Internet Engineering Task Force (IETF)   M. Komu
Request for Comments: 5770   HIIT
Category: Experimental   T. Henderson
ISSN: 2070-1721   The Boeing Company
                                                      H. Tschofenig
                                              Nokia Siemens Networks
                                                          J. Melen
                                                  A. Keranen, Ed.
                                         Ericsson Research Nomadiclab
                                                        April 2010

                Basic Host Identity Protocol (HIP) Extensions for
                   Traversal of Network Address Translators

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.

Status of This Memo

   This document is not an Internet Standards Track specification; it is
   published for examination, experimental implementation, and
   evaluation.

   This document defines an Experimental Protocol for the Internet
   community. This document is a product of the Internet Engineering
   Task Force (IETF). It represents the consensus of the IETF
   community. It has received public review and has been approved for
   publication by the Internet Engineering Steering Group (IESG). Not
   all documents approved by the IESG are a candidate for any level of
   Internet Standard; see Section 2 of RFC 5741.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   http://www.rfc-editor.org/info/rfc5770.

Table of Contents

Komu, et al.   Experimental   [Page 3]

---

[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 [[RFC5245](#)] 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 Traversal Using Relay NAT (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 that 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 [RFC5245]. [RFC5245] describes ICE as follows:

      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 [Simple
      Traversal of the UDP Protocol through NAT (STUN)] specification
      [RFC5389], now renamed to Session Traversal Utilities for NAT.

The standard ICE [RFC5245] is specified with SIP in mind and it has some features that are not necessary or suitable as such for other protocols. [MMUSIC-ICE] 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], [RFC5245], 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 [RFC5766].

   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

```
                           +-------+
                           | HIP   |
           +--------+      | Relay |       +--------+
           | TURN   |      +-------+       | STUN   |
           | Server |      /       \       | Server |
           +--------+     /         \      +--------+
                         /           \
                        /             \
                       /               \
                      /  <- Signaling ->  \
                     /                   \
                +-------+           +-------+
                | NAT   |           | NAT   |
                +-------+           +-------+
                   /                     \
                  /                       \
             +-------+                 +-------+
             | Init- |                 | Resp- |
             | iator |                 | onder |
             +-------+                 +-------+
```
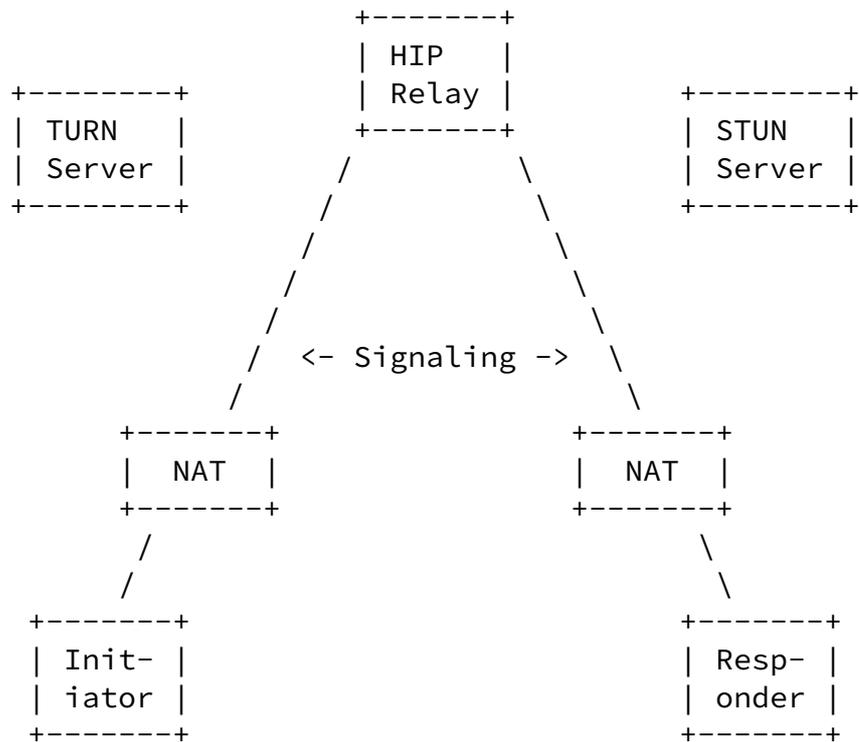
                Figure 1: Example Network Configuration

   In the 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 Host Identity Tag (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 [[RFC5766](#)].

