

HIP Working Group
Internet-Draft
Intended status: Experimental
Expires: August 28, 2008

M. Komu
HIIT
T. Henderson
The Boeing Company
P. Matthews
Avaya
H. Tschofenig
Nokia Siemens Networks
A. Keraenen
J. Melen
Ericsson Research Nomadiclab
M. Bagnulo
Huawei Lab at UC3M
February 25, 2008

Basic HIP Extensions for Traversal of Network Address Translators and
Firewalls

draft-ietf-hip-nat-traversal-03.txt

Status of this Memo

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 becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on August 28, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

Abstract

The Host Identity Protocol (HIP) provides a new namespace that can be used for uniquely identifying hosts. Existing HIP experimental specifications do not specify protocol operations across Network Address Translators (NATs).

This document specifies NAT traversal extensions for HIP. The HIP shim layer is located between the network and transport layer, the extensions can also provide a more general-purpose NAT traversal support for higher-layer networking applications. The extensions are based on the use of the The Interactive Connectivity Establishment (ICE) methodology to discover a working path between two end-hosts. Using the specified extensions, two HIP-capable hosts are able to communicate with each other even when both nodes are behind NATs or firewalls.

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

Table of Contents

1.	Introduction	4
2.	Terminology	5
3.	Protocol Description	7
3.1.	Relay Registration and NAT Detection	7
3.2.	Base Exchange via HIP Relay	9
4.	Connectivity Tests	11
4.1.	NAT Transformation Negotiation	11
4.2.	ICE Procedure	12
4.3.	NAT Keep-alives	12
5.	Packet Formats	13
5.1.	HIP Control Packets	13
5.2.	Keep-Alives	13
5.3.	Relay and Registration Parameters	14
5.4.	LOCATOR Parameter	14
5.5.	RELAY_HMAC	16
5.6.	Registration Types	16
5.7.	HIP ESP Data Packet Formats	17
6.	Security Considerations	17
6.1.	Privacy Considerations	17
6.2.	Opportunistic Mode	18
7.	IANA Considerations	18
8.	Contributors	18
9.	Acknowledgements	18
10.	References	19
10.1.	Normative References	19
10.2.	Informative References	20
Appendix A.	Firewall Traversal	21
Appendix B.	Base Exchange without ICE Connectivity Checks	22
Appendix C.	IPv4-IPv6 Interoperability	22
Appendix D.	Base Exchange through a Rendezvous Server	22
Appendix E.	Document Revision History	23
	Authors' Addresses	23
	Intellectual Property and Copyright Statements	26

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

1. Introduction

HIP [[I-D.ietf-hip-base](#)] is defined as a protocol that runs directly over IPv4 or IPv6. This approach is known to have problems traversing NATs. Several different types of NATs exist, see [[RFC2663](#)]. This document describes HIP extensions for the traversal of both Network Address Translator (NAT) and Network Address and Port Translator (NAPT) middleboxes. Additionally, it covers firewalls to a certain extent (see [Appendix A](#) for a more detailed discussion). The document generally uses the term NAT to refer to these types of middleboxes. A detailed description of HIP problems with traversing legacy middleboxes is documented in [[I-D.irtf-hiprg-nat](#)].

NAT devices do not operate consistently even though a recommended behavior is described in [[RFC4787](#)]. The HIP protocol extensions in this document make as few assumptions as possible about the behavior of the NAT devices so that NAT traversal will work even with legacy NAT devices. The purpose of these extensions is to allow two HIP-enabled hosts to communicate with each other even if one or both communicating hosts are in private address realms. With some legacy NAT devices, utilizing the shortest path between two end hosts located behind NATs is not possible without relaying the traffic through a relay, such as a TURN server [[I-D.ietf-behave-p2p-state](#)]. As a consequence, the TURN server increases the roundtrip delay and may become a point of network congestion. With the extensions described in this document, hosts try to avoid the use of such a relay when possible.

A distinction must be made between a HIP rendezvous server (defined

in [[I-D.ietf-hip-rvs](#)]) and a HIP Relay, defined herein. HIP rendezvous servers solve initial contact and mobility related problems in networks without NATs. HIP Relay solve the same problems, in addition to NAT traversal problems. HIP Relay servers can be used both in NATted and non-NATted networks.

Both rendezvous and relay services forward HIP control packets, but the main difference is that the rendezvous service forwards only the initial I1 packet of the base exchange while all other HIP control packets are sent directly between the communicating hosts. In contrast, the relay service relays all HIP control packets because p2p-unfriendly NAT devices drop the packets otherwise [[I-D.ietf-behave-p2p-state](#)]. The peers use the control channel to communicate their current locators to each other to find a direct path for carrying ESP encapsulated data traffic. A direct path between the hosts enables efficient delivery of data traffic without relaying of ESP packets through an intermediary TURN server. The direct path is searched using connectivity tests.

The basis for the connectivity tests is ICE [[I-D.ietf-mmusic-ice](#)].

[[I-D.ietf-mmusic-ice](#)] describes ICE as follows:

"The Interactive Connectivity Establishment (ICE) methodology is a technique for NAT traversal for UDP-based media streams (though ICE can be extended to handle other transport protocols, such as TCP) established by the offer/answer model. ICE is an extension to the offer/answer model, and works by including a multiplicity of IP addresses and ports in SDP offers and answers, which are then tested for connectivity by peer-to-peer connectivity checks. The IP addresses and ports included in the SDP and the connectivity checks are performed using the revised STUN specification [[I-D.ietf-behave-rfc3489bis](#)], now renamed to Session Traversal Utilities for NAT."

ICE for SIP is specified in [[I-D.ietf-mmusic-ice](#)] and ICE for non-SIP protocols is specified in [[I-D.rosenberg-mmusic-ice-nonsip](#)].

Two hosts communicate their peer address set (typically consisting of IP address and port number pairs) to each other in the HIP base exchange. They are then paired with the locally operational address

of the other end point and prioritized according to some policy. These address sets are then tested sequentially based on the procedure specified in ICE. Both sides participate in the connectivity tests. The tests also determine whether operational address pairs and select the preferred address pair to be used for subsequent communication.

