

Network Working Group
Request for Comments: 5207
Category: Informational

M. Stiernerling
J. Quittek
NEC
L. Eggert
Nokia
April 2008

NAT and Firewall Traversal Issues of Host Identity Protocol (HIP) Communication

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.

IESG Note

This RFC is a product of the Internet Research Task Force and is not a candidate for any level of Internet Standard. The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment.

Abstract

The Host Identity Protocol (HIP) changes the way in which two Internet hosts communicate. One key advantage over other schemes is that HIP does not require modifications to the traditional network-layer functionality of the Internet, i.e., its routers. In the current Internet, however, many devices other than routers modify the traditional network-layer behavior of the Internet. These "middleboxes" are intermediary devices that perform functions other than the standard functions of an IP router on the datagram path between source and destination hosts. Whereas some types of middleboxes may not interfere with HIP at all, others can affect some aspects of HIP communication, and others can render HIP communication impossible. This document discusses the problems associated with HIP communication across network paths that include specific types of middleboxes, namely, network address translators and firewalls. It identifies and discusses issues in the current HIP specifications that affect communication across these types of middleboxes. This document is a product of the IRTF HIP Research Group.

Table of Contents

1. Introduction	3
2. HIP across NATs	4
2.1. Phase 1: HIP Base Exchange	4
2.1.1. IPv4 HIP Base Exchange	4
2.1.2. IPv6 HIP Base Exchange	5
2.2. Phase 2: ESP Data Exchange	5
3. HIP Across Firewalls	6
3.1. Phase 1: HIP Base Exchange	6
3.1.1. IPv4 HIP Base Exchange	6
3.1.2. IPv6 HIP Base Exchange	6
3.2. Phase 2: ESP Data Exchange	7
4. HIP Extensions	7
5. NAT Extensions	8
6. Legacy NAT and Firewall Traversal	8
7. HIP across Other Middleboxes	9
8. Security Considerations	9
9. Acknowledgments	10
10. References	10
10.1. Normative References	10
10.2. Informative References	10

1. Introduction

The current specification of the Host Identity Protocol (HIP) [RFC4423] assumes simple Internet paths, where routers forward globally routable IP packets based on their destination address alone.

In the current Internet, such pure paths are becoming increasingly rare. For a number of reasons, several types of devices modify or extend the pure forwarding functionality the Internet's network layer used to deliver. [RFC3234] coins the term middleboxes for such devices: "A middlebox is (...) any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host".

Middleboxes affect communication in a number of ways. For example, they may inspect the flows of some transport protocols, such as TCP, and selectively drop, insert, or modify packets. If such devices encounter a higher-layer protocol they do not support, or even a variant of a supported protocol that they do not know how to handle, communication across the middlebox may become impossible for these kinds of traffic.

There are many different variants of middleboxes. The most common ones are network address translators and firewalls. [RFC3234] identifies many other types of middleboxes. One broad way of classifying them is by behavior. The first group operates on packets, does not modify application-layer payloads, and does not insert additional packets. This group includes NAT, NAT-PT, SOCKS gateways, IP tunnel endpoints, packet classifiers, markers, schedulers, transport relays, IP firewalls, application firewalls, involuntary packet redirectors, and anonymizers.

Other middleboxes exist (such as TCP performance-enhancing proxies, application-level gateways, gatekeepers, session control boxes, transcoders, proxies, caches, modified DNS servers, content and applications distribution boxes, and load balancers) that divert or modify URLs, application-level interceptors, and application-level multicast systems. However, NATs and firewalls are the most frequent middleboxes that HIP traffic can encounter in the Internet. Consequently, this memo focuses on how NAT and firewall middleboxes can interfere with HIP traffic.

Middleboxes can cause two different kinds of communication problems for HIP. They can interfere with the transmission of HIP control traffic or with the transmission of the HIP data traffic carried within the Encapsulating Security Payload (ESP) [RFC4303].

This document serves mainly as a problem description that solution proposals can reference. But it also discusses known approaches to solving the problem and gives recommendations for certain approaches depending on the specific scenario. It does not promote the use of any of the discussed types of middleboxes.

This memo was discussed and modified in the Host Identity Protocol Research Group, was reviewed by the Internet Research Steering Group (IRSG), and represents a consensus view of the research group at the time of its submission for publication.