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.

```
   HIP                                                          HIP
   Relay                                                        Relay
   Client                                                       Server
     |    1. UDP(I1)                                            |
     +-------------------------------------------------------->|
     |                                                          |
     |    2. UDP(R1(REG_INFO(RELAY_UDP_HIP)))                   |
     |<-------------------------------------------------------+
     |                                                          |
     |    3. UDP(I2(REG_REQ(RELAY_UDP_HIP)))                    |
     +-------------------------------------------------------->|
     |                                                          |
     |    4. UDP(R2(REG_RES(RELAY_UDP_HIP), REG_FROM))          |
     |<-------------------------------------------------------+
     |                                                          |
```
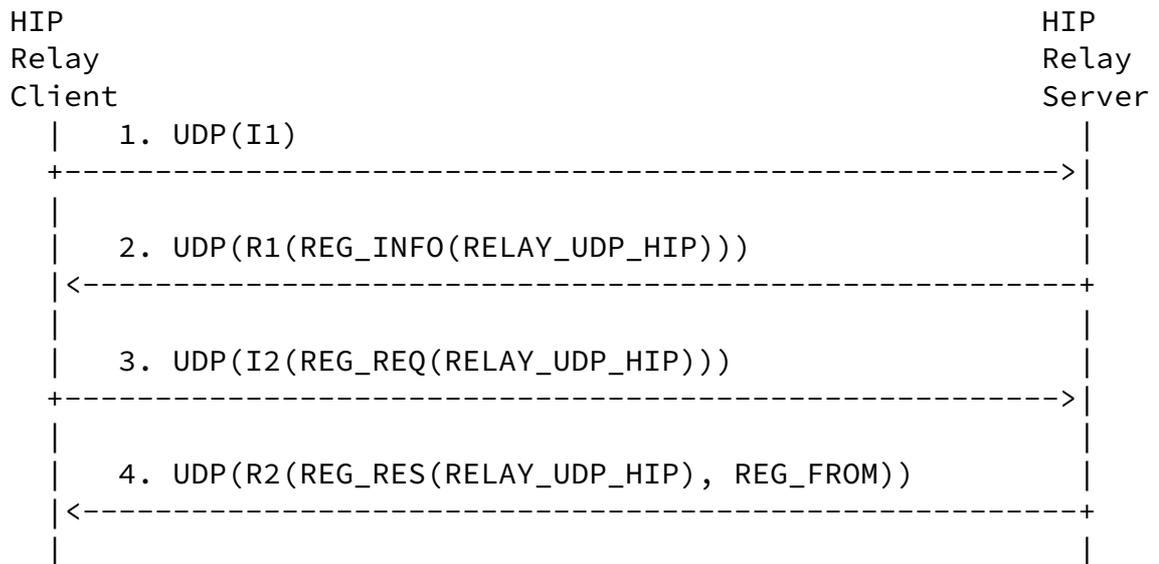
            Figure 2: Example Registration with a HIP Relay

