HIPRG
Internet-Draft
Expires: April 27, 2006

H. Tschofenig
Siemens
M. Shanmugam
TUHH

October 24, 2005

Traversing HIP-aware NATs and Firewalls: Problem Statement and Requirements

draft-tschofenig-hiprg-hip-natfw-traversal-03.txt

Status of this Memo

Copyright Notice

Abstract

The Host Identity Protocol is a signaling protocol which adds another layer to the Internet model and (optionally) establishes IPsec ESP SAs to protect subsequent data traffic. HIP is designed to be a middlebox friendly protocol, it allows the middleboxes (such as NATs and Firewalls) to participate in the base exchange messages in order to learn the flow identifier and thereby, relying the data traffic.

Adding authentication and authorization mechanisms can help the middlebox decide which end hosts are allowed to traverse a firewall. This can potentially prevent waste of network resources and limit unwanted traffic. This document gives a problem statement and requirements for HIP-aware middlebox traversal techniques.

Table of Contents

1. Introduction 3
2. Terminology 4
3. Problem Statement 5
4. Overview of HIP Base Exchange with Middleboxes 7
4.1. HIP Base Exchange with NAT 7
4.2. HIP Base Exchange with Firewall 8
5. Scenarios 10
5.1. Same Firewall at Initiator for both outgoing and incoming packets 10
5.2. Different Firewalls at Initiator for outgoing and incoming packets 11
5.3. Different Firewalls at Initiator and Receiver 12
6. Requirements 15
7. Security Considerations 16
8. Contributors 17
9. Acknowledgements 18
10. References 19
10.1. Normative References 19
10.2. Informative References 19
Authors' Addresses 21
Intellectual Property and Copyright Statements 22

1. Introduction

An IP address serves the dual role of a locator and an identifier for every host on the Internet. Since, the transport layer connections are bound to the IP address, end systems that use IP addresses as identifiers cannot support dynamic changes in the mapping between the identifier and the locator in case of multi-homing and mobility.

The Host Identity Protocol (HIP) [I-D.ietf-hip-base] proposes to separate the identifier from the locator by adding an additional layer between the transport layer and the network layer. The transport layer uses a new, mobility-unrelated identifier called as Host Identity Tags (HITs), in place of IP addresses, while the network layer uses conventional IP addresses for routing. IPsec security associations are bound to the HITs and are not modified with IP address changes. In other words, a host despite being mobile or multi-homed can use a single transport layer connection associated to one HIT and multiple IP addresses.

The Host Identity Protocol offers also the functionality to establish IPsec ESP SAs which are subsequently used to encrypt data traffic between the two end hosts. HIP is liable to all known incompatibilities of IPsec with middleboxes such as NATs [RFC3022] and firewalls. The problem statement for dealing with legacy NATs is described in [I-D.irtf-hiprg-nat]. The main goal of the draft is to present a problem statement and requirements in order to aim for a NAT/FW traversal solution using the Host Identity Protocol.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This draft used the terminology defined in [NATTerminology], [I-D.ietf-hip-base], [I-D.ietf-HIP-esp] and [draft-moskowitz-hip-arch] and [RFC2401].

The term SPI refers to the Security Parameter Index value used in IPsec packets. The initiator selects one SPI(I) that can be found in the ESP_info parameter, which is then used by the responder to create an IPsec packet (ESP packet in this case) for traffic sent to the initiator. The responder selects one SPI(R)(using ESP_info(R) parameter) which is used by the initiator to encrypt all data sent to the responder.

Other relevant abbreviations can be found in [I-D.ietf-hip-base] and [I-D.ietf-HIP-esp].

The concept of a flow identifier is described in [RFC4080].

We use the following notation throughout this document:

- [x] indicates that x is optional.
- {x} indicates that x is encrypted.

- o <x>y indicates that "x" is encrypted with the key "y".
- o --> signifies "Initiator to Responder" communication (requests).
- o <-- signifies "Responder to Initiator" communication (replies).