As a summary, the extensions in this document

- o illustrate how to encapsulate HIP packets in UDP
- o refer to the UDP encapsulation of IPsec ESP packets defined in [Section 2.1 of RFC 3948](#) [[RFC3948](#)]
- o define how a node interacts with a HIP rendezvous server (defined in [[I-D.ietf-hip-rvs](#)]) when middleboxes are present
- o describe a methodology to determine operational address pairs between two end hosts based on ICE.

[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 document borrows terminology from [[I-D.ietf-hip-base](#)], [[I-D.ietf-hip-mm](#)], [[RFC4423](#)], [[I-D.ietf-mmusic-ice](#)], and [[I-D.ietf-behave-rfc3489bis](#)]. Additionally, the following terms are used:

Rendezvous server:

A host that forwards I1 packets to the Responder

HIP Relay:

A host that forwards all HIP control packets between an Initiator and Responder

TURN server:

A server that forwards data traffic between two end-hosts

Locator:

A name that controls how the packet is routed through the network and demultiplexed by the end host. It may include a concatenation of traditional network addresses such as an IPv6 address and end-to-end identifiers such as an ESP SPI. It may also include transport port numbers or IPv6 Flow Labels as demultiplexing context, or it may simply be a network address. [[I-D.ietf-hip-mm](#)] "Address" is used in this document as a synonym for locator.

Transport address:

Transport layer port and the corresponding IPv4/v6 address

Candidate:

A transport address that has not been verified yet for reachability using ICE

Host candidate:

An IPv4 or IPv6 address of a network interface of a host

Server reflexive transport candidate:

A translated transport address of a host as observed by a HIP Relay or a STUN server

Peer reflexive transport candidate:

A translated transport address of a host as observed by its peer

Relayed transport candidate:

A transport address that exists on a TURN server. If a permission exists, packets that arrive at this address are relayed towards the TURN client.

3. Protocol Description

This section describes the normative behavior of the protocol extension. Examples of packet exchanges are provided for illustration purposes.

3.1. Relay Registration and NAT Detection

HIP rendezvous servers are used in non-NATted environments and their use is described in [[I-D.ietf-hip-rvs](#)]. This section specifies a new role for these rendezvous servers to act as HIP Relays. HIP Relays forward HIP control packets between the Initiator and the Responder. TURN servers [[I-D.ietf-behave-turn](#)] are used for relaying ESP traffic. A host SHOULD register to a TURN server before registering to a HIP Relay to guarantee that the host can accept ESP traffic immediately after HIP Relay registration.

A HIP relay forwards UDP-encapsulated HIP traffic, and in future extensions, a relay may also forward TCP-encapsulated traffic. The HIP Relay forwards HIP control packets. NAT traversal for HIP between two end-hosts may require the use of relays in certain scenarios. A successful NAT traversal therefore requires at least the Responder located behind a NAT to register with a HIP Relay.

A HIP Relay MUST silently drop packets to a HIP Relay Client that has not previously registered with the HIP Relay. The registration process follows the generic registration extensions defined in [[I-D.ietf-hip-registration](#)] and is illustrated in Figure 1.

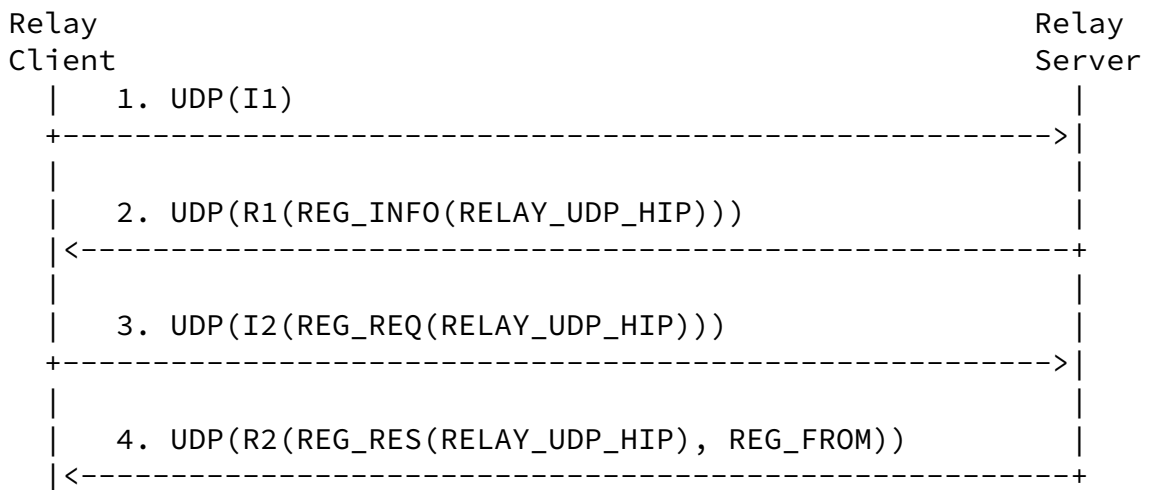


Figure 1: Example Registration to a HIP Relay

In step 1, the Initiator starts the registration procedure by sending an I1 packet over UDP. It is RECOMMENDED that the Initiator selects a random port number from the ephemeral port range 49152–65535 for initiating a base exchange. However, the allocated port MUST be maintained until all of the corresponding HIP Associations are closed. Alternatively, a host MAY also use a single fixed port for initiating all outgoing connections.

In step 2, the Responder lists the services that it supports in the R1 packet. The support for HIP-over-UDP relaying is denoted by the RELAY_UDP_HIP value. The R1 does not contain any NAT transform parameter (see [Section 4.1](#)) as discussed in [Appendix B](#).

In step 3, the Initiator selects the services it registers for and lists them in the REG_REQ parameter. In this example, the Initiator registers for HIP Relay service.

In step 4, the Responder concludes the registration procedure with an R2 packet and acknowledges the registered services in the REG_RES parameter. The Responder may also denote unsuccessful registrations in the REG_FAILED parameter in R2. The Responder also includes a REG_FROM parameter that contains the transport address of the client as observed by the Relay (Server Reflexive candidate). After the registration, the Initiator needs to send periodically NAT keep-alives.

