

Mobile IPv4 Traversal across IPsec-Based VPN Gateways

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.

Abstract

This document outlines a solution for the Mobile IPv4 (MIPv4) and IPsec coexistence problem for enterprise users. The solution consists of an applicability statement for using Mobile IPv4 and IPsec for session mobility in corporate remote access scenarios, and a required mechanism for detecting the trusted internal network securely.

Table of Contents

1.	Introduction	3
1.1.	Overview	3
1.2.	Scope	5
1.3.	Related Work	5
1.4.	Terms and Abbreviations	5
1.5.	Requirement Levels	6
1.6.	Assumptions and Rationale	7
1.7.	Why IPsec Lacks Mobility	8
2.	The Network Environment	9
2.1.	Access Mode: 'c'	12
2.2.	Access Mode: 'f'	13
2.3.	Access Mode: 'cvc'	13
2.4.	Access Mode: 'fvc'	14
2.5.	NAT Traversal	14
3.	Internal Network Detection	15
3.1.	Assumptions	16
3.2.	Implementation Requirements	16
3.2.1.	Separate Tracking of Network Interfaces	16
3.2.2.	Connection Status Change	16
3.2.3.	Registration-Based Internal Network Detection	17

3.2.4.	Registration-Based Internal Network Monitoring	17
3.3.	Proposed Algorithm	19
3.4.	Trusted Networks Configured (TNC) Extension	20
3.5.	Implementation Issues	20
3.6.	Rationale for Design Choices	21
3.6.1.	Firewall Configuration Requirements	21
3.6.2.	Registration-Based Internal Network Monitoring	22
3.6.3.	No Encryption When Inside	22
3.7.	Improvements	22
4.	Requirements	23
4.1.	Mobile Node Requirements	23
4.2.	VPN Device Requirements	23
4.3.	Home Agent Requirements	24
5.	Analysis	24
5.1.	Comparison against Guidelines	24
5.2.	Packet Overhead	26
5.3.	Latency Considerations	27
5.4.	Firewall State Considerations	27
5.5.	Intrusion Detection Systems (IDSs)	28
5.6.	Implementation of the Mobile Node	28
5.7.	Non-IPsec VPN Protocols	28
6.	Security Considerations	29
6.1.	Internal Network Detection	29
6.2.	Mobile IPv4 versus IPsec	30
7.	IANA Considerations	31
8.	Acknowledgements	31
9.	References	32
9.1.	Normative References	32
9.2.	Informative References	33
Appendix A.	Packet Flow Examples	34
A.1.	Connection Setup for Access Mode 'cvc'	34

1. Introduction

The Mobile IP working group set out to explore the problem and solution spaces of IPsec and Mobile IP coexistence. The problem statement and solution requirements for Mobile IPv4 case were first documented in [RFC4093]. This document outlines a solution for IPv4.

The document contains two parts:

- o a basic solution that is an applicability statement of Mobile IPv4 and IPsec to provide session mobility between enterprise intranets and external networks, intended for enterprise mobile users; and
- o a technical specification and a set of requirements for secure detection of the internal and the external networks, including a new extension that must be implemented by a mobile node and a home agent situated inside the enterprise network.

There are many useful ways to combine Mobile IPv4 and IPsec. The solution specified in this document is most applicable when the assumptions documented in the problem statement [RFC4093] are valid; among others that the solution:

- o must minimize changes to existing firewall/VPN/DMZ (DeMilitarized Zone) deployments;
- o must ensure that traffic is not routed through the DMZ when the mobile node is inside (to avoid scalability and management issues);
- o must support foreign networks with only foreign agent access;
- o should not require changes to existing IPsec or key exchange protocols;
- o must comply with the Mobile IPv4 protocol (but may require new extensions or multiple instances of Mobile IPv4); and
- o must propose a mechanism to avoid or minimize IPsec re-negotiation when the mobile node moves.

1.1. Overview

Typical corporate networks consist of three different domains: the Internet (untrusted external network), the intranet (trusted internal network), and the DMZ, which connects the two networks. Access to the internal network is guarded both by a firewall and a VPN device;

access is only allowed if both firewall and VPN security policies are respected.

Enterprise mobile users benefit from unrestricted seamless session mobility between subnets, regardless of whether the subnets are part of the internal or the external network. Unfortunately, the current Mobile IPv4 and IPsec standards alone do not provide such a service [tessier].

The solution is to use standard Mobile IPv4 (except for a new extension used by the home agent in the internal network to aid in network detection) when the mobile node is in the internal network, and to use the VPN tunnel endpoint address for the Mobile IPv4 registration when outside. IPsec-based VPN tunnels require re-negotiation after movement. To overcome this limitation, another layer of Mobile IPv4 is used underneath IPsec, in effect making IPsec unaware of movement. Thus, the mobile node can freely move in the external network without disrupting the VPN connection.

Briefly, when outside, the mobile node:

- o detects that it is outside (Section 3);
- o registers its co-located or foreign agent care-of address with the external home agent;
- o establishes a VPN tunnel using, e.g., Internet Key Exchange Protocol (IKE) (or IKEv2) if security associations are not already available;
- o registers the VPN tunnel address as its co-located care-of address with the internal home agent; this registration request is sent inside the IPsec tunnel.

The solution requires control over the protocol layers in the mobile node. It must be capable of (1) detecting whether it is inside or outside in a secure fashion, and (2) controlling the protocol layers accordingly. For instance, if the mobile node is inside, the IPsec layer needs to become dormant.

Except for the new Mobile IPv4 extension to improve security of internal network detection, current Mobile IPv4 and IPsec standards, when used in a suitable combination, are sufficient to implement the solution. No changes are required to existing VPN devices or foreign agents.

The solution described is compatible with different kinds of IPsec-based VPNs, and no particular kind of VPN is required. Because the

appropriate Security Policy Database (SPD) entries and other IKE and IPsec specifics differ between deployed IPsec-based VPN products, these details are not discussed in the document.

1.2. Scope

This document describes a solution for IPv4 only. The downside of the described approach is that an external home agent is required and that the packet overhead (see Section 5) and overall complexity increase. Optimizations would require significant changes to Mobile IPv4 and/or IPsec, and are out of scope of this document.

VPN, in this document, refers to an IPsec-based remote access VPN. Other types of VPNs are out of scope.

1.3. Related Work

Related work has been done on Mobile IPv6 in [RFC3776], which discusses the interaction of IPsec and Mobile IPv6 in protecting Mobile IPv6 signaling. The document also discusses dynamic updating of the IPsec endpoint based on Mobile IP signaling packets.

The "transient pseudo-NAT" attack, described in [pseudonat] and [mipnat], affects any approach that attempts to provide security of mobility signaling in conjunction with NAT devices. In many cases, one cannot assume any cooperation from NAT devices, which thus have to be treated as any other networking entity.

The IKEv2 Mobility and Multihoming Protocol (MOBIKE) [RFC4555] provides better mobility for IPsec. This would allow the external Mobile IPv4 layer described in this specification to be removed. However, deploying MOBIKE requires changes to VPN devices, and is thus out of scope of this specification.

1.4. Terms and Abbreviations

co-CoA: co-located care-of address.

DMZ: (DeMilitarized Zone) a small network inserted as a "neutral zone" between a company's private network and the outside public network to prevent outside users from getting direct access to the company's private network.

external network: the untrusted network (i.e., Internet). Note that a private network (e.g., another corporate network) other than the mobile node's internal network is considered an external network.

FA: mobile IPv4 foreign agent.

FA-CoA: foreign agent care-of address.

