

Mobile IP Traversal of Network Address Translation (NAT) Devices

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Mobile IP's datagram tunnelling is incompatible with Network Address Translation (NAT). This document presents extensions to the Mobile IP protocol and a tunnelling method which permits mobile nodes using Mobile IP to operate in private address networks which are separated from the public internet by NAT devices. The NAT traversal is based on using the Mobile IP Home Agent UDP port for encapsulated data traffic.

Table of Contents

1.	Introduction	2
1.1	Terminology	2
1.2	Problem description	3
1.3	Assumptions	4
2.	NAT Traversal Overview.	5
2.1	Basic Message Sequence.	5
3.	New Message Formats	6
3.1	UDP Tunnel Request Extension.	6
3.1.1	F (Force) Flag.	7
3.1.2	R (Registration through FA Required) flag	8
3.1.3	Reserved Fields	8
3.1.4	Encapsulation	8
3.1.5	Mobile IP Registration Bits	9
3.2	UDP Tunnel Reply Extension.	9
3.2.1	Reply Code.	10

3.3	MIP Tunnel Data Message	10
3.4	UDP Tunnelling Flag in Agent Advertisements	11
3.5	New Registration Reply Codes.	12
4.	Protocol Behaviour.	12
4.1	Relation to standard MIP tunnelling	12
4.2	Encapsulating IP Headers in UDP	13
4.3	Decapsulation	15
4.4	Mobile Node Considerations.	15
4.5	Foreign Agent Considerations.	16
4.6	Home Agent Considerations	18
4.6.1	Error Handling.	19
4.7	MIP signalling versus tunnelling.	20
4.8	Packet fragmentation.	21
4.9	Tunnel Keepalive.	21
4.10	Detecting and compensating for loss of NAT mapping.	22
4.11	Co-located registration through FA.	24
5.	Implementation Issues	24
5.1	Movement Detection and Private Address Aliasing	24
5.2	Mobility Binding Lifetime	25
6.	Security Considerations	26
6.1	Traffic Redirection Vulnerabilities	27
6.1.1	Manipulation of the Registration Request Message	27
6.1.2	Sending a Bogus Keepalive Message	27
6.2	Use of IPsec.	28
6.3	Firewall Considerations	28
7.	UNSAF Considerations.	28
8.	IANA Considerations	30
9.	Intellectual Property Rights.	30
10.	Acknowledgements.	31
11.	Normative References.	31
12.	Informative References.	32
13.	Authors' Addresses.	33
14.	Full Copyright Statement.	34

1. Introduction

1.1 Terminology

The Mobile IP related terminology described in RFC 3344 [10] is used in this document. In addition, the following terms are used:

Forward Tunnel

A tunnel that forwards packets towards the mobile node. It starts at the home agent, and ends at the mobile node's care-of address.

Reverse Tunnel

A tunnel that starts at the mobile node's care-of address and terminates at the home agent.

NAT

Network Address Translation. "Traditional NAT would allow hosts within a private network to transparently access hosts in the external network, in most cases. In a traditional NAT, sessions are uni-directional, outbound from the private network." -- RFC 2663 [11]. Basic NAT and NAPT are two varieties of NAT.

Basic NAT

"With Basic NAT, a block of external addresses are set aside for translating addresses of hosts in a private domain as they originate sessions to the external domain. For packets outbound from the private network, the source IP address and related fields such as IP, TCP, UDP and ICMP header checksums are translated. For inbound packets, the destination IP address and the checksums as listed above are translated." -- RFC 2663 [11].

NAPT

Network Address Port Translation. "NAPT extends the notion of translation one step further by also translating transport identifier (e.g., TCP and UDP port numbers, ICMP query identifiers). This allows the transport identifiers of a number of private hosts to be multiplexed into the transport identifiers of a single external address. NAPT allows a set of hosts to share a single external address. Note that NAPT can be combined with Basic NAT so that a pool of external addresses are used in conjunction with port translation." -- RFC 2663 [11].

In this document, the more general term NAT is used to cover both NATs and NAPT. In most deployment cases today, we believe that the NATs used are of the NAPT variety.

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 BCP 14, RFC 2119 [6].

1.2 Problem description

A basic assumption that Mobile IP [10] makes is that mobile nodes and foreign agents are uniquely identifiable by a globally routable IP address. This assumption breaks down when a mobile node attempts to communicate from behind a NAT.

Mobile IP relies on sending traffic from the home network to the mobile node or foreign agent through IP-in-IP tunnelling. IP nodes which communicate from behind a NAT are reachable only through the NAT's public address(es). IP-in-IP tunnelling does not generally contain enough information to permit unique translation from the common public address(es) to the particular care-of address of a mobile node or foreign agent which resides behind the NAT; in particular there are no TCP/UDP port numbers available for a NAT to work with. For this reason, IP-in-IP tunnels cannot in general pass through a NAT, and Mobile IP will not work across a NAT.

Mobile IP's Registration Request and Reply will on the other hand be able to pass through NATs and NAPT's on the mobile node or foreign agent side, as they are UDP datagrams originated from the inside of the NAT or NAPT. When passing out, they make the NAT set up an address/port mapping through which the Registration Reply will be able to pass in to the correct recipient. The current Mobile IP protocol does however not permit a registration where the mobile node's IP source address is not either the CoA, the Home Address, or 0.0.0.0.

What is needed is an alternative data tunnelling mechanism for Mobile IP which will provide the means needed for NAT devices to do unique mappings so that address translation will work, and a registration mechanism which will permit such an alternative tunnelling mechanism to be set up when appropriate.

This mechanism will address 3 different scenarios:

- A mobile node with co-located address behind a NAT; no FA
- A mobile node registered with an FA where both the mobile node and the FA are behind the same NAT
- A mobile node with co-located address using an FA which demands that registrations pass through the FA (sets the "R" bit) where both the mobile node and the FA are behind the same NAT

1.3 Assumptions

The primary assumption in this document is that the network allows communication between an UDP port chosen by the mobile node and the home agent UDP port 434. If this assumption does not hold, neither Mobile IP registration nor data tunnelling will work.

This document does NOT assume that mobility is constrained to a common IP address space. On the contrary, the routing fabric between the mobile node and the home agent may be partitioned into a

"private" and a "public" network, and the assumption is that some mechanism is needed in addition to vanilla Mobile IP according to RFC 3344 [10] in order to achieve mobility within disparate IP address spaces.

For a more extensive discussion of the problems with disparate address spaces, and how they may be solved, see RFC 3024 [9].

The reverse tunnels considered here are symmetric, that is, they use the same configuration (encapsulation method, IP address endpoints) as the forward tunnel.