There are different ways for an Initiator to learn it's publically visible IP address and port that are referred to as the "server reflexive transport candidate" in this document. This document makes use of two ways:

- o The Relay client may use STUN servers to detect the server reflexive locator, as described in [[I-D.ietf-behave-p2p-state](#)].
- o Alternatively, the Relay Client can learn it from the REG_FROM parameter when registering to a Relay.

3.2. Base Exchange via HIP Relay

It is RECOMMENDED that the Initiator sends an I1 packet encapsulated in UDP when it is destined to an IPv4 address of the Responder. Respectively, the Responder MUST respond to a such I1 packet with an R1 packet over the transport layer and using the same transport protocol. The rest of the base exchange, I2 and R2, MUST also use the same transport layer.

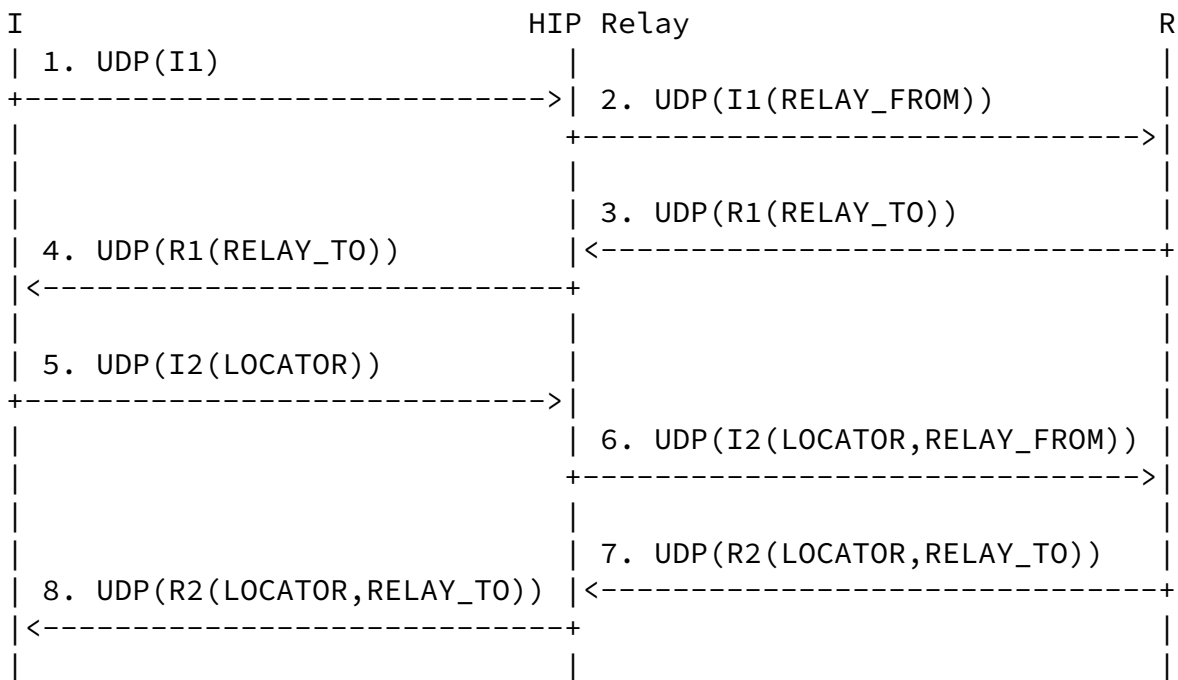


Figure 2: Base Exchange via a HIP Relay

In step 1 of Figure 2, the Initiator sends an I1 packet over the transport layer to the HIT of the Responder. The source address is one of the locators of the host. The locators of the end-hosts are referred as "host candidates" in this document.

In step 2, the HIP Relay receives the I1 packet at port HIPPORT. If the destination HIT belongs to a registered Responder, the Relay processes the packet. Otherwise, the Relay MUST drop the packet silently. The Relay appends a RELAY_FROM parameter to the I1 packet which contains the transport source address and port of the I1 as

observed by the Relay. The Relay protects the I1 packet with RELAY_HMAC as described in [[I-D.ietf-hip-rvs](#)], except that the

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

parameter type is different. The Relay changes the source and destination ports and IP addresses of the packet to match the values the Responder used when registering to the Relay, i.e., the reverse of the R2 used in the registration. The Relay MUST recalculate the transport checksum and forward the packet to the Responder.

In step 3, the Responder receives the I1 packet. The Responder processes it according to the rules in [[I-D.ietf-hip-base](#)]. In addition, the Responder validates the RELAY_HMAC according to [[I-D.ietf-hip-rvs](#)] and silently drops the packet if the validation fails. The Responder replies with an R1 packet to which it includes a RELAY_TO parameter. The RELAY_TO parameter contains same information as the RELAY_FROM parameter, i.e., Initiator transport address, but the type of the parameter is different. The RELAY_TO parameter is not integrity protected by the signature of the R1 to allow pre-created R1 packets at the Responder.

In step 4, the Relay receives the R1 packet. The Relay drops the packet silently if the source HIT belongs to an unregistered host. The Relay MAY verify the signature of the R1 packet and drop it if the signature is invalid. Otherwise, the Relay rewrites to source address and port, changes the destination address and port to match RELAY_TO information, recalculates transport checksum and forwards the packet.

In step 5, the Initiator receives the R1 packet and processes it according to [[I-D.ietf-hip-base](#)]. It replies with an I2 packet that uses the destination transport address of R1 as the source address and port. The I2 contains a LOCATOR parameter that lists all the ICE candidates (offer) of the Initiator. The candidates are encoded using the format defined in [Section 5.4](#).

In step 6, the Relay receives the I2 packet. The relay appends a RELAY_FROM and a RELAY_HMAC to the I2 packet as in the second step.