In step 1, the relay client (Initiator) starts the registration
procedure by sending an I1 packet over UDP. It is RECOMMENDED that
the Initiator select 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 listen to incoming connections
at UDP port 10500. If some other port number is used, it needs to be
known by potential 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 for which it registers
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, as described in
[Section 4.7](#), 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 [RFC5245]](#). 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 the UDP-ENCAPSULATION mode of NAT traversal since the relay is assumed to be in public address space. 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.

```
   Initiator                                                    Responder
   | 1. UDP(I1)                                                        |
   +----------------------------------------------------------------->|
   |                                                                   |
   | 2. UDP(R1(.., NAT_TRAVERSAL_MODE(list of modes), ..))             |
   |<----------------------------------------------------------------+
   |                                                                   |
   | 3. UDP(I2(.., NAT_TRAVERSAL_MODE(selected mode), LOCATOR, ..))   |
   +----------------------------------------------------------------->|
   |                                                                   |
   | 4. UDP(R2(.., LOCATOR, ..))                                       |
   |<----------------------------------------------------------------+
   |                                                                   |
```
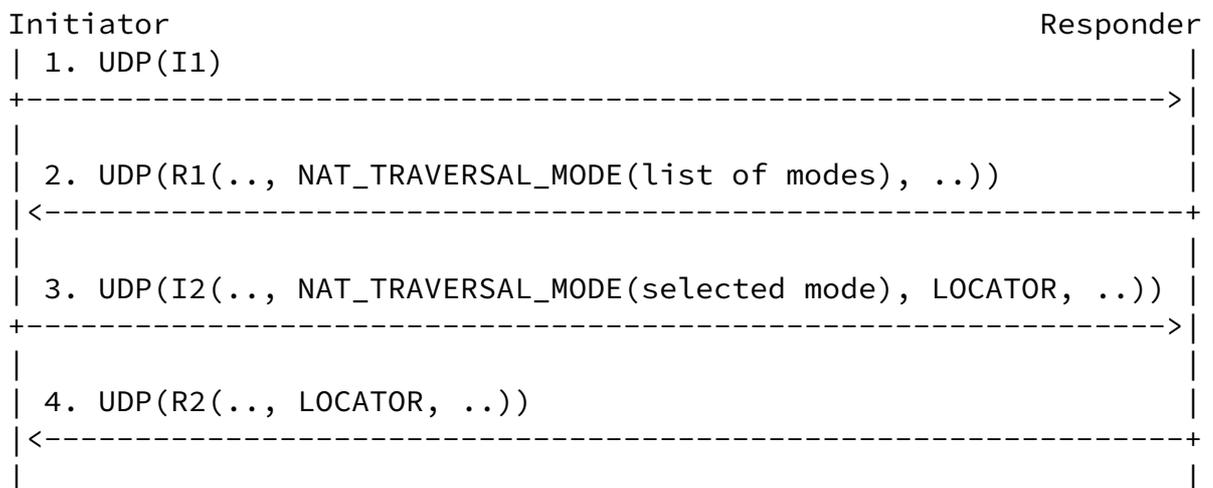
             Figure 3: Negotiation of NAT Traversal Mode

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 are specified 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 |
|                   | messages                                 |
+-------------------+------------------------------------------+
```

                   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 a NOTIFY packet with type
   NO_VALID_NAT_TRAVERSAL_MODE_PARAMETER and abort the base exchange.

4.4. Connectivity Check Pacing Negotiation

   As explained in [[RFC5245](#)], 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 500 ms 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 20 ms for the minimum Ta,

since such values may not work well with some NATs, as explained in
[RFC5245]. 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, a value of 500 ms 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 the 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 send 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.

```
   Initiator                     HIP relay                   Responder
   | 1. UDP(I1)                    |                            |
   +----------------------------->| 2. UDP(I1(RELAY_FROM))      |
   |                               +---------------------------->|
   |                               |                            |
   |                               | 3. UDP(R1(RELAY_TO, NAT_TM))   |
   | 4. UDP(R1(RELAY_TO, NAT_TM)) |<----------------------------+
   |<---------------------------+  |                            |
   |                               |                            |
   | 5. UDP(I2(LOCATOR, NAT_TM))   |                            |
   +----------------------------->| 6. UDP(I2(LOCATOR, RELAY_FROM, |
   |                               |             NAT_TM))        |
```