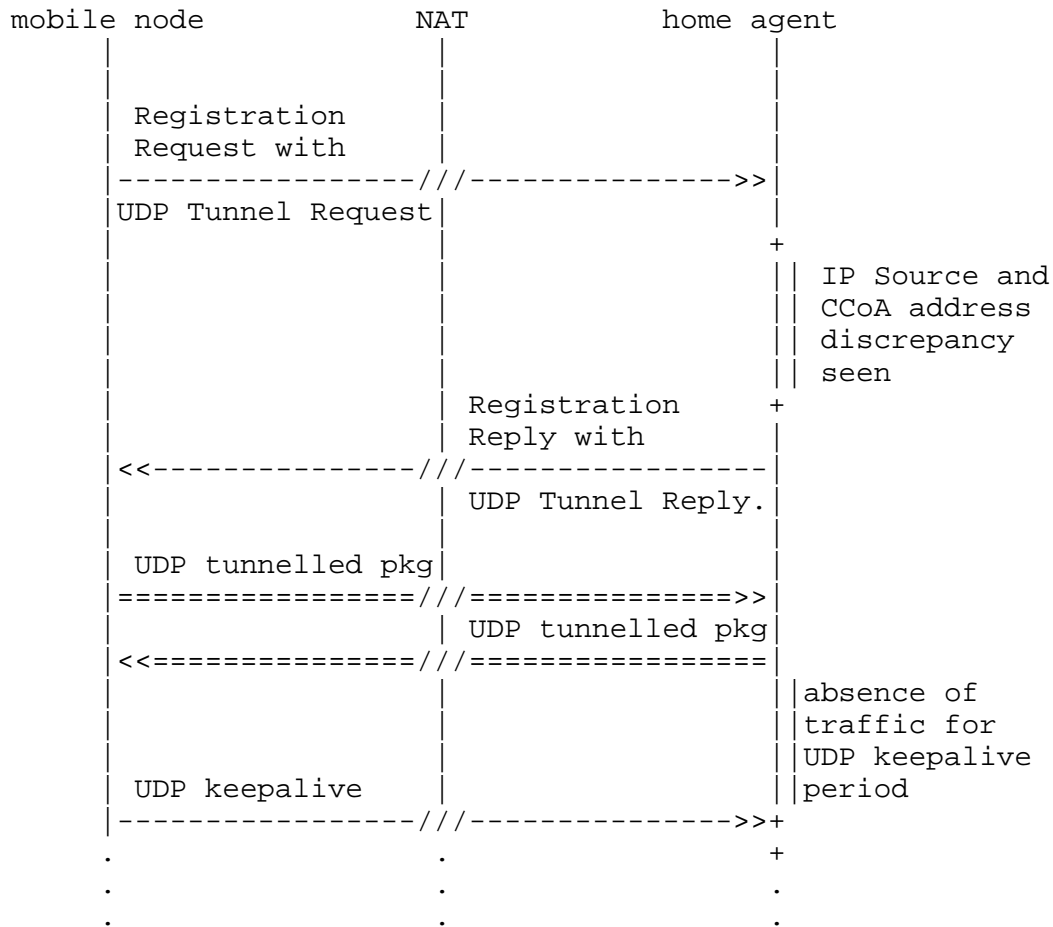
2. NAT Traversal Overview

This section gives a brief overview of the MIP UDP tunnelling mechanism which may be used to achieve NAT traversal for Mobile IP.

In MIP UDP tunnelling, the mobile node may use an extension (described below) in its Registration Request to indicate that it is able to use Mobile IP UDP tunnelling instead of standard Mobile IP tunnelling if the home agent sees that the Registration Request seems to have passed through a NAT. The home agent may then send a registration reply with an extension indicating acceptance (or denial). After assent from the home agent, MIP UDP tunnelling will be available for use for both forward and reverse tunnelling. UDP tunnelled packets sent by the mobile node use the same ports as the registration request message. In particular, the source port may vary between new registrations, but remains the same for all tunnelled data and re-registrations. The destination port is always 434. UDP tunnelled packets sent by the home agent uses the same ports, but in reverse.

2.1 Basic Message Sequence


The message sequence diagram below exemplifies setting up and using a Mobile IP UDP tunnel as described in this document. The tunnel is set up by the use of specific extensions in the initial Mobile IP Registration Request and Reply exchange. Thereafter, any traffic from the home agent to the mobile node is sent through the UDP tunnel. The mobile node may at its discretion use the UDP tunnel for reverse tunnelling or not, although in most cases where MIP UDP tunnelling is needed, reverse tunnelling will also be needed.



3. New Message Formats

3.1 UDP Tunnel Request Extension

This extension is a skippable extension. It signifies that the sender is capable of handling MIP UDP tunnelling, and optionally that a particular encapsulation format is requested in the MIP UDP tunnel. The format of this extension is as shown below. It adheres to the short extension format described in [10].



Type	144
Length	6. Length in bytes of this extension, not including the Type and Length bytes.
Sub-Type	0
Reserved 1	Reserved for future use. MUST be set to 0 on sending, MUST be ignored on reception.
F	F (Force) flag. Indicates that the mobile node wants to force MIP UDP tunnelling to be established.
R	R (Registration through FA Required) flag. Indicates that the R bit was set in the FA's Agent Advertisement. Registration is being made using a co-located care-of address, but through the FA.
Reserved 2	Reserved for future use. MUST be set to 0 on sending, MUST be ignored on reception.
Encapsulation	Indicates the type of tunnelled data, using the same numbering as the IP Header Protocol Field.
Reserved 3	Reserved for future use. MUST be set to 0 on sending, MUST be verified as 0 on receipt; otherwise the extension must be handled as not understood and silently skipped.

3.1.1 F (Force) Flag

Indicates that the mobile node wants to use traversal regardless of the outcome of NAT detection performed by the home agent. This is useful if the route between the mobile node and the home agent works for Mobile IP signalling packets, but not for generic data packets (e.g., because of firewall filtering rules). If the home agent supports this protocol, it SHOULD either accept the traversal and reply with a UDP Tunnel Reply Extension or reject the Registration Request. In case of the registration failing, the Home Agent SHOULD send a Registration Reply with Code field set to 129 ("administratively prohibited").

If the HA does not understand the UDP Tunnel Request Extension, it will silently discard it, and if everything else is fine, it will reply with a registration reply with reply code 0 (registration accepted), but without any UDP Tunnel Reply Extension. In this case, the mobile node MUST NOT use MIP UDP tunnelling.

3.1.2 R (Registration through FA Required) flag

This flag MUST be set by the mobile node when it is using a co-located address, but registering through an FA because it has received an Agent Advertisement with the 'R' bit set.

3.1.3 Reserved Fields

The 'Reserved 1' and 'Reserved 2' fields must be ignored on receipt, while the 'Reserved 3' field must be 0 on receipt, otherwise this extension MUST be handled as not understood and silently skipped. This permits future additions to this extension to be made which either can co-exist with old implementations, or will force a rejection of the extension from an old implementation.

3.1.4 Encapsulation

The 'Encapsulation' field defines the mode of encapsulation requested if MIP UDP tunnelling is accepted by the home agent. This field uses the same values as the IP header Protocol field:

- 4 IP header (IP-in-UDP tunnelling) RFC 2003 [4]
- 47 GRE Header (GRE-in-UDP tunnelling) RFC 2784 [8]
- 55 Minimal IP encapsulation header RFC 2004 [5]

If the home agent finds that UDP tunnelling is not needed, the encapsulation will be determined by the 'M' and 'G' flags of the registration request; but if the home agent finds that MIP UDP tunnelling should be done, the encapsulation is determined from the value of the Encapsulation field. If the value of this field is zero, it defaults to the value of 'M' and 'G' fields in the Registration Request message, but if it is non-zero, it indicates that a particular encapsulation is desired.