In step 7, the Responder receives the I2 packet and processes it according to [[I-D.ietf-hip-base](#)]. It replies with a R2 packet and includes a RELAY_TO parameter as in step three. The R2 packet includes a LOCATOR parameter that lists all the ICE candidates

(answer) of the Responder. The RELAY_TO parameter is protected by the HMAC.

In step 8, the Relay processes the R2 as described in step four. The Relay forwards the packet to the Responder.

4. Connectivity Tests

4.1. NAT Transformation Negotiation

This section describes usage of a new optional transform parameter type. The presence of the parameter in HIP base exchange means that the host supports all of the extensions defined in this document. If the transform parameter is used, hosts MUST use a password for STUN HMACs that is drawn from the DH keying material.

The transform parameter applies both to the registration to the HIP Relay as well as to a base exchange between end-hosts. The transform negotiation in base exchange is illustrated in Figure 3.

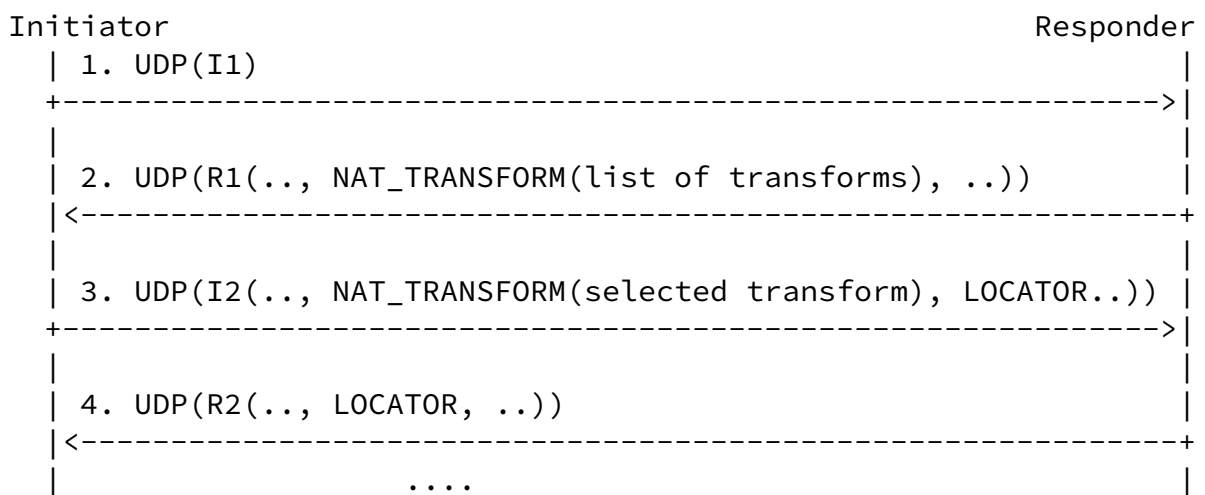


Figure 3: Negotiation of NAT Transforms

In step 1, the Initiator sends an I1 to the Responder. In step 2, the Responder responds with an R1. The R1 contains a list of

transforms the Responder supports in NAT_TRANSFORM parameter as shown in Table 1.

Transform Type	Purpose
RESERVED	Reserved for future use
ICE-STUN-UDP	UDP encapsulated control and data traffic with ICE-based connectivity tests using STUN messages

Table 1: Locator Transformations

In step 3, the Initiator sends an I2 that includes a NAT_TRANSFORM parameter. It contains the transform type selected by the Initiator


from the list of transforms offered by the Responder. The I2 also includes the locators of the Initiator in a LOCATOR parameter.

In step 4, the Responder concludes the base exchange with an R2 packet. The Responder includes a LOCATOR parameter in the R2 packet.

4.2. ICE Procedure

Hosts exchange HIP control packets through the HIP Relay. Connectivity tests are, however, directly exchanged between the address pairs to determine operational address pairs. If a working direct path between the hosts is found, also the HIP control traffic MAY start using it.

The base exchange is completed with an R2 packet. Then, the state of the HIP associations at both peers is ESTABLISHED, but the peers MUST NOT allow any ESP traffic until the connectivity tests are performed successfully. All of the locators, except the HIP Relay address, are in UNVERIFIED state. In the connectivity tests, the hosts test connectivity between different locator pairs in order to find a working one. The connectivity tests are illustrated in Figure 4. In this example, both hosts are behind NATs.



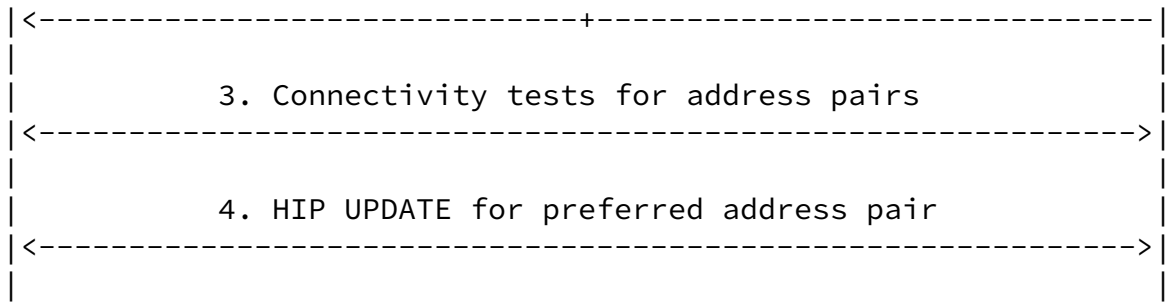


Figure 4: Connectivity tests

In steps 1 and 2, the R2 packet is relayed from the Responder through the Relay to the Initiator.

Afterwards, connectivity tests are started based on the procedure described in [[I-D.rosenberg-mmusic-ice-nonsip](#)] by using the candidates previously exchanged in the HIP base exchange.

4.3. NAT Keep-alives


Data channel keepalives are STUN Binding Indications. Keepalives MUST be sent every 20 seconds at the minimum when the channel is idle. To implement failure tolerance, a host SHOULD have smaller

keepalive period. When data traffic is exchanged between the end points then no further STUN keepalives need to be exchanged.

5. Packet Formats

The following subsections define the parameter and packet encodings. All values MUST be in network byte order.