3. Problem Statement

This document assumes that the data traffic following the HIP base exchange is IPsec protected using the mechanism described in [I-D.ietf-HIP-esp] for exchanging the IPsec parameters. A future version of this draft might also be extended to support other mechanisms for data traffic protection including no protection at all.

Besides the communicating hosts in the Internet, the entities such as NATs and Firewalls play a major role in the event of delivering packets to an appropriate host, and each meant for specific functionality. For instance, NATs are used to combat the IPv4 address depletion problem, and Firewalls are erected to protect unsolicited information flowing in and out of a corporate network. NATs use <src IP ,dst IP, src port, dst port, protocol> as an flow identifier to identify a particular traffic or connection. Because of this, protocols like IPsec suffers from well-known NAT related problems. The NAT traversal approach described in [RFC3947] and [I-D.ietf-ipsec-ikev2] allows the end hosts to detect one or more NATs in between them and [RFC3948] proposes to use the UDP encapsulation of IPsec ESP packets to traverse NATs.

Since HIP uses IPsec protection for the data traffic, the flow identifier takes the shape of a <destination IP address, SPI and ESP> (in order to support smooth traversal of the middleboxes) and the middleboxes should learn this flow id in order to relay the data packets. To achieve this, HIP aims to interact with middleboxes actively whereby these devices need to understand the HIP protocol and they need to be involved in the protocol exchange. HIP also provides a way to deal with legacy NATs, as described in [draft-nikander-hip-path-00.txt]. To support this functionality, it is necessary to provide UDP encapsulation for both HIP signaling and IPsec packets. Legacy NAT traversal does not require NATs to be HIP aware or to understand the HIP message exchange.

Even though HIP allows the middleboxes to participate in the base exchange, but scenarios like routing asymmetry poses a serious challenge for the HIP to traverse a middlebox. Section 5 explains some possible scenarios which have routing asymmetry. The inability of HIP to handle routing asymmetry motivates to use an explicit signaling mechanism for the HIP hosts in order to support secure and smooth traversal of the middleboxes.

Although HIP is described as a two-party protocol, middle boxes are supposed to intercept these messages in order to learn the flow identifier and to process them correctly. In other words, a multi party protocol is created such that the flow identifier is available

to middle boxes between the HIP hosts. To provide proper security, middleboxes should not be subject to denial of service attacks and might want to authenticate or authorize entities which create state. Note that the IPsec SA is unidirectional and therefore two IPsec SAs (with two different SPIs, ESP_info contains the SPI value) have to be established.

4. Overview of HIP Base Exchange with Middleboxes

This section explains the HIP base exchange together with the middleboxes and describes how the middleboxes behave during the base exchange.

4.1. HIP Base Exchange with NAT

Figure 1 shows the HIP base exchange traversing a NAT.