FW: firewall.

internal network: the trusted network; for instance, a physically secure corporate network where the i-HA is located.

i-FA: Mobile IPv4 foreign agent residing in the internal network.

i-HA: Mobile IPv4 home agent residing in the internal network; typically has a private address [privaddr].

i-HoA: home address of the mobile node in the internal home agent.

MN: mobile node.

NAI: Network Access Identifier [RFC4282].

R: router.

VPN: Virtual Private Network based on IPsec.

VPN-TIA: VPN tunnel inner address, the address(es) negotiated during IKE phase 2 (quick mode), assigned manually, using IPsec-DHCP [RFC3456], using the "de facto" standard Internet Security Association and Key Management Protocol (ISAKMP) configuration mode, or by some other means. Some VPN clients use their current care-of address as their Tunnel Inner Address (TIA) for architectural reasons.

VPN tunnel: an IPsec-based tunnel; for instance, IPsec tunnel mode IPsec connection, or Layer 2 Tunneling Protocol (L2TP) combined with IPsec transport connection.

x-FA: Mobile IPv4 foreign agent residing in the external network.

x-HA: Mobile IPv4 home agent residing in the external network.

x-HoA: home address of the mobile node in the external home agent.

1.5. Requirement Levels

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 [RFC2119].

1.6. Assumptions and Rationale

The solution is an attempt to solve the problem described in [RFC4093]. The major assumptions and their rationale is summarized below.

Changes to existing firewall and VPN deployments should be minimized:

- o The current deployment of firewalls and IPsec-based VPNs is much larger than corresponding Mobile IPv4 elements. Thus, a solution should work within the existing VPN infrastructure.
- o Current enterprise network deployments typically centralize management of security and network access into a compact DMZ.

When the mobile node is inside, traffic should not go through the DMZ network:

- o Routing all mobile node traffic through the DMZ is seen as a performance problem in existing deployments of firewalls. The more sophisticated firewall technology is used (e.g., content scanning), the more serious the performance problem is.
- o Current deployments of firewalls and DMZs in general have been optimized for the case where only a small minority of total enterprise traffic goes through the DMZ. Furthermore, users of current VPN remote access solutions do not route their traffic through the DMZ when connected to an internal network.

A home agent inside the enterprise cannot be reached directly from outside, even if the home agent contains IPsec functionality:

- o Deployment of current combined IPsec/MIPv4 solutions are not common in large installations.
- o Doing decryption in the home agents "deep inside" the enterprise effectively means having a security perimeter much larger than the typical, compact DMZ used by a majority of enterprises today.
- o In order to maintain a security level equal to current firewall/DMZ deployments, every home agent decapsulating IPsec would need to do the same firewalling as the current DMZ firewalls (content scanning, connection tracking, etc.).

Traffic cannot be encrypted when the mobile node is inside:

- o There is a considerable performance impact on home agents (which currently do rather light processing) and mobile nodes (especially for small devices). Note that traffic throughput inside the enterprise is typically an order (or more) of magnitude larger than the remote access traffic through a VPN.
- o Encryption consumes processing power and has a significant impact on device battery life.
- o There is also a usability issue involved; the user needs to authenticate the connection to the IPsec layer in the home agent to gain access. For interactive authentication mechanisms (e.g., SecurID), this always means user interaction.
- o Furthermore, if there is a separate VPN device in the DMZ for remote access, the user needs to authenticate to both devices, and might need to have separate credentials for both.
- o Current Mobile IPv4 home agents do not typically incorporate IPsec functionality, which is relevant for the solution when we assume zero or minimal changes to existing Mobile IPv4 nodes.
- o Note, however, that the assumption (no encryption when inside) does not necessarily apply to all solutions in the solution space; if the above mentioned problems were resolved, there is no fundamental reason why encryption could not be applied when inside.

1.7. Why IPsec Lacks Mobility

IPsec, as currently specified [RFC4301], requires that a new IKE negotiation be done whenever an IPsec peer moves, i.e., changes care-of address. The main reason is that a security association is unidirectional and identified by a triplet consisting of (1) the destination address (which is the outer address when tunnel mode is used), (2) the security protocol (Encapsulating Security Payload (ESP) or Authentication Header (AH)), and (3) the Security Parameter Index (SPI) ([RFC4301], Section 4.1). Although an implementation is not required to use all of these for its own Security Associations (SAs), an implementation cannot assume that a peer does not.

When a mobile IPsec peer sends packets to a stationary IPsec peer, there is no problem; the SA is "owned" by the stationary IPsec peer, and therefore the destination address does not need to change. The (outer) source address should be ignored by the stationary peer (although some implementations do check the source address as well).

The problem arises when packets are sent from the stationary peer to the mobile peer. The destination address of this SA (SAs are unidirectional) is established during IKE negotiation, and is effectively the care-of address of the mobile peer at time of negotiation. Therefore, the packets will be sent to the original care-of address, not a changed care-of address.

The IPsec NAT traversal mechanism can also be used for limited mobility, but UDP tunneling needs to be used even when there is no NAT in the route between the mobile and the stationary peers. Furthermore, support for changes in current NAT mapping is not required by the NAT traversal specification [RFC3947].

In summary, although the IPsec standard does not as such prevent mobility (in the sense of updating security associations on-the-fly), the standard does not include a built-in mechanism (explicit or implicit) for doing so. Therefore, it is assumed throughout this document that any change in the addresses comprising the identity of an SA requires IKE re-negotiation, which implies too heavy computation and too large latency for useful mobility.

The IKEv2 Mobility and Multihoming Protocol (MOBIKE) [RFC4555] provides better mobility for IPsec. This would allow the external Mobile IPv4 layer described in this specification to be removed. However, deploying MOBIKE requires changes to VPN devices, and is thus out of scope of this specification.

2. The Network Environment

Enterprise users will access both the internal and external networks using different networking technologies. In some networks, the MN will use FAs and in others it will anchor at the HA using co-located mode. The following figure describes an example network topology illustrating the relationship between the internal and external networks, and the possible locations of the mobile node (i.e., (MN)).

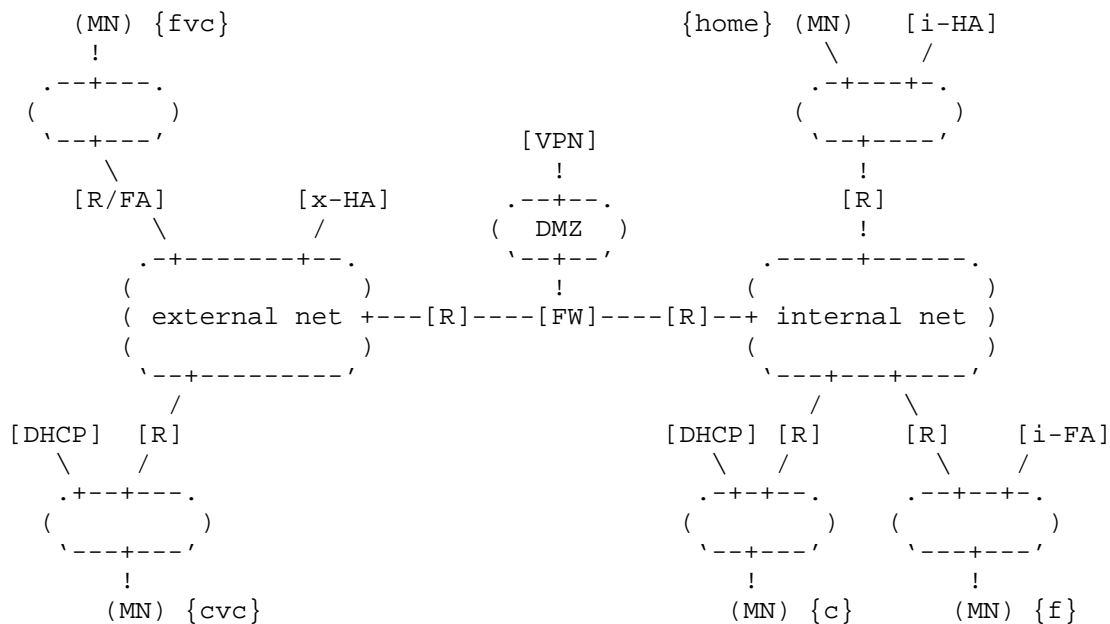


Figure 1: Basic topology, possible MN locations, and access modes

In every possible location described in the figure, the mobile node can establish a connection to the corresponding HA(s) by using a suitable "access mode". An access mode is here defined to consist of:

1. a composition of the mobile node networking stack (i-MIP or x-MIP/VPN/i-MIP); and
2. registration mode(s) of i-MIP and x-MIP (if used); i.e., co-located care-of address or foreign agent care-of address.

