A WWRF contribution to

- WG2 (enabling technologies and standards)
- WG3 (advanced network architectures),
- SIG2 (end-to-end security, identity management, new protocols), and
- SIG3 (potential tool for identity management, related to self-organization)

Contributing towards already identified research areas and reporting results from complementary research projects and activities.

Host Identity Protocol - Extended Abstract

Petri Jokela, Pekka Nikander, Jan Melen, Jukka Ylitalo, and Jorma Wall

Ericsson Research, NomadicLab
02420 Jorvas, Finland
email: <first>.<last>@ericsson.com

ABSTRACT

In the current Internet, hosts are identified using IP addresses that depend on the topological location of the hosts. In other words, the IP addresses are semantically overloaded since they identify both hosts and topological locations. The Host Identity Protocol (HIP) introduces a way of separating the location and host identity information. It introduces a new namespace, cryptographic in nature, for host identities. The IP addresses continue to be used for packet routing. In this paper we present the current status and latest development both at Ericsson Research and at the IETF.

I. INTRODUCTION

An IP address describes a topological location of a network interface, attached to the network. At the same time, the IP address is also used to identify the node hosting the interface, providing two mixed functions in a same thing. When mobility is added to the picture, the result is not pretty. Since IP addresses act as host identifiers, they must not be changed. However, since IP addresses describe topological locations, they must necessarily change when a host changes its location in the network. Obviously, it is impossible to achieve both stability and dynamic changes at the same time.

In the case of Mobile IP, the solution is to use a fixed home location providing a "home address" for the node. The home address both identifies the node and provides a stable location for it when it is at home. The current location information is available in the form of a care-of address, which is used for routing purposes when the node is away from home.

Another solution to the problem is to separate the identification and location functions from each other. One possible way is defined in the Host Identity Protocol (HIP) proposal [3][4]. HIP separates the location and identity roles of IP addresses by introducing a new name-space, the Host Identity. In HIP, the Host Identity is basically a public cryptographic key of a public-private key-pair. The public key identifies the party that holds the only copy of the private key. A host possessing the private key of the key-pair can directly prove that it ”owns” the public key that is used to identify it in the network. The separation also provides a means to handle mobility and multi-homing in a secure way.

A. Related work

There are other proposals that introduce similar ideas to introduce a better architecture in the Internet.

FARA [5] is a generalized model of ideas that provides a framework from which the actual architecture can be derived. The FARA model decouples the host identifier and location information without introducing a new global namespace. FARA could make use of the HIP when the node identifications are verified. Consequently, HIP could be a part of a particular FARA instantiation.

The PeerNet proposal [6] discusses the location and identity separation, but does not provide any solution for security. Each node has both identity and location information. The location information is not based on IP addresses, but it is defined to be a binary address tree. Routing is implemented using the bit-wise information in the locator. The host updates its location information to a suitable server from where the peer node can retrieve this information.

The Internet Indirection Infrastructure, I^3 [7], also defines a separation between the identity and routing information. The proposal concentrates on multicast environments, where the data is identified using an identifier. The receiver registers its IP address on the rendezvous server that is responsible for forwarding packets identified with the identifier to all parties that are registered to receive that particular data.

The rest of this paper is organized as follows. Section 2 gives a short introduction to the Host Identity Protocol: the new namespace, host identities, separating the identity from the location information and negotiating security associations between nodes. As the design of the HIP allows the routing information to be fully independent from the host identity, Section 3 shows the possibilities when HIP is used in combination with mobile and multi-homed hosts. Finally, Section 4 concludes the paper by discussing current status and recent developments.

II. Host Identity Protocol

The Host Identity Protocol introduces a separation between the location and identity information at the IP layer. In addition to the separation, a protocol is defined to negotiate security associations between HIP capable nodes.

A. The Separation Between the Identity and Location

If you are asked a question: "Who are you?" and you respond with your home *street address*, do you actually answer the question? However, the question is answered in an analogous way in the current Internet. When a host is identified, the IP address, providing the topological location of a node in the Internet, is given as the answer.

In real life, if you have to prove your identity and the asking person is unsure, you show your ID-card. Respectively, if you are asked to give your address, you will give the street address providing your (home) location. If this analogy is used in the Internet, the host identity and location information must be separated from each other. HIP provides one possible solution for decoupling the location from the identity.