3.1.5 Mobile IP Registration Bits


The Mobile IP registration bits S, B, D, M, G and T retain their meaning as described in RFC 3344 [10] and RFC 3024 [9] (except that the significance of the M and G bits may be modified by the Encapsulation field when MIP UDP tunnelling is used, as described above). The use of the M and G bits together with MIP UDP tunnelling is also touched upon in Section 4.1.

When the MN requests MIP UDP tunnelling, the 'D' bit (Decapsulation by the mobile node) MUST be set, otherwise UDP tunnelling would not be meaningful.

Both the MN and the FA SHOULD set the 'T' bit when requesting MIP UDP tunnelling, even if not all traffic will be reverse tunnelled. This ensures that a HA which is not prepared to accept reverse tunnelling will not accept a registration which may later turn out to be unusable. Also see the discussion of use of the 'T' bit in Foreign Agent Considerations (Section 4.5).

3.2 UDP Tunnel Reply Extension

This extension is a non-skippable extension. It is sent in reply to a UDP Tunnel Request extension, and indicates whether or not the HA will use MIP UDP tunnelling for the current mobility binding. The format of this extension is as shown below.



Type 44

Length 6. Length in bytes of this extension, not including the Type and Length bytes.

Sub-Type 0

Reply Code Indicates whether the HA assents or declines to use UDP tunnelling for the current mobility binding. See Section 3.2.1 below.

F (Forced) flag. Indicates that tunnelling is being forced because the F flag was set in the tunnelling request, irrespective of the detection of a NAT or not.

Keepalive Interval	Specifies the NAT keepalive interval that the mobile node SHOULD use. A keepalive packet SHOULD be sent if Keepalive Interval seconds have elapsed without any signalling or data traffic being sent. If this field is set to 0, the mobile node MUST use its default configured keepalive interval.
--------------------	--

Reserved	Reserved for future use. MUST be set to 0 on sending, MUST be ignored on reception.
----------	---

3.2.1 Reply Code

The Reply Code field of the UDP Tunnel Reply Extension indicates if UDP tunnelling have been accepted and will be used by the HA. Values in the range 0 .. 63 indicate assent, i.e., that tunnelling will be done, while values in the range 64 .. 255 indicate that tunnelling should not be done. More information may be given by the value of the response code.

The following response codes are defined for use in the code field of the UDP Tunnel Reply Extension:

```
0 Will do tunnelling
```

64 Tunnelling declined, reason unspecified

3.3 MIP Tunnel Data Message

This MIP message header serves to differentiate traffic tunneled through the well-known port 434 from other Mobile IP messages, e.g., Registration Requests and Registration Replies.

[illegible]

Type 4

Next Header	Indicates the type of tunnelled data, using the same numbering as the IP Header Protocol Field.
Reserved	Reserved for future use. MUST be set to 0 on sending, MUST be ignored on reception.

The Next Header field uses the same values as the IP header Protocol field. Immediately relevant for use with Mobile IP are the following values:

- 4 IP header (IP-in-UDP tunnelling) RFC 2003 [4]
- 47 GRE Header (GRE-in-UDP tunnelling) RFC 2784 [8]
- 55 Minimal IP encapsulation header RFC 2004 [5]

The receiver of a tunnelled packet MUST check that the Next Header value matches the tunnelling mode established for the mobility binding with which the packet was sent. If a discrepancy is detected, the packet MUST be dropped. A log entry MAY be written, but in this case the receiver SHOULD ensure that the amount of log entries written is not excessive.

In addition to the encapsulation forms listed above, the MIP UDP tunnelling can potentially support other encapsulations, by use of the Next Header field in the MIP Tunnel Data Header and the Encapsulation Header field of the UDP Tunnel Request Extension (Section 3.1).

3.4 UDP Tunnelling Flag in Agent Advertisements

The only change to the Mobility Agent Advertisement Extension defined in RFC 3344 [10] is a flag indicating that the foreign agent generating the Agent Advertisement supports MIP UDP Tunnelling. The flag is inserted after the flags defined in [10].

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	Sequence Number	
Lifetime		R B H F M G r T U	reserved
zero or more Care-of Addresses			
...			

The flag is defined as follows:

- U UDP Tunnelling support. This Agent supports MIP UDP Tunnelling as specified in this document. This flag SHOULD be set in advertisements sent by a foreign agent which supports MIP UDP tunnelling and is situated behind a NAT. It MUST NOT be set in advertisements from foreign agents which are not situated behind a NAT, and thus has no need to advertise the capability.

3.5 New Registration Reply Codes

One new registration reply code is defined:

ERROR_HA_UDP_ENCAP_UNAVAIL	Requested UDP tunnel encapsulation unavailable
----------------------------	--

This is used by the HA to distinguish the registration denial caused by an unavailable UDP tunnel encapsulation mode from a denial caused by unavailable standard tunnel encapsulation requested by use of the 'T' bit together with either 'M' or 'G' bit.

4. Protocol Behaviour

4.1 Relation to standard MIP tunnelling

The default encapsulation mode for MIP UDP tunnelling is IP-in-UDP encapsulation. The mobile node MAY request alternative forms of encapsulation to be used with UDP tunnelling by setting the 'M' bit and/or the 'G' bit of a Mobile IP registration request, or by explicitly requesting a particular encapsulation for the MIP UDP tunnel by using the Encapsulation field. The M and G bits retain the meaning as described in RFC 3344 [10] within the context of MIP UDP tunnelling. The UDP tunnelling version of the classic MIP encapsulation methods can be summarised as:

IP in UDP. When Mobile IP UDP tunnelling is used, this is the default encapsulation type. Any home agent and mobile node that implements Mobile IP UDP tunnelling MUST implement this encapsulation type.

GRE in UDP. If the 'G' bit is set in a registration request and the Encapsulation field is zero, the mobile node requests that its home agent use GRE encapsulation [3] for datagrams tunnelled to the mobile node. If MIP UDP tunnelling is also requested and accepted, GRE-in-UDP encapsulation SHALL be used in the same cases as GRE in IP encapsulation would be used if the MIP UDP tunnelling had not been requested.

Minimal encapsulation in UDP. If the 'M' bit is set and the Encapsulation field is zero, the mobile node requests that its home agent use minimal encapsulation [5] for datagrams tunneled to the mobile node. If MIP UDP tunnelling is also used, minimal encapsulation in UDP SHALL be used in the same cases as minimal encapsulation according to RFC 2004 [5] would be used if the MIP UDP tunnelling had not been requested.