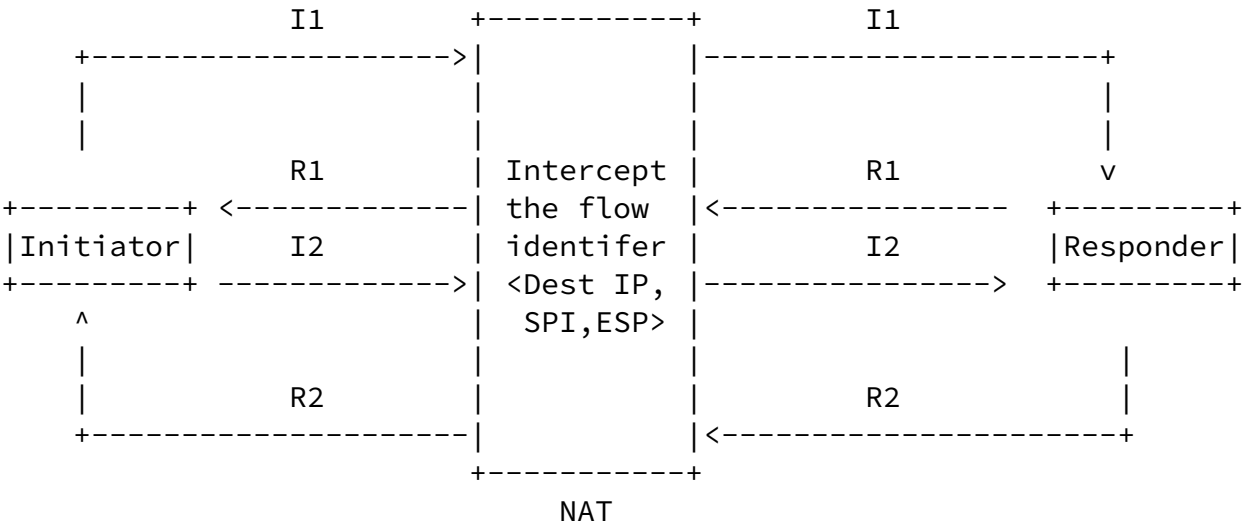

Figure 1: NAT and HIP Base Exchange

Subsequently, the HIP base exchange is described in more detail.

```
 I -> R: I1: Trigger exchange

 I <- R: R1: {Puzzle, D-H(R), HI(R), ESP Transform,
              HIP Transform }SIG

 I -> R: I2: {Solution, LSI(I), ESP_info(I), D-H(I),
              ESP_Transform, HIP Transform, {H(I)}SK }SIG

 I <- R: R2: {LSI(R), ESP_info(R), HMAC}SIG
```

A potential responsibility of the NAT, as shown in Figure 1, can be the following

- Intercept the signaling messages
- Authenticate and authorize the HIP nodes by verifying the signatures.

- Process the flow identifier information

- Perform actions according to the state machine

- Create state based on the content of message I2 with ESP_info(I) and R2 with ESP_info(R). Additionally, it might be necessary to include support for storing the respective HITs and host identities.

4.2. HIP Base Exchange with Firewall

In case of a firewall traversal, the routing asymmetry needs to be considered i.e., the fact that the messages I1 and I2 do not necessarily traverse the same devices as R1 and R2. The same is true with more complex network topologies with a mixture of NATs and Firewalls. This is an assumption made in the NSIS working group (and therefore also with NAT/Firewall traversal). Pure NAT traversal is therefore simpler to handle in comparison to middlebox traversal which also includes devices such as Firewalls. Figure 3 shows this circumstance graphically:

