Host Identity Protocol (HIP)
Research Group
Internet-Draft
Expires: April 18, 2005

L. Eggert
NEC
J. Laganier
LIP / Sun Microsystems
October 18, 2004

HIP Resolution and Rendezvous Problem Description
draft-eggert-hiprg-rr-prob-desc-00

Status of this Memo

This Internet-Draft will expire on April 18, 2005.

Copyright Notice

Abstract

This document investigates the design space for resolution and rendezvous mechanisms for the Host Identity Protocol (HIP.) It identifies and describes specific issues that HIP resolution and rendezvous mechanisms should address. These issues include

dependencies on the DNS, lack of support for direct communication based on host identities, lack of a reverse lookup mechanism for host identities, and DNS and node rendezvous.

This document does not propose specific resolution and rendezvous mechanisms. Different alternative solutions will be described and discussed in companion documents. These documents should analyze if and to what degree the specific proposals they present address the issues identified here.

Table of Contents

1. Introduction

The current Internet uses two global namespaces: domain names and IP addresses. The first namespace - domain names - has a single use. Domain names, usually simply called names, are symbolic identifiers for sets of numeric IP addresses, chosen for their mnemonic properties: humans need to interact with them.

IP addresses form the Internet's second global namespace. They have two uses. First, they are topological locators for network attachment points, addressing a specific location in the network topology. Their second use is as identifiers for the network interfaces - and thus nodes - that attach to the addressed locations. In this role as identifiers, IP addresses loose their topological meaning and become simple names. Routing and other network-layer mechanisms use the locator aspects of IP addresses. Transport-layer protocols and mechanisms typically use IP addresses in their role as names for communication endpoints. (Saltzer [6] discusses these naming concepts in detail.)

This dual use of IP addresses as both names and locators limits the flexibility of the Internet architecture. For example, the use of topology-dependent IP addresses as symbolic names for communication endpoints complicates node mobility. A mobile node changes its points of network attachment and hence its IP addresses dynamically. At the transport layer, this causes the logical endpoints of communication sessions - which are based on IP addresses - to change dynamically as well. Many of the Internet's transport protocols do not support changing the logical endpoints of established communication sessions. Arguably, they should not, because the identities of the communicating nodes have not changed, simply their points of network attachment.

The Host Identity Protocol (HIP) architecture defines a third global namespace [1]. The new host identity namespace decouples the name and locator roles currently filled by IP addresses. Host identities take over the naming role, while IP addresses become pure locators. With HIP, transport-layer mechanisms operate on host identities instead of using IP addresses as endpoint names. Network-layer mechanisms continue to use IP addresses as pure locators.

Due to the introduction of a new global namespace, HIP also affects the Internet's name resolution services. The Domain Name System

(DNS) is currently the Internet's only global resolution service [7]. The DNS provides a two-way lookup service between domain names and their set of corresponding IP addresses. HIP requires an additional resolution step. Domain names now map into sets of host identities, which in turn map into sets of IP addresses.

The additional HIP resolution step complicates the rendezvous procedure by which two nodes establish communication, i.e., the steps they need to perform until they obtain a peer's IP addresses. In the current Internet, the DNS maps the domain name of a target remote node directly into its set of IP addresses, which the local node may then use to address packets. The address of each node's DNS server is either manually configured or dynamically discovered (e.g., using DHCP [8]). When no DNS server is configured or has been discovered, nodes can still communicate by using IP addresses directly.

With HIP, the rendezvous procedure and resolution mechanisms are becoming more complex. The various alternatives for performing name and identity resolutions lead to rendezvous procedures that offer significantly different characteristics. This document discusses the limitations of the current HIP architecture and describes the general design space for alternative resolution and rendezvous mechanisms. It does not, however, present any specific resolution and rendezvous mechanisms. Specific alternatives will be described and discussed in companion documents.

This problem description document and its companion documents that describe specific resolution and rendezvous approaches obsolete prior contributions, i.e., [9].

2. HIP Resolution and Rendezvous

As mentioned in Section 1, HIP complicates the Internet's simple resolution and rendezvous procedures. Currently, nodes use DNS servers at well-known IP addresses to resolve domain names into IP addresses, which they can then use to address packets. The top illustration in Figure 1 shows this DNS resolution procedure. It also shows the reverse DNS resolution, which resolves an IP address back into its associated domain name.