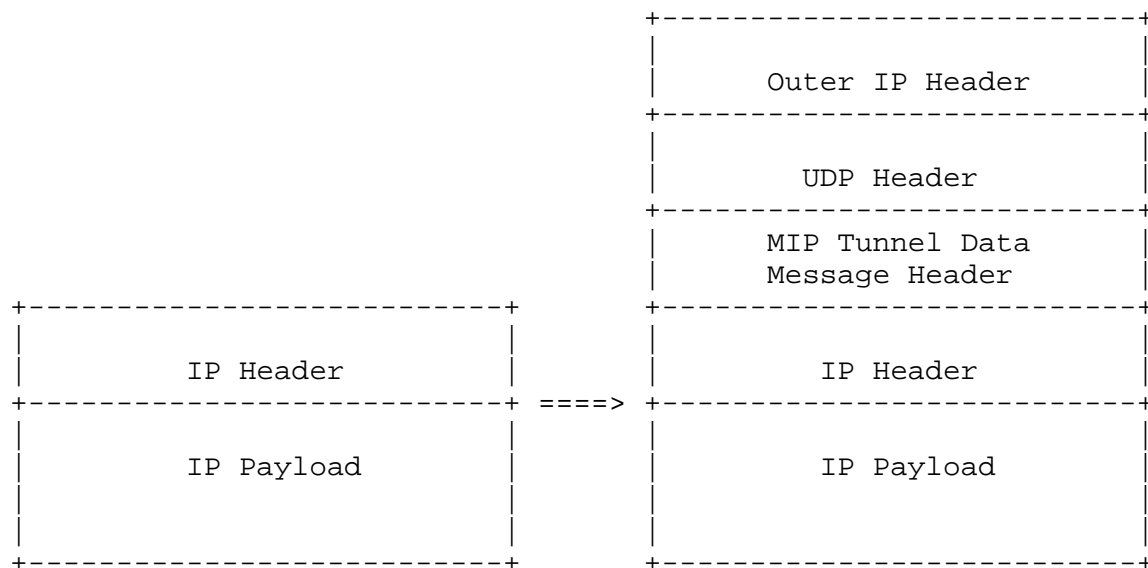
When the Encapsulation field is non-zero, a particular encapsulation format is requested for the MIP UDP tunnel. If tunnelling is indicated, the request MUST either be accepted using the requested encapsulation, or rejected with the error code `ERROR_HA_UDP_ENCAP_UNAVAIL`, "Requested UDP tunnel encapsulation unavailable" defined in Section 3.5. On receipt of this error, the mobile node MAY choose to send a new Registration Request with different requirements on MIP UDP tunnelling encapsulation.

4.2 Encapsulating IP Headers in UDP

MIP IP-in-UDP tunnelling, or UDP tunnelling for short, is done in a manner analogous to that described for IP-in-IP tunnelling in RFC 2003 [4], with the exception of the addition of an UDP header [1] and MIP Message header [10] between the outer and inner IP header.

Mobile IP Registration Requests and Registration Replies are already in the form of UDP messages, and SHALL NOT be tunneled even when MIP IP-in-UDP tunnelling is in force.

To encapsulate an IP datagram using MIP IP-in-UDP encapsulation, an outer IP header [2], UDP header [1] and MIP Message header [10] is inserted before the datagram's existing IP header, as follows:



The outer IP header Source Address and Destination Address, together with the UDP header Source Port and Destination Port, identify the "endpoints" of the tunnel. The inner IP header Source Address and Destination Addresses identify the original sender and the recipient of the datagram, respectively. The inner IP header is not changed by the encapsulator, except to decrement the TTL by one if the tunnelling is being done as part of forwarding the datagram as noted in RFC 2003 [4], and remains unchanged during its delivery to the tunnel exit point. No change to IP options in the inner header occurs during delivery of the encapsulated datagram through the tunnel. Note that the security options of the inner IP header MAY affect the choice of security options for the encapsulating (outer) IP header.

Minimal Encapsulation and GRE encapsulation is done in an analogous manner, following RFC 2004 [5] for Minimal Encapsulation and RFC 2784 [8] for GRE Encapsulation, but using outer IP, UDP and MIP headers in place of the outer IP header.

All other provisions and requirements of RFC 2003 [4] and RFC 3024 [9] are in force, except in one respect, as covered in Packet Fragmentation (Section 4.8).

4.3 Decapsulation

Before decapsulation is actually done, the decapsulating node **MUST** verify that the outer IP addresses and UDP port numbers exactly match the values used for the tunnel, with the exception of tunnels that are "half bound" (as described in Section 4.11) where the source UDP port can change.

IP-in-UDP encapsulated traffic is decapsulated simply by stripping off the outer IP, UDP and MIP header, which leaves the original IP packet which is forwarded as is.

Minimal IP encapsulation is processed by the receiver conceptually as follows. First, the UDP and the Mobile IP headers are removed from the packet, and the Protocol field of the IP header replaced with the Next Header field in the MIP Tunnel Data header. Second, the remaining IP header total length and checksum are adjusted to match the stripped packet. Third, ordinary minimal IP encapsulation processing is done.

GRE encapsulated traffic is processed according to RFC 2784 [8] and RFC 1701 [3], with the delivery header consisting of the outer IP, UDP and MIP headers.

4.4 Mobile Node Considerations

The UDP Tunnel Request Extension **MAY** be used in a Mobile IP Registration Request from the mobile node to the home agent when the mobile node uses a co-located care-of address. It **SHALL NOT** be used by the mobile node when it is registering with a foreign agent care-of address.

The purpose of this extension is to indicate to the home agent that the mobile node is able to accept MIP UDP tunnelling if the home agent has an indication that the mobile node resides behind a NAT or NAT. It thus functions as a conditional solicitation for the use of MIP UDP tunnelling.

As per Section 3.2 and 3.6.1.3 of RFC 3344 [10], the mobile node **MUST** place this Extension before the Mobile-Home Authentication Extension in registration messages, so that it is covered by the Mobile-Home Authentication Extension.

If the mobile node includes the UDP Tunnel Request extension in a registration request, but receives a registration reply without a UDP Tunnel Reply extension, it **MUST** assume that the HA does not

understand this extension, and it MUST NOT use UDP tunnelling. If the mobile node is in fact behind a NAT, the registration may then succeed, but traffic will not be able to traverse the NAT.

When the mobile node sends MIP UDP tunnelled data, it MUST use the same UDP source port as was used for the most recent registration request.

When the mobile node re-registers without having moved, it SHOULD take care to use the same source port as was used for the original registration of the current mobility binding. Otherwise, while the home agent would change destination port on acceptance of the new registration, and the mobile node would presumably start listening on the new port, the packets in flight from the home agent at the time of change will be dropped when arriving at the mobile node's old port. (This does not mean that the home agent should refuse a registration request using MIP UDP tunnelling where a new port have been used, as this might be the result of the NAT dropping state, the mobile node re-booting, changing interface, etc.)

If a mobile node is registering through a foreign agent but using a co-located care-of address, and the agent advertisement from the foreign agent had the 'U' bit set, the mobile node MUST set the 'R' flag in its UDP Tunnel Request Extension, in order to make the HA use MIP UDP tunnelling. In this case, the mobile node also MUST send a keepalive as soon as its registration has been accepted.

If a mobile node is registering through a foreign agent but using a co-located care-of address, and the agent advertisement from the foreign agent does not have the 'U' bit set, the mobile node MUST NOT include a UDP Tunnel Request Extension in the registration request.

