

HIP Working Group
Internet-Draft
Intended status: Experimental
Expires: April 26, 2010

M. Komu
HIIT
T. Henderson
The Boeing Company
H. Tschofenig
Nokia Siemens Networks
J. Melen
A. Keranen, Ed.
Ericsson Research Nomadiclab
October 23, 2009

Basic HIP Extensions for Traversal of Network Address Translators
draft-ietf-hip-nat-traversal-09.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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>.

This Internet-Draft will expire on April 26, 2010.

Internet-Draft

Basic NAT Traversal for HIP

October 2009

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document specifies extensions to the Host Identity Protocol (HIP) to facilitate Network Address Translator (NAT) traversal. The extensions are based on the use of the Interactive Connectivity Establishment (ICE) methodology to discover a working path between two end-hosts, and on standard techniques for encapsulating Encapsulating Security Payload (ESP) packets within the User Datagram Protocol (UDP). This document also defines elements of a procedure for NAT traversal, including the optional use of a HIP relay server. With these extensions HIP is able to work in environments that have NATs and provides a generic NAT traversal solution to higher-layer networking applications.

Internet-Draft

Basic NAT Traversal for HIP

October 2009

Table of Contents

1.	Introduction	4
2.	Terminology	6
3.	Overview of Operation	7
4.	Protocol Description	8
4.1.	Relay Registration	8
4.2.	ICE Candidate Gathering	10
4.3.	NAT Traversal Mode Negotiation	10
4.4.	Connectivity Check Pacing Negotiation	12
4.5.	Base Exchange via HIP Relay Server	12
4.6.	ICE Connectivity Checks	15
4.7.	NAT Keepalives	15
4.8.	Base Exchange without ICE Connectivity Checks	16
4.9.	Initiating a Base Exchange both with and without UDP Encapsulation	17
4.10.	Sending Control Packets after the Base Exchange	18
5.	Packet Formats	18
5.1.	HIP Control Packets	18
5.2.	Connectivity Checks	19
5.3.	Keepalives	20
5.4.	NAT Traversal Mode Parameter	20
5.5.	Connectivity Check Transaction Pacing Parameter	21
5.6.	Relay and Registration Parameters	22
5.7.	LOCATOR Parameter	23
5.8.	RELAY_HMAC Parameter	24
5.9.	Registration Types	24
5.10.	Notify Packet Types	25
5.11.	ESP Data Packets	25
6.	Security Considerations	26
6.1.	Privacy Considerations	26
6.2.	Opportunistic Mode	26
6.3.	Base Exchange Replay Protection for HIP Relay Server	26
6.4.	Demuxing Different HIP Associations	27
7.	IANA Considerations	27
8.	Contributors	28

9.	Acknowledgments	28
10.	References	28
10.1.	Normative References	28
10.2.	Informative References	29
Appendix A.	Selecting a Value for Check Pacing	30
Appendix B.	Base Exchange through a Rendezvous Server	31
Appendix C.	Document Revision History	31
	Authors' Addresses	31

[1.](#) Introduction

HIP [[RFC5201](#)] is defined as a protocol that runs directly over IPv4 or IPv6, and HIP coordinates the setup of ESP security associations [[RFC5202](#)] that are also specified to run over IPv4 or IPv6. This approach is known to have problems traversing NATs and other middleboxes [[RFC5207](#)]. This document defines HIP extensions for the traversal of both Network Address Translator (NAT) and Network Address and Port Translator (NAPT) middleboxes. The document generally uses the term NAT to refer to these types of middleboxes.

Currently deployed 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 of the communicating hosts are in a network that is behind one or more NATs.

Using the extensions defined in this document, HIP end-hosts use techniques drawn from the Interactive Connectivity Establishment (ICE) methodology [[I-D.ietf-mmusic-ice](#)] to find operational paths for the HIP control protocol and for ESP encapsulated data traffic. The hosts test connectivity between different locators and try to discover a direct end-to-end path between them. However, with some legacy NATs, 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 [[RFC5128](#)]. Because relaying

traffic increases the roundtrip delay and consumes resources from the relay, with the extensions described in this document, hosts try to avoid using the TURN server whenever possible.

HIP has defined a Rendezvous Server [[RFC5204](#)] to allow for mobile HIP hosts to establish a stable point-of-contact in the Internet. This document defines extensions to the Rendezvous Server that solve the same problems but for both NATed and non-NATed networks. The extended Rendezvous Server, called a "HIP relay server", forwards HIP control packets between an Initiator and a Responder, allowing hosts to be located behind NATs. This behavior is in contrast to the HIP rendezvous service that forwards only the initial I1 packet of the base exchange; an approach which is less likely to work in a NATed environment [[RFC5128](#)]. Therefore, when using relays to traverse NATs, HIP uses a HIP relay server for the control traffic and a TURN server for the data traffic.

The basis for the connectivity checks 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 [[RFC5389](#)], now renamed to Session Traversal Utilities for NAT."

The standard ICE [[I-D.ietf-mmusic-ice](#)] is specified with SIP in mind and it has some features that are not necessary or suitable as such for other protocols. [[I-D.rosenberg-mmusic-ice-nonsip](#)] gives instructions and recommendations on how ICE can be used for other protocols and this document follows those guidelines.

Two HIP hosts that implement this specification communicate their locators to each other in the HIP base exchange. The locators are then paired with the locators of the other endpoint and prioritized according to recommended and local policies. These locator pairs are

then tested sequentially by both of the end hosts. The tests may result in multiple operational pairs but ICE procedures determine a single preferred address pair to be used for subsequent communication.

In summary, the extensions in this document define:

- o UDP encapsulation of HIP packets
- o UDP encapsulation of IPsec ESP packets
- o registration extensions for HIP relay services
- o how the ICE "offer" and "answer" are carried in the base exchange
- o interaction with ICE connectivity check messages
- o backwards compatibility issues with rendezvous servers
- o a number of optimizations (such as when the ICE connectivity tests can be omitted)

[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 [[RFC5201](#)], [[RFC5206](#)], [[RFC4423](#)], [[I-D.ietf-mmusic-ice](#)], and [[RFC5389](#)]. Additionally, the following terms are used:

Rendezvous server:

A host that forwards I1 packets to the Responder.

HIP relay server:

A host that forwards any kind of HIP control packets between the

Initiator and the Responder.

TURN server:

A server that forwards data traffic between two end-hosts as defined in [[I-D.ietf-behave-turn](#)].

Locator:

As defined in [[RFC5206](#)]: "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."

LOCATOR (written in capital letters):

Denotes a HIP control packet parameter that bundles multiple locators together.

ICE offer:

The Initiator's LOCATOR parameter in a HIP I2 control packet.

ICE answer:

The Responder's LOCATOR parameter in a HIP R2 control packet.

Transport address:

Transport layer port and the corresponding IPv4/v6 address.

Candidate:

A transport address that is a potential point of contact for receiving data.

Host candidate:

A candidate obtained by binding to a specific port from an IP address on the host.

Server reflexive candidate:

A translated transport address of a host as observed by a HIP relay server or a STUN/TURN server.

Peer reflexive candidate:

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

Relayed candidate:

A transport address that exists on a TURN server. Packets that arrive at this address are relayed towards the TURN client.