```
|                                +---------------------------->|
|                                |                             |
|                                | 7. UDP(R2(LOCATOR, RELAY_TO))  |
| 8. UDP(R2(LOCATOR, RELAY_TO))|<----------------------------+
|<---------------------------+                               |
|                                |                             |
```
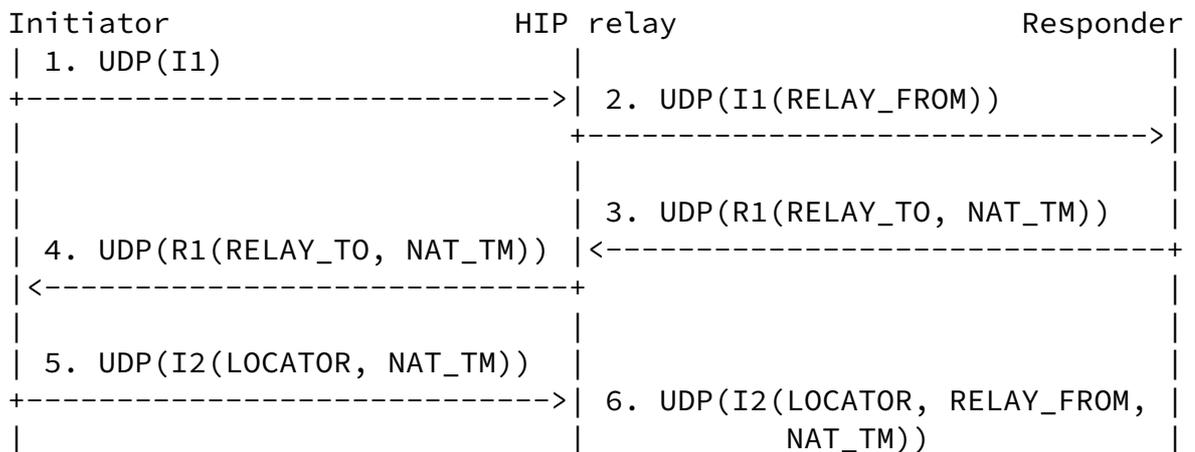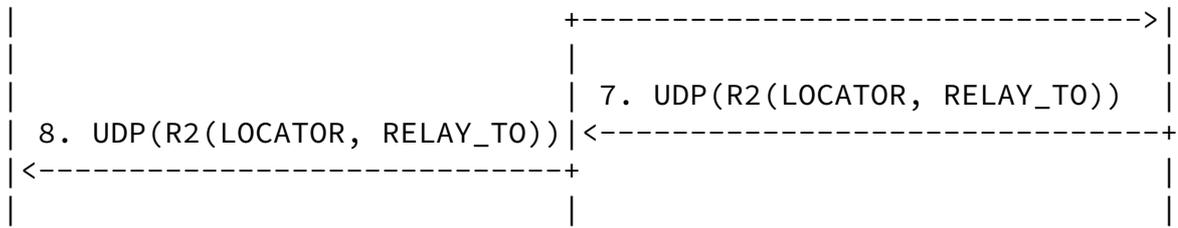
                Figure 4: Base Exchange via a HIP Relay Server

   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 checklist for the ICE connectivity checks should be
   performed as described in Section 5.7 of [RFC5245] 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 [RFC5245]. 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, a value of 500 ms MUST
   be used.

   For retransmissions, the retransmission timeout (RTO) value SHOULD be
   calculated as follows:

      RTO = MAX (500ms, Ta * (Num-Waiting + Num-In-Progress))

   In the RTO formula, Ta is the value used for the connectivity check
   pacing, Num-Waiting is the number of pairs in the checklist in the
   "Waiting" state, and Num-In-Progress is the number of pairs in the
   "In-Progress" state. This is identical to the formula in [RFC5245]
   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 a 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 the 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 a 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 that 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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |          Source Port          |       Destination Port        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Length            |            Checksum           |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                      32 bits of zeroes                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    ~                  HIP Header and Parameters                    ~
