

Host Identity Protocol (HIP) Architecture

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo describes a snapshot of the reasoning behind a proposed new namespace, the Host Identity namespace, and a new protocol layer, the Host Identity Protocol (HIP), between the internetworking and transport layers. Herein are presented the basics of the current namespaces, their strengths and weaknesses, and how a new namespace will add completeness to them. The roles of this new namespace in the protocols are defined. The memo describes the thinking of the authors as of Fall 2003. The architecture may have evolved since. This document represents one stable point in that evolution of understanding.

Table of Contents

1. Disclaimer	2
2. Introduction	2
3. Terminology	4
3.1. Terms Common to Other Documents	4
3.2. Terms Specific to This and Other HIP Documents	4
4. Background	6
4.1. A Desire for a Namespace for Computing Platforms	6
5. Host Identity Namespace	8
5.1. Host Identifiers	9
5.2. Storing Host Identifiers in DNS	9
5.3. Host Identity Tag (HIT)	10
5.4. Local Scope Identifier (LSI)	10
6. New Stack Architecture	11

6.1. Transport Associations and End-points	11
7. End-host Mobility and Multi-homing	12
7.1. Rendezvous Mechanism	13
7.2. Protection against Flooding Attacks	13
8. HIP and IPsec	14
9. HIP and NATs	15
9.1. HIP and TCP Checksums	15
10. Multicast	16
11. HIP Policies	16
12. Benefits of HIP	16
12.1. HIP's Answers to NSRG Questions	17
13. Security Considerations	19
13.1. HITs Used in ACLs	21
13.2. Non-security considerations	21
14. Acknowledgements	22
15. Informative References	22

1. Disclaimer

The purpose of this memo is to provide a stable reference point in the development of the Host Identity Protocol architecture. This memo describes the thinking of the authors as of Fall 2003; their thinking may have evolved since then. Occasionally, this memo may be confusing or self-contradicting. That is (partially) intentional, and it reflects the snapshot nature of this memo.

This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and notes that the decision to publish is not based on IETF review. However, the ideas put forth in this RFC have generated significant interest, including the formation of the IETF HIP Working Group and the IRTF HIP Research Group. These groups are expected to generate further documents, sharing their findings with the whole Internet community.

2. Introduction

The Internet has two important global namespaces: Internet Protocol (IP) addresses and Domain Name Service (DNS) names. These two namespaces have a set of features and abstractions that have powered the Internet to what it is today. They also have a number of weaknesses. Basically, since they are all we have, we try to do too much with them. Semantic overloading and functionality extensions have greatly complicated these namespaces.

The proposed Host Identity namespace fills an important gap between the IP and DNS namespaces. The Host Identity namespace consists of Host Identifiers (HIs). A Host Identifier is cryptographic in its

nature; it is the public key of an asymmetric key-pair. Each host will have at least one Host Identity, but it will typically have more than one. Each Host Identity uniquely identifies a single host; i.e., no two hosts have the same Host Identity. The Host Identity, and the corresponding Host Identifier, can be either public (e.g., published in the DNS) or unpublished. Client systems will tend to have both public and unpublished Identities.

There is a subtle but important difference between Host Identities and Host Identifiers. An Identity refers to the abstract entity that is identified. An Identifier, on the other hand, refers to the concrete bit pattern that is used in the identification process.

Although the Host Identifiers could be used in many authentication systems, such as the Internet Key Exchange (IKEv2) Protocol [9], the presented architecture introduces a new protocol, called the Host Identity Protocol (HIP), and a cryptographic exchange, called the HIP base exchange; see also Section 8. The HIP protocols provide for limited forms of trust between systems, enhance mobility, multi-homing, and dynamic IP renumbering; aid in protocol translation/transition; and reduce certain types of denial-of-service (DoS) attacks.

When HIP is used, the actual payload traffic between two HIP hosts is typically, but not necessarily, protected with IPsec. The Host Identities are used to create the needed IPsec Security Associations (SAs) and to authenticate the hosts. When IPsec is used, the actual payload IP packets do not differ in any way from standard IPsec-protected IP packets.

3. Terminology

3.1. Terms Common to Other Documents

Term	Explanation
public key	The public key of an asymmetric cryptographic key pair. Used as a publicly known identifier for cryptographic identity authentication.
Private key	The private or secret key of an asymmetric cryptographic key pair. Assumed to be known only to the party identified by the corresponding public key. Used by the identified party to authenticate its identity to other parties.
public key pair	An asymmetric cryptographic key pair consisting of public and private keys. For example, Rivest-Shamir-Adelman (RSA) and Digital Signature Algorithm (DSA) key pairs are such key pairs.
end-point	A communicating entity. For historical reasons, the term 'computing platform' is used in this document as a (rough) synonym for end-point.