Each possible access mode is encoded as "xyz", where:

- o "x" indicates whether the x-MIP layer is used, and if used, the mode ("f" indicates FA-CoA, "c" indicates co-CoA, absence indicates not used);
- o "y" indicates whether the VPN layer is used ("v" indicates VPN used, absence indicates not used); and
- o "z" indicates mode of i-MIP layer ("f" indicates FA-CoA, "c" indicates co-CoA).

This results in four access modes:

```
c: i-MIP with co-CoA
f: i-MIP with FA-CoA
cvc: x-MIP with co-CoA, VPN-TIA as i-MIP co-CoA
fvc: x-MIP with FA-CoA, VPN-TIA as i-MIP co-CoA
```

This notation is more useful when optimizations to protocol layers are considered. The notation is preserved here so that work on the optimizations can refer to a common notation.

The internal network is typically a multi-subnetted network using private addressing [privaddr]. Subnets may contain internal home agent(s), DHCP server(s), and/or foreign agent(s). Current IEEE 802.11 wireless LANs are typically deployed in the external network or the DMZ because of security concerns.

The figure leaves out a few details worth noticing:

- o There may be multiple NAT devices anywhere in the diagram.
 - * When the MN is outside, the NAT devices may be placed between the MN and the x-HA or the x-HA and the VPN.
 - * There may also be NAT(s) between the VPN and the i-HA, or a NAT integrated into the VPN. In essence, any router in the figure may be considered to represent zero or more routers, each possibly performing NAT and/or ingress filtering.
 - * When the MN is inside, there may be NAT devices between the MN and the i-HA.
- o Site-to-site VPN tunnels are not shown. Although mostly transparent, IPsec endpoints may perform ingress filtering as part of enforcing their policy.
- o The figure represents a topology where each functional entity is illustrated as a separate device. However, it is possible that several network functions are co-located in a single device. In fact, all three server components (x-HA, VPN, and i-HA) may be co-located in a single physical device.

The following issues are also important when considering enterprise mobile users:

- o Some firewalls are configured to block ICMP messages and/or fragments. Such firewalls (routers) cannot be detected reliably.

- o Some networks contain transparent application proxies, especially for HTTP. Like firewalls, such proxies cannot be detected reliably in general. IPsec and Mobile IPv4 are incompatible with such networks.

Whenever a mobile node obtains either a co-CoA or an FA-CoA, the following conceptual steps take place:

- o The mobile node detects whether the subnet where the care-of address was obtained belongs to the internal or the external network using the method described in Section 3 (or a vendor-specific mechanism fulfilling the requirements described).
- o The mobile node performs necessary registrations and other connection setup signaling for the protocol layers (in the following order):
 - * x-MIP (if used);
 - * VPN (if used); and
 - * i-MIP.

Note that these two tasks are intertwined to some extent: detection of the internal network results in a successful registration to the i-HA using the proposed network detection algorithm. An improved network detection mechanism not based on Mobile IPv4 registration messages might not have this side effect.

The following subsections describe the different access modes and the requirements for registration and connection setup phase.

2.1. Access Mode: 'c'

This access mode is standard Mobile IPv4 [RFC3344] with a co-located address, except that:

- o the mobile node MUST detect that it is in the internal network; and
- o the mobile node MUST re-register periodically (with a configurable interval) to ensure it is still inside the internal network (see Section 4).

2.2. Access Mode: 'f'

This access mode is standard Mobile IPv4 [RFC3344] with a foreign agent care-of address, except that

- o the mobile node MUST detect that it is in the internal network; and
- o the mobile node MUST re-register periodically (with a configurable interval) to ensure it is still inside the internal network (see Section 4).

2.3. Access Mode: 'cvc'

Steps:

- o The mobile node obtains a care-of address.
- o The mobile node detects it is not inside and registers with the x-HA, where
 - * T-bit MAY be set (reverse tunneling), which minimizes the probability of firewall-related connectivity problems
- o If the mobile node does not have an existing IPsec security association, it uses IKE to set up an IPsec security association with the VPN gateway, using the x-HoA as the IP address for IKE/IPsec communication. How the VPN-TIA is assigned is outside the scope of this document.
- o The mobile node sends a MIPv4 Registration Request (RRQ) to the i-HA, registering the VPN-TIA as a co-located care-of address, where
 - * T-bit SHOULD be set (reverse tunneling) (see discussion below)

Reverse tunneling in the inner Mobile IPv4 layer is often required because of IPsec security policy limitations. IPsec selectors define allowed IP addresses for packets sent inside the IPsec tunnel. Typical IPsec remote VPN selectors restrict the client address to be VPN-TIA (remote address is often unrestricted). If reverse tunneling is not used, the source address of a packet sent by the MN will be the MN's home address (registered with i-HA), which is different from the VPN-TIA, thus violating IPsec security policy. Consequently, the packet will be dropped, resulting in a connection black hole.

Some types of IPsec-based VPNs, in particular L2TP/IPsec VPNs (PPP-over-L2TP-over-IPsec), do not have this limitation and can use triangular routing.

Note that although the MN can use triangular routing, i.e., skip the inner MIPv4 layer, it MUST NOT skip the VPN layer for security reasons.

2.4. Access Mode: 'fvc'

Steps:

- o The mobile node obtains a foreign agent advertisement from the local network.
- o The mobile node detects it is outside and registers with the x-HA, where
 - * T-bit MAY be set (reverse tunneling), which minimizes the probability of firewall-related connectivity problems
- o If necessary, the mobile node uses IKE to set up an IPsec connection with the VPN gateway, using the x-HoA as the IP address for IKE/IPsec communication. How the VPN-TIA is assigned is outside the scope of this document.
- o The mobile node sends a MIPv4 RRQ to the i-HA, registering the VPN-TIA as a co-located care-of address, where
 - * T-bit SHOULD be set (reverse tunneling) (see discussion in Section 2.3)

Note that although the MN can use triangular routing, i.e., skip the inner MIPv4 layer, it MUST NOT skip the VPN layer for security reasons.

2.5. NAT Traversal

NAT devices may affect each layer independently. Mobile IPv4 NAT traversal [mipnat] SHOULD be supported for x-MIP and i-MIP layers, while IPsec NAT traversal [RFC3947][RFC3948] SHOULD be supported for the VPN layer.