3. Overview of Operation

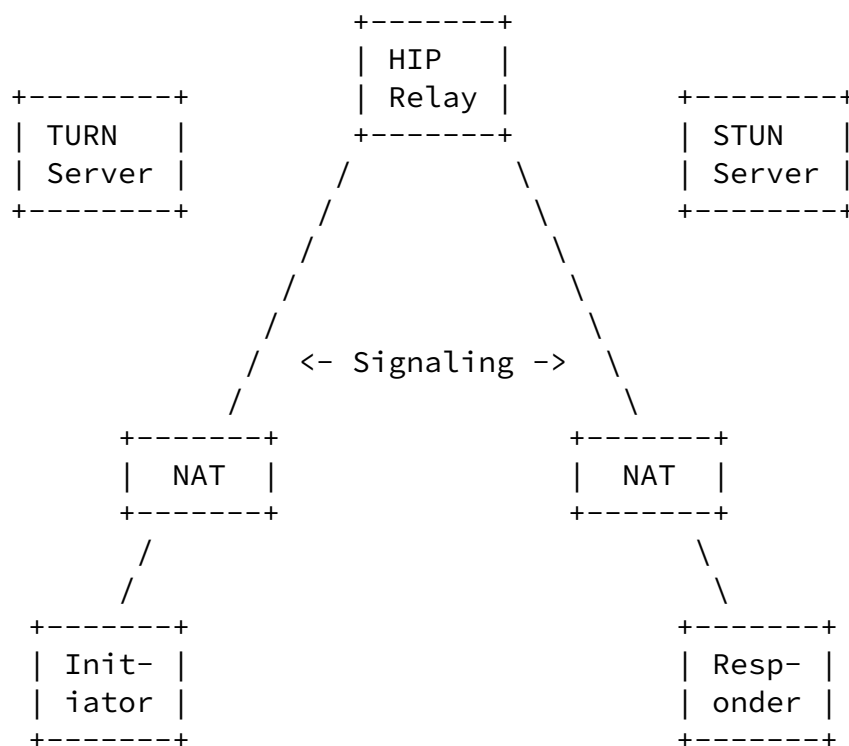


Figure 1: Example network configuration

In an example configuration depicted in Figure 1, both Initiator and Responder are behind one or more NATs, and both private networks are connected to the public Internet. To be contacted from behind a NAT, the Responder must be registered with a HIP relay server reachable on the public Internet, and we assume as a starting point that the Initiator knows both the Responder's HIT and the address of one of

its relay servers (how the Initiator learns of the Responder's relay

server is outside of the scope of this document, but may be through DNS or another name service).

The first steps are for both the Initiator and Responder to register with a relay server (need not be the same one) and gather a set of address candidates. The hosts may use TURN and STUN servers for gathering the candidates. Next, the HIP base exchange is carried out by encapsulating the HIP control packets in UDP datagrams and sending them through the Responder's relay server. As part of the base exchange, each HIP host learns of the peer's candidate addresses through the ICE offer/answer procedure embedded in the base exchange.

Once the base exchange is completed, HIP has established a working communication session (for signaling) via a relay server, but the hosts still work to find a better path, preferably without a relay, for the ESP data flow. For this, ICE connectivity checks are carried out until a working pair of addresses is discovered. At the end of the procedure, if successful, the hosts will have enabled a UDP-based flow that traverses both NATs, with the data flowing directly from NAT to NAT or via a TURN server. Further HIP signaling can be sent over the same address/port pair and is demultiplexed from data traffic via a marker in the payload. Finally, NAT keepalives will be sent as needed.

If either one of the hosts knows that it is not behind a NAT, hosts can negotiate during the base exchange a different mode of NAT traversal that does not use ICE connectivity checks, but only UDP encapsulation of HIP and ESP. Also, it is possible for the Initiator to simultaneously try a base exchange with and without UDP encapsulation. If a base exchange without UDP encapsulation succeeds, no ICE connectivity checks or UDP encapsulation of ESP are needed.

[4.](#) Protocol Description

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

[4.1.](#) Relay Registration

HIP rendezvous servers operate in non-NATed environments and their use is described in [\[RFC5204\]](#). This section specifies a new middlebox extension, called the HIP relay server, for operating in NATed environments. A HIP relay server forwards HIP control packets between the Initiator and the Responder.

End-hosts cannot use the HIP relay service for forwarding the ESP data plane. Instead, they use TURN servers [[I-D.ietf-behave-turn](#)] for that.

A HIP relay server **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 [[RFC5203](#)] and is illustrated in Figure 2.

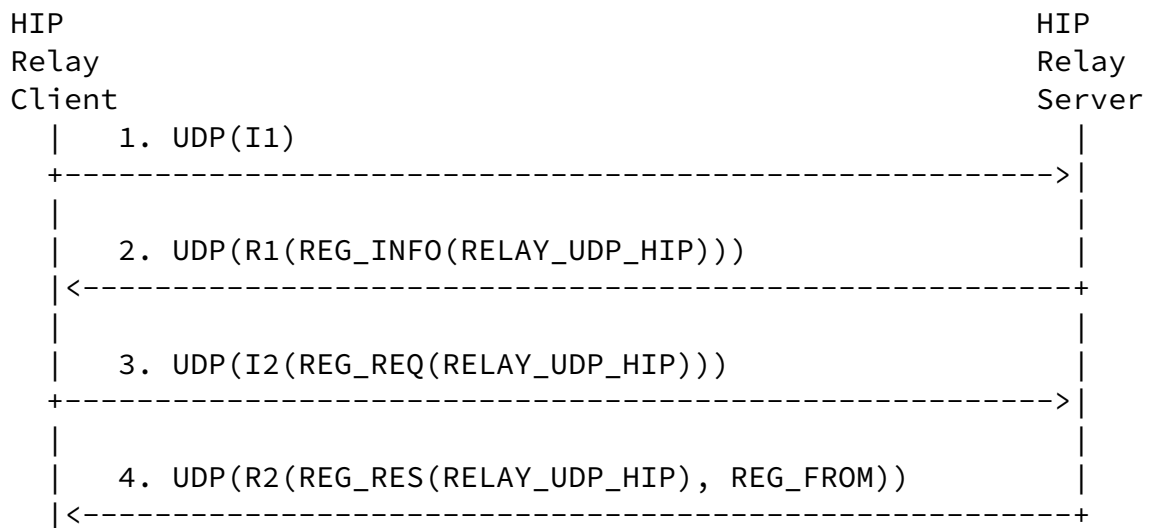


Figure 2: Example Registration with a HIP Relay

In step 1, the relay client (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. Alternatively, a host MAY also use a single fixed port for initiating all outgoing connections. However, the allocated port **MUST** be maintained until all of the corresponding HIP Associations are closed. It is RECOMMENDED that the HIP relay server listens to incoming connections at UDP port HIPPORT. If some other port number is used, it needs to be communicated to possible Initiators.

In step 2, the HIP relay server (Responder) lists the services that it supports in the R1 packet. The support for HIP-over-UDP relaying is denoted by the Registration Type value RELAY_UDP_HIP (see [Section 5.9](#)).

In step 3, the Initiator selects the services it registers for and lists them in the REG_REQ parameter. The Initiator registers for HIP relay service by listing the RELAY_UDP_HIP value in the request

parameter.

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 denotes unsuccessful registrations (if any) in the REG_FAILED parameter of 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 client sends NAT keepalives periodically to the relay to keep possible NAT bindings between the client and the relay alive. The relay client maintains the HIP association with the relay server as long as it requires relaying service from it.

[4.2.](#) ICE Candidate Gathering

If a host is going to use ICE, it needs to gather a set of address candidates. The candidate gathering SHOULD be done as defined in Section 4.1 of [[I-D.ietf-mmusic-ice](#)]. Candidates need to be gathered for the UDP encapsulated flow of HIP and ESP traffic. This flow corresponds to one ICE media stream and component. Since ICE component IDs are not needed, they are not explicitly signaled and ID value of 1 SHOULD be used for ICE processing, where needed. The Initiator takes the role of the ICE controlling agent.

The candidate gathering can be done at any time, but it needs to be done before sending an I2 or R2 in the base exchange if ICE is to be used for the connectivity checks. It is RECOMMENDED that all three types of candidates (host, server reflexive and relayed) are gathered to maximize the probability of successful NAT traversal. However, if no TURN server is used, and the host has only a single local IP address to use, the host MAY use the local address as the only host candidate and the address from the REG_FROM parameter discovered during the relay registration as a server reflexive candidate. In this case, no further candidate gathering is needed.