3.2. Terms Specific to This and Other HIP Documents

It should be noted that many of the terms defined herein are tautologous, self-referential, or defined through circular reference to other terms. This is due to the succinct nature of the definitions. See the text elsewhere in this document for more elaborate explanations.

Term	Explanation
computing platform	An entity capable of communicating and computing, for example, a computer. See the definition of 'end-point', above.
HIP base exchange	A cryptographic protocol; see also Section 8.
HIP packet	An IP packet that carries a 'Host Identity Protocol' message.
Host Identity	An abstract concept assigned to a 'computing platform'. See 'Host Identifier', below.
Host Identity namespace	A namespace formed by all possible Host Identifiers.
Host Identity Protocol	A protocol used to carry and authenticate Host Identifiers and other information.
Host Identity Tag	A 128-bit datum created by taking a cryptographic hash over a Host Identifier.
Host Identifier	A public key used as a name for a Host Identity.
Local Scope Identifier	A 32-bit datum denoting a Host Identity.
Public Host Identifier and Identity	A published or publicly known Host Identifier used as a public name for a Host Identity, and the corresponding Identity.
Unpublished Host Identifier and Identity	A Host Identifier that is not placed in any public directory, and the corresponding Host Identity. Unpublished Host Identities are typically shortlived in nature, being often replaced and possibly used just once.
Rendezvous Mechanism	A mechanism used to locate mobile hosts based on their Host Identity Tag (HIT).

4. Background

The Internet is built from three principal components: computing platforms (end-points), packet transport (i.e., internetworking) infrastructure, and services (applications). The Internet exists to service two principal components: people and robotic services (silicon-based people, if you will). All these components need to be named in order to interact in a scalable manner. Here we concentrate on naming computing platforms and packet transport elements.

There are two principal namespaces in use in the Internet for these components: IP numbers and Domain Names. Domain Names provide hierarchically assigned names for some computing platforms and some services. Each hierarchy is delegated from the level above; there is no anonymity in Domain Names. Email, HTTP, and SIP addresses all reference Domain Names.

IP numbers are a confounding of two namespaces, the names of a host's networking interfaces and the names of the locations ('confounding' is a term used in statistics to discuss metrics that are merged into one with a gain in indexing, but a loss in informational value). The names of locations should be understood as denoting routing direction vectors, i.e., information that is used to deliver packets to their destinations.

IP numbers name networking interfaces, and typically only when the interface is connected to the network. Originally, IP numbers had long-term significance. Today, the vast number of interfaces use ephemeral and/or non-unique IP numbers. That is, every time an interface is connected to the network, it is assigned an IP number.

In the current Internet, the transport layers are coupled to the IP addresses. Neither can evolve separately from the other. IPng deliberations were strongly shaped by the decision that a corresponding TCPng would not be created.

There are three critical deficiencies with the current namespaces. First, dynamic readdressing cannot be directly managed. Second, anonymity is not provided in a consistent, trustable manner. Finally, authentication for systems and datagrams is not provided. All of these deficiencies arise because computing platforms are not well named with the current namespaces.

4.1. A Desire for a Namespace for Computing Platforms

An independent namespace for computing platforms could be used in end-to-end operations independent of the evolution of the internetworking layer and across the many internetworking layers.

This could support rapid readdressing of the internetworking layer because of mobility, rehomining, or renumbering.

If the namespace for computing platforms is based on public key cryptography, it can also provide authentication services. If this namespace is locally created without requiring registration, it can provide anonymity.

Such a namespace (for computing platforms) and the names in it should have the following characteristics:

- o The namespace should be applied to the IP 'kernel'. The IP kernel is the 'component' between applications and the packet transport infrastructure.
- o The namespace should fully decouple the internetworking layer from the higher layers. The names should replace all occurrences of IP addresses within applications (like in the Transport Control Block, TCB). This may require changes to the current APIs. In the long run, it is probable that some new APIs are needed.
- o The introduction of the namespace should not mandate any administrative infrastructure. Deployment must come from the bottom up, in a pairwise deployment.
- o The names should have a fixed-length representation, for easy inclusion in datagram headers and existing programming interfaces (e.g., the TCB).
- o Using the namespace should be affordable when used in protocols. This is primarily a packet size issue. There is also a computational concern in affordability.
- o Name collisions should be avoided as much as possible. The mathematics of the birthday paradox can be used to estimate the chance of a collision in a given population and hash space. In general, for a random hash space of size n bits, we would expect to obtain a collision after approximately $1.2 \cdot \sqrt{2^n}$ hashes were obtained. For 64 bits, this number is roughly 4 billion. A hash size of 64 bits may be too small to avoid collisions in a large population; for example, there is a 1% chance of collision in a population of 640M. For 100 bits (or more), we would not expect a collision until approximately 2^{50} (1 quadrillion) hashes were generated.
- o The names should have a localized abstraction that can be used in existing protocols and APIs.