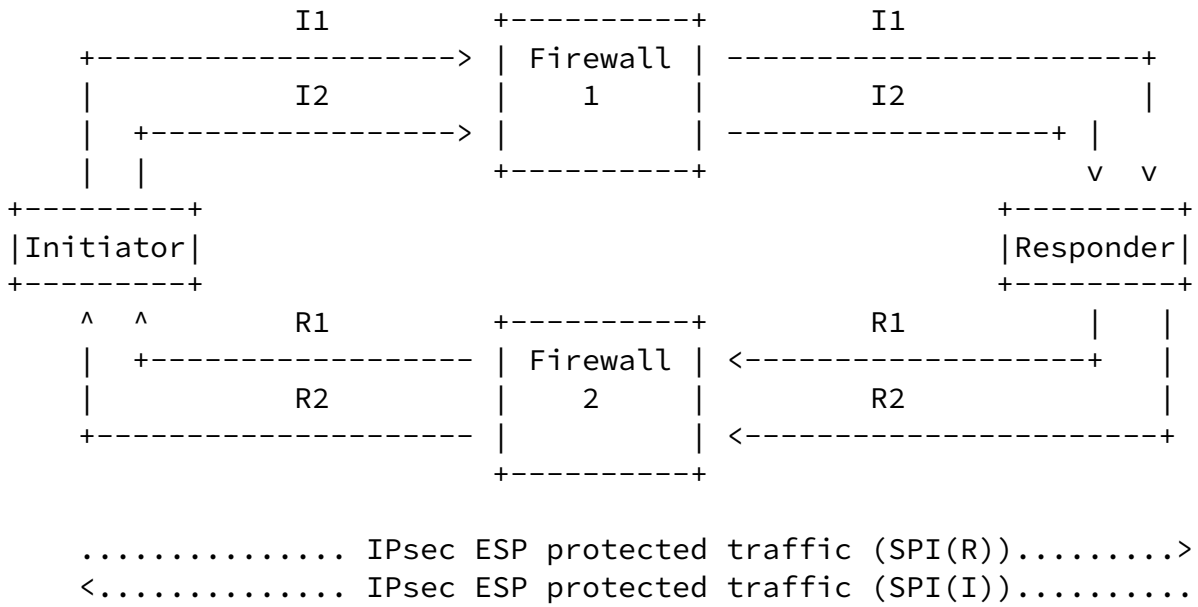

Figure 3: Firewall and HIP Base Exchange

With one single NAT between the HIP nodes, all messages of the base exchange are forced through it. With firewalls, it becomes obvious that the nice property of a NAT with respect to the symmetric

forwarding path is lost and the individual firewalls (Firewall 1 and Firewall 2) are unable to create the necessary firewall pinholes. SPI(I) is exchanged in I2 message (ESP_info(I)) through firewall 1, however firewall 2 only needs it. Similarly firewall 2 needs SPI (R) which is sent in message R2 (ESP_info(I)) through firewall 1.

5. Scenarios

The following section describes some sample scenarios, in the context of involving middleboxes, to learn the flow identifier:

5.1. Same Firewall at Initiator for both outgoing and incoming packets

This scenario assumes that the initiator I alone is behind a firewall named FW(I). This firewall is both for the outgoing and incoming packets and hence can look into all the base exchange messages. This scenario is also applicable for NATs as well. This is illustrated in Figure 4