When HIP is used, each host has identities, one or more, long-term or short-term, that can be used to identify it in the network. In HIP, the identifier is the public key of a public-private key pair. When the host possesses the private key, it can prove that it actually "owns" this identity that the public key represents. It is like showing an ID-card.

Each host can generate short-term public keys to be used only for a short time. These are handy when it is not necessary for the node to be identified with the same identity later. For example, buying books from a bookstore may be a long-term relationship, while once contacting a server that may collect user profiles may be considered to be a short-term action where the long-term identity is not wanted to be revealed.

The HIP Host Identity (HI), being a public key, is not practical in all actions; it is somewhat long. In HIP, the HI is represented with a 128-bit long Host Identity Tag (HIT) that is generated from the HI by hashing it. Thus, the HIT identifies a HI. Since the HIT is 128-bits long, it can be used for IPv6 applications directly as it is exactly the same length as IPv6 addresses.

When HIP is used, the upper layers, including the applications, do not see the IP address any longer. Instead, they see the HIT as the "address" of the destination host. The location information is hidden at a new layer, described in the next subsection.

B. A New Layer

Applications are typically not interested in location information but want to know the identity of their peers. To achieve this, HIP insulates the upper layers from IP addresses. Each host is represented by its HI, and the upper layers never see the actual IP addresses.

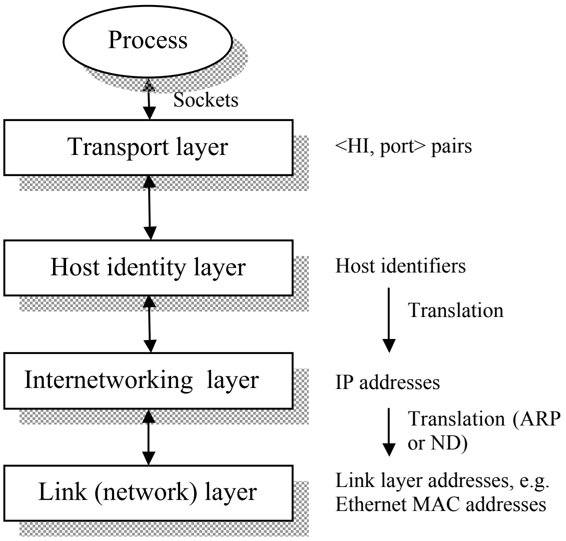

Figure 1. The proposed new architecture

Locally, each HI is mapped to the IP addresses of the node. When packets are leaving the host, the correct route is chosen and corresponding IP addresses are put into the packet as the source and destination addresses. How the path is chosen is a policy question and out of the scope of this paper.

HIP defines a base message exchange containing four messages, a four-way handshake. During this message exchange, a Diffie-Hellman authenticated key exchange is used to create a session key and to establish a pair of IPsec ESP Security Associations (SA) between the nodes. See Figure 2 for an overview of the four-way handshake; the details go beyond the scope of this paper.

The ESP SAs between the hosts are bound to the Host Identities. However, the packets travelling in the network do not contain the actual HI information, but the arriving packet is identified and mapped to the correct SA using the Security Parameter Index (SPI) value in the IPsec header. Figure 3 shows the logical and actual packet structures for packets in the network.

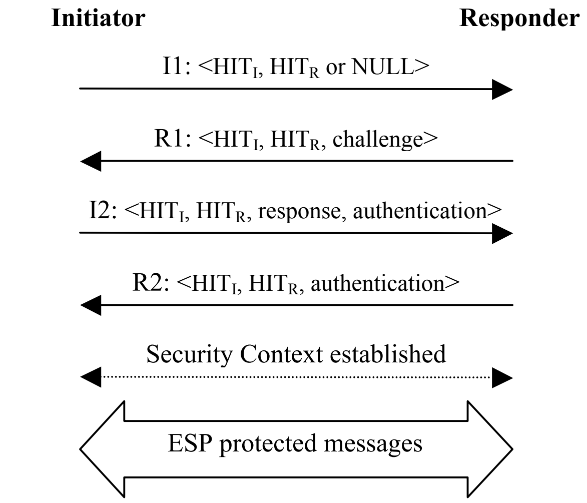
Initiator
Responder
I1: <HIT_I, HIT_R or NULL>
R1: <HIT_I, HIT_R, challenge>
I2: <HIT_I, HIT_R, response, authentication>
R2: <HIT_I, HIT_R, authentication>
Security Context established
ESP protected messages

Figure 2. A HIP session

From the previous it is clear that changing the location information in the packet does not generate any problems for the IPsec processing. The packet is still correctly identified using the SPI. If, for some reason, the packet is routed to a wrong destination, the receiver is not able to open the packet as it does not have the correct key.