Note that NAT traversal for the internal MIPv4 layer may be necessary even when there is no separate NAT device between the VPN gateway and the internal network. Some VPN implementations NAT VPN tunnel inner addresses before routing traffic to the intranet. Sometimes this is done to make a deployment easier, but in some cases this approach

makes VPN client implementation easier. Mobile IPv4 NAT traversal is required to establish a MIPv4 session in this case.

3. Internal Network Detection

Secure detection of the internal network is critical to prevent plaintext traffic from being sent over an untrusted network. In other words, the overall security (confidentiality and integrity of user data) relies on the security of the internal network detection mechanism in addition to IPsec. For this reason, security requirements are described in this section.

In addition to detecting entry into the internal network, the mobile node must also detect when it has left the internal network. Entry into the internal network is easier security-wise: the mobile node can ensure that it is inside the internal network before sending any plaintext traffic. Exit from the internal network is more difficult to detect, and the MN may accidentally leak plaintext packets if the event is not detected in time.

Several events can cause the mobile node to leave the internal network, including:

- o a routing change upstream;
- o a reassociation of 802.11 on layer 2 that the mobile node software does not detect;
- o a physical cable disconnect and reconnect that the mobile node software does not detect.

Whether the mobile node can detect such changes in the current connection reliably depends on the implementation and the networking technology. For instance, some mobile nodes may be implemented as pure layer three entities. Even if the mobile node software has access to layer 2 information, such information is not trustworthy security-wise, and depends on the network interface driver.

If the mobile node does not detect these events properly, it may leak plaintext traffic into an untrusted network. A number of approaches can be used to detect exit from the internal network, ranging from frequent re-registration to the use of layer two information.

A mobile node MUST implement a detection mechanism fulfilling the requirements described in Section 3.2; this ensures that basic security requirements are fulfilled. The basic algorithm described in Section 3.3 is one way to do that, but alternative methods may be used instead or in conjunction. The assumptions that the

requirements and the proposed mechanism rely upon are described in Section 3.1.

3.1. Assumptions

The enterprise firewall MUST be configured to block traffic originating from external networks going to the i-HA. In other words, the mobile node MUST NOT be able to perform a successful Registration Request/Registration Reply (RRQ/RRP) exchange (without using IPsec) unless it is connected to the trusted internal network; the mobile node can then stop using IPsec without compromising data confidentiality.

If this assumption does not hold, data confidentiality is compromised in a potentially silent and thus dangerous manner. To minimize the impact of this scenario, the i-HA is also required to check the source address of any RRQ to determine whether it comes from a trusted (internal network) address. The i-HA needs to indicate to the MN that it supports the checking of trusted source addresses by including a Trusted Networks Configured extension in its registration reply. This new extension, which needs to be implemented by both i-HA and the MN, is described in Section 3.4.

The firewall MAY be configured to block registration traffic to the x-HA originating from within the internal network, which makes the network detection algorithm simpler and more robust. However, as the registration request is basically UDP traffic, an ordinary firewall (even a stateful one) would typically allow the registration request to be sent and a registration reply to be received through the firewall.

3.2. Implementation Requirements

Any mechanism used to detect the internal network MUST fulfill the requirements described in this section. An example of a network detection mechanism fulfilling these requirements is given in Section 3.3.

3.2.1. Separate Tracking of Network Interfaces

The mobile node implementation MUST track each network interface separately. Successful registration with the i-HA through interface X does not imply anything about the status of interface Y.

3.2.2. Connection Status Change

When the mobile node detects that its connection status on a certain network interface changes, the mobile node MUST:

- o immediately stop relaying user data packets;
- o detect whether this interface is connected to the internal or the external network; and
- o resume data traffic only after the internal network detection and necessary registrations and VPN tunnel establishment have been completed.

The mechanisms used to detect a connection status change depends on the mobile node implementation, the networking technology, and the access mode.

3.2.3. Registration-Based Internal Network Detection

The mobile node MUST NOT infer that an interface is connected to the internal network unless a successful registration has been completed through that particular interface to the i-HA, the i-HA registration reply contained a Trusted Networks Configured extension (Section 3.4), and the connection status of the interface has not changed since.

3.2.4. Registration-Based Internal Network Monitoring

Some leak of plaintext packets to a (potentially) untrusted network cannot always be completely prevented; this depends heavily on the client implementation. In some cases, the client cannot detect such a change, e.g., if upstream routing is changed.

More frequent re-registrations when the MN is inside is a simple way to ensure that the MN is still inside. The MN SHOULD start re-registration every (T_MONITOR - N) seconds when inside, where N is a grace period that ensures that re-registration is completed before T_MONITOR seconds are up. To bound the maximum amount of time that a plaintext leak may persist, the mobile node must fulfill the following security requirements when inside:

- o The mobile node MUST NOT send or receive a user data packet if more than T_MONITOR seconds have elapsed since the last successful (re-)registration with the i-HA.
- o If more than T_MONITOR seconds have elapsed, data packets MUST be either dropped or queued. If the packets are queued, the queues MUST NOT be processed until the re-registration has been successfully completed without a connection status change.

- o The T_MONITOR parameter MUST be configurable, and have the default value of 60 seconds. This default is a trade-off between traffic overhead and a reasonable bound to exposure.

This approach is reasonable for a wide range of mobile nodes (e.g., laptops), but has unnecessary overhead when the mobile node is idle (not sending or receiving packets). If re-registration does not complete before T_MONITOR seconds are up, data packets must be queued or dropped as specified above. Note that re-registration packets MUST be sent even if bidirectional user data traffic is being relayed: data packets are no substitute for an authenticated re-registration.

To minimize traffic overhead when the mobile node is idle, re-registrations can be stopped when no traffic is being sent or received. If the mobile node subsequently receives or needs to send a packet, the packet must be dropped or queued (as specified above) until a re-registration with the i-HA has been successfully completed. Although this approach adds packet processing complexity, it may be appropriate for small, battery-powered devices, which may be idle much of the time. (Note that ordinary re-registration before the mobility binding lifetime is exhausted should still be done to keep the MN reachable.)

T_MONITOR is required to be configurable so that an administrator can determine the required security level for the particular deployment. Configuring T_MONITOR in the order of a few seconds is not practical; alternative mechanisms need to be considered if such confidence is required.

The re-registration mechanism is a worst-case fallback mechanism. If additional information (such as layer two triggers) is available to the mobile node, the mobile node SHOULD use the triggers to detect MN movement and restart the detection process to minimize exposure.

Note that re-registration is required by Mobile IPv4 by default (except for the atypical case of an infinite binding lifetime); however, the re-registration interval may be much larger when using an ordinary Mobile IPv4 client. A shorter re-registration interval is usually not an issue, because the internal network is typically a fast, wired network, and the shortened re-registration interval applies only when the mobile node is inside the internal network. When outside, the ordinary Mobile IPv4 re-registration process (based on binding lifetime) is used.

3.3. Proposed Algorithm

When the MN detects that it has changed its point of network attachment on a certain interface, it issues two simultaneous registration requests, one to the i-HA and another to the x-HA. These registration requests are periodically retransmitted if reply messages are not received.

Registration replies are processed as follows:

- o If a response from the x-HA is received, the MN stops retransmitting its registration request to the x-HA and tentatively determines it is outside. However, the MN MUST keep on retransmitting its registration to the i-HA for a period of time. The MN MAY postpone the IPsec connection setup for some period of time while it waits for a (possible) response from the i-HA.
- o If a response from the i-HA is received and the response contains the Trusted Networks Configured extension (Section 3.4), the MN SHOULD determine that it is inside. In any case, the MN MUST stop retransmitting its registration requests to both i-HA and x-HA.
- o When successfully registered with the i-HA directly, MN SHOULD de-register with the x-HA.