4.5 Foreign Agent Considerations

The UDP Tunnel Request Extension MAY be used by a foreign agent when it is forwarding a Mobile IP Registration Request to a home agent, when the foreign agent is situated behind a NAT or has some other compelling reason to require MIP UDP tunnelling.

The purpose of this extension is to indicate to the home agent that the foreign agent is able to accept MIP UDP tunnelling if the home agent has an indication that the foreign agent resides behind a NAT or NAPT. It thus functions as a conditional solicitation for the use of MIP UDP tunnelling.

A foreign agent which requires the mobile node to register through a foreign agent by setting the 'R' bit in the agent advertisement, MUST NOT add the UDP Tunnel Request Extension when forwarding a

registration request which uses a co-located care-of address, as this will lead to a UDP tunnel being set up from the home agent to the foreign agent instead of to the mobile node.

As per Section 3.2 and 3.7.2.2 of RFC 3344 [10], the foreign agent when using this extension MUST place it after the Mobile-Home Authentication Extension in registration messages. If the foreign agent shares a mobility security association with the home agent and therefore appends a Foreign-Home Authentication Extension, the UDP Tunnel Request Extension MUST be placed before the Foreign-Home Authentication Extension.

As the home agent detects the presence of a NAT in the path between the sender and itself by seeing a mismatch between the IP source address and the care-of address given in the registration request, it is REQUIRED that the foreign agent, when using this extension, sends the registration request with an IP source address matching the care-of address.

A foreign agent using MIP UDP tunnelling to a home agent because the FA is situated behind a NAT may be configured to encourage reverse tunnelling, or be neutral about it, depending on the characteristics of the NAT. If the NAT translates all source addresses of outgoing packets to its own public address, it will not be possible to maintain sessions when moving away from this network if the mobile node has used triangular routing instead of reverse tunnelling. On the other hand, if it is known that the NAT is smart enough to not translate publicly routable source addresses, AND does not do ingress filtering, triangular routing may succeed. The leg from the home agent to the foreign agent will still use MIP UDP tunnelling to pass through the NAT.

Therefore, if it is known when configuring a foreign agent behind a NAT that the NAT will translate public as well as private addresses, or it is known that ingress filtering is being done between the private and public networks, the foreign agent SHOULD reply to registration requests which don't have the 'T' bit set with a reply code 75, "reverse tunnel is mandatory and 'T' bit not set".

Conversely, if it is known that the NAT is smart about not translating public addresses, and no ingress filtering is done, so it is reasonable to believe that a mobile node with a publicly routable address may be able to keep up sessions when moving to or from this network, the foreign agent MAY be configured to forward registration requests even if they don't have the 'T' bit set.

If the behaviour of the NAT is unknown in this respect, it SHOULD be assumed that it will translate all addresses, thus the foreign agent SHOULD be configured to reply to registration requests which don't have the 'T' bit set with a reply code 75, "reverse tunnel is mandatory and 'T' bit not set".

4.6 Home Agent Considerations

The purpose of the MIP UDP Tunnel Reply Extension is to indicate whether or not the home agent accepts the use of MIP UDP tunnelling for this mobility binding, and to inform the mobile node or foreign agent of the suggested tunnel keepalive interval to be used.

The UDP Tunnel Reply Extension MUST be used in a Mobile IP Registration Reply from the home agent to the mobile node when it has received and accepted a UDP Tunnel Request (Section 3.1) from a mobile node.

The home agent MUST use a mismatch between source IP address and care-of address in the Mobile IP Registration Request message as the indication that a mobile node may reside behind a NAT. If the Registration Request also contains the UDP Tunnel Request extension without the 'R' flag set, and the home agent is capable of, and permits MIP UDP tunnelling, the home agent SHALL respond with a registration reply containing an assenting UDP Tunnel Reply Extension as described in Section 3.2. If the 'R' flag is set, special considerations apply, as described below.

If the home agent receives a Registration Request with matching source IP address and co-located care-of address which contains a MIP UDP Tunnel Request Extension, the home agent SHOULD respond with a Registration Reply containing a declining UDP Tunnel Reply - unless tunnelling has been explicitly requested by the mobile node using the 'F' flag as described in Section 3.1.

If the home agent assents to UDP tunnelling, it MUST use the source address of the registration request as the effective care-of address, rather than the care-of address given in the registration request, except in the case where the 'R' flag is set in the UDP Tunnel Request Extension.

If the home agent receives a Registration Request with the 'R' flag set in the UDP Tunnel Request Extension, it SHOULD reply with an assenting UDP Tunnel Reply Extension if it is capable of, and permits MIP UDP tunnelling. In this case, however, the source address and port of the registration request may be a NAT'ed version of the foreign agent source address and port. In order to direct tunnelled traffic correctly to the mobile node, the home agent MUST wait for

the first keepalive packet from the mobile node to arrive, before it can send traffic back to the correct NAT port (the one which is mapped to the MN). In this case, the home agent MUST check that the outer source address (but not the port) of this keepalive packet is identical with the source address of the corresponding registration request. The inner source address (that of the encapsulated ICMP echo request) MUST be the home address of the mobile node, and the inner destination address MUST be that of the home agent. If all this holds, the outer source address and port of this keepalive packet SHALL be used by the HA as the outer destination address and port of the MIP UDP tunnel when forwarding traffic to the mobile node.

The home agent SHOULD be consistent in acknowledging support for UDP tunnelling or not. A home agent which understands the UDP Tunnel Request Extension and is prepared to respond positively to such a request SHOULD also respond with a UDP Tunnel Reply Extension containing a declining reply code if use of MIP UDP tunnelling is not indicated for a session. The mobile node MUST NOT assume such behaviour from the home agent, since the home agent may undergo a software change with reboot, a policy change or a replacement; and consequently a change of behaviour.

4.6.1 Error Handling

The following actions take place when things go wrong.

The HA does not support the UDP Tunnel Request extension:

The home agent ignores the extension and proceeds normally, which would be to refuse the registration if the IP source address does not match the care-of address, the home address or 0.0.0.0. Even if the HA mistakenly does accept the registration, the mobile node will not be able to receive forward tunnelled data if it is behind a NAT.

(It would be beneficial to have the mobile node de-register in this case. The mobile node, however, normally has no way of telling that it is behind a NAT if it does not receive a UDP Tunnelling Reply.)

NAT detected by home agent, but traversal not allowed:

In some cases the home agent may disable NAT traversal even though it supports the UDP Tunnel Request extension and a NAT is detected. In this case, the home agent SHOULD send a Registration Reply with the Code field set to 129, "administratively prohibited".

NAT not detected, 'F' flag set, but home agent does not allow forced use of MIP UDP tunnelling:

The home agent SHOULD send a Registration Reply with the Code field set to 129, "administratively prohibited".

UDP Tunnel Request extension sent by the mobile node (placed before the MN-HA authentication extension), but 'D' bit in registration request header not set:

The home agent SHOULD send a Registration Reply with the Code field set to 134, "poorly formed Request".

UDP Tunnel Request extension sent by the foreign agent (placed after the MN-HA authentication extension), but 'D' bit is set:

The home agent SHOULD send a Registration Reply with the Code field set to 134, "poorly formed Request".