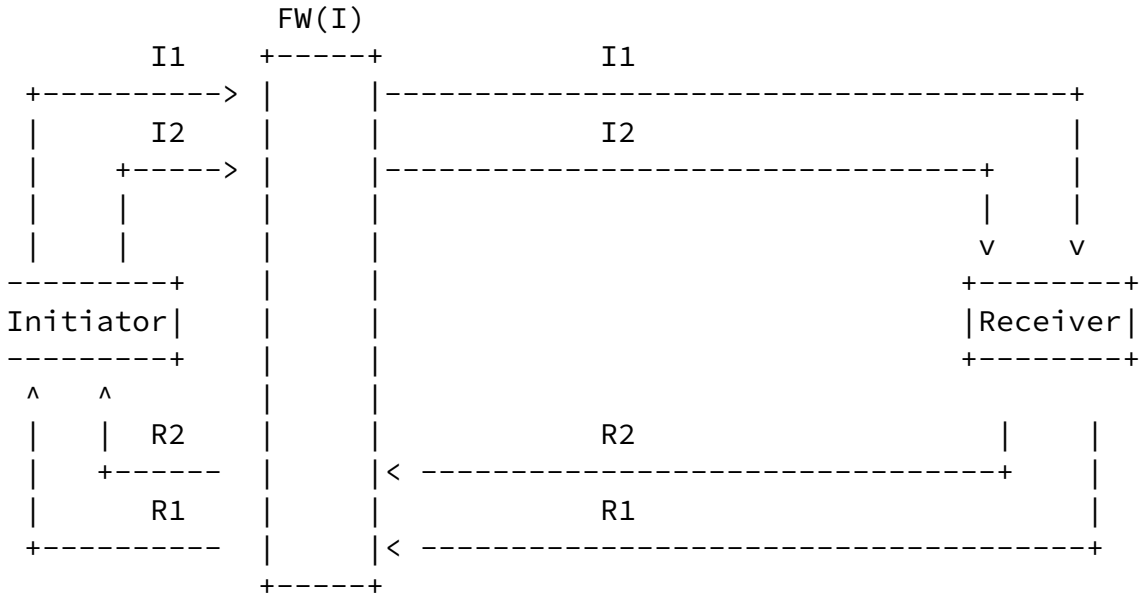

Figure 4: One FW only at initiator end

1. I1 packet is sent from the initiator I to receiver R.

2. FW(I) forward the packet to the Receiver.

3. R sends R1 message with puzzle,D-H key protected with the signature of R.

4. FW(I) forward the packet to the Initiator.

5. On receiving I2, FW(I) verifies the signature of I and learns the SPI value form the ESP_info parameter and forwards it to the Receiver

6. R sends the message R2 to I.

7. On receiving R2, FW(I) verifies the signature of R. Accordingly, it earns the SPI value form the ESP_info parameter and forwards it to the Initiator.

8. The base exchange continues until complete. Since all messages I1,R1,I2 and R2 traverse through FW(I), it can look into these messages to learn the flow identifier information.

5.2. Different Firewalls at Initiator for outgoing and incoming packets

This scenario assumes that both the initiator I and the receiver R is situated behind firewalls named FW(I) and FW(R) respectively. FW(I) is for the incoming packets to I and FW(R) is for the incoming packets to R. It is necessary that both the firewalls must learn the flow identifier information and should store the state <SPI,IP,HIT> to forward IPsec protected payload packets. This scenario is illustrated in Figure 5

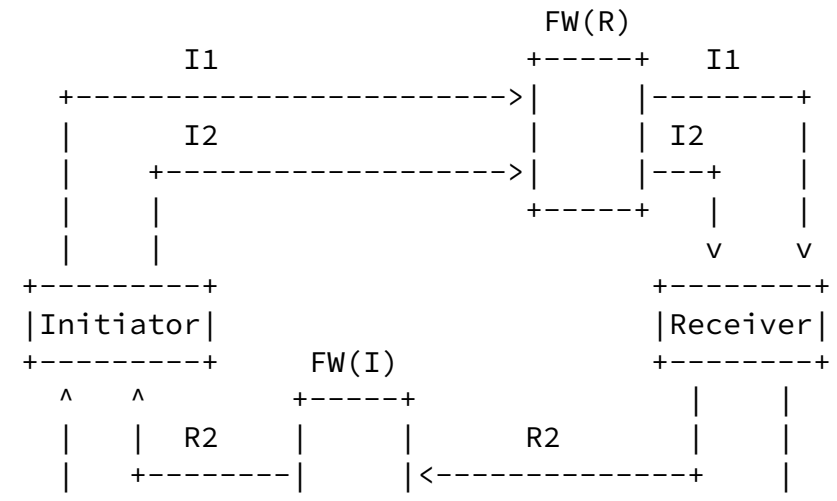

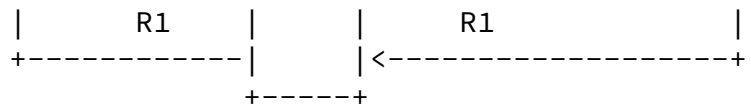

Figure 5: End hosts behind FWs

1. I1 packet is sent from the initiator I to receiver R.

2. FW(R) forwards the packet to the Receiver.

3. Then, R sends R1 message with puzzle,D-H key protected with the signature of R.

4. FW(I) forward the packet to the Initiator.

5. Now, I sends the I2 packet, on receiving I2, FW(R) verifies the signature of I and learns the SPI value form the ESP_info parameter and forwards it to the Receiver

6. To complete the base exchange, R sends the message R2 to I.

7. On receiving R2, FW(I) verifies the signature of R. Accordingly, it earns the SPI value from the ESP_info parameter and forwards it to the Initiator.