2. HIP across NATs

This section focuses on the traversal of HIP across network address translator (NAT) middleboxes. This document uses the term NAT for a basic translation of IP addresses, whereas it uses the term NATP for NATs that additionally perform port translation [RFC2663], if a differentiation between the two is important.

HIP operates in two phases. It first performs a HIP "base exchange" handshake before starting to exchange application data in the second phase. This section describes the problems that occur in each of the two phases when NATs are present along the path from the HIP initiator to the responder.

2.1. Phase 1: HIP Base Exchange

The HIP base exchange uses different transport mechanisms for IPv6 and IPv4. With IPv6, it uses a HIP-specific IPv6 extension header, whereas it uses the IP payload with IPv4 [RFC5201].

2.1.1. IPv4 HIP Base Exchange

The HIP protocol specification [RFC5201] suggests encapsulating the IPv4 HIP base exchange in a new IP payload type. The chances of NAT traversal for this traffic are different, depending on the type of NAT in the path. The IPv4 HIP base exchange traverses basic NATs (that translate IP addresses only) without problems, if the NAT only interprets and modifies the IP header, i.e., it does not inspect the IP payload.

However, basic NATs are rare. NATP devices that inspect and translate transport-layer port numbers are much more common. Because the IP payload used for the IPv4 base exchange does not contain port numbers or other demultiplexing fields, NATPs cannot relay it.

A second issue is the well-known "data receiver behind a NAT" problem. HIP nodes behind a NAT are not reachable unless they initiate the communication themselves, because the necessary translation state is otherwise not present at the NAT.

2.1.2. IPv6 HIP Base Exchange

The IPv6 HIP base exchange uses empty IPv6 packets (without a payload). New HIP extension headers carry the base exchange information. This approach can cause problems if NAT middleboxes translate or multiplex IP addresses.

At this time, IPv6 NATs are rare. However, when they exist, IPv6 NATs operate similarly to IPv4 NATs. Consequently, they will likely block IP payloads other than the "well-known" transport protocols. This includes the IPv6 HIP base exchange, which does not contain any IP payload.

2.2. Phase 2: ESP Data Exchange

HIP uses ESP to secure the data transmission between two HIP nodes after the base exchange completes. Thus, HIP faces the same challenges as IPsec with regard to NAT traversal. [RFC3715] discusses these issues for IPsec and describes three distinct problem categories: NAT-intrinsic issues, NAT implementation issues, and helper incompatibilities.

This section focuses on the first category, i.e., NAT-intrinsic issues. The two other problem categories are out of this document's scope. They are addressed in the BEHAVE working group or in [RFC3489].

With ESP-encrypted data traffic, all upper-layer headers are invisible to a NAT. Thus, changes of the IP header during NAT traversal can invalidate upper-layer checksums contained within the ESP-protected payload. HIP hosts already avoid this problem by substituting Host Identity Tags (HITs) for IP addresses during checksum calculations [RFC5201].

Although the traversal of ESP-encrypted packets across NATs is possible, [RFC3715] notes that the Security Parameter Index (SPI) values of such traffic have only one-way significance. NATs can use SPI values to demultiplex different IPsec flows, similar to how they use port number pairs to demultiplex unencrypted transport flows. Furthermore, NATs may modify the SPIs, similar to how they modify port numbers, when multiple IPsec nodes behind them happen to choose

identical SPIs. However, NATs can only observe the SPIs of outgoing IPsec flows and cannot determine the SPIs of the corresponding return traffic.

3. HIP Across Firewalls

This section focuses on the traversal of HIP across IP firewalls and packet filters. These types of middleboxes inspect individual packets and decide whether to forward, discard, or process them in some special way, based on a set of filter rules and associated actions.

Firewalls are not inherently problematic for HIP, as long as their policy rules permit HIP base exchange and IPsec traffic to traverse. The next sections discuss specific issues for HIP in typical firewall configurations.

3.1. Phase 1: HIP Base Exchange

3.1.1. IPv4 HIP Base Exchange

A common and recommended configuration for IPv4 firewalls is to block all unknown traffic by default and to allow well-known transport protocols only and often just on specific ports and with specific characteristics ("scrubbed" traffic). This common configuration blocks the HIP base exchange.

3.1.2. IPv6 HIP Base Exchange

The configuration of IPv6 firewalls is similar to IPv4 firewalls. Many IPv4 firewalls discard any IP packet that includes an IP option. With IPv6, the expectation is that firewalls will block IPv6 extension headers in general or will at least block unknown extension headers. Furthermore, payloads other than specific, well-known transport protocols are likely to be blocked as well. Like IPv4, this behavior blocks the HIP base exchange.