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
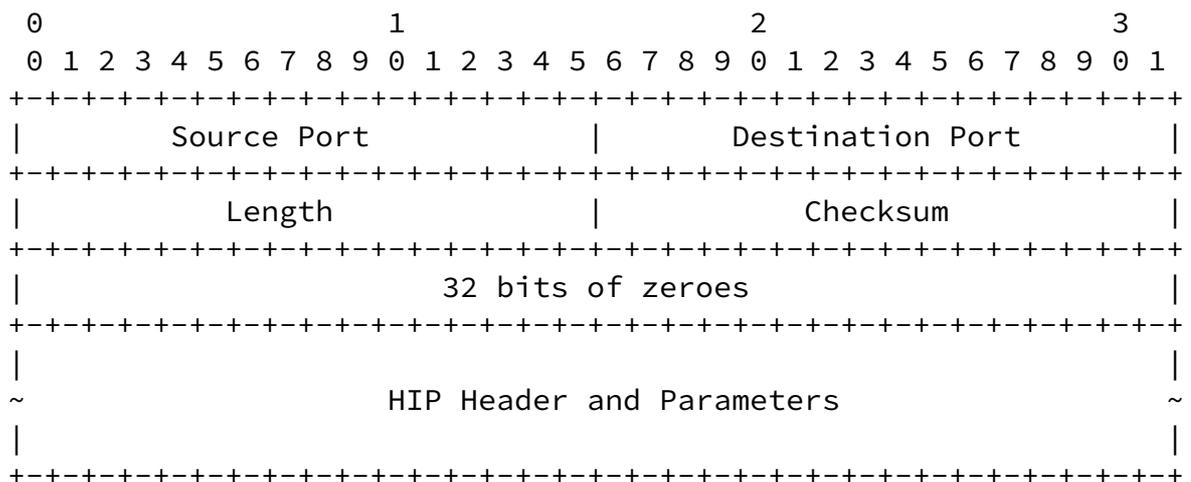
        Figure 5: Format of UDP-Encapsulated HIP Control Packets

HIP control packets are encapsulated in UDP packets as defined in
Section 2.2 of [RFC3948], "IKE Header Format for Port 4500", 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 10500 for incoming UDP-encapsulated HIP control packets. If
some other port number is used, it needs to be known by potential
Initiators.

5.2. Connectivity Checks

The connectivity checks are performed using STUN Binding requests as
defined in [RFC5245]. This section describes the details of the
parameters in the STUN messages.

The Binding requests MUST use STUN short-term credentials with the
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 [RFC5245]. 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 Diffie-Hellman (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 the PRIORITY attribute. They
   MAY contain the USE-CANDIDATE attribute as defined in Section 7.1.1.1
   of [RFC5245].

   The Initiator is always in the controlling role during a base
   exchange. When two hosts are initiating a connection to each other
   simultaneously, the 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

   The 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.
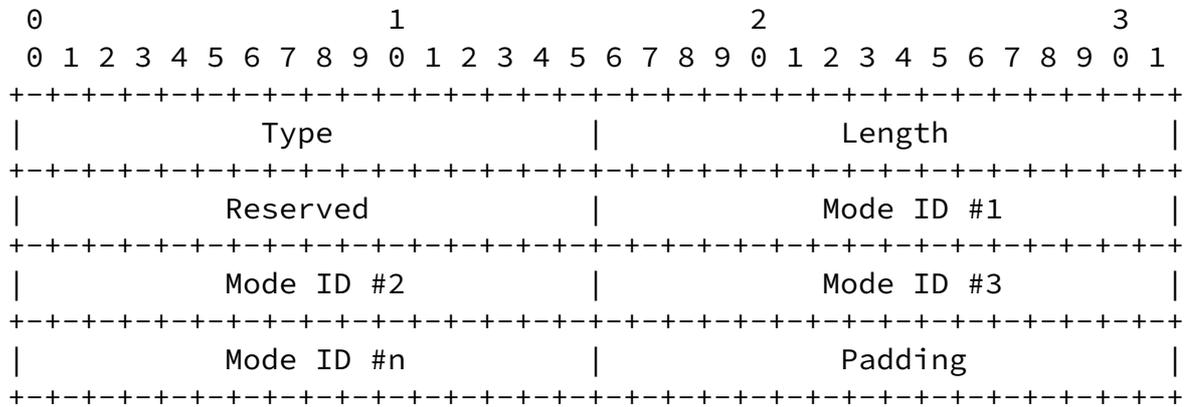
```
         0                   1                   2                   3
         0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        |              Type             |             Length            |
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        |            Reserved           |           Mode ID #1          |
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        |           Mode ID #2          |           Mode ID #3          |
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
        |           Mode ID #n          |            Padding            |
        +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