With HIP, domain names map into sets of host identities, each of which maps into sets of IP addresses. This results in a logical two-step resolution process before a node knows the IP addresses associated with target domain name. The middle illustration in Figure 1 shows this two-step process. To maintain application compatibility, the first mapping - from names into host identities - should remain in the DNS. For the second mapping - from host identities into IP addresses - various alternatives are possible. Logically, this HIP lookup is a completely separate operation from

the initial DNS lookup. A reverse HIP lookup is also useful; it maps IP addresses back into their associated host identities.

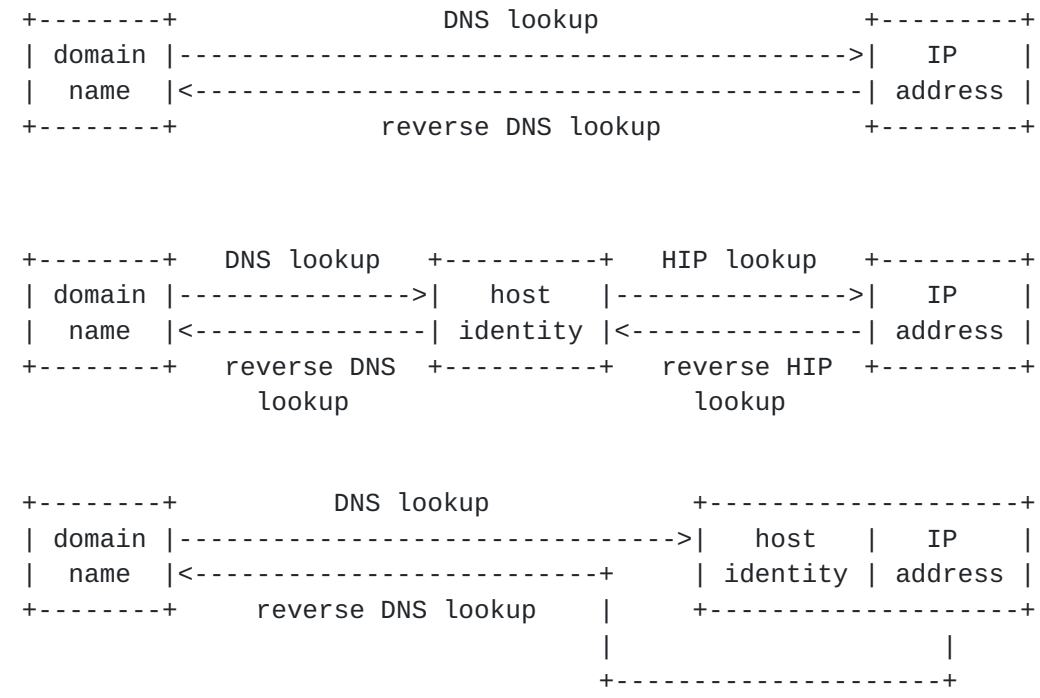

Figure 1: Domain name resolution without HIP (top), logical lookups with HIP (middle) and with the current HIP architecture (bottom)

Current HIP prototypes choose to maintain the second mapping between host identities and IP addresses in the DNS as well. One proposal simply stores a node's host identities alongside its IP addresses in the node's DNS record [2]. A DNS resolution of a domain name thus returns a pair of host identities and IP addresses, as shown in the bottom illustration of Figure 1. This simplistic approach creates several problems that the following sections discuss in more detail.

Companion documents to this document that propose specific mechanisms for resolution and rendezvous should investigate if and to what degree the individual proposals address these issues. [Comment.1]

2.1 Issue 1: DNS Dependency

One critical problem of storing host identities in a node's DNS record, as shown in the bottom of Figure 1, is that this approach creates a dependency between HIP and the DNS. To communicate with HIP under this approach, DNS resolution of a domain name is required to obtain a peer's host identities and IP addresses. It is not

possible to communicate with HIP based on host identities alone - no direct resolution mechanism exists to map host identities into IP addresses.

This is a drastic change from the current Internet, where the DNS is an optional component and communication can occur based on IP addresses alone.

2.2 Issue 2: Direct Communication

Storing host identities in a node's DNS record causes a second, related issue: even if a node already knows the host identity of a peer, it cannot use this host identity to initiate communication. It needs to know and resolve the peer's domain name to obtain the peer's IP addresses.

HIP allows protocols and services above the network layer to use host identities instead of IP-address-based names. Consequently, with HIP, host identities should replace many uses of IP addresses above the network layer. For example, applications and users should be able to substitute host identities wherever they now use IP addresses. A direct mechanism to resolve host identities into IP addresses, i.e., one that does not depend on knowledge of the corresponding domain name, is required to enable this transparency.