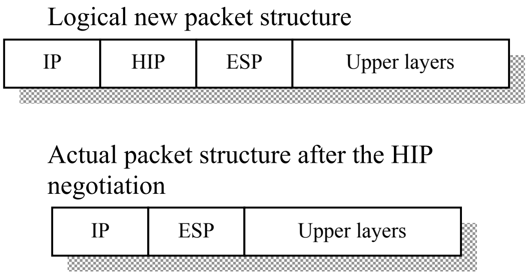
Logical new packet structure
IP
HIP
ESP
Upper layers
Actual packet structure after the HIP negotiation
IP
ESP
Upper layers

Figure 3. The packet structure

III. Mobility and Multi-homing

In this paper, we discuss the mobility and multi-homing from the end-host point of view. There are some similarities, but also differences, when the mobility concerns a whole network, i.e., network mobility. Network mobility, multi-homed hosts in a mobile network, and multi-homed mobile networks are, however, out of the scope of this paper.

A. Mobility

A mobile host can change the location inside one access network, between different access technologies, or even between different IP address realms. The most interesting handover happens in the latter case, when the host moves between IPv4 and IPv6 networks. In HIP, the application doesn't notice the change in the IP address version. The HI layer hides the change completely from upper layers. Of course, the peer node must be able to handle the location update that changes the IP version and packets must be routable using some compatible address. If a node does not have both IPv4 and IPv6 connectivity, it may use a proxy node that performs the address version conversion and provides connectivity on behalf of the node.

B. Multihoming

Multi-homing refers to a situation where an end-point has several parallel communication paths that it can use. Usually multi-homing is a result of either the host having several network interfaces (end-host multi-homing) or due to a network between the host and the rest of the network having redundant paths (site multi-homing). As said, in this paper we concentrate on the end-host multi-homing.

C. *Mobility with HIP*

With HIP, the separation between the location and identity information makes it clear that the packet identification and routing can be cleanly separated from each other. The host receiving a packet identifies the sender by first getting the correct key and then decrypting the packet. Thus, the actual IP addresses that were used for routing the packet are irrelevant.

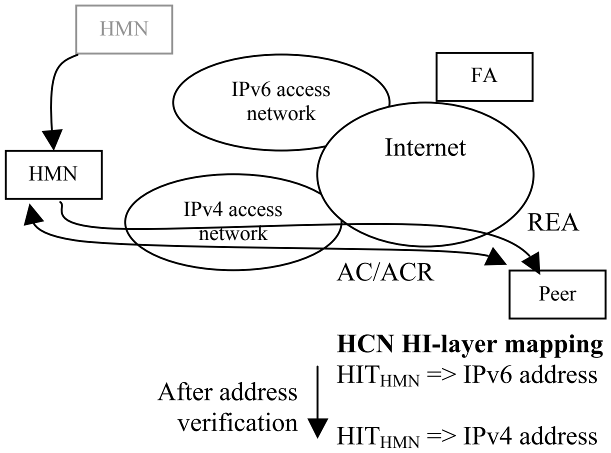

Figure 4. IPv6 - IPv4 handover

A HIP Mobile Host (HMN), moving in the network, may change the point of attachment to the Internet constantly. When the connection point is changed, also the IP address changes. This changed location information must be sent to the peer nodes (see Figure 4.). The same address can also be sent to a Forwarding Agent (FA) of the HMN, so that the HMN can be reached also via a more stable point. The DNS system is too slow to be used for constantly changing location information. Therefore, there must be a more stable address that can be used to contact the HMN. This address is the address provided by the FA.

The HIP mobility and multi-homing protocol defines a readdress (REA) parameter that contains the current IP address(es) of the HMN. When the HMN changes location and IP address, it generates an update packet with a REA parameter, signs the packet with the private key matching to the used HI, and sends the packet to the peer node and to the FA.