Here, the problem with this asymmetric base exchange is that the SPI needed for the FW(I) is sent through the I2 message, which flows through the FW(R) and the SPI needed for for the FW(I) is sent to FW(R).

5.3. Different Firewalls at Initiator and Receiver

This scenario looks into a more complicated situation. Initiator I is behind multiple firewalls FW1(I) for outgoing packets and FW2(I) and FW3(I) are for incoming packets. Similarly, receiver R is behind FW1(R) and FW2(R) for incoming packets and FW3(R) for outgoing packets. The incoming firewalls are chosen based on the type of the application and the hosts are unaware from which firewall they receive packets. Here, however for our scenario we assume that FW2(R) and FW2(I) are chosen about which also the hosts are unaware

of. This scenario is illustrated in Figure 6

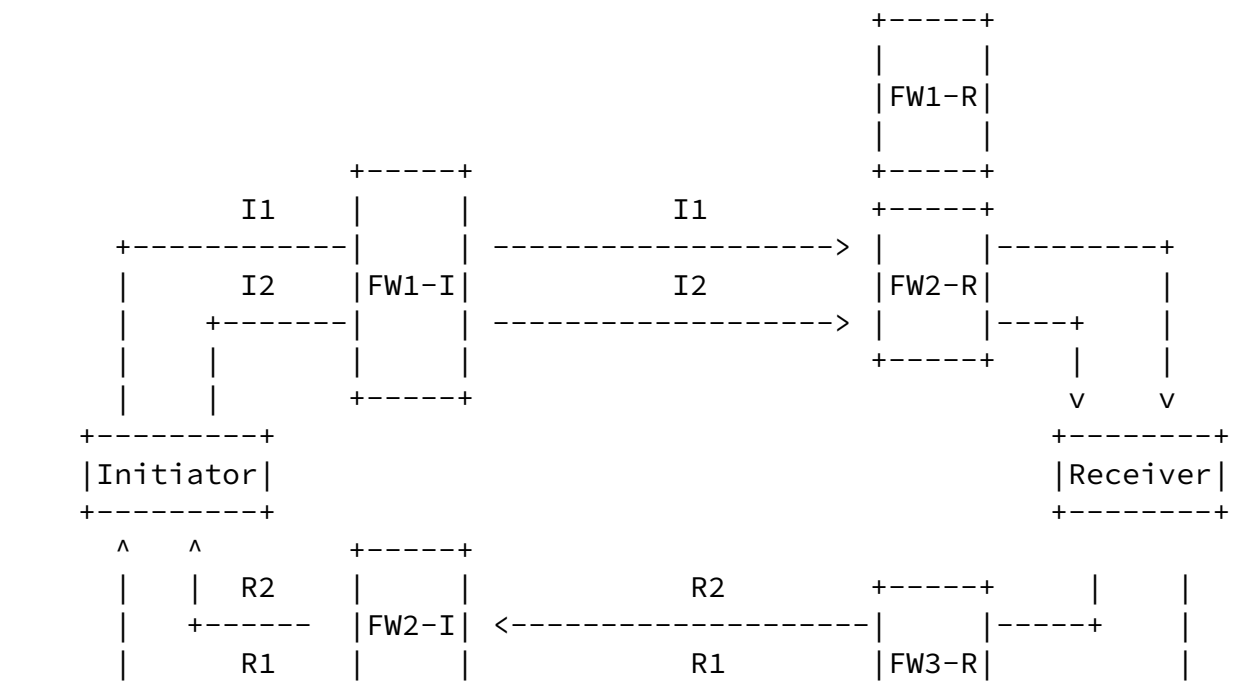

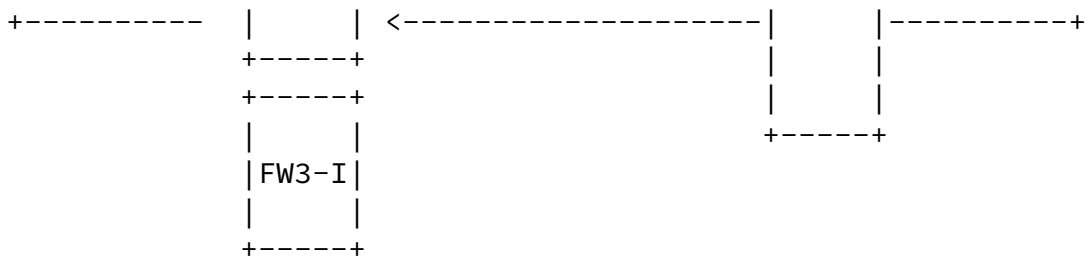

Figure 6: Multiple FWs at initiator's and receiver's end

1. I1 packet is sent from the initiator I to receiver R.

2. FW1(I) and FW2(R) forwards the packet to the Receiver.

3. Then, R sends R1 message with puzzle,D-H key protected with the signature of R.

4. Now, FW3(R) and FW2(I) forward the packet to the Initiator.

5. Now, the I sends the I2 packet, on receiving I2, FW1(I) and FW2(R) can verify the signature of I and can learn the SPI value from the ESP_info parameter and forward it to the Receiver.

6. To complete the base exchange, R sends the message R2 to I.

7. On receiving R2, FW3(R) and FW2(I) can verify the signature of R. Accordingly, they learn the SPI value form the ESP_info parameter and forwards it to the Initiator.

Here, the problems are :

1. With this asymmetric base exchange is that the SPI needed for the Firewall(s) on the receiver side is sent through the I2 message, Which is actually sent through FW1(I) and FW2(R) and the SPI needed for for the Firewall(s) on the Initiator side is sent to FW3(R) and FW2(I).

2. When hosts are behind multiple incoming firewalls, there are unable to decide to which firewall they have to inform their SPI values to.

3. The second problem is to secure the SPI signalling message from