If the MN ends up detecting that it is inside, it MUST re-register periodically (regardless of binding lifetime); see Section 3.2.4. If the re-registration fails, the MN MUST stop sending and receiving plaintext traffic, and MUST restart the detection algorithm.

Plaintext re-registration messages are always addressed either to the x-HA or the i-HA, not to both. This is because the MN knows, after initial registration, whether it is inside or outside. (However, when the mobile node is outside, it re-registers independently with the x-HA using plaintext, and with the i-HA through the VPN tunnel.)

Postponing the IPsec connection setup could prevent aborted IKE sessions. Aborting IKE sessions may be a problem in some cases because IKE does not provide a reliable, standardized, and mandatory-to-implement mechanism for terminating a session cleanly.

If the x-HA is not reachable from inside (i.e., the firewall configuration is known), a detection period of zero is preferred, as it minimizes connection setup overhead and causes no timing problems. Should the assumption have been invalid and a response from the i-HA received after a response from the x-HA, the MN SHOULD re-register with the i-HA directly.

- o a domain name in a DHCPDISCOVER / DHCPPOFFER message
- o an NAI in a foreign agent advertisement
- o a list of default gateway MAC addresses that are known to reside in the internal network (i.e., configured as such, or have been previously verified to be inside)

For instance, if the MN has reason to believe it is inside, it MAY postpone sending a registration request to the x-HA for some time. Similarly, if the MN has reason to believe it is outside, it may start IPsec connection setup immediately after receiving a registration reply from the x-HA. However, should the MN receive a registration reply from the i-HA after IPsec connection setup has been started, the MN SHOULD still switch to using the i-HA directly.

3.6. Rationale for Design Choices

3.6.1. Firewall Configuration Requirements

The requirement that the i-HA cannot be reached from the external network is necessary. If not, a successful registration with the i-HA (without IPsec) cannot be used as a secure indication that the mobile node is inside. A possible solution to the obvious security problem would be to define and deploy a secure internal network detection mechanism based on, e.g., signed FA advertisement or signed DHCP messages.

However, unless the mechanism is defined for both FA and DHCP messages and is deployed in every internal network, it has limited applicability. In other words, the mobile node MUST NOT assume it is in the internal network unless it receives a signed FA or DHCP message (regardless of whether or not it can register directly with the i-HA). If it receives an unsigned FA or DHCP message, it MUST use IPsec; otherwise, the mobile node can be easily tricked into using plaintext.

Assuming that all FA and DHCP servers in the internal network are upgraded to support such a feature does not seem realistic; it is highly desirable to be able to take advantage of existing DHCP and FA deployments. Similar analysis seems to apply regardless of what kind of additional security mechanism is defined.

Because a firewall configuration error can have catastrophic data security consequences (silent exposure of user data to external attackers), a separate protection mechanism is provided by the i-HA. The i-HA must be configured, by the administrator, with a list of trusted networks. The i-HA advertises that it knows which

registration request source addresses are trusted, using a registration reply extension (Trusted Networks Configured extension, Section 3.4). Without this extension, an MN may not rely on a successful registration to indicate that it is connected to the internal network. This ensures that user data compromise does not occur unless both the firewall and the i-HA are configured incorrectly. Further, occurrences of registration requests from untrusted addresses should be logged by the i-HA, exposing them to administrator review.

3.6.2. Registration-Based Internal Network Monitoring

This issue also affects IPsec client security. However, as IPsec specifications take no stand on how and when client IPsec policies are configured or changed (for instance, in response to a change in network connectivity), the issue is out of scope for IPsec. Because this document describes an algorithm and requirements for (secure) internal network detection, the issue is in scope of the document.

The current requirement for internal network monitoring was added as a fallback mechanism.

3.6.3. No Encryption When Inside

If encryption was applied also when MN was inside, there would be no security reason to monitor the internal network periodically.

The main rationale for why encryption cannot be applied when the MN is inside was given in Section 1.6. In short, the main issues are (1) power consumption; (2) extra CPU load, especially because internal networks are typically switched networks and a lot of data may be routinely transferred; (3) existing HA devices do not typically integrate IPsec functionality; (4) (IPsec) encryption requires user authentication, which may be interactive in some cases (e.g., SecurID) and thus a usability issue; and (5) user may need to have separate credentials for VPN devices in the DMZ and the HA.

3.7. Improvements

The registration process can be improved in many ways. One simple way is to make the x-HA detect whether a registration request came from inside or outside the enterprise network. If it came from inside the enterprise network, the x-HA can simply drop the registration request.

This approach is feasible without protocol changes in scenarios where a corporation owns both the VPN and the x-HA. The x-HA can simply determine based on the incoming interface identifier (or the router

that relayed the packet) whether or not the registration request came from inside.

In other scenarios, protocol changes may be needed. Such changes are out of scope of this document.

4. Requirements

4.1. Mobile Node Requirements

The mobile node MUST implement an internal network detection algorithm fulfilling the requirements set forth in Section 3.2. A new configurable MN parameter, T_MONITOR, is required. The value of this parameter reflects a balance between security and the amount of signaling overhead, and thus needs to be configurable. In addition, when doing internal network detection, the MN MUST NOT disable IPsec protection unless the registration reply from the i-HA contains a Trusted Networks Configured extension (Section 3.4).

The mobile node MUST support access modes c, f, cvc, fvc (Section 2).

The mobile node SHOULD support Mobile IPv4 NAT traversal [mipnat] for both internal and external Mobile IP.

The mobile node SHOULD support IPsec NAT traversal [RFC3947] [RFC3948].

When the mobile node has direct access to the i-HA, it SHOULD use only the inner Mobile IPv4 layer to minimize firewall and VPN impact.

When the mobile node is outside and using the VPN connection, IPsec policies MUST be configured to encrypt all traffic sent to and from the enterprise network. The particular Security Policy Database (SPD) entries depend on the type and configuration of the particular VPN (e.g., plain IPsec vs. L2TP/IPsec, full tunneling or split tunneling).

4.2. VPN Device Requirements

The VPN security policy MUST allow communication using UDP to the internal home agent(s), with home agent port 434 and any remote port. The security policy SHOULD allow IP-IP to internal home agent(s) in addition to UDP port 434.

The VPN device SHOULD implement the IPsec NAT traversal mechanism described in [RFC3947] and [RFC3948].

4.3. Home Agent Requirements

The home agent SHOULD implement the Mobile IPv4 NAT traversal mechanism described in [mipnat]. (This also refers to the i-HA: NAT traversal is required to support VPNs that NAT VPN tunnel addresses or block IP-IP traffic.)

To protect user data confidentiality against firewall configuration errors, the i-HA:

- o MUST be configured with a list of trusted IP subnets (containing only addresses from the internal network), with no subnets being trusted by default.
- o MUST reject (drop silently) any registration request coming from a source address that is not inside any of the configured trusted subnets. These dropped registration requests SHOULD be logged.
- o MUST include a Trusted Networks Configured extension (Section 3.4) in a registration reply sent in response to a registration request coming from a trusted address.

5. Analysis

This section provides a comparison against guidelines described in Section 6 of the problem statement [RFC4093] and additional analysis of packet overhead with and without the optional mechanisms.

5.1. Comparison against Guidelines

Preservation of existing VPN infrastructure

- o The solution does not mandate any changes to existing VPN infrastructure, other than possibly changes in configuration to avoid stateful filtering of traffic.

Software upgrades to existing VPN clients and gateways

- o The solution described does not require any changes to VPN gateways or Mobile IPv4 foreign agents.

