

Internet Draft
Document: [draft-stiemerling-hip-nat-01.txt](#)
Expires: January 2005

M. Stiemerling
J. Quittek
NEC Europe Ltd.

July 2004

Problem Statement: HIP operation over Network Address Translators

<[draft-stiemerling-hip-nat-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This memo investigates issues for Host Identity Protocol (HIP) nodes that communicate over a network path that includes Network Address Translators (NATs). There are two groups of issues: Operating HIP itself across NATs and operating the IPsec-based data transmission initiated by HIP across NATs. For both groups problems are summarized.

Internet-Draft

HIP and NAT

July 2004

Table of Contents

1	Introduction	3
2	Terminology	3
3	HIP Network Environment	4
4	Problems with operating HIP across NATs	5
4.1	HIP Base Exchange	5
4.1.1	HIP over IPv6	6
4.1.2	HIP over IPv4	6
4.2	IPsec Data Exchange	7
5	Extensions to HIP	7
6	Extension to NATs	8
7	HIP unaware NATs	8
8	Security Considerations	9
9	Acknowledgements	9
10	References	9
11	Authors' Addresses	11

1. Introduction

The Host Identity Protocol (HIP) architecture [[HIP-ARCH](#)] introduces a new kind of identifier for each host. Instead of using the IP address as host identity, a new Host Identifier is introduced. This additional namespace, separates the host identity from the routing information contained in the IP address.

The current version of the HIP Architecture is based on the assumption of full reachability between the IP addresses assigned to the participating hosts. This is usually given if both hosts have IP addresses belonging to the same address realm, for example, if both have public addresses. But if the host's addresses belong to different realms, for example one has a private IP address and the other one a public one, then network address translation must be provided between the realms. So far, network address translation is not integrated into the HIP architecture and HIP [[HIP-PROTO](#)] does not have particular means to support it.

Problems arising in conjunction with Network Address Translators (NATs) are two-fold. First the HIP base exchange needs to traverse NATs and then the IPsec encoded transport initiated by HIP needs to traverse it (see [[RFC2401](#)] for an introduction to IPsec). In general, both protocols have the risk of being blocked by NATs.

This document discusses the problems encountered when using HIP in the presence of NATs in [Section 4](#). Impacts on HIP itself and the used IPsec mechanism are elaborated. [Section 5](#) discusses extensions to HIP that potentially can solve some of the problems. [Section 6](#) suggests extensions to the functions provided by today's NATs complementing the HIP extensions.

It is not the intend of this memo to promoted the use of NATs in the context of HIP, but since NATs are widely deployed it is necessary to

take them into account while developing HIP.

2. Terminology

Terms used throughout this document are consistent with IPsec terminology [[RFC2401](#)], HIP terminology [[HIP-ARCH](#)] [[HIP-PROTO](#)] [[HIP-MM](#)], and NAT terminology [[RFC2663](#)]. For an introduction to NATs readers are recommended to read [[RFC2663](#)]. In this memo the term NAT refers to NAT and NAPT. So-called twice-NATs are not considered, yet.

3. HIP Network Environment

The HIP Architecture described in [[HIP-ARCH](#)] assumes full reachability between the IP addresses assigned to the participating hosts.




Figure 1: Assumed network environment for HIP

This is usually given if both hosts have IP addresses belonging to the same address realm, for instance, if both have public addresses or if both have addresses belonging to the same private address realm as shown by Figure 1. Within one realm, IP addresses provide end-to-end routing information between hosts A and B. If A wants to send a packet to B, it is sufficient for A to know the IP address assigned to B.

The scenario sketched by Figure 1 is an ideal one. It does not (anymore) match today's current scenarios where an increasing number of hosts use private IP addresses, mainly, because of lack of IP

addresses in IPv4. Furthermore, a separation of an organization's network from the public Internet realm hides the organization's network structure and enables more flexible internal renumbering of IP addresses.

If the host's addresses belong to different realms, then communication between these hosts requires network address translation between the realms. Network Address Translators (NATs, [\[RFC2663\]](#)) provide this service. Figure 2 shows an example of a host B with a private IP address and a host A with a public IP address. The scenario in Figure 2 is limited to a NAT on peer B's side only, but in general both peers might be located behind NATs.