the end host to the FW. Since the end hosts authenticate and authorize to the FW that lets outgoing packets, they share keys only with them. However, as mentioned earlier, they, somehow, need to signal the SPI value to the FW on the other end which forwards incoming packets.

6. Requirements

In the context of middlebox signaling, a few high-level requirements have to be accomplished:

- Add some authentication and authorization capabilities to NAT traversal. Many NAT/Firewall traversal solutions do not allow the end host to interact with the middlebox. As a consequence, some security vulnerabilities are introduced.

- Add secure firewall traversal functionality as another type of middlebox signaling by using <destination IP address, SPI and protocol> triplet. as a substitute for the typical < source IP, destination IP, source port, destination port, transport protocol> information.

Such a solution for HIP-based middlebox signaling needs to have the following properties:

- A HIP-aware NAT/FW MUST be able to authenticate the entity requesting a NAT binding or a firewall pinhole.

- A HIP-aware NAT/FW MUST authorize the entity requesting a NAT binding or a firewall pinhole before storing state information. This requirement might be accomplished by identity based authorization or an identity independent authorization mechanism.

- A HIP-aware NAT/FW MUST be able to intercept HIP messages in order to extract the flow identifier information and other related information. HIP messages are base exchange messages during context establishment and readdressing messages during IP address changes. A NAT/FW MUST be able to distinguish these messages.

- A NAT/FW node MUST NOT introduce new denial of service attacks based on authentication or key management mechanisms.

- A potential solution MUST respect the property of some middleboxes which do not allow traffic (data and signaling traffic) to traverse this middlebox without proper authorization.

Some requirements are taken from [I-D.ietf-nsis-nslp-natfw].

7. Security Considerations

This document analyzes the traversal of HIP-aware middleboxes. A problem statement is given and scenarios are described that lead to a number of requirements.