IPsec protocol

- o The solution does not require any changes to existing IPsec or key exchange standard protocols, and does not require implementation of new protocols in the VPN device.

Multi-vendor interoperability

- o The solution provides easy multi-vendor interoperability between server components (VPN device, foreign agents, and home agents). Indeed, these components need not be aware of each other.
- o The mobile node networking stack is somewhat complex to implement, which may be an issue for multi-vendor interoperability. However, this is a purely software architecture issue, and there are no known protocol limitations for multi-vendor interoperability.

MIPv4 protocol

- o The solution adheres to the MIPv4 protocol, but requires the new Trusted Networks Configured extension to improve the trustworthiness of internal network detection.
- o The solution requires the use of two parallel MIPv4 layers.

Handoff overhead

- o The solution provides a mechanism to avoid VPN tunnel SA renegotiation upon movement by using the external MIPv4 layer.

Scalability, availability, reliability, and performance

- o The solution complexity is linear with the number of MNs registered and accessing resources inside the intranet.
- o Additional overhead is imposed by the solution.

Functional entities

- o The solution does not impose any new types of functional entities or required changes to existing entities. However, an external HA device is required.

Implications of intervening NAT gateways

- o The solution leverages existing MIPv4 NAT traversal [mipnat] and IPsec NAT traversal [RFC3947] [RFC3948] solutions and does not require any new functionality to deal with NATs.

Security implications

- o The solution requires a new mechanism to detect whether the mobile node is in the internal or the external network. The security of this mechanism is critical in ensuring that the security level

provided by IPsec is not compromised by a faulty detection mechanism.

- o When the mobile node is outside, the external Mobile IPv4 layer may allow some traffic redirection attacks that plain IPsec does not allow. Other than that, IPsec security is unchanged.
- o More security considerations are described in Section 6.

5.2. Packet Overhead

The maximum packet overhead depends on access mode as follows:

- o f: 0 octets
- o c: 20 octets
- o fvc: 77 octets
- o cvc: 97 octets

The maximum overhead of 97 octets in the 'cvc' access mode consists of the following:

- o IP-IP for i-MIPv4: 20 octets
- o IPsec ESP: 57 octets total, consisting of 20 (new IP header), 4+4+8 = 16 (SPI, sequence number, cipher initialization vector), 7+2 = 9 (padding, padding length field, next header field), 12 (ESP authentication trailer)
- o IP-IP for x-MIPv4: 20 octets

When IPsec is used, a variable amount of padding is present in each ESP packet. The figures were computed for a cipher with 64-bit block size, padding overhead of 9 octets (next header field, padding length field, and 7 octets of padding; see Section 2.4 of [RFC4303]), and ESP authentication field of 12 octets (HMAC-SHA1-96 or HMAC-MD5-96). Note that an IPsec implementation MAY pad with more than a minimum amount of octets.

NAT traversal overhead is not included, and adds 8 octets when IPsec NAT traversal [RFC3947] [RFC3948] is used and 12 octets when MIP NAT traversal [mipnat] is used. For instance, when using access mode cvc, the maximum NAT traversal overhead is 12+8+12 = 32 octets. Thus, the worst case scenario (with the above mentioned ESP assumptions) is 129 octets for cvc.

5.3. Latency Considerations

When the MN is inside, connection setup latency does not increase compared to standard MIPv4 if the MN implements the suggested parallel registration sequence (see Section 3.3). Exchange of RRQ/RRP messages with the i-HA confirms the MN is inside, and the MN may start sending and receiving user traffic immediately. For the same reason, handovers in the internal network have no overhead relative to standard MIPv4.

When the MN is outside, the situation is slightly different. Initial connection setup latency essentially consists of (1) registration with the x-HA, (2) optional detection delay (waiting for i-HA response), (3) IPsec connection setup (IKE), and (4) registration with the i-HA. All but (4) are in addition to standard MIPv4.

However, handovers in the external network have performance comparable to standard MIPv4. The MN simply re-registers with the x-HA and starts to send IPsec traffic to the VPN gateway from the new address.

The MN may minimize latency by (1) not waiting for an i-HA response before triggering IKE if the x-HA registration succeeds and (2) sending first the RRQ most likely to succeed (e.g., if the MN is most likely outside). These can be done based on heuristics about the network, e.g., addresses, MAC address of the default gateway (which the mobile node may remember from previous access); based on the previous access network (i.e., optimize for inside-inside and outside-outside movement); etc.

5.4. Firewall State Considerations

A separate firewall device or an integrated firewall in the VPN gateway typically performs stateful inspection of user traffic. The firewall may, for instance, track TCP session status and block TCP segments not related to open connections. Other stateful inspection mechanisms also exist.

Firewall state poses a problem when the mobile node moves between the internal and external networks. The mobile node may, for instance, initiate a TCP connection while inside, and later go outside while expecting to keep the connection alive. From the point of view of the firewall, the TCP connection has not been initiated, as it has not witnessed the TCP connection setup packets, thus potentially resulting in connectivity problems.

When the VPN-TIA is registered as a co-located care-of address with the i-HA, all mobile node traffic appears as IP-IP for the firewall.

Typically, firewalls do not continue inspection beyond the IP-IP tunnel, but support for deeper inspection is available in many products. In particular, an administrator can configure traffic policies in many firewall products even for IP-IP encapsulated traffic. If this is done, similar statefulness issues may arise.

In summary, the firewall must allow traffic coming from and going into the IPsec connection to be routed, even though they may not have successfully tracked the connection state. How this is done is out of scope of this document.

5.5. Intrusion Detection Systems (IDSs)

Many firewalls incorporate intrusion detection systems monitoring network traffic for unusual patterns and clear signs of attack. Since traffic from a mobile node implementing this specification is UDP to i-HA port 434, and possibly IP-IP traffic to the i-HA address, existing IDSs may treat the traffic differently than ordinary VPN remote access traffic. Like firewalls, IDSs are not standardized, so it is impossible to guarantee interoperability with any particular IDS system.

5.6. Implementation of the Mobile Node

Implementation of the mobile node requires the use of three tunneling layers, which may be used in various configurations depending on whether that particular interface is inside or outside. Note that it is possible that one interface is inside and another interface is outside, which requires a different layering for each interface at the same time.

For multi-vendor implementation, the IPsec and MIPv4 layers need to interoperate in the same mobile node. This implies that a flexible framework for protocol layering (or protocol-specific APIs) is required.

5.7. Non-IPsec VPN Protocols

The solution also works for VPN tunneling protocols that are not IPsec-based, provided that the mobile node is provided IPv4 connectivity with an address suitable for registration. However, such VPN protocols are not explicitly considered.

6. Security Considerations

6.1. Internal Network Detection

If the mobile node by mistake believes it is in the internal network and sends plaintext packets, it compromises IPsec security. For this reason, the overall security (confidentiality and integrity) of user data is a minimum of (1) IPsec security and (2) security of the internal network detection mechanism.

Security of the internal network detection relies on a successful registration with the i-HA. For standard Mobile IPv4 [RFC3344], this means HMAC-MD5 and Mobile IPv4 replay protection. The solution also assumes that the i-HA is not directly reachable from the external network, requiring careful enterprise firewall configuration. To minimize the impact of a firewall configuration problem, the i-HA is separately required to be configured with trusted source addresses (i.e., addresses belonging to the internal network), and to include an indication of this in a new Trusted Networks Configured extension. The MN is required not to trust a registration as an indication of being connected to the internal network, unless this extension is present in the registration reply. Thus, to actually compromise user data confidentiality, both the enterprise firewall and the i-HA have to be configured incorrectly, which reduces the likelihood of the scenario.