- o It must be possible to create names locally. This can provide anonymity at the cost of making resolvability very difficult.
 - * Sometimes the names may contain a delegation component. This is the cost of resolvability.
- o The namespace should provide authentication services.
- o The names should be long-lived, but replaceable at any time. This impacts access control lists; short lifetimes will tend to result in tedious list maintenance or require a namespace infrastructure for central control of access lists.

In this document, a new namespace approaching these ideas is called the Host Identity namespace. Using Host Identities requires its own protocol layer, the Host Identity Protocol, between the internetworking and transport layers. The names are based on public key cryptography to supply authentication services. Properly designed, it can deliver all of the above-stated requirements.

5. Host Identity Namespace

A name in the Host Identity namespace, a Host Identifier (HI), represents a statistically globally unique name for naming any system with an IP stack. This identity is normally associated with, but not limited to, an IP stack. A system can have multiple identities, some 'well known', some unpublished or 'anonymous'. A system may self-assert its own identity, or may use a third-party authenticator like DNS Security (DNSSEC) [2], Pretty Good Privacy (PGP), or X.509 to 'notarize' the identity assertion. It is expected that the Host Identifiers will initially be authenticated with DNSSEC and that all implementations will support DNSSEC as a minimal baseline.

In theory, any name that can claim to be 'statistically globally unique' may serve as a Host Identifier. However, in the authors' opinion, a public key of a 'public key pair' makes the best Host Identifier. As will be specified in the Host Identity Protocol specification, a public-key-based HI can authenticate the HIP packets and protect them from man-in-the-middle attacks. Since authenticated datagrams are mandatory to provide much of HIP's DoS protection, the Diffie-Hellman exchange in HIP has to be authenticated. Thus, only public key HI and authenticated HIP messages are supported in practice. In this document, the non-cryptographic forms of HI and HIP are presented to complete the theory of HI, but they should not be implemented as they could produce worse DoS attacks than the Internet has without Host Identity.

5.1. Host Identifiers

Host Identity adds two main features to Internet protocols. The first is a decoupling of the internetworking and transport layers; see Section 6. This decoupling will allow for independent evolution of the two layers. In addition, it can provide end-to-end services over multiple internetworking realms. The second feature is host authentication. Because the Host Identifier is a public key, this key can be used for authentication in security protocols like IPsec.

The only completely defined structure of the Host Identity is that of a public/private key pair. In this case, the Host Identity is referred to by its public component, the public key. Thus, the name representing a Host Identity in the Host Identity namespace, i.e., the Host Identifier, is the public key. In a way, the possession of the private key defines the Identity itself. If the private key is possessed by more than one node, the Identity can be considered to be a distributed one.

Architecturally, any other Internet naming convention might form a usable base for Host Identifiers. However, non-cryptographic names should only be used in situations of high trust / low risk, that is, any place where host authentication is not needed (no risk of host spoofing and no use of IPsec). However, at least for interconnected networks spanning several operational domains, the set of environments where the risk of host spoofing allowed by non-cryptographic Host Identifiers is acceptable is the null set. Hence, the current HIP documents do not specify how to use any other types of Host Identifiers but public keys.

The actual Host Identities are never directly used in any Internet protocols. The corresponding Host Identifiers (public keys) may be stored in various DNS or Lightweight Directory Access Protocol (LDAP) directories as identified elsewhere in this document, and they are passed in the HIP base exchange. A Host Identity Tag (HIT) is used in other protocols to represent the Host Identity. Another representation of the Host Identities, the Local Scope Identifier (LSI), can also be used in protocols and APIs.

5.2. Storing Host Identifiers in DNS

The public Host Identifiers should be stored in DNS; the unpublished Host Identifiers should not be stored anywhere (besides the communicating hosts themselves). The (public) HI is stored in a new Resource Record (RR) type, to be defined. This RR type is likely to be quite similar to the IPSECKEY RR [6].

Alternatively, or in addition to storing Host Identifiers in the DNS, they may be stored in various kinds of Public Key Infrastructure (PKI). Such a practice may allow them to be used for purposes other than pure host identification.

5.3. Host Identity Tag (HIT)

A Host Identity Tag is a 128-bit representation for a Host Identity. It is created by taking a cryptographic hash over the corresponding Host Identifier. There are two advantages of using a hash over using the Host Identifier in protocols. First, its fixed length makes for easier protocol coding and also better manages the packet size cost of this technology. Second, it presents the identity in a consistent format to the protocol independent of the cryptographic algorithms used.