(Communication based on IP addresses alone is still possible with the simplistic HIP lookup, as shown on the bottom of Figure 1, but obviously will not incur the benefits of using HIP.)

2.3 Issue 3: Reverse Lookup

A third problem with the simplistic HIP resolution shown in the bottom of Figure 1 is that the DNS currently provides no mechanism to perform a reverse HIP lookup, i.e., determine the domain name of a node based on its host identity. Only the traditional reverse DNS lookup exists, which operates on IP addresses, not host identities. A reverse HIP lookup would need to be emulated by performing a reverse lookup on an IP address, and a forward DNS lookup on the resulting name to obtain the host identities. This is cumbersome.

Alternatively, reverse lookup capability could be added to the DNS through a new root, similar to how it provides reverse lookups on IP addresses. This approach, however, is problematic, because it maps resolution of a flat namespace into an hierarchical name system and also creates additional dependencies between HIP and the DNS [2].

2.4 Issue 4: DNS Rendezvous

Rendezvous with the DNS infrastructure is another issue with the simplistic HIP lookup. It may be useful to communicate with DNS servers using HIP instead of IP, i.e., access a DNS server through its well-known host identity instead of its well-known IP address. This would enable DNS servers to benefit from HIP's mobility, multihoming and security mechanisms.

The simplistic HIP lookup (bottom of Figure 1) requires a deployed

DNS infrastructure.

2.5 Issue 5: Node Rendezvous

Arguably the largest issue with the current HIP approach of using the DNS to store both the host identities and IP addresses associated with a name is rendezvous between nodes. As mentioned before, the rendezvous procedure encompasses all required steps for two nodes to obtain enough information about the other peer to be able to address packets to it. In most cases, this involves determining the set of a peer's IP addresses.

Without HIP, node rendezvous involves a DNS lookup of a peer's domain name to obtain its IP addresses. If the peer's addresses are known already, a host may skip the rendezvous procedure and address packets to it directly.

Consequently, different resolution mechanisms can have significantly different effects on node rendezvous. The following sections describe specific issues in detail.

2.5.1 Issue 5.1: Middlebox Traversal

A different document discusses middlebox traversal of HIP traffic [3], focusing on the current specification of the HIP protocols. This document discusses two separate problems: middlebox traversal of the HIP base exchange - i.e., the current rendezvous procedure - and traversal of HIP data traffic, which is currently carried inside IPsec.

New resolution and rendezvous solutions, which may consequently change the current base exchange, must consider how they interact with various middleboxes [10].

2.5.2 Issue 5.2: Location Privacy

Internet users are becoming more sensitive to privacy concerns. For example, the introduction of IPv6 already caused concern because of the possibility to trace users based on the unique EUI48 NIC identifiers included in their global IPv6 addresses. HIP may

potentially worsen the situation through its use of cryptographic, semi-permanent identifiers.

One approach to mitigating these concerns is through the periodic regeneration of host identities. Instead of reusing the same identity, nodes will generate new identities on the fly, similar to RFC 3041 [11]. This approach makes it more difficult to correlate a node's HIP associations and may thus reduce traceability concerns. A

second approach to increasing location privacy is concealing the IP addresses of two communicating nodes from one another. The SPI-multiplexed NAT (SPINAT) described as part of the BLIND framework offers this ability [12].

New resolution and rendezvous solutions should consider if and how they may provide location privacy or integrate with other mechanisms that do.

2.5.3 Issue 5.3: Mobility and Multihoming

HIP already includes mechanisms that allow a peer to signal its peers when its IP addresses have changed [4]. These mechanisms, however, require an already-established HIP association. For establishing new HIP associations after a move, the regular rendezvous procedure must complete. Consequently, a node must update its IP addresses after a move in whatever mechanism provides rendezvous service.

When host identities and IP addresses are stored in the DNS, dynamic DNS may provide a simple update mechanism [13][14]. However, the caching mechanisms of the DNS and to a lesser degree implementation issues with some current DNS servers make frequent dynamic DNS updates, i.e., support for highly mobile nodes, problematic.

A two-step resolution procedure, as described in Section 2 or dedicated external rendezvous servers [5] can improve HIP operation in such cases by offering a range of different update/lookup operations with different performance trade-offs.