This document therefore lists a number of security aspects throughout the document. Care should be taken when solutions are developed and the security solution must not introduce new vulnerabilities to the middlebox.

8. Contributors

We would like to thank Aarthi Nagarajan, Vesa Torvinen, Jochen Grimminger and Jukka Ylitalo for their help with initial versions of this document.

9. Acknowledgements

The authors would like to thank Pekka Nikander, Dieter Gollmann and Thomas Aura for their feedback to this document.

This document is a byproduct of the Ambient Networks Project, partially funded 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.

10. References

10.1. Normative References

[I-D.ietf-HIP-esp]
Moskowitz, R., Nikander, P., and P. Jokela, "Using ESP transport format with HIP", draft-ietf-hip-esp-00 (work in progress), June 2005.

[I-D.ietf-hip-base]
Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol", draft-ietf-hip-base-03 (work in progress), June 2005.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

10.2. Informative References

[I-D.ietf-ipsec-ikev2]
Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", draft-ietf-ipsec-ikev2-17 (work in progress), September 2004.

[I-D.ietf-nsis-nslp-natfw]
Stiemerling, M., Tschofenig, H., and C. Aoun, "A NAT/

Firewall NSIS Signaling Layer Protocol (NSLP)",
draft-ietf-nsis-nslp-natfw-07 (work in progress),
July 2005.

[I-D.irtf-hiprg-nat]
Stiemerling, M., Quittek, J., and L. Eggert, "Middlebox
Traversal Issues of Host Identity Protocol (HIP)
Communication", draft-irtf-hiprg-nat-00.txt (work in
progress) (work in progress), October 2005.

[NATTerminology]
Srisuresh, P. and M. Holdrege, "IP Network Address
Translator (NAT) Terminology and Considerations", Request
For Comments RFC 2663, August 1999.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the
Internet Protocol", RFC 2401, November 1998.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network
Address Translator (Traditional NAT)", RFC 3022,
January 2001.

[RFC3947] Kivinen, T., A. Huttunen, A., Swander, B., and V. Volpe,
"Negotiation of NAT-Traversal in the IKE", RFC 3947,
January 2005.

[RFC3948] A. Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and
M. Stenberg, "UDP Encapsulation of IPsec Packets",
RFC 3948, January 2005.

[RFC4080] Hancock, R., "Next Steps in Signaling: Framework",
RFC 4080, November 2004.

[draft-ietf-hip-mm]
Henderson (Editor), T., "End-Host Mobility and Multi-
Homing with Host Identity Protocol",
draft-nikander-hip-mm-02.txt (work in progress) (work in
progress), July 2005.

[draft-ietf-ipsec-esp-v3-08]
Kent, S., "IP Encapsulating Security Payload (ESP)",

draft-ietf-ipsec-esp-v3-10 (work in progress) (work in progress), March 2005.

[draft-moskowitz-hip-arch]
Moskowitz, R. and P. Nikander, "Host Identity Protocol Architecture", draft-ietf-hip-arch-03 (work in progress) (work in progress), August 2005.

[draft-nikander-hip-path-00.txt]
Nikander, P., Tschofenig, H., Henderson, T., Eggert, L., and J. Laganier, "Preferred Alternatives for Tunnelling HIP (PATH)", draft-nikander-hip-path-00.txt (work in progress) (work in progress), February 2005.

[rfc3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", Request For Comments RFC 3022, January 2001.

Authors' Addresses

Hannes Tschofenig
Siemens
Otto-Hahn-Ring 6
Munich, Bayern 81739
Germany

Email: Hannes.Tschofenig@siemens.com

Murugaraj Shanmugam

Technical Univeristy Hamburg-Harburg
Schwarzenbergstrasse 95
Harburg, Hamburg 21073
Germany

Email: murugaraj.shanmugam@tuhh.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

Copyright Statement

Acknowledgment

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