In the HIP packets, the HITs identify the sender and recipient of a packet. Consequently, a HIT should be unique in the whole IP universe as long as it is being used. In the extremely rare case of a single HIT mapping to more than one Host Identity, the Host Identifiers (public keys) will make the final difference. If there is more than one public key for a given node, the HIT acts as a hint for the correct public key to use.

5.4. Local Scope Identifier (LSI)

A Local Scope Identifier (LSI) is a 32-bit localized representation for a Host Identity. The purpose of an LSI is to facilitate using Host Identities in existing protocols and APIs. LSI's advantage over HIT is its size; its disadvantage is its local scope.

Examples of how LSIs can be used include: as the address in an FTP command and as the address in a socket call. Thus, LSIs act as a bridge for Host Identities into IPv4-based protocols and APIs.

6. New Stack Architecture

One way to characterize Host Identity is to compare the proposed new architecture with the current one. As discussed above, the IP addresses can be seen to be a confounding of routing direction vectors and interface names. Using the terminology from the IRTF Name Space Research Group Report [7] and, e.g., the unpublished Internet Draft "Endpoints and Endpoint Names" [10] by Noel Chiappa, the IP addresses currently embody the dual role of locators and end-point identifiers. That is, each IP address names a topological location in the Internet, thereby acting as a routing direction vector, or locator. At the same time, the IP address names the physical network interface currently located at the point-of-attachment, thereby acting as an end-point name.

In the HIP architecture, the end-point names and locators are separated from each other. IP addresses continue to act as locators. The Host Identifiers take the role of end-point identifiers. It is important to understand that the end-point names based on Host Identities are slightly different from interface names; a Host Identity can be simultaneously reachable through several interfaces.

The difference between the bindings of the logical entities is illustrated in Figure 1.




Figure 1

6.1. Transport Associations and End-points

Architecturally, HIP provides for a different binding of transport-layer protocols. That is, the transport-layer associations, i.e., TCP connections and UDP associations, are no longer bound to IP addresses but to Host Identities.

It is possible that a single physical computer hosts several logical end-points. With HIP, each of these end-points would have a distinct Host Identity. Furthermore, since the transport associations are bound to Host Identities, HIP provides for process migration and clustered servers. That is, if a Host Identity is moved from one physical computer to another, it is also possible to simultaneously move all the transport associations without breaking them. Similarly, if it is possible to distribute the processing of a single Host Identity over several physical computers, HIP provides for cluster-based services without any changes at the client end-point.

7. End-host Mobility and Multi-homing

HIP decouples the transport from the internetworking layer, and binds the transport associations to the Host Identities (through actually either the HIT or LSI). Consequently, HIP can provide for a degree of internetworking mobility and multi-homing at a low infrastructure cost. HIP mobility includes IP address changes (via any method) to either party. Thus, a system is considered mobile if its IP address can change dynamically for any reason like PPP, Dynamic Host Configuration Protocol (DHCP), IPv6 prefix reassignments, or a Network Address Translation (NAT) device remapping its translation. Likewise, a system is considered multi-homed if it has more than one globally routable IP address at the same time. HIP links IP addresses together, when multiple IP addresses correspond to the same Host Identity, and if one address becomes unusable, or a more preferred address becomes available, existing transport associations can easily be moved to another address.

When a node moves while communication is already ongoing, address changes are rather straightforward. The peer of the mobile node can just accept a HIP or an integrity protected IPsec packet from any address and ignore the source address. However, as discussed in Section 7.2 below, a mobile node must send a HIP readdress packet to inform the peer of the new address(es), and the peer must verify that the mobile node is reachable through these addresses. This is especially helpful for those situations where the peer node is sending data periodically to the mobile node (that is restarting a connection after the initial connection).

7.1. Rendezvous Mechanism

Making a contact to a mobile node is slightly more involved. In order to start the HIP exchange, the initiator node has to know how to reach the mobile node. Although infrequently moving HIP nodes could use Dynamic DNS [1] to update their reachability information in the DNS, an alternative to using DNS in this fashion is to use a piece of new static infrastructure to facilitate rendezvous between HIP nodes.

The mobile node keeps the rendezvous infrastructure continuously updated with its current IP address(es). The mobile nodes must trust the rendezvous mechanism to properly maintain their HIT and IP address mappings.

The rendezvous mechanism is also needed if both of the nodes happen to change their address at the same time, either because they are mobile and happen to move at the same time, because one of them is off-line for a while, or because of some other reason. In such a case, the HIP readdress packets will cross each other in the network and never reach the peer node.

A separate document will specify the details of the HIP rendezvous mechanism.