One drawback of rendezvous servers is that they may introduce triangle routing: packets no longer follow the direct path between two peers, but instead flow through the rendezvous server. Besides increasing the end-to-end latency, this may decrease communication reliability. The two step resolution process, however, may be unable to support very high-rate mobility due to caching - it is impractical to translate from host identity to current IP address for every single packet.

New resolution and rendezvous solutions should discuss the trade-offs between update and lookup performance and routing inefficiencies versus move frequency.

2.5.4 Issue 5.4: Interoperation with Legacy Nodes

With HIP, node rendezvous breaks down into two distinct cases: rendezvous between two HIP nodes and rendezvous between a HIP node and a legacy non-HIP node. Specific resolution mechanisms for host identities may affect the two rendezvous cases in different ways.

For example, the current rendezvous server extensions to the base HIP protocol [5] do not support communication with legacy nodes, because they store the addresses of rendezvous servers in new "RVS" DNS record types that legacy nodes do not support [2]. For communication between two HIP nodes, this approach is successful, because they can use the RVS records to establish communication with the peer. A legacy node, however, expects to receive A or AAAA records that contain IP addresses the peer is directly reachable at.

New resolution and rendezvous solutions for HIP nodes should consider how they may provide rendezvous functionality with legacy non-HIP nodes. There are two aspects to this question. First, whether resolution mechanisms allow rendezvous with non-HIP nodes at all, and second, whether they can offer some of the benefits of HIP communication to non-HIP nodes as well.

3. Conclusion

This document described the design space of HIP resolution and rendezvous mechanisms and described a number of issues that the current architecture fails to support adequately. These issues include dependencies on the DNS, lack of support for direct communication based on host identities, lack of a reverse lookup mechanism for host identities, and DNS and node rendezvous.

This document does not propose specific resolution and rendezvous solutions; this will occur in future companion documents that should also describe how the solutions they propose address the issues identified here.

4. Security Considerations

The security aspects of HIP resolution and rendezvous mechanisms are currently being investigated and will be included in a future revision of this document.

5. Acknowledgments

Part of this work is a product of the Ambient Networks project, partially supported by the European Commission under its Sixth

Framework Programme. It is provided "as is" and without any express or implied warranties, including, without limitation, the implied warranties of fitness for a particular purpose. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Ambient Networks project or the European Commission.

6. References

6.1 Normative References

[1] Moskowitz, R., "Host Identity Protocol Architecture", draft-moskowitz-hip-arch-06 (work in progress), June 2004.

[2] Nikander, P. and J. Laganier, "Host Identity Protocol (HIP) Domain Name System (DNS) Extensions", draft-nikander-hip-dns-00 (work in progress), July 2004.

[3] Stiemerling, M., "Problem Statement: HIP operation over Network Address Translators", draft-stiemerling-hip-nat-01 (work in progress), July 2004.

[4] Nikander, P., "End-Host Mobility and Multi-Homing with Host Identity Protocol", draft-nikander-hip-mm-02 (work in progress), July 2004.

[5] Eggert, L., "Host Identity Protocol (HIP) Rendezvous Extensions", draft-eggert-hip-rvs-00 (work in progress), July 2004.

6.2 Informative References

[6] Saltzer, J., "On the Naming and Binding of Network Destinations", RFC 1498, August 1993.

[7] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[8] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[9] Eggert, L., "Host Identity Protocol (HIP) Rendezvous Mechanisms", draft-eggert-hip-rendezvous-01 (work in progress), July 2004.

[10] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.

[11] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[12] Ylitalo, J. and P. Nikander, "BLIND: A Complete Identity Protection Framework for End-points", Proc. Twelfth International Workshop on Security Protocols, Cambridge, England, April 2004.

[13] Vixie, P., Thomson, S., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.

[14] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.

Editorial Comments

[Comment.1] LE: The authors would appreciate feedback on the list of issues; specifically, whether it is complete and whether all of the currently identified issues are valid.

Authors' Addresses

Lars Eggert
NEC Network Laboratories
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 90511 43
Fax: +49 6221 90511 55
EMail: lars.eggert@netlab.nec.de
URI: http://www.netlab.nec.de/

Julien Laganier
Sun Labs (Sun Microsystems) & LIP (CNRS/INRIA/ENSL/UCBL)
180, Avenue de l'Europe
Saint Ismier CEDEX 38334
France

Phone: +33 476 188 815
EMail: ju@sun.com
URI: http://research.sun.com/

Appendix A. Document Revision History

Revision	Comments
00	Initial version. This document and its future companion documents that will propose and analyze specific solutions obsolete prior contributions [9].

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.