Reserved 3 field of UDP Tunnel Request extension is nonzero:

The home agent SHOULD send a Registration Reply with the Code field set to 134, "poorly formed Request".

Encapsulation type requested in UDP Tunnel Request extension is unsupported:

The home agent SHOULD send a Registration Reply with the Code field set to ERROR_HA_UDP_ENCAP_UNAVAIL, "Requested UDP tunnel encapsulation unavailable" defined in Section 3.5.

4.7 MIP signalling versus tunnelling

UDP tunnelling SHALL be used only for data packets, and only when the mobility binding used for sending was established using the UDP Tunnel Request, and accepted by an UDP Tunnel Reply from the home agent. After MIP UDP tunnelling has been established for a mobility binding, data packets that are forward or reverse tunneled using this mobility binding MUST be tunneled using MIP UDP tunnelling, not IP-in-IP tunnelling or some other tunnelling method.

As a consequence:

- Mobile IP signalling is never tunneled.
- When using simultaneous bindings, each binding may have a different type (i.e., UDP tunnelling bindings may be mixed with non-UDP tunnelling bindings).

- Tunnelling is only allowed for the duration of the binding lifetime.

4.8 Packet fragmentation

From RFC 3022 [12]:

"Translation of outbound TCP/UDP fragments (i.e., those originating from private hosts) in NAPT set-up are doomed to fail. The reason is as follows. Only the first fragment contains the TCP/UDP header that would be necessary to associate the packet to a session for translation purposes. Subsequent fragments do not contain TCP/UDP port information, but simply carry the same fragmentation identifier specified in the first fragment. Say, two private hosts originated fragmented TCP/UDP packets to the same destination host. And, they happened to use the same fragmentation identifier. When the target host receives the two unrelated datagrams, carrying same fragmentation id, and from the same assigned host address, it is unable to determine which of the two sessions the datagrams belong to. Consequently, both sessions will be corrupted."

Because of this, if the mobile node or foreign agent for any reason needs to send fragmented packets, the fragmentation MUST be done prior to the encapsulation. This differs from the case for IP-in-IP tunnelling, where fragmentation may be done before or after encapsulation, although RFC 2003 [4] recommends doing it before encapsulation.

A similar issue exists with some firewalls, which may have rules that only permit traffic on certain TCP and UDP ports, and not arbitrary outbound (or inbound) IP traffic. If this is the case and the firewall is not set to do packet reassembly, a home agent behind a firewall will also have to do packet fragmentation before MIP UDP encapsulation. Otherwise, only the first fragment (which contains the UDP header) will be allowed to pass from the home agent out through the firewall.

For this reason, the home agent SHOULD do packet fragmentation before it does MIP UDP encapsulation.

4.9 Tunnel Keepalive

As the existence of the bi-directional UDP tunnel through the NAT is dependent on the NAT keeping state information associated with a session, as described in RFC 2663 [11], and as the NAT may decide that the session has terminated after a certain time, keepalive messages may be needed to keep the tunnel open. The keepalives should be sent more often than the timeout value used by the NAT.

This timeout may be assumed to be a couple of minutes, according to RFC 2663 [11], but it is conceivable that shorter timeouts may exist in some NATs.

For this reason the extension used to set up the UDP tunnel has the option of setting the keepalive message interval to another value than the default value, see Section 3.2.

The keepalive message sent MUST consist of a properly UDP encapsulated ICMP echo request directed to the home agent.

For each mobility binding which has UDP tunnelling established, the non-HA endpoint of the Mobile-IP UDP tunnel MUST send a keepalive packet if no other packet to the HA has been sent in K seconds. Here K is a parameter with a default value of 110 seconds. K may be set to another value by the HA as described in the UDP tunnelling reply extension (Section 3.2).

Except for the case where the mobile node registers with a co-located address through an FA (see Section 4.11) MIP UDP tunnelling is done using the same ports that have already been used for the registration request / reply exchange. The MN or FA will send its first keepalive message at the earliest K seconds after the registration request was sent. The same UDP source port MUST be used for the keepalive messages as was used for the original Registration Messages and for data messages.

The remote UDP tunnel endpoint MUST use two-way keepalives consisting of UDP encapsulated ICMP Echo Request/Reply messages. The rationale for using two-way keepalives is two-fold:

1. Two-way keepalives allow the mobile node to detect loss of a NAT mapping. Detection of NAT mapping loss in turn allows the MN to compensate by re-registering and using a shorter keepalive to avoid loss of NAT mappings in the future.
2. One-way keepalives (keepalives sent by MN or FA, but without any reply from the home agent) actually cause more keepalive traffic overhead; the keepalive messages have to be sent more frequently to compensate for occasional loss of keepalive messages. In contrast, two-way keepalives are acknowledged, and retransmissions occur only when a response is not received for a keepalive request within a reasonable time.

4.10 Detecting and compensating for loss of NAT mapping

When a mobile node is using UDP encapsulated ICMP Echo Request/Reply messages as keepalives, it will have to deal with the possibility

that a NAT mapping is lost by a NAT device. The crucial thing here is of course not the loss of the NAT mapping in itself; but rather that the home agent, in the absence of a Registration Request through the new mapping, will continue to send traffic to the NAT port associated with the old mapping.

If the mobile node does not get a reply to its UDP encapsulated ICMP Echo Request even after a number of retransmissions, and is still connected to the same router that was used to establish the current mobility binding, the mobile node SHOULD re-register with the home agent by sending an Registration Request. This lets the HA know about the new NAT mapping and restores connectivity between mobile node and home agent.

Having established a new mobility binding, the mobile node MAY use a shorter keepalive interval than before the NAT mapping was lost; in particular, the mobile node MAY deviate from the keepalive interval assigned by the home agent. If the binding loss continues to occur, the mobile node may shorten the keepalive interval each time it re-registers, in order to end up with a keepalive interval that is sufficient to keep the NAT mapping alive. The strategy used to arrive at a keepalive interval when a NAT mapping is lost is implementation dependent. However, the mobile node MUST NOT use a keepalive less than 10 seconds.

Note that the above discussion only applies when the mobile node is re-registering through the same router, and thus presumably through the same NAT device that lost a NAT mapping earlier. If the MN moves and still finds itself behind a NAT, it SHOULD return to its original keepalive interval (the default value, or as assigned by the home agent) and it SHOULD NOT do any keepalive interval compensation unless it discovers a loss of NAT mapping in the new environment.

The home agent MUST NOT attempt to detect or compensate for NAT binding loss by dynamically changing the keepalive interval assigned in the Registration Reply; the home agent does not have enough information to do this reliably and should thus not do it at all. The mobile node is in a much better position to determine when a NAT mapping has actually been lost. Note also that a MN is allowed to let a NAT mapping expire if the MN no longer needs connectivity.

The discussion above does only in a limited sense apply to a foreign agent which is situated behind a NAT and using MIP UDP tunnelling. In this case, it is a matter of permanently configuring the FA to use a keepalive interval which is lower than the NAT mapping lifetime, rather than trying to dynamically adapt to the binding lifetimes of different NATs.