        Type      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 a 32-
   bit unsigned integer.
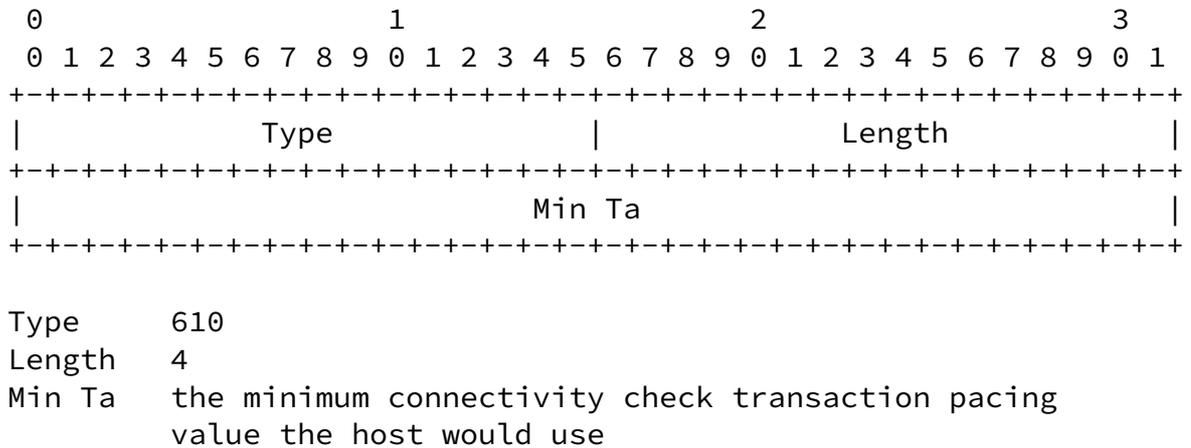
```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Type              |            Length             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                            Min Ta                             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

    Type     610
    Length   4
    Min Ta   the minimum connectivity check transaction pacing
             value the host would use

          Figure 7: Format of the TRANSACTION_PACING Parameter
```