7.2. Protection against Flooding Attacks

Although the idea of informing about address changes by simply sending packets with a new source address appears appealing, it is not secure enough. That is, even if HIP does not rely on the source address for anything (once the base exchange has been completed), it appears to be necessary to check a mobile node's reachability at the new address before actually sending any larger amounts of traffic to the new address.

Blindly accepting new addresses would potentially lead to flooding DoS attacks against third parties [8]. In a distributed flooding attack, an attacker opens high-volume HIP connections with a large number of hosts (using unpublished HIs), and then claims to all of these hosts that it has moved to a target node's IP address. If the peer hosts were to simply accept the move, the result would be a packet flood to the target node's address. To close this attack, HIP includes an address check mechanism where the reachability of a node is separately checked at each address before using the address for larger amounts of traffic.

Whenever HIP is used between two hosts that fully trust each other, the hosts may optionally decide to skip the address tests. However,

such performance optimization must be restricted to peers that are known to be trustworthy and capable of protecting themselves from malicious software.

8. HIP and IPsec

The preferred way of implementing HIP is to use IPsec to carry the actual data traffic. As of today, the only completely defined method is to use IPsec Encapsulating Security Payload (ESP) to carry the data packets. In the future, other ways of transporting payload data may be developed, including ones that do not use cryptographic protection.

In practice, the HIP base exchange uses the cryptographic Host Identifiers to set up a pair of ESP Security Associations (SAs) to enable ESP in an end-to-end manner. This is implemented in a way that can span addressing realms.

While it would be possible, at least in theory, to use some existing cryptographic protocol, such as IKEv2 together with Host Identifiers, to establish the needed SAs, HIP defines a new protocol. There are a number of historical reasons for this, and there are also a few architectural reasons. First, IKE and IKEv2 were not designed with middle boxes in mind. As adding a new naming layer allows one to potentially add a new forwarding layer (see Section 9, below), it is very important that the HIP protocols are friendly toward any middle boxes.

Second, from a conceptual point of view, the IPsec Security Parameter Index (SPI) in ESP provides a simple compression of the HITs. This does require per-HIT-pair SAs (and SPIs), and a decrease of policy granularity over other Key Management Protocols, such as IKE and IKEv2. In particular, the current thinking is limited to a situation where, conceptually, there is only one pair of SAs between any given pair of HITs. In other words, from an architectural point of view, HIP only supports host-to-host (or endpoint-to-endpoint) Security Associations. If two hosts need more pairs of parallel SAs, they should use separate HITs for that. However, future HIP extensions may provide for more granularity and creation of several ESP SAs between a pair of HITs.

Since HIP is designed for host usage, not for gateways or so-called Bump-in-the-Wire (BITW) implementations, only ESP transport mode is supported. An ESP SA pair is indexed by the SPIs and the two HITs (both HITs since a system can have more than one HIT). The SAs need not be bound to IP addresses; all internal control of the SA is by the HITs. Thus, a host can easily change its address using Mobile IP, DHCP, PPP, or IPv6 readdressing and still maintain the SAs.

Since the transports are bound to the SA (via an LSI or a HIT), any active transport is also maintained. Thus, real-world conditions like loss of a PPP connection and its re-establishment or a mobile handover will not require a HIP negotiation or disruption of transport services [12].

Since HIP does not negotiate any SA lifetimes, all lifetimes are local policy. The only lifetimes a HIP implementation must support are sequence number rollover (for replay protection) and SA timeout. An SA times out if no packets are received using that SA. Implementations may support lifetimes for the various ESP transforms.

9. HIP and NATs

Passing packets between different IP addressing realms requires changing IP addresses in the packet header. This may happen, for example, when a packet is passed between the public Internet and a private address space, or between IPv4 and IPv6 networks. The address translation is usually implemented as Network Address Translation (NAT) [4] or NAT Protocol Translation (NAT-PT) [3].

In a network environment where identification is based on the IP addresses, identifying the communicating nodes is difficult when NAT is used. With HIP, the transport-layer end-points are bound to the Host Identities. Thus, a connection between two hosts can traverse many addressing realm boundaries. The IP addresses are used only for routing purposes; they may be changed freely during packet traversal.

For a HIP-based flow, a HIP-aware NAT or NAT-PT system tracks the mapping of HITs, and the corresponding IPsec SPIs, to an IP address. The NAT system has to learn mappings both from HITs and from SPIs to IP addresses. Many HITs (and SPIs) can map to a single IP address on a NAT, simplifying connections on address-poor NAT interfaces. The NAT can gain much of its knowledge from the HIP packets themselves; however, some NAT configuration may be necessary.

NAT systems cannot touch the datagrams within the IPsec envelope; thus, application-specific address translation must be done in the end systems. HIP provides for 'Distributed NAT', and uses the HIT or the LSI as a placeholder for embedded IP addresses.