[4.3.](#) NAT Traversal Mode Negotiation

This section describes the usage of a new non-critical parameter type. The presence of the parameter in a HIP base exchange means that the end-host supports NAT traversal extensions described in this document. As the parameter is non-critical (as defined in [Section](#)

[5.2.1 of \[RFC5201\]](#)), it can be ignored by an end-host which means that the host does not support or is not willing to use these extensions.

With registration with a HIP relay it is usually sufficient to use UDP-ENCAPSULATION mode of NAT traversal since the relay should not be behind a NAT. Thus, the relay SHOULD propose the UDP-ENCAPSULATION mode as the preferred or only mode. The NAT traversal mode negotiation in a HIP base exchange is illustrated in Figure 3.

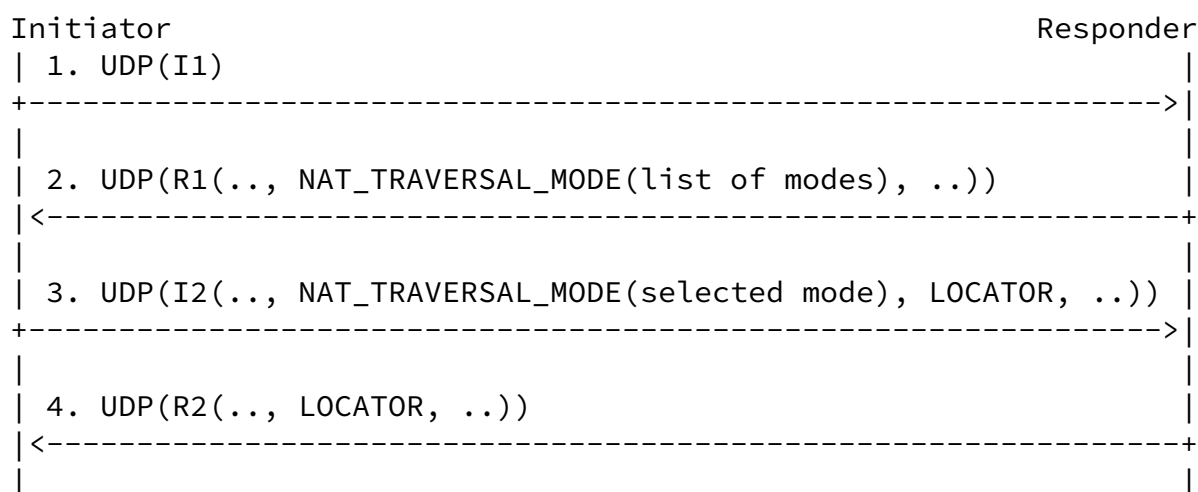


Figure 3: Negotiation of NAT Traversal Mode

In step 1, the Initiator sends an I1 to the Responder. In step 2, the Responder responds with an R1. The NAT_TRAVERSAL_MODE parameter in R1 contains a list of NAT traversal modes the Responder supports. The modes specified in this document are shown in Table 1 and their values in [Section 5.4](#).

Type	Purpose
RESERVED	Reserved for future use
UDP-ENCAPSULATION	Use only UDP encapsulation of the HIP signaling traffic and ESP (no ICE connectivity checks)
ICE-STUN-UDP	UDP-encapsulated control and data traffic with ICE-based connectivity checks using STUN

Table 1: NAT Traversal Modes

In step 3, the Initiator sends an I2 that includes a NAT_TRAVERSAL_MODE parameter. It contains the mode selected by the Initiator from the list of modes offered by the Responder. If ICE mode was selected, the I2 also includes the "Transport address" locators (as defined in [Section 5.7](#)) of the Initiator in a LOCATOR parameter. The locators in I2 are the "ICE offer".

In step 4, the Responder concludes the base exchange with an R2 packet. If the Initiator chose ICE NAT traversal mode, the Responder includes a LOCATOR parameter in the R2 packet. The locators in R2, encoded like the locators in I2, are the "ICE answer". If the NAT

traversal mode selected by the Initiator is not supported by the Responder, the Responder SHOULD reply with NOTIFY packet with type NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER and abort the base exchange.

4.4. Connectivity Check Pacing Negotiation

As explained in [\[I-D.ietf-mmusic-ice\]](#), when a NAT traversal mode with connectivity checks is used, new transactions should not be started too fast to avoid congestion and overwhelming the NATs.

For this purpose, during the base exchange, hosts can negotiate a transaction pacing value, Ta, using a TRANSACTION_PACING parameter in R1 and I2 packets. The parameter contains the minimum time (expressed in milliseconds) the host would wait between two NAT traversal transactions, such as starting a new connectivity check or retrying a previous check. If a host does not include this parameter in the base exchange, a Ta value of 500ms MUST be used as that host's minimum value. The value that is used by both of the hosts is the higher out of the two offered values.

Hosts SHOULD NOT use values smaller than 20ms for the minimum Ta, since such values may not work well with some NATs, as explained in [\[I-D.ietf-mmusic-ice\]](#). The Initiator MUST NOT propose a smaller value than what the Responder offered.

The minimum Ta value SHOULD be configurable, and if no value is configured, value of 500ms MUST be used. Guidelines for selecting a Ta value are given in [Appendix A](#). Currently this feature applies only to the ICE-STUN-UDP NAT traversal mode, but any other mode using connectivity checks SHOULD utilize this feature.

4.5. Base Exchange via HIP Relay Server

This section describes how Initiator and Responder perform a base exchange through a HIP relay server. The NAT traversal mode negotiation (denoted as NAT_TM in the example) was described in [Section 4.3](#) and is not repeated here. If a relay receives an R1 or I2 packet without the NAT traversal mode parameter, it MUST drop it and SHOULD send a NOTIFY error packet with type NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER to the sender of the R1/I2.

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 such an I1 packet with a UDP-encapsulated R1 packet and the rest of the base exchange, I2 and R2, MUST also use UDP encapsulation.

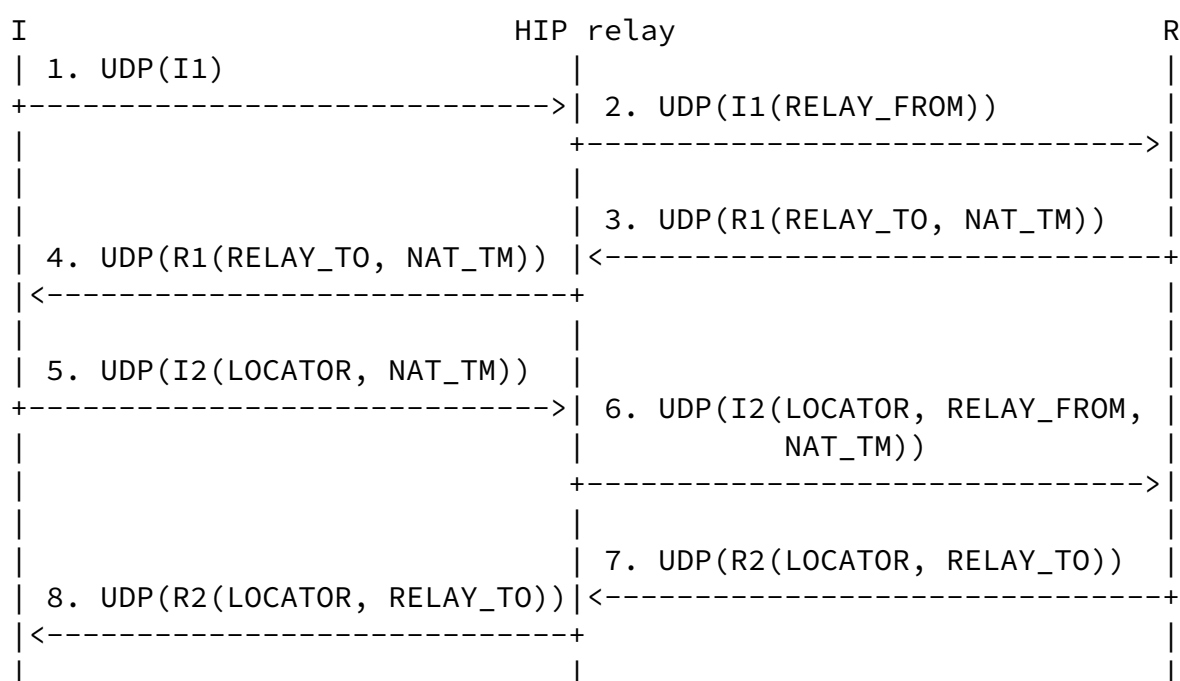


Figure 4: Base Exchange via a HIP Relay Server

In step 1 of Figure 4, the Initiator sends an I1 packet over the transport layer to the HIT of the Responder and IP address and port of the HIP relay server. The source address is one of the locators of the Initiator.