A problem similar to the "data receiver behind a NAT" issue (see Section 2.1.1) applies to both IPv4 and IPv6 firewalls. Typically, firewalls block all traffic into the protected network that is not identifiable return traffic of a prior outbound communication. This means that HIP peers are not reachable outside the protected network, because firewalls block base exchange attempts from outside peers.

3.2. Phase 2: ESP Data Exchange

Firewalls are less problematic than NATs with regard to passing ESP traffic. The largest concern is commonly used firewall configurations that block ESP traffic, because it is not a well-known transport protocol and ports cannot be used to identify return flows. However, firewalls could use mechanisms similar to Security Parameter Index (SPI) multiplexed NAT (SPINAT) to use SPIs as flow identifiers [YLITALO].

4. HIP Extensions

This section identifies possible changes to HIP that attempt to improve NAT and firewall traversal, specifically, the reachability of HIP peers behind those middleboxes and traversal of the HIP base exchange. Sections 2 and 3 describe several problems related to encapsulation schemes for the HIP base exchange in IPv4 and IPv6.

UDP may improve HIP operation in the presence of NATs and firewalls. It may also aid traversal of other middleboxes. For example, load balancers that use IP- and transport-layer information can correctly operate with UDP-encapsulated HIP traffic.

HIP nodes located behind a NAT must notify their communication peers about the contact information. The contact information is the NAT's public IP address and a specific UDP port number. This measure enables the peers to send return traffic to HIP nodes behind the NAT. This would require a new HIP mechanism.

To be reachable behind a NAT, a rendezvous point is required that lets HIP nodes behind a NAT register an IP address and port number that can be used to contact them. Depending on the type of NAT, use of this rendezvous point may be required only during the base exchange or throughout the duration of a communication instance. A rendezvous point is also useful for HIP nodes behind firewalls, because they suffer from an analogous problem, as described in Section 3.

The proposed mobility management packet exchange [RFC5206] [NIKANDER] can support this method of NAT traversal. The original intention of this extension is to support host mobility and multihoming. This mechanism is similar to the Alternate Network Address Types (ANAT) described in [RFC4091]. However, HIP peers use mobility management messages to notify peers about rendezvous points, similar to [RFC4091]. HIP peers must determine their contact address before they can announce it to their peers.

5. NAT Extensions

IPsec SPIs have only one-way significance, as described in Section 2.2. Consequently, NATs and firewalls can observe the SPI values of outgoing packets, but they cannot learn the SPI values of the corresponding inbound return traffic in the same way. Two methods exist:

First, NATs can observe the HIP base exchange and learn the SPI values that HIP peers agree to use. Afterwards, NATs can map outgoing and incoming IPsec flows accordingly. This approach is called architected NAT, or SPINAT [YLITALO], and can be used by firewalls as well. It requires HIP-specific NAT modifications.

Second, HIP peers can use a generic NAT or firewall signaling protocol to explicitly signal appropriate SPI values to their NATs and firewalls. This approach does not require HIP-specific changes at the middlebox, but does require integration of HIP with the signaling protocol at the end systems.

Possible solutions for signaling SPI values are the mechanisms proposed in the IETF NSIS WG (NATFW NSLP) and MIDCOM MIB module [RFC5190]. Using MIDCOM in the context of HIP requires additional knowledge about network topology. For example, in multihomed environments with different border NATs or firewalls, a host must know which of the multiple NATs/firewalls to signal. Therefore, this solution can be problematic.

By using the NSIS NAT/FW traversal (NATFW NSLP) mechanism HIP nodes can signal the used SPI values for both directions. NATFW NSLP ensures that signaling messages will reach all NATs and firewalls along the data path (path-coupled signaling). Although NSIS is generally supported at both peers, the NATFW NSLP offers a "proxy mode" for scenarios where only one end supports NSIS. This has deployment advantages.

6. Legacy NAT and Firewall Traversal

The solutions outlined in Section 5 require that NATs and firewalls are updated to support new functions, such as HIP itself or NSIS NATFW signaling. NATs and firewalls are already widely deployed. It will be impossible to upgrade or replace all such middleboxes with HIP support. This section explores how HIP operates in the presence of legacy NATs and firewalls that are not HIP-aware. Because the vast majority of deployed NATs currently support IPv4 only, this section focuses on them.