9.1. HIP and TCP Checksums

There is no way for a host to know if any of the IP addresses in an IP header are the addresses used to calculate the TCP checksum. That is, it is not feasible to calculate the TCP checksum using the actual IP addresses in the pseudo header; the addresses received in the incoming packet are not necessarily the same as they were on the

sending host. Furthermore, it is not possible to recompute the upper-layer checksums in the NAT/NAT-PT system, since the traffic is IPsec protected. Consequently, the TCP and UDP checksums are calculated using the HITs in the place of the IP addresses in the pseudo header. Furthermore, only the IPv6 pseudo header format is used. This provides for IPv4/IPv6 protocol translation.

10. Multicast

Back in the Fall of 2003, there were little if any concrete thoughts about how HIP might affect IP-layer or application-layer multicast.

11. HIP Policies

There are a number of variables that will influence the HIP exchanges that each host must support. All HIP implementations should support at least 2 HIs, one to publish in DNS and an unpublished one for anonymous usage. Although unpublished HIs will be rarely used as responder HIs, they are likely to be common for initiators. Support for multiple HIs is recommended.

Many initiators would want to use a different HI for different responders. The implementations should provide for a policy of initiator HIT to responder HIT. This policy should also include preferred transforms and local lifetimes.

Responders would need a similar policy, describing the hosts allowed to participate in HIP exchanges, and the preferred transforms and local lifetimes.

12. Benefits of HIP

In the beginning, the network layer protocol (i.e., IP) had the following four "classic" invariants:

- o Non-mutable: The address sent is the address received.
- o Non-mobile: The address does not change during the course of an "association".
- o Reversible: A return header can always be formed by reversing the source and destination addresses.
- o Omniscient: Each host knows what address a partner host can use to send packets to it.

Actually, the fourth can be inferred from 1 and 3, but it is worth mentioning for reasons that will be obvious soon if not already.

In the current "post-classic" world, we are intentionally trying to get rid of the second invariant (both for mobility and for multi-homing), and we have been forced to give up the first and the fourth. Realm Specific IP [5] is an attempt to reinstate the fourth invariant without the first invariant. IPv6 is an attempt to reinstate the first invariant.

Few systems on the Internet have DNS names that are meaningful. That is, if they have a Fully Qualified Domain Name (FQDN), that name typically belongs to a NAT device or a dial-up server, and does not really identify the system itself but its current connectivity. FQDNs (and their extensions as email names) are application-layer names, more frequently naming services than a particular system. This is why many systems on the Internet are not registered in the DNS; they do not have services of interest to other Internet hosts.

DNS names are references to IP addresses. This only demonstrates the interrelationship of the networking and application layers. DNS, as the Internet's only deployed, distributed database, is also the repository of other namespaces, due in part to DNSSEC-specific and application-specific key records. Although each namespace can be stretched (IP with v6, DNS with KEY records), neither can adequately provide for host authentication or act as a separation between internetworking and transport layers.

The Host Identity (HI) namespace fills an important gap between the IP and DNS namespaces. An interesting thing about the HI is that it actually allows one to give up all but the 3rd network-layer invariant. That is to say, as long as the source and destination addresses in the network-layer protocol are reversible, then things work OK because HIP takes care of host identification, and reversibility allows one to get a packet back to one's partner host. You do not care if the network-layer address changes in transit (mutable), and you do not care what network-layer address the partner is using (non-omniscient).

12.1. HIP's Answers to NSRG Questions

The IRTF Name Space Research Group has posed a number of evaluating questions in its report [7]. In this section, we provide answers to these questions.

1. How would a stack name improve the overall functionality of the Internet?

HIP decouples the internetworking layer from the transport layer, allowing each to evolve separately. The decoupling makes end-host mobility and multi-homing easier, also across

IPv4 and IPv6 networks. HIs make network renumbering easier, and they also make process migration and clustered servers easier to implement. Furthermore, being cryptographic in nature, they provide the basis for solving the security problems related to end-host mobility and multi-homing.

2. What does a stack name look like?

A HI is a cryptographic public key. However, instead of using the keys directly, most protocols use a fixed-size hash of the public key.

3. What is its lifetime?

HIP provides both stable and temporary Host Identifiers. Stable HIs are typically long-lived, with a lifetime of years or more. The lifetime of temporary HIs depends on how long the upper-layer connections and applications need them, and can range from a few seconds to years.

4. Where does it live in the stack?

The HIs live between the transport and internetworking layers.

5. How is it used on the end-points?

The Host Identifiers may be used directly or indirectly (in the form of HITs or LSIs) by applications when they access network services. In addition, the Host Identifiers, as public keys, are used in the built-in key agreement protocol, called the HIP base exchange, to authenticate the hosts to each other.