In step 2, the HIP relay server receives the I1 packet. 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 [RFC5204], except that the parameter type is different (see [Section 5.8](#)). 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 [RFC5201]. In addition, the Responder validates the RELAY_HMAC according to [RFC5204] and silently drops the packet if the validation fails. The Responder replies with an R1 packet to which it includes RELAY_TO and NAT traversal mode parameters. The RELAY_TO parameter **MUST** contain the same information as the RELAY_FROM parameter, i.e., the Initiator's 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 the source address and port, and changes the destination address and port to match RELAY_TO information. Finally, the relay recalculates transport checksum and forwards the packet.

In step 5, the Initiator receives the R1 packet and processes it according to [RFC5201]. The Initiator MAY use the address in the RELAY_TO parameter as a local peer-reflexive candidate for this HIP association if it is different from all known local candidates. The Initiator replies with an I2 packet that uses the destination transport address of R1 as the source address and port. The I2 packet contains a LOCATOR parameter that lists all the ICE candidates (ICE offer) of the Initiator. The candidates are encoded using the format defined in [Section 5.7](#). The I2 packet MUST also contain a NAT traversal mode parameter with the mode the Initiator selected.

In step 6, the relay receives the I2 packet. The relay appends a RELAY_FROM and a RELAY_HMAC to the I2 packet as explained in step 2.

In step 7, the Responder receives the I2 packet and processes it according to [RFC5201]. It replies with an R2 packet and includes a RELAY_TO parameter as explained in step 3. The R2 packet includes a LOCATOR parameter that lists all the ICE candidates (ICE 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 4. The relay forwards the packet to the Initiator. After the Initiator has received the R2 and processed it successfully, the base exchange is completed.

Hosts MUST include the address of one or more HIP relay servers (including the one that is being used for the initial signaling) in the LOCATOR parameter in I2/R2 if they intend to use such servers for relaying HIP signaling immediately after the base exchange completes. The traffic type of these addresses MUST be "HIP signaling" and they MUST NOT be used as ICE candidates. If the HIP relay server locator used for the base exchange is not included in I2/R2 LOCATOR parameters, it SHOULD NOT be used after the base exchange, but further HIP signaling SHOULD use the same path as the data traffic.

[4.6.](#) ICE Connectivity Checks

If a HIP relay server was used, the Responder completes the base exchange with the R2 packet through the relay. However, the

destination address the Initiator and Responder used for the base exchange packets belongs to the HIP relay server. Therefore, that address MUST NOT be used as a destination for ESP traffic. Instead, if a NAT traversal mode with ICE connectivity checks was selected, the Initiator and Responder MUST start the connectivity checks.

Creating the check list for the ICE connectivity checks should be performed as described in Section 5.7 of [[I-D.ietf-mmusic-ice](#)] bearing in mind that only one media stream and component is needed (so there will be only a single checklist and all candidates should have the same component ID value). The actual connectivity checks MUST be performed as described in Section 7 of [[I-D.ietf-mmusic-ice](#)]. Regular mode SHOULD be used for the candidate nomination. [Section 5.2](#) defines the details of the STUN control packets. As a result of the ICE connectivity checks, ICE nominates a single transport address pair to be used if an operational address pair was found. The end-hosts MUST use this address pair for the ESP traffic.

The connectivity check messages MUST be paced by the value negotiated during the base exchange as described in [Section 4.4](#). If neither one of the hosts announced a minimum pacing value, value of 500ms MUST be used.

For retransmissions, the RT0 value should be calculated as follows:

$$RT0 = \text{MAX} (500\text{ms}, T_a * P)$$

In the RT0 formula, T_a is the value used for the connectivity check pacing and P is the number of pairs in the checklist when the connectivity checks begin. This is identical to the formula in [[I-D.ietf-mmusic-ice](#)] if there is only one checklist.

If the ICE connectivity checks failed, the hosts MUST NOT send ESP traffic to each other but MAY continue communicating using HIP packets and the locators used for the base exchange. Also, the hosts SHOULD notify each other about the failure with a CONNECTIVITY_CHECKS_FAILED NOTIFY packet (see [Section 5.10](#)).

[4.7](#). NAT Keepalives

To prevent NAT states from expiring, communicating hosts send periodic keepalives to each other. HIP relay servers MAY refrain from sending keepalives if it's known that they are not behind a middlebox that requires keepalives. An end-host MUST send keepalives

every 15 seconds to refresh the UDP port mapping at the NAT(s) when the control or data channel is idle. To implement failure tolerance, an end-host SHOULD have shorter keepalive period.

The keepalives are STUN Binding Indications if the hosts have agreed on ICE-STUN-UDP NAT traversal mode during the base exchange. Otherwise, HIP NOTIFY packets MAY be used as keepalives.

The communicating hosts MUST send keepalives to each other using the transport locators they agreed to use for data and signaling when they are in ESTABLISHED state. Also, the Initiator MUST send a NOTIFY packet to the relay to keep the NAT states alive on the path between the Initiator and relay when the Initiator has not received any response to its I1 or I2 from the Responder in 15 seconds.

[4.8.](#) Base Exchange without ICE Connectivity Checks

In certain network environments the ICE connectivity checks can be omitted to reduce initial connection set up latency because a base exchange acts as an implicit connectivity test itself. For this to work, the Initiator MUST be able to reach the Responder by simply UDP encapsulating HIP and ESP packets sent to the Responder's address. Detecting and configuring this particular scenario is prone to failure unless carefully planned.

In such a scenario, the Responder MAY include UDP-ENCAPSULATION NAT traversal mode as one of the supported modes in the R1 packet. If the Responder has registered to a HIP relay server, it MUST also include a LOCATOR parameter in R1 that contains a preferred address where the Responder is able to receive UDP-encapsulated ESP and HIP packets. This locator MUST be of type "Transport address", its Traffic type MUST be "both" and it MUST have the "Preferred bit" set (see Table 2). If there is no such locator in R1, the source address of R1 is used as the Responder's preferred address.

The Initiator MAY choose the UDP-ENCAPSULATION mode if the Responder listed it in the supported modes and the Initiator does not wish to use ICE for searching for a more optimal path. In this case, the Initiator sends the I2 with UDP-ENCAPSULATION mode in the NAT traversal mode parameter directly to the Responder's preferred address (i.e., to the preferred locator in R1 or to the address where R1 was received from if there was no preferred locator in R1). The Initiator MAY include locators in I2 but they MUST NOT be taken as ICE candidates, since ICE will not be used for connections with UDP-ENCAPSULATION NAT traversal mode. Instead, if R2 and I2 are received and processed successfully, a security association can be created and UDP-encapsulated ESP can be exchanged between the hosts after the

base exchange completes. However, the Responder SHOULD NOT send any

ESP to the Initiator's address before it has received data from the Initiator, as specified in Sections 4.4.2. and 6.9 of [RFC5201] and in Sections 3.2.9 and 5.4 of [RFC5206].