4.11 Co-located registration through FA

This section summarizes the protocol details which have been necessary in order to handle and support the case when a mobile node registers with a co-located address through a foreign agent, due to the FA advertisements having the 'R' bit set. It gives background information, but lists no new requirements.

When a mobile registers a co-located care-of address through an FA, the registration request which reaches the HA will have a different care-of address in the registration request compared to the source address in the registration request IP-header. If the registration request also contains a UDP Tunnel Request Extension, the HA will erroneously set up a UDP tunnel, which will go to the FA instead of the MN. For this reason, as mentioned in Section 4.4, the mobile node must not include a UDP Tunnel Request Extension in the registration if it is registering a co-located address through an FA which does not have the 'U' bit set in its advertisements.

In order to still be able to use MIP UDP tunnelling in this case, foreign agents which are situated behind a NAT are encouraged to send advertisements which have the 'U' bit set, as described in Section 3.4.

If the FA advertisement has the 'U' bit set, indicating that it is behind a NAT, and also the 'R' bit set, and the mobile node wishes to use a co-located care-of address, it MUST set the 'R' flag in the UDP Tunnel Request Extension, in order to inform the HA of the situation so that it may act appropriately, as described in Section 4.4.

Because the UDP tunnel is now taking another path than the registration requests, the home agent, when handling registrations of this type, must wait till the arrival of the first keepalive packet before it can set up the tunnel to the correct address and port. To reduce the possibility of tunnel hijacking by sending a keepalive with a phony source address, it is required that only the port of the keepalive packet may be different from that of the registration request; the source address must be the same. This means that if the FA and MN are communicating with the HA through different NATs, the connection will fail.

5. Implementation Issues

5.1 Movement Detection and Private Address Aliasing

In providing a mobile node with a mechanism for NAT traversal of Mobile IP traffic, we expand the address space where a mobile node may function and acquire care-of addresses. With this comes a new

problem of movement detection and address aliasing. We here have a case which may not occur frequently, but is mentioned for completeness:

Since private networks use overlapping address spaces, they may be mistaken for one another in some situations; this is referred to as private address aliasing in this document. For this reason, it may be necessary for mobile nodes implementing this specification to monitor the link layer address(es) of the gateway(s) used for sending packets. A change in the link layer address indicates probable movement to a new network, even if the IP address remains reachable using the new link layer address.

For instance, a mobile node may obtain the co-located care-of address 10.0.0.1, netmask 255.0.0.0, and gateway 10.255.255.254 using DHCP from network #1. It then moves to network #2, which uses an identical addressing scheme. The only difference for the mobile node is the gateway's link layer address. The mobile node should store the link layer address it initially obtains for the gateway (using ARP, for instance). The mobile node may then detect changes in the link layer address in successive ARP exchanges as part of its ordinary movement detection mechanism.

In rare cases the mobile nodes may not be able to monitor the link layer address of the gateway(s) it is using, and may thus confuse one point of attachment with another. This specification does not explicitly address this issue. The potential traffic blackout caused by this situation may be limited by ensuring that the mobility binding lifetime is short enough; the re-registration caused by expiration of the mobility binding fixes the problem (see Section 5.2).

5.2 Mobility Binding Lifetime

When responding to a registration request with a registration reply, the home agent is allowed to decrease the lifetime indicated in the registration request, as covered in RFC 3344 [10]. When using UDP tunnelling, there are some cases where a short lifetime is beneficial.

First, if the NAT mapping maintained by the NAT device is dropped, a connection blackout will arise. New packets sent by the mobile node (or the foreign agent) will establish a new NAT mapping, which the home agent will not recognize until a new mobility binding is established by a new registration request.

A second case where a short lifetime is useful is related to the aliasing of private network addresses. In case the mobile node is

not able to detect mobility and ends up behind a new NAT device (as described in Section 5.1), a short lifetime will ensure that the traffic blackout will not be exceedingly long, and is terminated by a re-registration.

The definition of "short lifetime" in this context is dependent on the requirements of the usage scenario. Suggested maximum lifetime returned by the home agent is 60 seconds, but in case the abovementioned scenarios are not considered a problem, longer lifetimes may of course be used.

6. Security Considerations

The ordinary Mobile IP security mechanisms are also used with the NAT traversal mechanism described in this document. However, there is one noticeable change: the NAT traversal mechanism requires that the HA trust unauthenticated address (and port) fields possibly modified by NATs.

Relying on unauthenticated address information when forming or updating a mobility binding leads to several redirection attack vulnerabilities. In essence, an attacker may do what NATs do, i.e., modify addresses and ports and thus cause traffic to be redirected to a chosen address. The same vulnerabilities apply to both MN-HA and FA-HA NAT traversal.

In more detail: without a NAT, the care-of address in the registration request will be directly used by the HA to send traffic back to the MN (or the FA), and the care-of address is protected by the MN-HA (or FA-HA) authentication extension. When communicating across a NAT, the effective care-of address from the HA point of view is that of the NAT, which is not protected by any authentication extension, but inferred from the apparent IP source address of received packets. This means that by using the mobile IP registration extensions described in this document to enable traversal of NATs, one is opening oneself up to having the care-of address of a MN (or a FA) maliciously changed by an attacker.

Some, but not all, of the attacks could be alleviated to some extent by using a simple routability check. However, this document does not specify such a mechanism for simplicity reasons and because the mechanism would not protect against all redirection attacks. To limit the duration of such redirection attacks, it is RECOMMENDED to use a conservative (that is, short) mobility binding lifetime when using the NAT traversal mechanism specified in this document.

The known security issues are described in the sections that follow.

6.1 Traffic Redirection Vulnerabilities

6.1.1 Manipulation of the Registration Request Message

An attacker on the route between the mobile node (or foreign agent) and the home agent may redirect mobility bindings to a desired address simply by modifying the IP and UDP headers of the Registration Request message. Having modified the binding, the attacker no longer needs to listen to (or manipulate) the traffic. The redirection is in force until the mobility binding expires or the mobile node re-registers.

This vulnerability may be used by an attacker to read traffic destined to a mobile node, and to send traffic impersonating the mobile node. The vulnerability may also be used to redirect traffic to a victim host in order to cause denial-of-service on the victim.

The only defense against this vulnerability is to have a short time between re-registrations, which limits the duration of the redirection attack after the attacker has stopped modifying registration messages.

6.1.2 Sending a Bogus Keepalive Message

When registering through an FA using a co-located care-of address, another redirection vulnerability opens up. Having exchanged Registration Request/Reply messages with the HA through the FA, the MN is expected to send the first keepalive message to the HA, thus finalizing the mobility binding (the binding will remain in a "half bound" state until the keepalive is received).

Having observed a Registration Request/Reply exchange, an attacker may send a bogus keepalive message assuming that the mobility binding is in the "half bound" state. This opens up a similar redirection attack as discussed in Section 6.1.1. Note, however, that the attacker does not need to be able to modify packets in flight; simply being able to observe the Registration Request/Reply message exchange is sufficient to mount the attack.

With this in mind, the home agent MUST NOT accept a keepalive message from a different source IP address than where the Registration Request came from, as specified in Section 4.6. This requirement limits the extent of the attack to redirecting the traffic to a bogus UDP port, while the IP address must remain the same as in the initial Registration Request.