For HIP over IPv4, UDP encapsulation of HIP traffic already solves some NAT traversal issues. Usually, UDP packets can traverse NATs and firewalls when communication was initiated from the inside. However, traffic initiated outside a NAT is typically dropped, because it cannot be demultiplexed to the final destination (for NATs) or is prohibited by policy (for firewalls).

Even when UDP encapsulation enables the HIP base exchange to succeed, ESP still causes problems [RFC3715]. Some NAT implementations offer "VPN pass-through", where the NAT learns about IPsec flows and tries to correlate outgoing and incoming SPI values. This often works reliably only for a small number of nodes behind a single NAT, due to the possibility of SPI collisions.

A better solution may be to use UDP encapsulation for ESP [RFC3948], enabled through a new parameter in the base exchange. It is for further study whether to mandate UDP encapsulation for all HIP traffic to reduce the complexity of the protocol.

HIP may also consider other NAT/firewall traversal mechanisms, such as the widely deployed Universal Plug and Play (UPNP) [UPNP]. UPNP can be used to configure middleboxes on the same link as a HIP node.

7. HIP across Other Middleboxes

This document focuses on NAT and firewall middleboxes and does not discuss other types identified in [RFC3234]. NATs and firewalls are the most frequently deployed middleboxes at the time of writing. However, future versions of this document may describe how HIP interacts with other types of middleboxes.

8. Security Considerations

Opening pinholes in firewalls (i.e., loading firewall rules allowing packets to traverse) and creating NAT bindings are highly security-sensitive actions. Any mechanism that does so in order to support HIP traversal across middleboxes should be well protected. Detailed discussion of the related security issues can be found in the security considerations sections of the corresponding standards documents, such as [RFC3715] and [RFC5190].

This document has not considered whether some of the options listed above pose additional threats to security of the HIP protocol itself.

9. Acknowledgments

The following people have helped to improve this document through thoughtful suggestions and feedback: Pekka Nikander, Tom Henderson, and the HIP research group. The authors would like to thank the final reviewers, Kevin Fall, Mark Allman, and Karen Sollins.

Lars Eggert and Martin Stiernerling are partly funded by Ambient Networks, a research project supported by the European Commission under its Sixth Framework Program. 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.

10. References

10.1. Normative References

- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets", RFC 3948, January 2005.
- [RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.
- [RFC5201] Moskowitz, R., Nikander, P., Jokela, P., Ed., and T. Henderson, "Host Identity Protocol", RFC 5201, April 2008.

10.2. Informative References

- [NIKANDER] Nikander, P., Ylitalo, J., and J. Wall, "Integrating Security, Mobility, and Multi-Homing in a HIP Way", Proc. Network and Distributed Systems Security Symposium (NDSS) 2003, February 2003.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

- [RFC3715] Aboba, B. and W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements", RFC 3715, March 2004.
- [RFC4091] Camarillo, G. and J. Rosenberg, "The Alternative Network Address Types (ANAT) Semantics for the Session Description Protocol (SDP) Grouping Framework", RFC 4091, June 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [RFC5190] Quittek, J., Stiemerling, M., and P. Srisuresh, "Definitions of Managed Objects for Middlebox Communication", RFC 5190, March 2008.
- [RFC5206] Henderson, T., Ed., "End-Host Mobility and Multihoming with the Host Identity Protocol", RFC 5206, April 2008.
- [UPNP] UPNP Web Site, "Universal Plug and Play Web Site", Web Site <http://www.upnp.org/>, July 2005.
- [YLITALO] Ylitalo, J., Melen, J., Nikander, P., and V. Torvinen, "Re-Thinking Security in IP-Based Micro-Mobility", Proc. 7th Information Security Conference (ISC) 2004, September 2004.

Authors' Addresses

Martin Stiernerling
NEC Europe Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 113
Fax: +49 6221 4342 155
EMail: stiernerling@nw.neclab.eu
URI: <http://www.nw.neclab.eu/>

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

Phone: +49 6221 4342 115
Fax: +49 6221 4342 155
EMail: quittek@nw.neclab.eu
URI: <http://www.nw.neclab.eu/>

Lars Eggert
Nokia Research Center
P.O. Box 407
Nokia Group 00045
Finland

Phone: +358 50 48 24461
EMail: lars.eggert@nokia.com
URI: http://research.nokia.com/people/lars_eggert/

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at <http://www.rfc-editor.org/copyright.html>, 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, THE IETF TRUST 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.