When the mobile node sends a registration request to the i-HA from an untrusted network that does not go through the IPsec tunnel, it will reveal the i-HA's address, its own identity including the NAI and the home address, and the Authenticator value in the authentication extensions to the untrusted network. This may be a concern in some deployments.

When the connection status of an interface changes, an interface previously connected to the trusted internal network may suddenly be connected to an untrusted network. Although the same problem is also relevant to IPsec-based VPN implementations, the problem is especially relevant in the scope of this specification.

In most cases, mobile node implementations are expected to have layer 2 information available, making connection change detection both fast and robust. To cover cases where such information is not available (or fails for some reason), the mobile node is required to periodically re-register with the internal home agent to verify that it is still connected to the trusted network. It is also required that this re-registration interval be configurable, thus giving the administrator a parameter by which potential exposure may be controlled.

6.2. Mobile IPv4 versus IPsec

MIPv4 and IPsec have different goals and approaches for providing security services. MIPv4 typically uses a shared secret for authentication of signaling traffic, while IPsec typically uses IKE (an authenticated Diffie-Hellman exchange) to set up session keys. Thus, the overall security properties of a combined MIPv4 and IPsec system depend on both mechanisms.

In the solution outlined in this document, the external MIPv4 layer provides mobility for IPsec traffic. If the security of MIPv4 is broken in this context, traffic redirection attacks against the IPsec traffic are possible. However, such routing attacks do not affect other IPsec properties (confidentiality, integrity, replay protection, etc.), because IPsec does not consider the network between two IPsec endpoints to be secure in any way.

Because MIPv4 shared secrets are usually configured manually, they may be weak if easily memorizable secrets are chosen, thus opening up redirection attacks described above. Note especially that a weak secret in the i-HA is fatal to security, as the mobile node can be fooled into dropping encryption if the i-HA secret is broken.

Assuming the MIPv4 shared secrets have sufficient entropy, there are still at least the following differences and similarities between MIPv4 and IPsec worth considering:

- o Both IPsec and MIPv4 are susceptible to the "transient pseudo NAT" attack described in [pseudonat] and [mipnat], assuming that NAT traversal is enabled (which is typically the case). "Pseudo NAT" attacks allow an attacker to redirect traffic flows, resulting in resource consumption, lack of connectivity, and denial of service. However, such attacks cannot compromise the confidentiality of user data protected using IPsec.
- o When considering a "pseudo NAT" attack against standard IPsec and standard MIP (with NAT traversal), redirection attacks against MIP may be easier because:
 - * MIPv4 re-registrations typically occur more frequently than IPsec SA setups (although this may not be the case for mobile hosts).
 - * It suffices to catch and modify a single registration request, whereas attacking IKE requires that multiple IKE packets are caught and modified.

- o There may be concerns about mixing of algorithms. For instance, IPsec may be using HMAC-SHA1-96, while MIP is always using HMAC-MD5 (RFC 3344) or prefix+suffix MD5 (RFC 2002). Furthermore, while IPsec algorithms are typically configurable, MIPv4 clients typically use only HMAC-MD5 or prefix+suffix MD5. Although this is probably not a security problem as such, it is more difficult to communicate to users.
- o When IPsec is used with a Public Key Infrastructure (PKI), the key management properties are superior to those of basic MIPv4. Thus, adding MIPv4 to the system makes key management more complex.
- o In general, adding new security mechanisms increases overall complexity and makes the system more difficult to understand.

7. IANA Considerations

This document specifies a new skippable extension (in the short format) in Section 3.4, whose Type and Sub-Type values have been assigned.

Allocation of new Sub-Type values can be made via Expert Review and Specification Required [RFC5226].

8. Acknowledgements

This document is a joint work of the contributing authors (in alphabetical order):

- Farid Adrangi (Intel Corporation)
- Nitsan Baider (Check Point Software Technologies, Inc.)
- Gopal Dommety (Cisco Systems)
- Eli Gelasco (Cisco Systems)
- Dorothy Gellert (Nokia Corporation)
- Espen Klovning (Birdstep)
- Milind Kulkarni (Cisco Systems)
- Henrik Levkowetz (ipUnplugged AB)
- Frode Nielsen (Birdstep)
- Sami Vaarala (Codebay)
- Qiang Zhang (Liqwid Networks, Inc.)

The authors would like to thank the MIP/VPN design team, especially Mike Andrews, Gaetan Feige, Prakash Iyer, Brijesh Kumar, Joe Lau, Kent Leung, Gabriel Montenegro, Ranjit Narjala, Antti Nuopponen, Alan O'Neill, Alpesh Patel, Ilkka Pietikainen, Phil Roberts, Hans Sjostrand, and Serge Tessier for their continuous feedback and helping us improve this document. Special thanks to Radia Perlman for giving the document a thorough read and a security review. Tom

Hiller pointed out issues with battery-powered devices. We would also like to thank the previous Mobile IP working group chairs (Gabriel Montenegro, Basavaraj Patil, and Phil Roberts) for important feedback and guidance.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3344] Perkins, C., Ed., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3947] Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, "Negotiation of NAT-Traversal in the IKE", RFC 3947, January 2005.
- [RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec packets", RFC 3948, January 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [mipnat] Levkowitz, H. and S. Vaarala, "Mobile IP Traversal of Network Address Translation (NAT) Devices", RFC 3519, April 2003.
- [privaddr] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