5.6. Relay and Registration Parameters

   The 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.

```
     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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Type              |            Length             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Port              |    Protocol   |    Reserved   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |                            Address                            |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

    Type       REG_FROM:   950
               RELAY_FROM: 63998
               RELAY_TO:   64002
    Length     20
    Port       transport port number; zero when plain IP is used
    Protocol   IANA assigned, Internet Protocol number.
               17 for UDP, 0 for plain IP
```
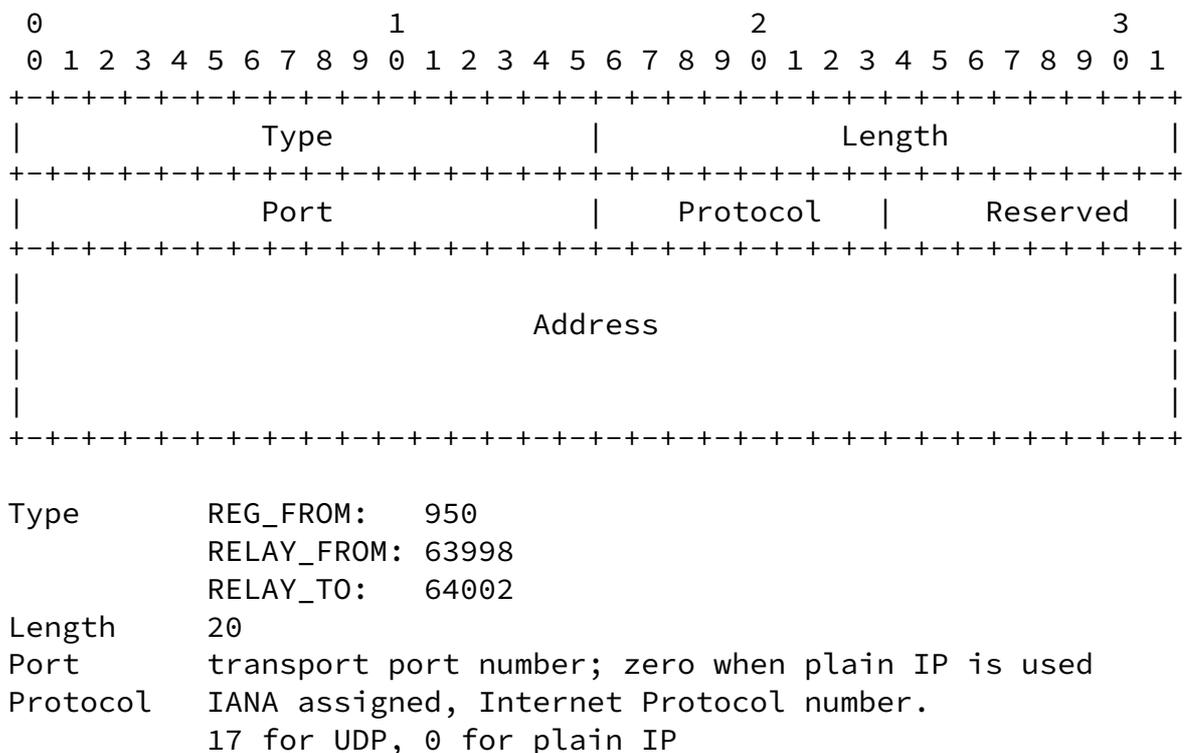
```
      Reserved    reserved for future use; zero when sent, ignored
                  when received
      Address     an IPv6 address or an IPv4 address in "IPv4-Mapped
                  IPv6 address" format
```