6. What administrative infrastructure is needed to support it?

In some environments, it is possible to use HIP opportunistically, without any infrastructure. However, to gain full benefit from HIP, the HIs must be stored in the DNS or a PKI, and a new rendezvous mechanism is needed. Such a new rendezvous mechanism may need new infrastructure to be deployed.

7. If we add an additional layer, would it make the address list in Stream Control Transmission Protocol (SCTP) unnecessary?

Yes.

8. What additional security benefits would a new naming scheme offer?

HIP reduces dependency on IP addresses, making the so-called address ownership [11] problems easier to solve. In practice, HIP provides security for end-host mobility and multi-homing. Furthermore, since HIP Host Identifiers are public keys, standard public key certificate infrastructures can be applied on the top of HIP.

9. What would the resolution mechanisms be, or what characteristics of a resolution mechanisms would be required?

For most purposes, an approach where DNS names are resolved simultaneously to HIs and IP addresses is sufficient. However, if it becomes necessary to resolve HIs into IP addresses or back to DNS names, a flat resolution infrastructure is needed. Such an infrastructure could be based on the ideas of Distributed Hash Tables, but would require significant new development and deployment.

13. Security Considerations

HIP takes advantage of the new Host Identity paradigm to provide secure authentication of hosts and to provide a fast key exchange for IPsec. HIP also attempts to limit the exposure of the host to various Denial-of-Service (DoS) and Man-in-the-Middle (MitM) attacks. In so doing, HIP itself is subject to its own DoS and MitM attacks that potentially could be more damaging to a host's ability to conduct business as usual.

Resource-exhausting DoS attacks take advantage of the cost of setting up a state for a protocol on the responder compared to the 'cheapness' on the initiator. HIP allows a responder to increase the cost of the start of state on the initiator and makes an effort to reduce the cost to the responder. This is done by having the responder start the authenticated Diffie-Hellman exchange instead of the initiator, making the HIP base exchange 4 packets long. There are more details on this process in the Host Identity Protocol.

HIP optionally supports opportunistic negotiation. That is, if a host receives a start of transport without a HIP negotiation, it can attempt to force a HIP exchange before accepting the connection. This has the potential for DoS attacks against both hosts. If the method to force the start of HIP is expensive on either host, the attacker need only spoof a TCP SYN. This would put both systems into the expensive operations. HIP avoids this attack by having the responder send a simple HIP packet that it can pre-build. Since this

packet is fixed and easily replayed, the initiator reacts to it only if it has just started a connection to the responder.

MitM attacks are difficult to defend against, without third-party authentication. A skillful MitM could easily handle all parts of the HIP base exchange, but HIP indirectly provides the following protection from an MitM attack. If the responder's HI is retrieved from a signed DNS zone or secured by some other means, the initiator can use this to authenticate the signed HIP packets. Likewise, if the initiator's HI is in a secure DNS zone, the responder can retrieve it and validate the signed HIP packets. However, since an initiator may choose to use an unpublished HI, it knowingly risks an MitM attack. The responder may choose not to accept a HIP exchange with an initiator using an unknown HI.

In HIP, the Security Association for IPsec is indexed by the SPI; the source address is always ignored, and the destination address may be ignored as well. Therefore, HIP-enabled IPsec Encapsulated Security Payload (ESP) is IP address independent. This might seem to make it easier for an attacker, but ESP with replay protection is already as well protected as possible, and the removal of the IP address as a check should not increase the exposure of IPsec ESP to DoS attacks.

Since not all hosts will ever support HIP, ICMPv4 'Destination Unreachable, Protocol Unreachable' and ICMPv6 'Parameter Problem, Unrecognized Next Header' messages are to be expected and present a DoS attack. Against an initiator, the attack would look like the responder does not support HIP, but shortly after receiving the ICMP message, the initiator would receive a valid HIP packet. Thus, to protect against this attack, an initiator should not react to an ICMP message until a reasonable time has passed, allowing it to get the real responder's HIP packet. A similar attack against the responder is more involved.

Another MitM attack is simulating a responder's administrative rejection of a HIP initiation. This is a simple ICMP 'Destination Unreachable, Administratively Prohibited' message. A HIP packet is not used because it would have to either have unique content, and thus difficult to generate, resulting in yet another DoS attack, or be just as spoofable as the ICMP message. Like in the previous case, the defense against this attack is for the initiator to wait a reasonable time period to get a valid HIP packet. If one does not come, then the initiator has to assume that the ICMP message is valid. Since this is the only point in the HIP base exchange where this ICMP message is appropriate, it can be ignored at any other point in the exchange.