Since an I2 packet with UDP-ENCAPSULATION NAT traversal mode selected MUST NOT be sent via a relay, the Responder SHOULD reject such I2 packets and reply with NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER NOTIFY packet (see Section 5.10).

If there is no answer for the I2 packet sent directly to the Responder's preferred address, the Initiator MAY send another I2 via the HIP relay server, but it MUST NOT choose UDP-ENCAPSULATION NAT traversal mode for that I2.

[4.9.](#) Initiating a Base Exchange both with and without UDP Encapsulation

The Initiator MAY also try to simultaneously perform a base exchange with the Responder without UDP encapsulation. In such a case, the Initiator sends two I1 packets, one without and one with UDP encapsulation, to the Responder. The Initiator MAY wait for a while before sending the other I1. How long to wait and in which order to send the I1 packets can be decided based on local policy. For retransmissions, the procedure is repeated.

The I1 packet without UDP encapsulation may arrive directly, without any relays, at the Responder. When this happens, the procedures in [RFC5201] are followed for the rest of the base exchange. The Initiator may receive multiple R1 packets, with and without UDP encapsulation, from the Responder. However, after receiving a valid R1 and answering it with an I2, further R1 packets that are not retransmits of the original R1 MUST be ignored.

The I1 packet without UDP encapsulation may also arrive at a HIP-capable middlebox. When the middlebox is a HIP rendezvous server and the Responder has successfully registered with the rendezvous service, the middlebox follows rendezvous procedures in [RFC5204].

If the Initiator receives a NAT traversal mode parameter in R1 without UDP encapsulation, the Initiator MAY ignore this parameter and send an I2 without UDP encapsulation and without any selected NAT

traversal mode. When the Responder receives the I2 without UDP encapsulation and without NAT traversal mode, it will assume that no NAT traversal mechanism is needed. The packet processing will be done as described in [\[RFC5201\]](#). The Initiator MAY store the NAT traversal modes for future use e.g., in case of a mobility or multihoming event which causes NAT traversal to be used during the lifetime of the HIP association.

[4.10.](#) Sending Control Packets after the Base Exchange

After the base exchange, the end-hosts MAY send HIP control packets directly to each other using the transport address pair established for a data channel without sending the control packets through the HIP relay server. When a host does not get acknowledgments, e.g., to an UPDATE or CLOSE packet after a timeout based on local policies, the host SHOULD resend the packet through the relay, if it was listed in the LOCATOR parameter in the base exchange.

If control packets are sent through a HIP relay server, the host registered with the relay MUST utilize the RELAY_TO parameter as in the base exchange. The HIP relay server SHOULD forward HIP packets to the registered hosts and forward packets from a registered host to the address in the RELAY_TO parameter. The relay MUST add a RELAY_FROM parameter to the control packets it relays to the registered hosts.

If the HIP relay server is not willing or able to relay a HIP packet, it MAY notify the sender of the packet with MESSAGE_NOT_RELAYED error notification (see [Section 5.10](#)).

[5.](#) Packet Formats

The following subsections define the parameter and packet encodings for the HIP, ESP and ICE connectivity check packets. All values MUST be in network byte order.

[5.1.](#) HIP Control Packets

0		1		2		3																									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1

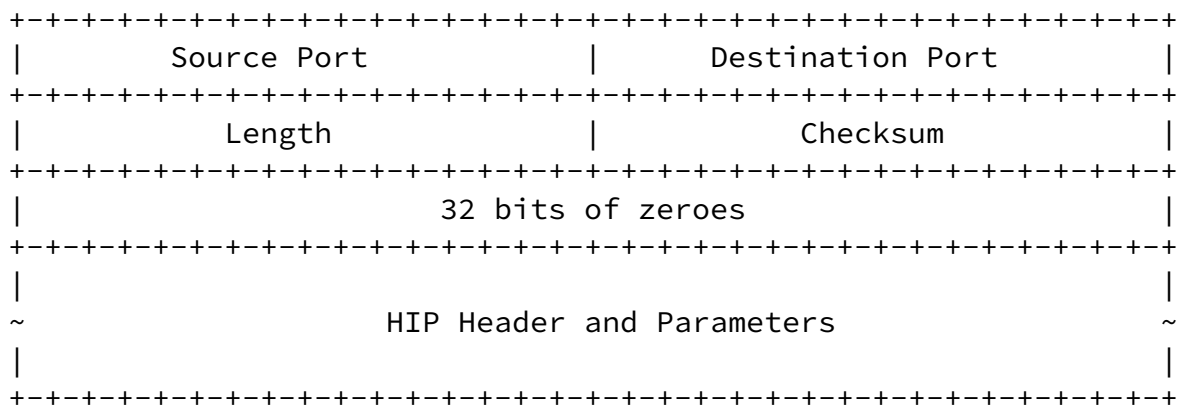


Figure 5: Format of UDP-encapsulated HIP Control Packets

HIP control packets are encapsulated in UDP packets as defined in

[Section 2.2 of \[RFC3948\]](#), "rules for encapsulating IKE messages", except a different port number is used. Figure 5 illustrates the encapsulation. The 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 [\[RFC5201\]](#) with the exception that the HIP header checksum MUST be zero. The HIP header checksum is zero for two reasons. First, the UDP header already contains a checksum. Second, the checksum definition in [\[RFC5201\]](#) 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 server or a Responder without a relay SHOULD listen at UDP port HIPPORT for incoming UDP-encapsulated HIP control packets.

5.2. Connectivity Checks

The connectivity checks are performed using STUN Binding Requests as defined in [\[I-D.ietf-mmusic-ice\]](#). This section describes the details of the parameters in the STUN messages.

The Binding Requests MUST use STUN short term credentials with last 32 bits of the HITs of the Initiator and Responder as the username fragments. The username is formed from the username fragments as defined in Section 7.1.1.3 of [\[I-D.ietf-mmusic-ice\]](#). The 32 bit username fragments are expressed using lowercase hexadecimal ASCII

characters. The leading zeroes MUST NOT be omitted so that the username's size is fixed (8 characters): for example, if the local HIT is 2001:15:8ebe:1aa7:42f5:b413:7237:6c0a and the remote HIT is 2001:18:46fa:97c0:ba5:cd77:51:47b, the local username would be 72376c0a and the remote username 0051047b.

The STUN password is drawn from the DH keying material. Drawing of HIP keys is defined in [\[RFC5201\] Section 6.5](#) and drawing of ESP keys in [\[RFC5202\] Section 7](#). Correspondingly, the hosts MUST draw symmetric keys for STUN according to [\[RFC5201\] Section 6.5](#). The hosts draw the STUN key after HIP keys, or after ESP keys if ESP transform was successfully negotiated in the base exchange. Both hosts draw a 128 bit key from the DH keying material, express that in hexadecimal ASCII format using only lowercase letters (resulting in 32 numbers or lowercase letters), and use that as both the local and peer password. [\[RFC5389\]](#) describes how hosts use the password for message integrity of STUN messages.

Both the username and password are expressed in ASCII hexadecimal format to prevent the need to run them through SASLPrep as defined in [\[RFC5389\]](#).

The connectivity checks MUST contain PRIORITY attribute. They MAY contain USE-CANDIDATE attribute as defined in Section 7.1.1.1 of [\[I-D.ietf-mmusic-ice\]](#).

The Initiator is always in the controlling role during a base exchange. When two hosts are initiating a connection to each other simultaneously, HIP state machine detects it and assigns the host with the larger HIT as the Responder as explained in Sections [4.4.2](#) and [6.7](#) in [\[RFC5201\]](#). Hence, the ICE-CONTROLLED and ICE-CONTROLLING attributes are not needed to resolve role conflicts. However, the attributes SHOULD be added to the connectivity check messages to ensure interoperability with different ICE stacks and they can be safely ignored on received connectivity checks.