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

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

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
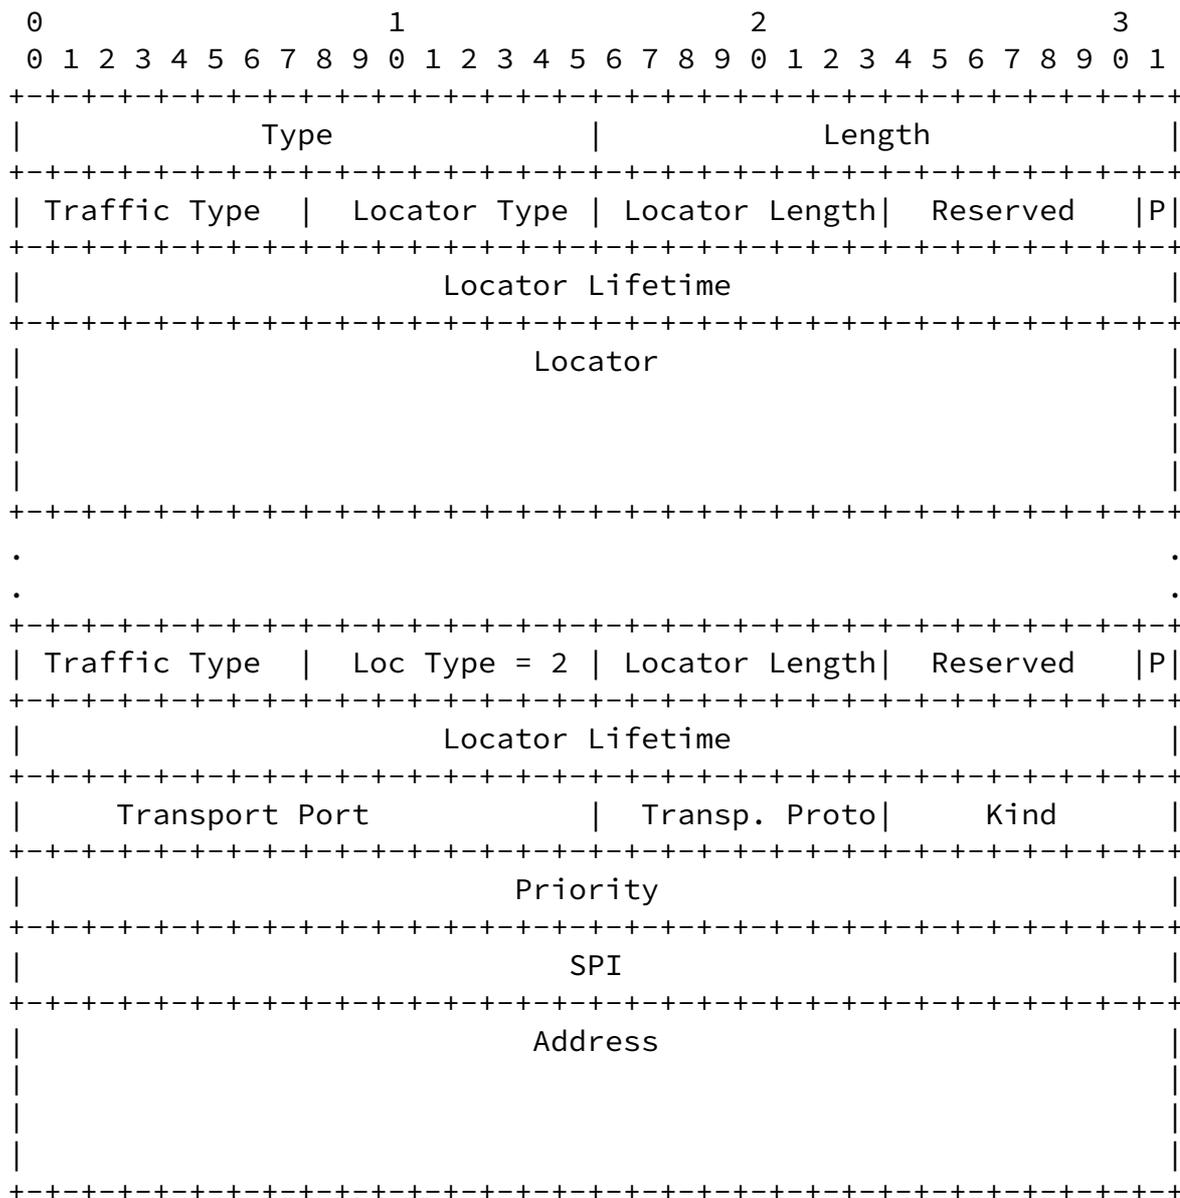   in Figure 9.

```
     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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Type              |            Length             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    | Traffic Type  |  Locator Type | Locator Length|  Reserved   |P|
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                       Locator Lifetime                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                            Locator                            |
    |                                                               |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    .                                                               .
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    | Traffic Type  | Loc Type = 2  | Locator Length|  Reserved   |P|
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                       Locator Lifetime                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |      Transport Port           | Transp. Proto|     Kind       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                            Priority                           |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                              SPI                              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                            Address                            |
    |                                                               |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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  | 0-2      | Is the locator for HIP signaling (1), for|
| Type     |          | ESP (2), or for both (0)                 |
| Locator  | 2        | "Transport address" locator type         |
| Type     |          |                                          |
| Locator  | 7        | Length of the fields after Locator       |
| Length   |          | Lifetime in 4-octet units                |
| Reserved | 0        | Reserved for future extensions           |
| Preferred| 0 or 1   | Set to 1 for a Locator in R1 if the      |
| (P) bit  |          | Responder can use it for the rest of the |
|          |          | base exchange, otherwise set to zero     |
| Locator  | Variable | Locator lifetime in seconds              |
| Lifetime |          |                                          |
| Transport| Variable | Transport layer port number              |
| Port     |          |                                          |
| Transport| Variable | IANA assigned, transport layer Internet  |
| Protocol |          | 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       |
|          |          | [RFC5245]                                |
| SPI      | Variable | Security Parameter Index (SPI) value that|
|          |          | 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. 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. 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 is was not able or
        willing to relay a HIP packet.

5.11. ESP Data Packets

    [RFC3948] describes the 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 Bound End-to-End Tunnel (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 document 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 that 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 (DoS) 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 client 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 [HIP-MIDDLE] 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].

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 has created and maintains 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.

8. Contributors

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

9. Acknowledgments

   Thanks to 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, and 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 has been working in the Networking Research group at
   Helsinki Institute for Information Technology (HIIT). The work has
   been funded by Tekes, Telia-Sonera, Elisa, Nokia, the Finnish Defence
   Forces, Ericsson and Birdstep in InfraHIP I and II projects.

10. References

10.1. Normative References

   [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.

   [RFC5245]   Rosenberg, J., "Interactive Connectivity Establishment
                (ICE): A Protocol for Network Address Translator (NAT)
                Traversal for Offer/Answer Protocols", RFC 5245,
                April 2010.

   [RFC5389]   Rosenberg, J., Mahy, R., Matthews, P., and D. Wing,
                "Session Traversal Utilities for NAT (STUN)", RFC 5389,
                October 2008.

   [RFC5766]   Rosenberg, J., Mahy, R., and P. Matthews, "Traversal
                Using Relays around NAT (TURN): Relay Extensions to
                Session Traversal Utilities for NAT (STUN)", RFC 5766,
                April 2010.

10.2. Informative References

   [HIP-MIDDLE] Heer, T., Wehrle, K., and M. Komu, "End-Host
                Authentication for HIP Middleboxes", Work in Progress,
                February 2009.

   [MMUSIC-ICE] Rosenberg, J., "Guidelines for Usage of Interactive
                Connectivity Establishment (ICE) by non Session
                Initiation Protocol (SIP) Protocols", 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]   Stiemerling, 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 Ta 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 round-trip time (RTT) between them and set the minimum
   Ta 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 a rendezvous server (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 are reachable, it MAY try
   them out using its own HIP relay server as described above.

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.priv.at/

Jan Melen
Ericsson Research Nomadiclab
Hirsalantie 11
02420 Jorvas
Finland
Phone: +358 9 2991
EMail: jan.melen@ericsson.com

Ari Keranen (editor)
Ericsson Research Nomadiclab
Hirsalantie 11
02420 Jorvas
Finland
Phone: +358 9 2991
EMail: ari.keranen@ericsson.com