13.1. HITs Used in ACLs

It is expected that HITs will be used in Access Control Lists (ACLs). Future firewalls can use HITs to control egress and ingress to networks, with an assurance level difficult to achieve today. As discussed above in Section 8, once a HIP session has been established, the SPI value in an IPsec packet may be used as an index, indicating the HITs. In practice, firewalls can inspect HIP packets to learn of the bindings between HITs, SPI values, and IP addresses. They can even explicitly control IPsec usage, dynamically opening IPsec ESP only for specific SPI values and IP addresses. The signatures in HIP packets allow a capable firewall to ensure that the HIP exchange is indeed happening between two known hosts. This may increase firewall security.

There has been considerable bad experience with distributed ACLs that contain public-key-related material, for example, with Secure Shell Protocol (SSH). If the owner of a key needs to revoke it for any reason, the task of finding all locations where the key is held in an ACL may be impossible. If the reason for the revocation is due to private key theft, this could be a serious issue.

A host can keep track of all of its partners that might use its HIT in an ACL by logging all remote HITs. It should only be necessary to log responder hosts. With this information, the host can notify the various hosts about the change to the HIT. There has been no attempt to develop a secure method to issue the HIT revocation notice.

HIP-aware NATs, however, are transparent to the HIP-aware systems by design. Thus, the host may find it difficult to notify any NAT that is using a HIT in an ACL. Since most systems will know of the NATs for their network, there should be a process by which they can notify these NATs of the change of the HIT. This is mandatory for systems that function as responders behind a NAT. In a similar vein, if a host is notified of a change in a HIT of an initiator, it should notify its NAT of the change. In this manner, NATs will get updated with the HIT change.

13.2. Non-security considerations

The definition of the Host Identifier states that the HI need not be a public key. It implies that the HI could be any value; for example, an FQDN. This document does not describe how to support such a non-cryptographic HI. A non-cryptographic HI would still offer the services of the HIT or LSI for NAT traversal. It would be possible to carry HITs in HIP packets that had neither privacy nor authentication. Since such a mode would offer so little additional functionality for so much addition to the IP kernel, it has not been

defined. Given how little public key cryptography HIP requires, HIP should only be implemented using public key Host Identities.

If it is desirable to use HIP in a low-security situation where public key computations are considered expensive, HIP can be used with very short Diffie-Hellman and Host Identity keys. Such use makes the participating hosts vulnerable to MitM and connection hijacking attacks. However, it does not cause flooding dangers, since the address check mechanism relies on the routing system and not on cryptographic strength.

14. Acknowledgements

For the people historically involved in the early stages of HIP, see the Acknowledgements section in the Host Identity Protocol specification.

During the later stages of this document, when the editing baton was transferred to Pekka Nikander, the comments from the early implementors and others, including Jari Arkko, Tom Henderson, Petri Jokela, Miika Komu, Mika Kousa, Andrew McGregor, Jan Melen, Tim Shepard, Jukka Ylitalo, and Jorma Wall, were invaluable. Finally, Lars Eggert, Spencer Dawkins, and Dave Crocker provided valuable input during the final stages of publication, most of which was incorporated but some of which the authors decided to ignore in order to get this document published in the first place.

15. Informative References

- [1] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [2] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005
- [3] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.

- [4] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [5] Borella, M., Lo, J., Grabelsky, D., and G. Montenegro, "Realm Specific IP: Framework", RFC 3102, October 2001.
- [6] Richardson, M., "A Method for Storing IPsec Keying Material in DNS", RFC 4025, March 2005.
- [7] Lear, E. and R. Droms, "What's In A Name: Thoughts from the NSRG", Work in Progress, September 2003.
- [8] Nikander, P., et al, "Mobile IP Version 6 Route Optimization Security Design Background", RFC 4225, December 2005.
- [9] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [10] Chiappa, J., "Endpoints and Endpoint Names: A Proposed Enhancement to the Internet Architecture", URL <http://users.exis.net/~jnc/tech/endpoints.txt>, 1999.
- [11] Nikander, P., "Denial-of-Service, Address Ownership, and Early Authentication in the IPv6 World", in Security Protocols, 9th International Workshop, Cambridge, UK, April 25-27 2001, LNCS 2467, pp. 12-26, Springer, 2002.
- [12] Bellovin, S., "EIDs, IPsec, and HostNAT", in Proceedings of the 41st IETF, Los Angeles, CA, March 1998.

Authors' Addresses

Robert Moskowitz
ICSALabs, a Division of Cybertrust Corporation
1000 Bent Creek Blvd, Suite 200
Mechanicsburg, PA
USA

Email: rgm@icsalabs.com

Pekka Nikander
Ericsson Research Nomadic Lab
JORVAS FIN-02420
FINLAND

Phone: +358 9 299 1
Email: pekka.nikander@nomadiclab.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