[5.3](#). Keepalives


The keepalives for HIP associations that are created with ICE are STUN Binding Indications, as defined in [\[RFC5389\]](#). In contrast to the UDP-encapsulated HIP header, the non-ESP-marker between the UDP

header and the STUN header is excluded. Keepalives MUST contain the FINGERPRINT STUN attribute but SHOULD NOT contain any other STUN attributes and SHOULD NOT utilize any authentication mechanism. STUN messages are demultiplexed from ESP and HIP control packets using the STUN markers, such as the magic cookie value and the FINGERPRINT attribute.

Keepalives for HIP associations created without ICE are HIP control packets that have NOTIFY as the packet type. The keepalive NOTIFY packets do not contain any parameters.

5.4. NAT Traversal Mode Parameter

Format of the NAT_TRAVERSAL_MODE parameter is similar to the format of the ESP_TRANSFORM parameter in [RFC5202] and is shown in Figure 6. This specification defines traversal mode identifiers UDP-ENCAPSULATION and ICE-STUN-UDP. The identifier RESERVED is reserved for future use. Future specifications may define more traversal modes.



Type	[TBD by IANA: 608]
Length	length in octets, excluding Type, Length, and padding
Reserved	zero when sent, ignored when received
Mode ID	defines the proposed or selected NAT traversal mode(s)

The following NAT traversal mode IDs are defined:


ID name	Value
RESERVED	0
UDP-ENCAPSULATION	1
ICE-STUN-UDP	2

Figure 6: Format of the NAT_TRAVERSAL_MODE parameter

The sender of a NAT_TRAVERSAL_MODE parameter MUST make sure that there are no more than six (6) Mode IDs in one NAT_TRAVERSAL_MODE parameter. Conversely, a recipient MUST be prepared to handle received NAT traversal mode parameters that contain more than six Mode IDs by accepting the first six Mode IDs and dropping the rest. The limited number of Mode IDs sets the maximum size of the NAT_TRAVERSAL_MODE parameter. The modes MUST be in preference order, most preferred mode(s) first.

5.5. Connectivity Check Transaction Pacing Parameter

The TRANSACTION_PACING parameter shown in Figure 7 contains only the connectivity check pacing value, expressed in milliseconds, as 32 bit unsigned integer.



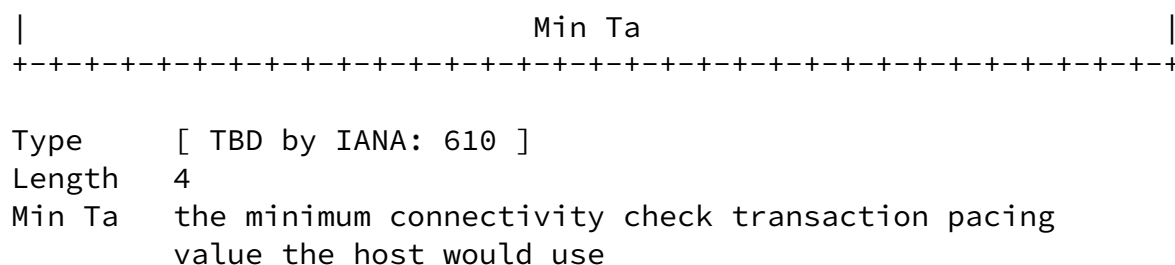


Figure 7: Format of the TRANSACTION_PACING parameter

5.6. Relay and Registration Parameters

Format of the REG_FROM, RELAY_FROM and RELAY_TO parameters is shown in Figure 8. All parameters are identical except for the type. REG_FROM is the only parameter covered with the signature.

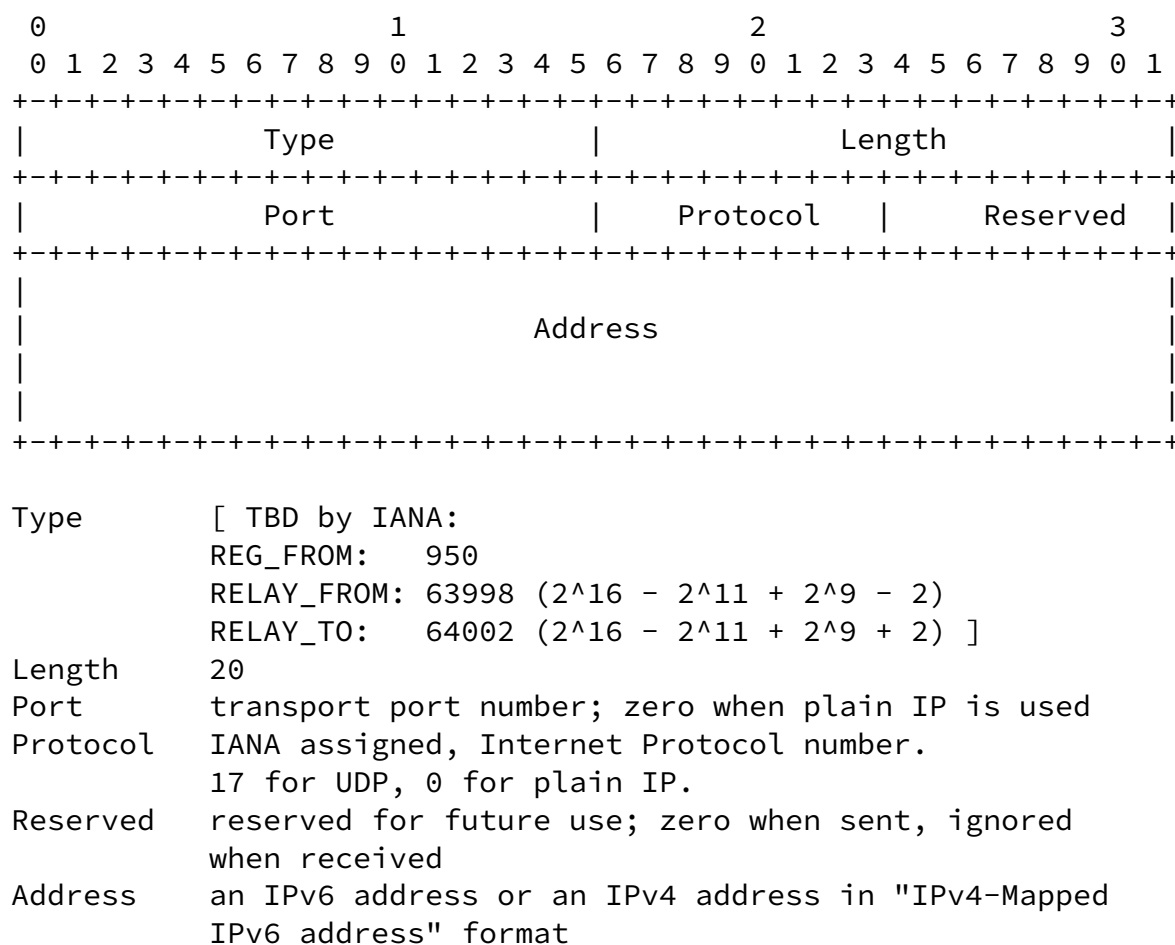


Figure 8: Format of the REG_FROM, RELAY_FROM and RELAY_TO parameters

REG_FROM contains the transport address and protocol where the HIP relay server sees the registration coming from. RELAY_FROM contains the address where the relayed packet was received from by the relay server and the protocol that was used. RELAY_TO contains same information about the address where a packet should be forwarded to.

5.7. LOCATOR Parameter

The generic LOCATOR parameter format is the same as in [\[RFC5206\]](#). However, presenting ICE candidates requires a new locator type. The generic and NAT traversal specific locator parameters are illustrated in Figure 9.