Figure 2: Network environment with NAT

The translation performed by NATs affects the IP header and (in most existing cases) also the transport header. IP addresses and transport port numbers are translated between the address realms. In the scenario illustrated by Figure 2, host A would address all HIP messages that it intends to send to host B to a public IP address that is provided by the NAT. The NAT receives the message for this address (typically in combination with a dynamically chosen port number) and replaces the destination IP address (and typically also the TCP or UDP port number) before it forwards the packet into the private realm.

The public IP address (and port number) that the NAT uses for this translation service is in general not known to host A. Therefore, applications that carry address information in their payload or applications that rely on static protocol header fields experience severe problems. Often, the application-level signaling can traverse NATs, but subsequent data exchange fails (see [\[NATP2P\]](#)).

NATs are heavily used in IPv4 networks when there is a significant lack of address space. Initially, they were expected not to be present in IPv6 networks, but the authors believe that NATs will also be used in IPv6 networks, despite the availability of a much larger address space. Today, even organizations owning a sufficiently large IPv4 address space tend to use private Internet realms, because this hides their internal network structure and allows to renumber the internal network addresses without communicating this to the public network.

[4.](#) Problems with operating HIP across NATs

HIP operation can be divided into two phases. The first one is the HIP base exchange using HIP only. The second one is the actual application data exchange via IPsec. This section separately describes the problems that occur in each of the phases when network address translation is integrated into the HIP architecture.

[4.1.](#) HIP Base Exchange

The HIP base exchange uses different transport mechanisms for IPv6 and IPv4.

- o IPv6
When IPv6 is used, a HIP-specific IPv6 extension header carries all information necessary (see Section 6 of [[HIP-PROTO](#)]).
- o IPv4
When IPv4 is used, HIP messages are transmitted as UDP payload. (see [Appendix E](#) of [[HIP-PROTO](#)]).

[4.1.1.](#) HIP over IPv6

Current implementations use plain IPv6 packets without any payload other than the HIP extension header. This causes problems in combination with NATs that multiplex many private IP addresses into a few public addresses. Usually, this kind of multiplexing is performed on a combination of IP address and TCP or UDP port number. In the absence of a TCP or UDP transport header, typical network address and port translation will fail. Currently there is no specific scheme defined for translating IPv6 HIP extension headers in

the absence of TCP or UDP headers. This all is a minor concern right now, since the likelihood of NATs in IPv6 is small.

Another problem (that also applies to IPv4) is that HIP provides no explicit means for transmitting the IP address used by a host other than using the source IP address of the packet carrying HIP messages. Even if a host B in Figure 2 would know the public IP address used by the NAT for translation, it has no means of providing this information to host A.

[4.1.2.](#) HIP over IPv4

The use of UDP encapsulation for IPv4 transmission enables the HIP base exchange to traverse NATs and to reach their final destination. This encapsulation scheme replaces the IPv6 extension header with a UDP header followed by a HIP header. All HITs and HIP parameters are appended to this new HIP UDP header. UDP transport suffers from many well known problems when traversing NATs (compare [Section 2.2 of \[RFC3715\]](#) also for some issues on UDP traversal). One problem is that often NATs are not able or do not allow to traverse UDP packets, they are blocked and discarded. If NATs do allow UDP packets to traverse, it is not determinable which UDP port number and IP address is used for outgoing UDP packets. HIP over UDP is mandating a fixed port number of 272 for source and destination ports. NATs may change the source port number to any possible port number, for instance, a source port of 272 may be changed by the NAT to 34657.

The above paragraph implies that outgoing packets (packets origin in a private address realm and traversing towards the public address realm) are allowed to traverse, but for the opposite way, incoming packets, this is not necessary given. Only incoming packets that are replies to outgoing packets are able to traverse the NAT. For incoming packets, from flows originating in the public address realm, there is no way of traversing. The destination address in the private address realm (more precise the IP address NAT binding) is not given, thus no address mapping is possible. So transmissions originating in the public Internet are unable to traverse (or even reach the NAT), unless the NAT has some pre-configured configurations.