5.1. HIP Control Packets



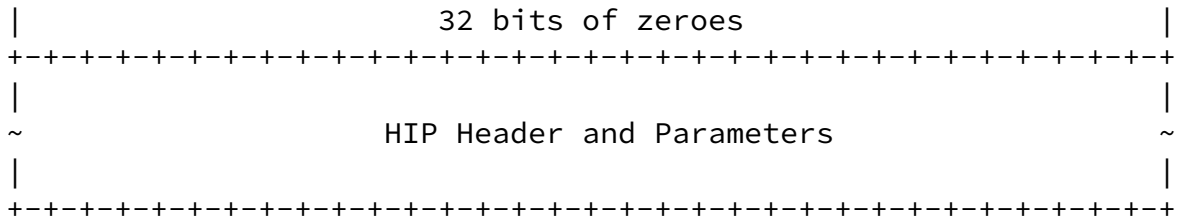


Figure 5: Format for UDP-encapsulated HIP Control Packets

HIP control packets are encapsulated in UDP packets like in [Section 2.2 of \[RFC3948\]](#), "rules for encapsulating IKE messages", except that a different port number is used. Figure 5 shows the encapsulation: UDP header is followed by 32 zero bits that can be used to differentiate HIP control packets from ESP packets. The HIP header and parameters follow the conventions of [\[I-D.ietf-hip-base\]](#) with the exception that the HIP header checksum MUST be zero. The HIP header checksum is zero for two reasons. First, the UDP header contains already a checksum. Second, the checksum definition in [\[I-D.ietf-hip-base\]](#) includes the IP addresses in the checksum calculation. The NATs unaware of HIP cannot recompute the HIP checksum after changing IP addresses.


A HIP Relay or a Responder without a relay MUST listen at transport port HIPPORT for incoming UDP-encapsulated HIP control packets.

5.2. Keep-Alives

Control and data channel keep-alives are STUN Binding Indications, as defined in [\[I-D.ietf-behave-rfc3489bis\]](#). They use the same UDP header as the HIP control packets but there is no non-ESP-marker

between the UDP header and the STUN header. STUN messages are demultiplexed from ESP and HIP control messages using the STUN markers, such as the magic cookie value.

5.3. Relay and Registration Parameters



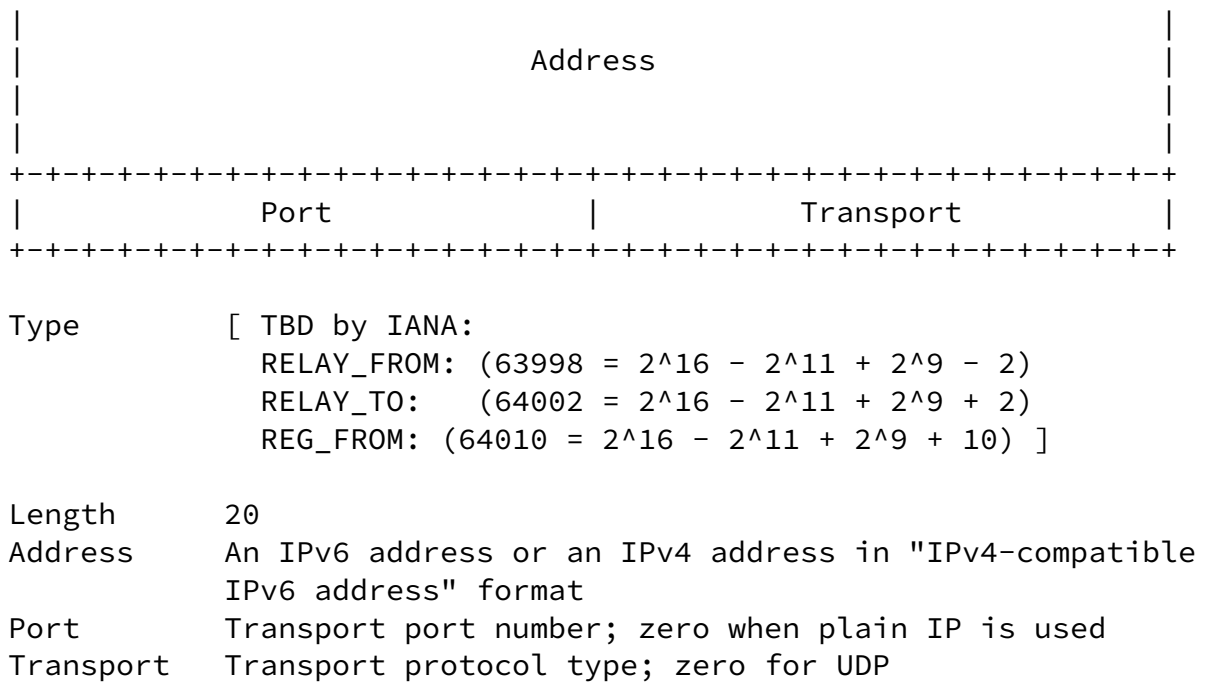



Figure 6: Format for the RELAY_FROM, RELAY_TO and REG_FROM parameters

Format of the RELAY_FROM, RELAY_TO and REG_FROM parameters is shown in Figure 6. Parameters are identical except for the type field.

5.4. LOCATOR Parameter

The generic LOCATOR parameter format is the same as in [I-D.ietf-hip-mm]. However, presenting ICE candidates requires a new locator type. The generic and NAT traversal specific locator parameters are illustrated in Figure 7.