The only defenses against this vulnerability are: (1) to have a short time between re-registrations, which limits the duration of the redirection attack after the attacker has stopped sending bogus keepalive messages, and (2) to minimize the time the binding is in a "half bound" state by having the mobile node send the first keepalive message immediately after receiving an affirmative registration reply.

6.2 Use of IPsec

If the intermediate network is considered insecure, it is recommended that IPsec be used to protect user data traffic. However, IPsec does not protect against the redirection attacks described previously, other than to protect confidentiality of hijacked user data traffic.

The NAT traversal mechanism described in this document allows all IPsec-related traffic to go through NATs without any modifications to IPsec. In addition, the IPsec security associations do not need to be re-established when the mobile node moves.

6.3 Firewall Considerations

This document does not specify a general firewall traversal mechanism. However, the mechanism makes it possible to use only a single address and a port for all MN-HA (or FA-HA) communication. Furthermore, using the same port for the MIP UDP tunnelled traffic as for control messages makes it quite probable that if a MIP registration can reach the home agent, MIP tunnelling and reverse tunnelling using the described mechanism will also work.

7. UNSAF Considerations

The mechanism described in this document is not an "UNilateral Self-Address Fixing" (UNSAF) mechanism. Although the mobile node makes no attempt to determine or use the NAT translated address, the mobile node through the registration process does attempt to keep the NAT mapping alive through refresh messages. This section attempts to address issues that may be raised through this usage through the framework of the unsaf considerations IAB document [13].

1. Precise definition.

This proposal extends the Mobile IP v4 registration process to work across intervening NATs. The Home Agent detects the presence of the NAT by examining the source address in the packet header and comparing it with the address contained in the registration message.

The NAT address and port detected by the home agent are not exported or communicated to any other node anywhere.

2. Exit strategy.

This mechanism will go out of use as IPv6 and Mobile IP v6 is deployed, obviating the need for MIPv4 NAT traversal.

It can also be noted that this mechanism makes no changes to the base MIPv4 protocol which makes it dependent on the presence of NATs or the current extensions - i.e., no additional protocol changes would be needed if NATs were to go away.

3. Issues making systems more brittle.

The specific issue which is relevant here is that the effective care-of address (being the source address in the IP header received by the HA) is not protected by the Mobile IP authentication extension, and therefore may be spoofed. This is discussed in some detail in Section 6, Security Considerations.

4. Requirements for longer term solutions.

The trivial long term solution is a transition to an environment where NATs are not required. The most obvious such environment would be an IPv6 based internet.

In the presence of NATs, an improved solution would require

- * the ability to discover the translations done by each NAT along the route
- * the ability to validate the authority of each NAT to do those translations
- * communicating as part of the signed registration request the address of the NAT closest to the HA, for use as the effective care-of address from the viewpoint of the HA.
- * configuration of all intermediate NATs to accept only packets from the neighbour NATs.

5. Impact on existing, deployed NATs.

One precursor of the mechanism described here has been used successfully across deployed NATs in Sweden, Germany, England, Japan and the USA, without necessitating neither adjustments of the NATs in question, nor adjustment of any protocol parameters. At the time of writing, little experience exist with the exact implementation proposed in this document, but effort has been put into making this mechanism even more robust and adaptive than its precursors.

With respect to the base Mobile IP specification, the impact of this document is that an increased frequency of registration requests is recommended from a security perspective when the NAT traversal mechanism is used.

8. IANA Considerations

The numbers for the extensions defined in this document have been taken from the numbering space defined for Mobile IP messages, registration extensions and error codes defined in RFC 3344 [10]. This document proposes one new message, two new extensions and a new error code that require type numbers and an error code value that have been assigned by IANA. The two new extensions also introduce two new sub-type numbering spaces to be managed by IANA.

Section 3.1 defines a new Mobile IP extension, the UDP Tunnel Request Extension. The type number for this extension is 144. This extension introduces a new sub-type numbering space where the value 0 has been assigned to this extension. Approval of new Tunnel Request Extension sub-type numbers is subject to Expert Review, and a specification is required [7].

Section 3.2 defines a new Mobile IP extension, the UDP Tunnel Reply Extension. The type value for this extension is 44. This extension introduces a new sub-type numbering space where the value 0 has been assigned to this extension. Approval of new Tunnel Reply Extension sub-type numbers is subject to Expert Review, and a specification is required [7].

Section 3.3 defines a new Mobile IP message, the Tunnel Data message. The type value for this message is 4.

Section 3.5 defines a new error code, `ERROR_HA_UDP_ENCAP_UNAVAIL`: "Requested UDP tunnel encapsulation unavailable", from the numbering space for values defined for use with the Code field of Mobile IP Registration Reply Messages. Code number 142 has been assigned from the subset "Error Codes from the Home Agent".

The values for the Next Header field in the MIP Tunnel Data Message (Section 3.3) shall be the same as those used for the Protocol field of the IP header [2], and requires no new number assignment.

9. Intellectual Property Rights

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights (www.ietf.org/ipr.html).

10. Acknowledgements

Much of the text in Section 4.2 has been taken almost verbatim from RFC 2003, IP Encapsulation within IP [4].

Adding support for the FA case was suggested by George Tsirtsis and Frode B. Nilsen. Roy Jose pointed out a problem with binding updates, and Frode, Roy and George pointed out that there are cases where triangular routes may work, and suggested that reverse tunnelling need not be mandatory. Roy and Qiang Zhang drew attention to a number of sections which needed to be corrected or stated more clearly.

Phil Roberts helped remove a number of rough edges. Farid Adrangi pointed out the DoS issue now covered in Security Considerations (Section 6). Francis Dupont's helpful comments made us extend the Security Considerations section to make it more comprehensive and clear. Milind Kulkarni and Madhavi Chandra pointed out the required match between the FA source and care-of addresses when the mechanism is used by an FA, and also contributed a number of clarifications to the text.

Thanks also to our co-workers, Ilkka Pietikainen, Antti Nuopponen and Timo Aalto at Netseal and Hans Sjostrand, Fredrik Johansson and Erik Liden at ipUnplugged. They have read and re-read the text, and contributed many valuable corrections and insights.

11. Normative References

- [1] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.
- [2] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [3] Hanks, S., Li, T., Farinacci, D. and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 1701, October 1994.
- [4] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [5] Perkins, C., "Minimal Encapsulation within IP", RFC 2004, October 1996.
- [6] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [7] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

- [8] Farinacci, D., Li, T., Hanks, S., Meyer, D. and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [9] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [10] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.

12. Informative References

- [11] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [12] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [13] Daigle, L., Editor, and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF)", RFC 3424, November 2002.

13. Authors' Addresses

Henrik Levkowetz
ipUnplugged AB
Arenavagen 23
Stockholm S-121 28
SWEDEN

Phone: +46 708 32 16 08
EMail: henrik@levkowetz.com

Sami Vaarala
Netseal
Niittykatu 6
Espoo 02201
FINLAND

Phone: +358 9 435 310
EMail: sami.vaarala@iki.fi

14. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