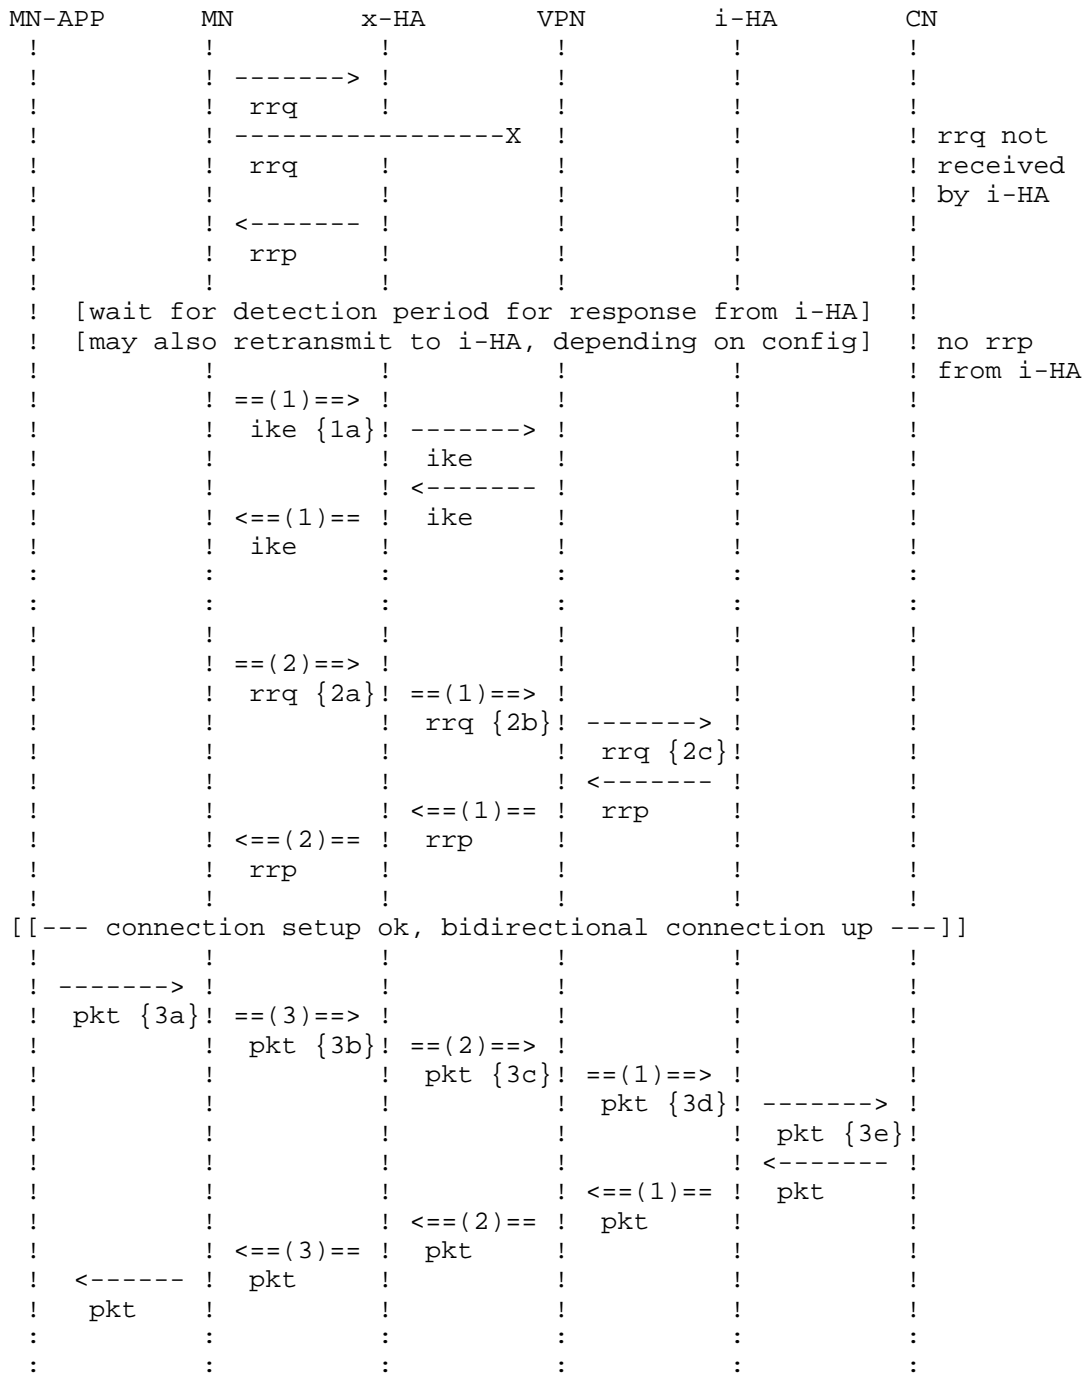
9.2. Informative References

- [RFC2002] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [RFC3456] Patel, B., Aboba, B., Kelly, S., and V. Gupta, "Dynamic Host Configuration Protocol (DHCPv4) Configuration of IPsec Tunnel Mode", RFC 3456, January 2003.
- [RFC3776] Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", RFC 3776, June 2004.
- [RFC4093] Adrangi, F. and H. Levkowitz, "Problem Statement: Mobile IPv4 Traversal of Virtual Private Network (VPN) Gateways", RFC 4093, August 2005.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.
- [RFC4555] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, June 2006.
- [pseudonat] Dupont, F. and J. Bernard, "Transient pseudo-NAT attacks or how NATs are even more evil than you believed", Work in Progress, June 2004.
- [tessier] Tessier, S., "Guidelines for Mobile IP and IPsec VPN Usage", Work in Progress, December 2002.

Appendix A. Packet Flow Examples

A.1. Connection Setup for Access Mode 'cvc'

The following figure illustrates connection setup when the mobile node is outside and using a co-located care-of address. IKE connection setup is not shown in full, and involves multiple round trips (4.5 round trips when using main mode followed by quick mode).



The notation "==(N)==>" or "<==(N)==>" indicates that the innermost packet has been encapsulated N times, using IP-IP, ESP, or MIP NAT traversal.

Packets marked with {xx} are shown in more detail below. Each area represents a protocol header (labeled). Source and destination addresses or ports are shown underneath the protocol name when applicable. Note that there are no NAT traversal headers in the example packets.

Packet {1a}

```

.------.
! IP      ! IP      ! UDP  ! IKE  !
! co-CoA ! x-HoA  ! 500  !     !
! x-HA   ! VPN-GW ! 500  !     !
'-----'

```

Packet {2a}

```

.------.
! IP      ! IP      ! ESP  ! IP      ! UDP  ! MIP RRQ !
! co-CoA ! x-HoA  !     ! VPN-TIA ! ANY  !     !
! x-HA   ! VPN-GW !     ! i-HA   ! 434  !     !
'-----'

```

Packet {2b}

```

.------.
! IP      ! ESP  ! IP      ! UDP  ! MIP RRQ !
! x-HoA  !     ! VPN-TIA ! ANY  !     !
! VPN-GW !     ! i-HA   ! 434  !     !
'-----'

```

Packet {2c}

```

.------.
! IP      ! UDP  ! MIP RRQ !
! VPN-TIA ! ANY  !     !
! i-HA   ! 434  !     !
'-----'

```

Packet {3a}

```

.------.
! IP      ! user  !
! i-HoA  ! protocol !
! CN     !     !
'-----'

```

Packet {3b}

```

-----
! IP      ! IP      ! ESP ! IP      ! IP      ! user   \
! co-CoA ! x-HoA  !     ! VPN-TIA ! i-HoA  ! protocol.. /
! x-HA   ! VPN-GW !     ! i-HA   ! CN     !         \
-----
      .
      |
      | \..user      ! ESP
      | / protocol ! trailer
      | \          !
      |  .-----

```

Packet {3c}

```

-----
! IP      ! ESP ! IP      ! IP      ! user      ! ESP !
! x-HoA  !     ! VPN-TIA ! i-HoA  ! protocol ! trailer !
! VPN-GW !     ! i-HA   ! CN     !         !
-----

```

Packet {3d}

```

-----
! IP      ! IP      ! user      !
! VPN-TIA ! i-HoA  ! protocol !
! i-HA   ! CN     !
-----

```

Packet {3e}

```

-----
! IP      ! user      !
! i-HoA  ! protocol !
! CN     !
-----

```

Packet {3b} with all NAT traversal headers (x-MIP, ESP, and i-MIP) is shown below for comparison.

Packet {3b} (with NAT traversal headers)

```

-----
! IP      ! UDP  ! MIP   ! IP      ! UDP  ! ESP.. \
! co-CoA ! ANY ! tunnel ! x-HoA  ! 4500 !      /
! x-HA   ! 434 ! data  ! VPN-GW ! 4500 !      \
-----
<=== external MIPv4 =====> <=== IPsec ESP ===== = =

-----
\..ESP ! IP      ! UDP  ! MIP   ! IP      ! user   \
/      ! VPN-TIA ! ANY ! tunnel ! i-HoA  ! protocol../
\      ! i-HA   ! 434 ! data  ! CN     !       \
-----
= =====> <===== internal MIPv4 =====> <== user packet == =

-----
\..user      ! ESP
/ protocol ! trailer !
\           !       !
-----
= = =====> <= ESP =>

```

Authors' Addresses

Sami Vaarala
 Codebay
 P.O. Box 63
 Espoo 02601
 FINLAND

Phone: +358 (0)50 5733 862
 EMail: sami.vaarala@iki.fi

Espen Klovning
 Birdstep
 Bryggegata 7
 Oslo 0250
 NORWAY

Phone: +47 95 20 26 29
 EMail: espen@birdstep.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.