When the peer node receives a REA parameter, it must start an address verification process for the IP address(es) that are included in the parameter. Reachability verification is needed to avoid accepting false updates from the HMN.

Because the HMN can move between networks using different IP address versions, the address received by the peer may also be from different address family than the previous address. The peer may support only one IP address version. In this case, the peer node must use some other proxy node that can be used for routing packets over to the other IP address version network.

D. Host Multi-homing

A multi-homed HIP host, having multiple IP addresses configured on different interfaces connected to different access networks, has much more possibilities to handle the traffic towards a peer node. As it has multiple IP addresses presenting its current locations in the network, it may want to tell all of these addresses to its peer nodes. To do so, the multi-homed HIP node creates one or more REA parameters that contains all the addresses that it is able to use towards that particular node. This set of addresses may contain all addresses it has, or some subset of these addresses. When the peer node receives the REA with the multiple addresses, it must make verify the reachability of each of these addresses to avoid possible false updates.

False, or non-routable, addresses in the REA may be caused either because the HMN is malicious node, it has an error in the stack implementation, or the HMN node may be inside a network that uses private addresses that are not routable in the Internet.

Basically, a multi-homed HIP node is able to use all of the available connections, but efficient usage of the connections requires a policy system that has knowledge of the underlying access networks and can control the usage of them. Such a policy system can use different kinds of information: user preferences, operator preferences, input from the network connections, such as QoS, and so on. While we acknowledge the need for such a system, further considerations are out of the scope of this paper.

IV. Current status

A. Standardization and other collaborative activities

At the 58th IETF meeting, held in November 2003 in Minneapolis, it was decided to form a HIP working group. The working group will focus on finishing the current HIP protocol proposals and publishing them as experimental RFCs. This will make it possible to experiment HIP in a wider scale to see how HIP works in practice.

It has also been proposed to form a parallel research group at the Internet Research Task Force (IRTF), the research "branch" of the IETF. The research group would focus on studying how HIP and other similar alternatives would affect the Internet in the large.

HIP has been proposed as one starting point for the Ambient Networks project, a large collaborative European research project funded by the CEC under 6th framework. HIP is also an active research item at the Finnish Vertical Handover project, funded by Ericsson, Sonera, Tekes and others.

B. Prototyping at Ericsson Research

We have implemented the proposed HIP protocol, including the mobility and multi-homing functions. Our implementation uses the FreeBSD 5.2 operating system as the platform. Currently, the prototype implements the four-way handshake, IPsec ESP protection of all communications, mobility management, IPv4 – IPv6 interoperability, and multi-homing. As the design allows, our implementation does not care if the underlying IP address is from IPv4 or IPv6 realms. It can make handovers between access networks even when the IP address realm changes. The prototype is available for download at `http://www.hip4inter.net`

There are four other publicly known implementations. Our prototype has been tested against these other prototypes and the concept has been proven to work. Hosts are able to negotiate security associations and use the SA for secure communication.

REFERENCES

[1] C. Perkins, "IP Mobility Support for IPv4", RFC 3220, IETF, 2002.

[2] D. Johnson, C. Perkins, J. Arkko, "Mobility Support in IPv6", Internet Draft, work in progress, draft-ietf-mobileip-ipv6-24.txt, IETF, 2003.

[3] R. Moskowitz, P. Nikander, P. Jokela, "Host Identity Protocol", Internet Draft, work in progress, draft-moskowitz-hip-09.txt, IETF, 2004.

[4] Pekka Nikander, Jukka Ylitalo, and Jorma Wall, "Integrating Security, Mobility, and Multi-Homing in a HIP Way," NDSS'03, February 6-7, 2003, San Diego, CA, pp. 87-99, Internet Society, February, 2003.

[5] D. Clark, R. Braden, A. Falk, V. Pingali, "FARA: Reorganizing the Addressing Architecture", ACM SIGCOMM 2003 Workshops, August 25 & 27, 2003.

[6] J. Eriksson, M. Faloutsos, S. Krishnamurthy, "PeerNet: Pushing Peer-to-Peer Down the Stack", In IPTPS '03, February 20 - 21, 2003.

[7] I. Stoica, et.al., "Internet Indirection Infrastructure", ACM SIGCOMM '02, August 19-23, 2002.