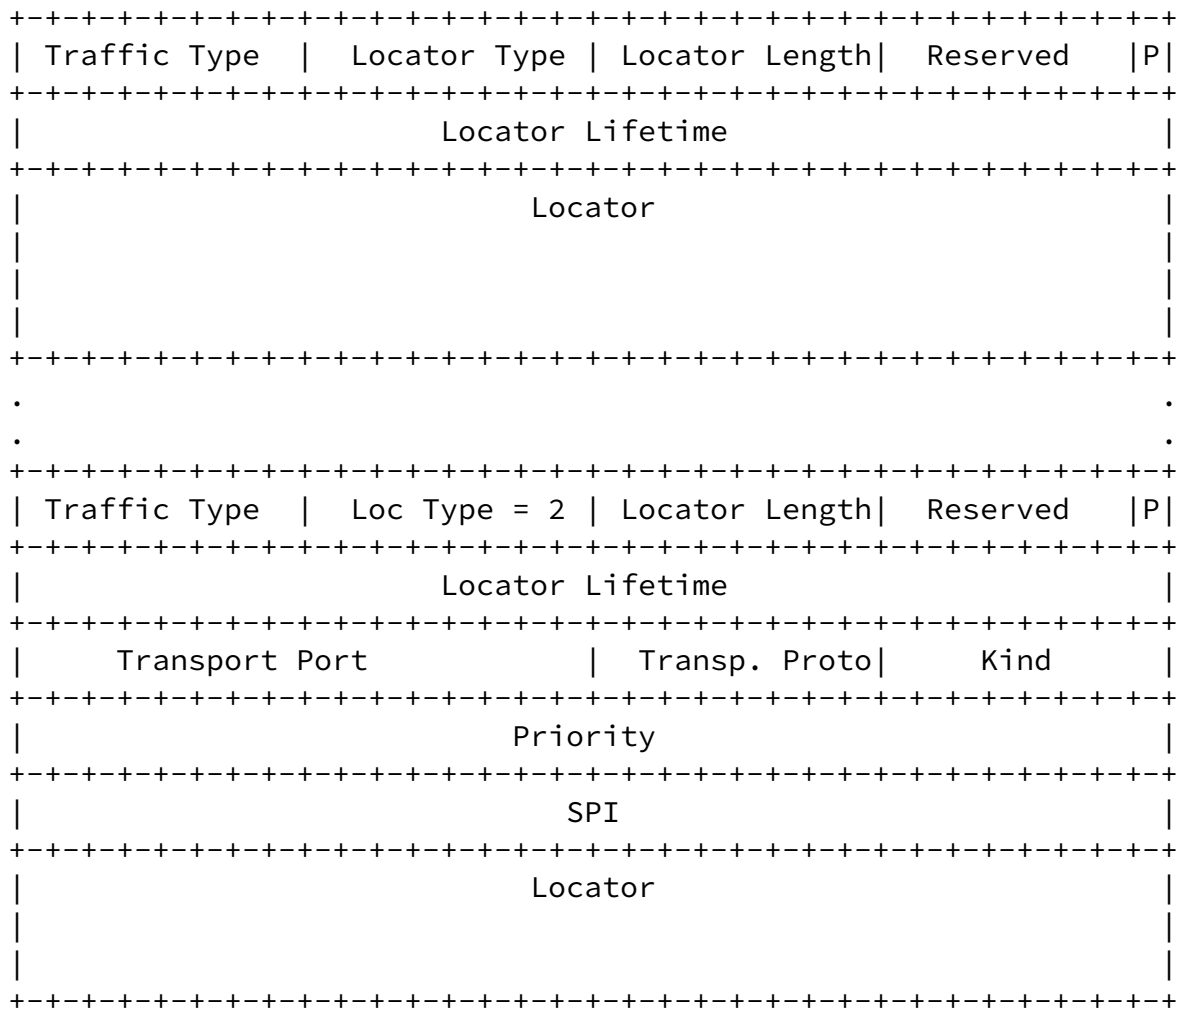


Figure 7: LOCATOR parameter

The individual fields in the LOCATOR parameter are described in Table 2.

Field	Value(s)	Purpose
Type	193	Parameter type
Length	Variable	Length in octets, excluding Type and Length fields and padding
Traffic Type	0-2	Is the locator for HIP signaling (1), for ESP (2), or for both (0)
Locator Type	2	"Transport address" locator type
Locator Length	7	Length of the Locator field in 4-octet units
Reserved	0	Reserved for future extensions
Preferred (P) bit	0	Not used for transport address locators; MUST be ignored by the receiver.
Locator Lifetime	Variable	Locator lifetime in seconds
Transport Port	Variable	Transport layer port number
Transport Protocol	0	0 for UDP
Kind	Variable	0 for host, 1 for server reflexive, 2 for peer reflexive or 3 for relayed address
Priority	Variable	Locator's priority as described in [I-D.ietf-mmusic-ice]
SPI	Variable	SPI value which the host expects to see in incoming ESP packets that use this locator
Locator	Variable	IPv6 address or an "IPv4-compatible IPv6 address" format IPv4 address [RFC3513] , obfuscated by XORring it with the owner's HIT

Table 2: Fields of the LOCATOR parameter

5.5. RELAY_HMAC

The RELAY_HMAC parameter value has the TLV type 65520 ($2^{16} - 2^5 + 2^4$). It has the same semantics as RVS_HMAC [\[I-D.ietf-hip-rvs\]](#).

5.6. Registration Types

The REG_INFO, REQ_REQ, REG_RESP and REG_FAILED parameters contain values for HIP Relay registration. The value for RELAY_UDP_HIP is 2. The value for RELAY_UDP_ESP is 3.

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

[5.7.](#) HIP ESP Data Packet Formats

[RFC3948] describes UDP encapsulation of the IPsec ESP transport and tunnel mode. On the wire the HIP ESP packets do not differ from the transport mode ESP and thus the encapsulation of the HIP ESP packets is same as the UDP encapsulation transport mode ESP.

During the HIP base exchange, the two peers exchange parameters that enable them to define a pair of IPsec ESP security associations (SAs), as described in [[I-D.ietf-hip-esp](#)]. When two peers perform a UDP-encapsulated base exchange, they MUST define a pair of IPsec SAs that produces UDP-encapsulated ESP data traffic.

The management of encryption/authentication protocols and security parameter indices (SPIs) is defined in [[I-D.ietf-hip-esp](#)]. The UDP encapsulation format and processing of HIP ESP traffic is described in Section 6.1 of [[I-D.ietf-hip-esp](#)].