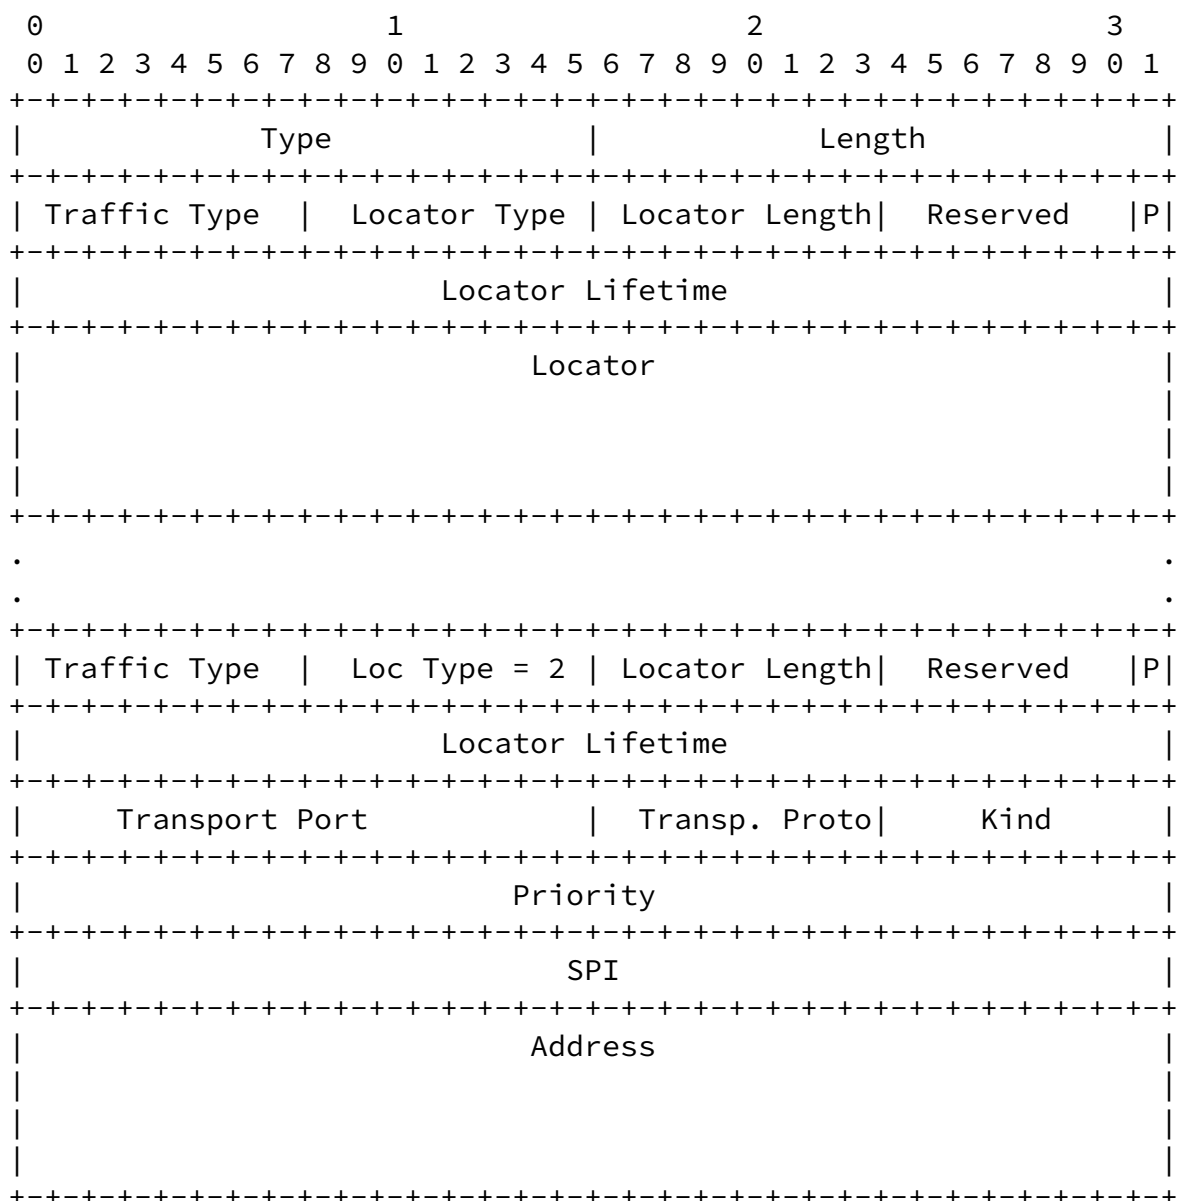


Figure 9: LOCATOR parameter

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 fields after Locator Lifetime in 4-octet units
Reserved	0	Reserved for future extensions
Preferred (P) bit	0 or 1	Set to 1 for a Locator in R1 if the Responder can use it for the rest of the base exchange, otherwise set to zero
Locator Lifetime	Variable	Locator lifetime in seconds
Transport Port	Variable	Transport layer port number
Transport Protocol	Variable	IANA Assigned, transport layer Internet Protocol number. Currently only UDP (17) is supported.
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
Address	Variable	IPv6 address or an "IPv4-Mapped IPv6 address" format IPv4 address [RFC4291]

Table 2: Fields of the LOCATOR parameter

[5.8.](#) RELAY_HMAC Parameter

The RELAY_HMAC parameter value has the TLV type 65520 ($2^{16} - 2^5 + 2^4$). It has the same semantics as RVS_HMAC [\[RFC5204\]](#).

[5.9.](#) Registration Types

The REG_INFO, REG_REQ, REG_RESP and REG_FAILED parameters contain Registration Type [[RFC5203](#)] values for HIP relay server registration. The value for RELAY_UDP_HIP is 2.

[5.10.](#) Notify Packet Types

A HIP relay server and end hosts can use NOTIFY packets to signal different error conditions. The new Notify Packet Types [[RFC5201](#)] defined in this document are shown below [values TBD by IANA]. The Notification Data field for the error notifications SHOULD contain the HIP header of the rejected packet and SHOULD be empty for the CONNECTIVITY_CHECKS_FAILED type.

NOTIFICATION PARAMETER - ERROR TYPES -----	Value -----
NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER	60

If a HIP relay server does not forward a base exchange packet due to missing NAT traversal mode parameter, or the Initiator selects a NAT traversal mode that the Responder did not expect, the relay or the Responder may send back a NOTIFY error packet with this type.

CONNECTIVITY_CHECKS_FAILED	61
----------------------------	----

Used by the end hosts to signal that NAT traversal connectivity checks failed and did not produce a working path.

MESSAGE_NOT_RELAYED	62
---------------------	----

Used by a HIP relay server to signal that it was not able or willing to relay a HIP packet.

[5.11.](#) ESP Data Packets

[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. However, the (semantic) difference to BEET mode ESP packets used by HIP is that IP header is not used in BEET integrity protection calculation.

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 [[RFC5202](#)]. 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 SPIs is

defined in [[RFC5202](#)]. The UDP encapsulation format and processing of HIP ESP traffic is described in [Section 6.1 of \[RFC5202\]](#).

[6.](#) Security Considerations

[6.1.](#) Privacy Considerations

The locators are in plain text format in favor 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 to avoid disclosing them to middleboxes.

It is also possible that end-users may not want to reveal all locators to each other. For example, tracking the physical location of a multihoming end-host may become easier if it reveals all locators to its peer during a base exchange. Also, revealing host addresses exposes information about the local topology which may not be allowed in all corporate environments. For these two reasons, an end-host may exclude certain host addresses from its LOCATOR parameter. However, such behavior creates non-optimal paths when the hosts are located behind the same NAT. Especially, this could be problematic 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 [[RFC5128](#)]. With such a legacy NAT, the only option left would be to use a relayed transport address from a TURN server.

The use of HIP relay servers and TURN relays can be also useful for privacy purposes. For example, a privacy concerned Responder may reveal only its HIP relay server and Relayed candidates to Initiators. This same mechanism also protects the Responder against Denial-of-Service attacks by allowing the Responder to initiate new connections even if its relays would be unavailable due to a DoS attack.