[4.2.](#) IPsec Data Exchange

IPsec secures data transmission between two HIP nodes after their base exchange is completed. Well-known issues with IPsec and NAT are listed in [[RFC3715](#)] and apply to the IPsec use of HIP as well. Those issues are separated into NAT intrinsic issues, NAT implementation issues and helper incompatibilities. The NAT intrinsic issues related to IPsec (IKE issues are omitted) are:

1. By using ESP all upper layer headers are invisible to the NAT. So changes of the IP header during NAT traversal invalidate any upper layer checksums that are protected within ESP after ESP decryption. HIP takes care about not invalidating upper layer checksums, by using HITs instead of IP addresses for checksum calculations (see Section 3.5 of [[HIP-PROTO](#)]). So this is not applicable here.
2. It is possible to use ESP encrypted packets through a NAT, but as [[RFC3715](#)] remarks, the used SPI values have only one-way significance. Furthermore, SPI values may collide at the NAT, meaning that two different peers behind the NAT are using the same SPI value.
- 3 SPI could be use for multiplexing different IPsec flows at the NAT. But since SPI have only this one-way significance, NATs can only learn the SPI value of outgoing ESP flows, but not the SPI value of the response ESP flow.

A possible way of carrying IPsec traffic through NATs has been proposed in [[IPSEC-UDP](#)]. IPsec traffic is encapsulated in UDP packets.

5. Extensions to HIP

[[HIP-MM](#)] [[HIP-PA](#)] propose new extensions to HIP to make it usable for end host mobility and multi-homing.

Out of those proposals, one extension to HIP can be used for HIP and NAT traversal: HIP peers are able to notify other HIP peers about new addresses they have obtained with the REA packet exchange (see Section 4 of [[HIP-MM](#)]). With REA a HIP node can notify other HIP nodes about new addresses, for example, about its new address at a NAT. This functionality has been already described and can be used for NAT traversal.

The issue is how to determine the publicly reachable IP address, so that it can be announced within the REA packet to the HIP peer.

[6.](#) Extension to NATs

As described in [Section 4.2](#), IPsec SPIs have only one-way significance, meaning that NATs can learn SPI values of outgoing packets, but they cannot learn the corresponding SPI values of incoming packets. Therefore, a matching between SPI values used for outgoing flows to SPI values used for incoming flows is not possible.

NATs must learn the corresponding SPI values for outgoing and incoming IPsec flows. There are two ways in doing so. First, NATs can monitor the HIP base exchange and learn the SPI values agreed by both HIP peers. Afterwards, NATs can remember these values and map outgoing and incoming IPsec flows accordingly.

Second, HIP peers can use a NAT signaling protocol and signal the appropriate SPI values with IP addresses. NATs can so learn the SPI values out of the signaling protocol.

Both approaches require changes to NATs. The first approach would require changes that are only beneficial to HIP, the second one would be beneficial for other protocols as well. Possible solutions for signaling NATs SPI values are NSIS NAT/FW traversal [[NSIS-NATFW](#)] and MIDCOM.

Using MIDCOM in the context of HIP would require some additional knowledge about network topology, for instance, in multihomed environments with different border NATs, host must know which of the multiple NATs to signal for. Therefore, this solution is hard to deploy.

By using the NSIS NAT/FW traversal (NATFW NSLP) mechanism HIP nodes could signal the later on used SPI values for both directions. NATFW NSLP always ensures that signaling messages will reach an appropriate NAT and those messages follow exactly the data path (path-coupled signaling). NSIS requires usually both ends, i.e., both HIP peers, to support this new signaling protocol. However, NATFW NSLP offers support for only one end supporting this protocol. HIP peers behind a NATFW NSLP enabled NAT could so configure the local NATs without impacting other networks. An add-on is given through NATFW traversal, too, since on-path firewalls are configured as well.

[7.](#) HIP unaware NATs

The solutions outlined in [Section 6](#) require that NATs are updated to support new functions, such as HIP itself or NSIS NATFW signaling.

By today's measures, NATs are widely deployed and are currently getting a push through low-cost devices rolled-out to broadband connection users, for instance, DSL lines. NATs are deployed in various places, enterprise borders, mobile phone networks offering

IP-based services, and many more.

Since there many NATs deployed, it will be impossible to upgrade or replace all of them to support HIP or HIP-related extensions. It is the intent of this section to explore how HIP works even in the presence of these HIP unaware NATs. Deployed NATs are currently IPv4 only, so that this section takes them into account only.

