomhend@u.washington.edu

Christian Vogt
Independent
3473 North First Street
San Jose, CA 95134
United States of America

Email: mail@christianvogt.net

Jari Arkko
Ericsson
Jorvas, FIN-02420
Finland

Phone: +358 40 5079256
Email: jari.arkko@piuha.net