[6.2.](#) Opportunistic Mode

A HIP relay server should have one address per relay client when a HIP relay is serving more than one relay clients and supports opportunistic mode. Otherwise, it cannot be guaranteed that the HIP relay server can deliver the I1 packet to the intended recipient.

[6.3.](#) Base Exchange Replay Protection for HIP Relay Server

In certain scenarios, it is possible that an attacker, or two attackers, can replay an earlier base exchange through a HIP relay server by masquerading as the original Initiator and Responder. The

attack does not require the attacker(s) to compromise the private key(s) of the attacked host(s). However, for this attack to succeed, the Responder has to be disconnected from the HIP relay server.

The relay can protect itself against replay attacks by becoming involved in the base exchange by introducing nonces that the end-hosts (Initiator and Responder) are required to sign. One way to do this is to add ECHO_REQUEST_M parameters to the R1 and I2 packets as described in [[I-D.heer-hip-middle-auth](#)] and drop the I2 or R2 packets if the corresponding ECHO_RESPONSE_M parameters are not present.

[6.4.](#) Demuxing Different HIP Associations

[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 uses HITs to distinguish between different Initiators.

7. IANA Considerations

This section is to be interpreted according to [[RFC5226](#)].

[TO BE REMOVED: Upon publication of this document, IANA is requested to register a UDP port from the Registered Ports range 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 [[RFC5201](#)] by assigning new HIP Parameter Type values for the new HIP Parameters: RELAY_FROM, RELAY_TO and REG_FROM (defined in [Section 5.6](#)), RELAY_HMAC (defined in [Section 5.8](#)), TRANSACTION_PACING (defined in [Section 5.5](#)), and NAT_TRAVERSAL_MODE (defined in [Section 5.4](#)).

This document defines an additional registration type for the HIP Registration Extension [[RFC5203](#)] that allows registering with a HIP relay server for relaying service: RELAY_UDP_HIP (defined in [Section 5.9](#)).

This document also defines NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER,

CONNECTIVITY_CHECKS_FAILED and MESSAGE_NOT_RELAYED Notify Packet Types [[RFC5201](#)] in [Section 5.10](#).

The NAT_TRAVERSAL_MODE parameter has 16-bit unsigned integer fields for different modes, for which IANA is to create and maintain a new sub-registry entitled "HIP NAT traversal modes" under the "Host Identity Protocol (HIP) Parameters". Initial values for the NAT traversal mode registry are given in [Section 5.4](#); future assignments are to be made through IETF Review [[RFC5226](#)]. Assignments consist of a NAT traversal mode identifier name and its associated value. [TO BE REMOVED: This registration should take place at the following location: <http://www.iana.org/assignments/hip-parameters/>]

8. Contributors

This draft is a product of a design team which also included Marcelo Bagnulo and Philip Matthews who both have made major contributions to this document.

9. Acknowledgments

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 thank Andrei Gurtov, Simon Schuetz, Martin Stiemerling, Lars Eggert, Vivien Schmitt, Abhinav Pathak for their contributions and Tobias Heer, Teemu Koponen, Juhana Mattila, Jeffrey M. Ahrenholz, Kristian Slavov, Janne Lindqvist, Pekka Nikander, Lauri Silvennoinen, Jukka Ylitalo, Juha Heinanen, Joakim Koskela, Samu Varjonen, Dan Wing and Jani Hautakorpi 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, Ericsson, and Birdstep.

10. References

10.1. Normative References

[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-16](#) (work in progress), July 2009.

[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.

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

Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", [RFC 4423](#), May 2006.
- [RFC5201] Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol", [RFC 5201](#), April 2008.
- [RFC5202] Jokela, P., Moskowitz, R., and P. Nikander, "Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP)", [RFC 5202](#), April 2008.
- [RFC5203] Laganier, J., Koponen, T., and L. Eggert, "Host Identity Protocol (HIP) Registration Extension", [RFC 5203](#), April 2008.
- [RFC5204] Laganier, J. and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extension", [RFC 5204](#), April 2008.
- [RFC5206] Nikander, P., Henderson, T., Vogt, C., and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol", [RFC 5206](#), April 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", [RFC 5389](#), October 2008.

[10.2.](#) Informative References

- [I-D.heer-hip-middle-auth]
Heer, T., Wehrle, K., and M. Komu, "End-Host Authentication for HIP Middleboxes",
[draft-heer-hip-middle-auth-02](#) (work in progress),
February 2009.

- [I-D.rosenberg-mmusic-ice-nonsip]
Rosenberg, J., "Guidelines for Usage of Interactive Connectivity Establishment (ICE) by non Session Initiation Protocol (SIP) Protocols", [draft-rosenberg-mmusic-ice-nonsip-01](#) (work in progress), July 2008.
- [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.
- [RFC5128] Srisuresh, P., Ford, B., and D. Kegel, "State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)", [RFC 5128](#), March 2008.
- [RFC5207] Stiernerling, M., Quittek, J., and L. Eggert, "NAT and Firewall Traversal Issues of Host Identity Protocol (HIP) Communication", [RFC 5207](#), April 2008.

[Appendix A](#). Selecting a Value for Check Pacing

Selecting a suitable value for the connectivity check transaction pacing is essential for the performance of connectivity check-based NAT traversal. The value should not be so small that the checks cause network congestion or overwhelm the NATs. On the other hand, a pacing value that is too high makes the checks last for a long time, thus increasing the connection setup delay.

The T_a value may be configured by the user in environments where the network characteristics are known beforehand. However, if the characteristics are not known, it is recommended that the value is adjusted dynamically. In this case it's recommended that the hosts estimate the RTT between them and set the minimum T_a value so that only two connectivity check messages are sent on every RTT.

One way to estimate the RTT is to use the time it takes for the HIP relay server registration exchange to complete; this would give an estimate on the registering host's access link's RTT. Also the I1/R1 exchange could be used for estimating the RTT, but since the R1 can be cached in the network, or the relaying service can increase the delay notably, it is not recommended.

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

When the Initiator looks up the information of the Responder from DNS, it's possible that it discovers an RVS record [[RFC5204](#)]. In this case, if the Initiator uses NAT traversal methods described in this document, it MAY use its own HIP relay server to forward HIP traffic to the Rendezvous server. The Initiator will send the I1 packet using its HIP relay server which will then forward it to the RVS server of the Responder. In this case, the value of the protocol field in the RELAY_TO parameter MUST be IP since RVS does not support UDP-encapsulated base exchange packets. 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 Responder's address records (e.g., if the DNS server did not support HIP extensions), the Initiator SHOULD first try to contact the Responder directly, without using a HIP relay server. If none of the addresses is reachable, it MAY try out them using its own HIP relay server as described above.

[Appendix C](#). Document Revision History

To be removed upon publication

+-----+-----+	
Revision	Comments
+-----+-----+	
-00	Initial version.
-01	Draft based on RVS.
-02	Draft based on Relay proxies and ICE concepts.
-03	Draft based on STUN/ICE formats.
-04	Issues 25-27,29-36
-05	Issues 28,40-43,47,49,51
-06	New copyright boilerplate and STUN username encoding
-07	New NOTIFY error packet parameters, changed handling
	of I2/R2 via relay with UDP-ENCAPSULATION mode
-08	Small editorial fixes regarding WGLC comments
-09	Small fixes regarding IESG and Gen-ART review comments
+-----+-----+	

Internet-Draft

Basic NAT Traversal for HIP

October 2009

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/>

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

Email: thomas.r.henderson@boeing.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

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

Jan Melen
Ericsson Research Nomadiclab

Hirsalantie 11
02420 Jorvas
Finland

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

Komu, et al.

Expires April 26, 2010

[Page 32]

Internet-Draft

Basic NAT Traversal for HIP

October 2009

Ari Keranen (editor)
Ericsson Research Nomadiclab
Hirsalantie 11
02420 Jorvas
Finland

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