For HIP over IPv4, UDP encapsulated HIP messages solve already some problems in traversing NATs. Usually, UDP packets can traverse NATs from inside (private Network) to outside (public network) and vice versa when they have been initiated from inside. The other way around, initiated from outside, will be blocked or impossible since the destination IP address of the NAT is not known a priory. In the case that UDP encapsulation works fine for the HIP message exchange, IPsec is still troublesome (see [[RFC3715](#)]). Some NAT implementations offer some sort of so-called VPN passthrough, where the NAT learns about IPsec flows and tries to correlate outgoing and incoming SPI values (see [[AIREXT](#)] for an example). This works probably only for a small numbers of nodes behind a single NAT, until there will be SPI collisions.

The solution for running IPsec through NATs is documented in [IPSEC-UDP] and applicable here, too. HIP should support IPsec over UDP transport through an extension to the signaling. This extension would indicate when to use IPsec over UDP, instead of plain IPsec.

[8.](#) Security Considerations

To be done.

[9.](#) Acknowledgements

The authors like to thank Lars Eggert for his valuable comments.

10. References

- [RFC2401] Kent, S., and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [RFC2402] Kent, S., and R. Atkinson, "IP Authentication Header", [RFC 2402](#), November 1998.
- [RFC2406] Kent, S., and R. Atkinson, "IP Encapsulating Security Payload (ESP)", [RFC 2406](#), November 1998.

Stiemerling, Quittek

[Page 9]

Internet-Draft

HIP and NAT

July 2004

- [IPSEC-UDP] Huttunen, A., et al, "UDP Encapsulation of IPsec Packets ", [draft-ietf-ipsec-udp-encaps-09.txt](#), Internet Draft (work in progress), May 2004.
- [RFC3715] Aboba, B., and W. Dixon, "IPsec-NAT Compatibility Requirements ", [RFC 3715](#), March 2004.
- [HIP-ARCH] Moskowitz, R., and P. Nikander, "Host Identity Protocol Architecture", [draft-moskowitz-hip-arch-06.txt](#), Internet Draft (work in progress), June 2004.
- [HIP-PROTO] Moskowitz, R., Nikander, P., and T. Henderson, "Host Identity Protocol", [draft-moskowitz-hip-09.txt](#), Internet Draft (work in progress), February 2004.
- [HIP-MM] Nikander, P., and P. Jokela, "End-Host Mobility and Multi-Homing with Host Identity Protocol", [draft-nikander-hip-mm02.txt](#), Internet Draft (work in progress), July 2004.
- [HIP-PA] Nikander, P., Ylitalo, J., and J. Wall, "Integrating Security, Mobility, and Multi-Homing in a HIP way", NDSS 2003, 2003.
- [NATP2P] Ford, B., Srisuresh, P., and D. Kegel, "(Peer-to-Peer (P2P) communication across middleboxes)", [draft-ford-midcom-p2p-03.txt](#), Internet Draft (work in progress), June 2004
- [RFC2663] Srisuresh, P., and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#),

August 1999

- [RFC793] Postel, J., "Transmission Control Protocol", [RFC 793](#), September 1981
- [RFC768] Postel, J., "User Datagram Protocol", [RFC 768](#), August 1980
- [NSIS-NATFW] Stiemerling, M., Tschofenig, H., Martin, M., and C. Aoun, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", [draft-ietf-nsis-nslp-natfw-03.txt](#), Internet Draft (work in progress), July 2004
- [AIREXT] Apple Airport stations with VPN passthrough support, <http://www.apple.com/airportexpress/specs.html>, July 2004,

Stiemerling, Quittek

[Page 10]

Internet-Draft

HIP and NAT

July 2004

11. Authors' Addresses

Martin Stiemerling
NEC Europe Ltd.
Network Laboratories
Kurfuersten-Anlage 36
69115 Heidelberg
Germany

Phone: +49 6221 90511-13
Email: stiemerling@netlab.nec.de

Juergen Quittek
NEC Europe Ltd.
Network Laboratories
Kurfuersten-Anlage 36
69115 Heidelberg
Germany

Phone: +49 6221 90511-15
Email: quittek@netlab.nec.de

Intellectual Property Statement

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.

Disclaimer of Validity

This document and the information contained herein are provided on an

Stiemerling, Quittek

[Page 11]

Internet-Draft

HIP and NAT

July 2004

"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.

Full Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Acknowledgment

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