[Section 5.1 of \[RFC3948\]](#) describes a security issue for the UDP encapsulation in the standard IP tunnel mode when two hosts behind different NATs have the same private IP address and initiate communication to the same Responder in the public Internet. The Responder cannot distinguish between two hosts, because security associations are based on the same inner IP addresses.

This issue does not exist with the UDP encapsulation of HIP ESP transport format because the Responder use HITs to distinguish between different communication instances.

[6.](#) Security Considerations

[6.1.](#) Privacy Considerations

The LOCATORs are sent XORed format in plain text in favour of inspection at HIP-aware middleboxes in the future. The current draft does not specify encrypted versions of LOCATORs even though it could be beneficial for privacy reasons.

It is possible that an Initiator or Responder may not want to reveal

all of its locators to its peer. For example, a host may not want to reveal the internal topology of the private address realm and it discards host addresses. Such behavior creates non-optimal paths when the hosts are located behind the same NAT. Especially, this could be a problem with a legacy NAT that does not support routing from the private address realm back to itself through the outer address of the NAT. This scenario is referred to as the hairpin problem [[I-D.ietf-behave-p2p-state](#)]. With such a legacy NAT, the

only option left would be to use a relayed transport address from a TURN server. As a consequence, a host may support locator-based privacy by leaving out the reflexive candidates. Using only host candidates can produce suboptimal paths possibly causing congestion.

The use of HIP Relays or TURN Relays can be useful for protection against Denial-of-Service attacks. If a Responder reveals only its HIP and ESP relay candidates to malign Initiators, the Initiators can only attack the relays that does not prevent the Responder from initiating new outgoing connections if a path around the relay exists.

[6.2.](#) Opportunistic Mode

A HIP Relay should have one address per Relay Client when a HIP Relay is serving more than one Relay Clients and is willing to support opportunistic mode. Otherwise, it cannot be guaranteed that the Relay can deliver the I1 packet to the intended recipient.

[7.](#) IANA Considerations

This section is to be interpreted according to [[RFC2434](#)].

This draft currently uses a UDP port in the "Dynamic and/or Private Port" and HIPPORT. Upon publication of this document, IANA is requested to register a UDP port and the RFC editor is requested to change all occurrences of port HIPPORT to the port IANA has registered. The HIPPORT number 50500 should be used for initial experimentation.

This document updates the IANA Registry for HIP Parameter Types by assigning new HIP Parameter Type values for the new HIP Parameters:

RELAY_FROM, RELAY_TO and REG_FROM (defined in [Section 5.3](#)) and RELAY_HMAC (defined in [Section 5.5](#)).

[8.](#) Contributors

Marcelo Bagnulo, Jan Melen, Simon Schuetz, Martin Stiernerling, Lars Eggert, Vivien Schmitt, Abhinav Pathak and Andrei Gurtov have contributed to the initial versions of this draft.

[9.](#) Acknowledgements

Thanks for Jonathan Rosenberg and the rest of the MMUSIC WG folks for the excellent work on ICE. In addition, the authors would like to

Komu, et al.

Expires August 28, 2008

[Page 18]

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

thank Andrei Gurtov, Tobias Heer, Teemu Koponen, Juhana Mattila, Jeffrey M. Ahrenholz, Thomas Henderson, Kristian Slavov, Janne Lindqvist, Pekka Nikander, Lauri Silvennoinen, Jukka Ylitalo, Juha Heinanen, Joakim Koskela, Samu Varjonen, Dan Wing, Hannes Tschofenig, Jan Melen, Jani Hautakorpi and Ari Keraenen For their comments on this document.

Miika Komu is working in the Networking Research group at Helsinki Institute for Information Technology (HIIT). The InfraHIP project was funded by Tekes, Telia-Sonera, Elisa, Nokia, the Finnish Defence Forces, and Ericsson and Birdstep.

[10.](#) References

[10.1.](#) Normative References

[I-D.ietf-behave-rfc3489bis]

Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)", [draft-ietf-behave-rfc3489bis-15](#) (work in progress), February 2008.

[I-D.ietf-behave-turn]

Rosenberg, J., Mahy, R., and P. Matthews, "Traversal Using Relays around NAT (TURN): Relay Extensions to Session

Traversal Utilities for NAT (STUN)",
[draft-ietf-behave-turn-06](#) (work in progress),
January 2008.

[I-D.ietf-hip-base]

Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson,
"Host Identity Protocol", [draft-ietf-hip-base-10](#) (work in
progress), October 2007.

[I-D.ietf-hip-esp]

Jokela, P., "Using ESP transport format with HIP",
[draft-ietf-hip-esp-06](#) (work in progress), June 2007.

[I-D.ietf-hip-mm]

Henderson, T., "End-Host Mobility and Multihoming with the
Host Identity Protocol", [draft-ietf-hip-mm-05](#) (work in
progress), March 2007.

[I-D.ietf-hip-registration]

Laganier, J., "Host Identity Protocol (HIP) Registration
Extension", [draft-ietf-hip-registration-02](#) (work in
progress), June 2006.

[I-D.ietf-hip-rvs]

Laganier, J. and L. Eggert, "Host Identity Protocol (HIP)
Rendezvous Extension", [draft-ietf-hip-rvs-05](#) (work in
progress), June 2006.

[I-D.ietf-mmusic-ice]

Rosenberg, J., "Interactive Connectivity Establishment
(ICE): A Protocol for Network Address Translator (NAT)
Traversal for Offer/Answer Protocols",
[draft-ietf-mmusic-ice-19](#) (work in progress), October 2007.

[I-D.rosenberg-mmusic-ice-nonsip]

Rosenberg, J., "NICE: Non Session Initiation Protocol
(SIP) usage of Interactive Connectivity Establishment
(ICE)", [draft-rosenberg-mmusic-ice-nonsip-00](#) (work in
progress), February 2008.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", [RFC 3513](#), April 2003.
- [RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", [RFC 4423](#), May 2006.

10.2. Informative References

- [I-D.ietf-behave-p2p-state]
Srisuresh, P., Ford, B., and D. Kegel, "State of Peer-to-Peer (P2P) Communication Across Network Address Translators (NATs)", [draft-ietf-behave-p2p-state-06](#) (work in progress), November 2007.
- [I-D.irtf-hiprg-nat]
Stiemerling, M., "NAT and Firewall Traversal Issues of Host Identity Protocol (HIP) Communication", [draft-irtf-hiprg-nat-04](#) (work in progress), March 2007.
- [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.

- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", [BCP 127](#), [RFC 4787](#), January 2007.

Appendix A. Firewall Traversal

This section describes firewall traversal issues separately from NAT issues. When the Initiator or the Responder of a HIP association is

behind a firewall, additional issues arise. The firewall discussion applies both to IPv4 and IPv6 addressing.

The NAT traversal mechanisms described in this document require that the firewall - stateful or not - allows UDP traffic. At the minimum, successful firewall control packet traversal requires that the host behind the firewall is allowed to communicate packets with a HIP Relay (or a Responder without HIP Relay) that is listening on UDP port HIPPORT. Successful ESP data packet traversal requires the same for the TURN server. For unrelayed traffic, the destination port HIPPORT should be open at the firewall to all hosts behind the firewall.

Most firewall implementations support "UDP connection tracking", i.e., after a host behind a firewall has initiated UDP communication to the public Internet, the firewall accepts UDP response traffic in the return direction. If no such return traffic arrives for a specific period of time, the firewall stops accepting the given IP address and port pair. The mechanisms described in this document already enable traversal of such firewalls, if the keep-alive interval used is less than the refresh interval of the firewall.

When the Initiator is behind a firewall, the NAT traversal mechanisms described in this document depend on the ability to initiate communication via UDP to the destination port HIPPORT from arbitrary source ports and to receive UDP response traffic from that port to the chosen source port. If the Initiator is behind a firewall that does not support "UDP connection tracking", the NAT traversal mechanisms described in this document can still be supported, if the firewall allows permanently inbound UDP traffic from the port HIPPORT and destined to arbitrary source IP addresses and UDP ports.

When the Responder is behind a firewall, the NAT traversal mechanisms described in this document depend on the ability to send and receive UDP traffic originating from HIPPORT of the HIP Relays and TURN servers. When end-to-end traffic is preferred, arbitrary source IP

addresses and ports are required.

In certain network environments, the ICE connectivity tests can be omitted to reduce initial connection set up latency because base exchange acts as an implicit connectivity test itself. There are three assumptions about such as environments. First, the Responder should have a long-term, fixed locator in the network. Second, the Responder should not have a HIP Relay configured for itself. Third, the Initiator can reach the Responder by simply UDP encapsulating HIP and ESP packets to the host. Detecting and configuring this particular scenario is prone to administrative failure unless carefully planned.

In such a scenario, the Initiator sends an I1 packet over UDP to the Responder. The Responder replies with a R1 packet that does not contain the transform parameter as explained in [Section 4.1](#). The Initiator receives the R1 packet and determines from the absence of the transform and RELAY_TO parameters that ICE connectivity tests can be omitted with the Responder. Finally, the hosts set up IPsec security associations using the locators observed from the concluding I2 and R2 packets of the base exchange without ICE connectivity tests.

[Appendix C](#). IPv4-IPv6 Interoperability

Currently Relay Client and Server do not have to run any ICE connectivity tests as described in [Appendix B](#). However, it could be useful for IPv4-IPv6 interoperability when the Relay Server actually includes both the NAT transform parameter and multiple locators in R2. The interoperability benefit is that the Relay could support IPv4-based Initiators and IPv6-based Responders by converting the network headers and recalculating UDP checksums.

Such an approach is underspecified in this document currently. It is not yet recommended because it may consume resources at the Relay and requires also similar conversion support at the TURN relay for data packets.

[Appendix D](#). Base Exchange through a Rendezvous Server

This section describes handling for a scenario where Initiator looks up the information of the Responder from DNS and discovers a RVS record [[I-D.ietf-hip-rvs](#)]. In such a case, the Initiator uses its

own HIP Relay to forward HIP traffic to the Rendezvous server. The Initiator will send the I1 message using the its HIP Relay server which will then forward it to the RVS server of the responder. The responder will send the R1 packet directly to the Initiator's HIP Relay server and the following I2 and R2 packets are also sent directly using the Relay.

In case the Initiator is not able to distinguish which records are RVS address records and which are Responders address records, then the Initiator SHOULD first try to contact the Responder directly and if none of the addresses is reachable it MAY try out them using its own HIP Relay as described in the above.

[Appendix E](#). Document Revision History

To be removed upon publication

Revision	Comments
draft-ietf-nat-traversal-00	Initial version.
draft-ietf-nat-traversal-01	Draft based on RVS.
draft-ietf-nat-traversal-02	Draft based on Relay proxies and ICE concepts.
draft-ietf-nat-traversal-03	Draft based on STUN/ICE formats.

Authors' Addresses

Miika Komu
Helsinki Institute for Information Technology
Metsanneidonkuja 4
Espoo
Finland

Phone: +358503841531
Fax: +35896949768
Email: miika@iki.fi
URI: <http://www.hiit.fi/>

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

Thomas Henderson
The Boeing Company
P.O. Box 3707
Seattle, WA
USA

Email: thomas.r.henderson@boeing.com

Philip Matthews
Avaya
100 Innovation Drive
Ottawa, Ontario K2K 3G7
Canada

Phone: +1 613 592 4343 224
Email: philip_matthews@magma.ca

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.com>

Ari Keraenen
Ericsson Research Nomadiclab
Hirsalantie 11
02420 Jorvas
Finland

Phone: +358 9 2991
Email: ari.keranen@ericsson.com

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

Jan Melen
Ericsson Research Nomadiclab
Hirsalantie 11
02420 Jorvas
Finland

Phone: +358 9 2991
Email: jan.melen@ericsson.com

Marcelo Bagnulo
Huawei Lab at UC3M
Av. Universidad 30
Leganes, Madrid 28911
Spain

Phone: 34 91 6249500
Email: marcelo@it.uc3m.es
URI: <http://www.it.uc3m.es/>

Internet-Draft Basic NAT and Firewall Traversal for HIP February 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).