

Mobile IP Working Group  
Internet-Draft  
Expires: December 22, 2003

F. Adrangi, Ed.  
intel  
H. Levkowetz, Ed.  
ipUnplugged  
June 23, 2003

**Problem Statement: Mobile IPv4 Traversal of VPN Gateways**  
**draft-ietf-mobileip-vpn-problem-statement-req-03**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on December 22, 2003.

Copyright Notice

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

Abstract

Deploying Mobile-IP v4 in networks which are connected to the internet through a VPN (Virtual Private Network) gateway presents some problems which do not currently have well-described solutions. This document aims to describe and illustrate these problems, and propose some guidelines for possible solutions.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.1</a>	Overview of the Problem . . . . .	<a href="#">3</a>
<a href="#">1.2</a>	Terminology . . . . .	<a href="#">4</a>
<a href="#">2.</a>	MIP and VPN Deployment Scenarios . . . . .	<a href="#">5</a>
<a href="#">2.1</a>	MIPv4 HA(s) Inside the Intranet behind a VPN Gateway . . . . .	<a href="#">5</a>
<a href="#">2.2</a>	VPN Gateway and MIPv4 HA(s) in parallel . . . . .	<a href="#">6</a>
<a href="#">2.3</a>	Combined VPN Gateway and MIPv4 HA . . . . .	<a href="#">7</a>
<a href="#">2.4</a>	MIPv4 HA(s) Outside the VPN domain . . . . .	<a href="#">8</a>
<a href="#">2.5</a>	Combined VPN Gateway and MIPv4 HA(s) on the Local Link . . . . .	<a href="#">8</a>
<a href="#">3.</a>	Deployment Scenarios Selection . . . . .	<a href="#">9</a>
<a href="#">4.</a>	Problem statement . . . . .	<a href="#">10</a>
<a href="#">4.1</a>	Registering in co-located mode . . . . .	<a href="#">11</a>
<a href="#">4.2</a>	Registering via an FA . . . . .	<a href="#">11</a>
<a href="#">4.3</a>	Summary: MIP Incompatibilities with IPsec-based VPN Gateways . . . . .	<a href="#">13</a>
<a href="#">5.</a>	Solution Guidelines . . . . .	<a href="#">13</a>
<a href="#">5.1</a>	Preservation of Existing VPN Infrastructure . . . . .	<a href="#">13</a>
<a href="#">5.2</a>	Software Upgrades to Existing VPN Client and Gateways . . . . .	<a href="#">14</a>
<a href="#">5.3</a>	IPsec Protocol . . . . .	<a href="#">14</a>
<a href="#">5.4</a>	Multi-Vendor Interoperability . . . . .	<a href="#">14</a>
<a href="#">5.5</a>	MIPv4 Protocol . . . . .	<a href="#">14</a>
<a href="#">5.6</a>	Handoff Overhead . . . . .	<a href="#">15</a>
<a href="#">5.7</a>	Scalability, Availability, Reliability, and Performance . . . . .	<a href="#">15</a>
<a href="#">5.8</a>	Functional Entities . . . . .	<a href="#">15</a>
<a href="#">5.9</a>	Implications of Intervening NAT Gateways . . . . .	<a href="#">15</a>
<a href="#">5.10</a>	Security Implications . . . . .	<a href="#">15</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">15</a>
	Normative References . . . . .	<a href="#">16</a>
	Informative References . . . . .	<a href="#">16</a>
	Authors' Addresses . . . . .	<a href="#">17</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">19</a>



## **1. Introduction**

Mobile IP [[1](#)] agents are being deployed in enterprise networks to enable mobility across wired and wireless LANs while roaming inside the enterprise intranet. With the growing deployment of IEEE 802.11 access points ("hot spots") in public places such as hotels, airports, and convention centers, and wireless WAN data networks such as GPRS, the need for enabling mobile users to maintain their transport connections and constant reachability while connecting back to their target "home" networks protected by Virtual Private Network (VPN) technology is increasing. This implies that Mobile IP and VPN technologies have to coexist and function together in order to provide mobility and security to the enterprise mobile users.

The goal of this draft is to:

- o Identify and describe practical deployment scenarios for Mobile IP and VPN in enterprise and operator environments.
- o Identify example usage scenarios for remote users roaming outside the "home" network protected by a VPN gateway.
- o Articulate the problems resulting from Mobile IP and VPN coexistence. Specify a set of framework guidelines to evaluate proposed solutions, supporting multi-vendor seamless IPv4 mobility across IPsec-based VPN gateways.

### **1.1 Overview of the Problem**

Real life networks typically consist of three different domains from a corporate point of view. The first domain is the Internet (i.e., the untrusted external network). The second domain is the trusted intranet (also referred to as VPN Domain in this document). The third domain is the DMZ, which is between the Internet and the intranet.

Access to the intranet is typically guarded by both a firewall and a VPN device. The intranet can only be accessed by respecting the security policies in the firewall and the VPN device.

When MIP is deployed in a corporate network behind a VPN device, roaming between these two different domains (i.e., the untrusted Internet and the trusted intranet) becomes problematic. It would be desirable to have seamless session mobility between the two domains, because MIP was designed for session mobility regardless of the network point of attachment. Unfortunately, the current MIP standards fall short of this promise for an important customer



segment, corporate users behind VPN gateways.

Because current standards do not provide for session mobility across these two domains the possibility of finding a solution to this problem has been investigated. The goal is to provide seamless session mobility when the mobile node moves between these two domains or between subnets in either domain.

From the beginning it was also assumed that VPNs and firewalls were to be taken as more or less granted because they have much wider deployments than MIP at the present. Therefore any solutions would need to minimize impact on existing VPN and firewall deployments, related standards and "de facto" standards.

## **1.2 Terminology**

MIPv4      Mobile IP for IPv4 [[1](#)]

MIPv6      Mobile IP for IPv6

VPN          Virtual Private Network

GW          Gateway

VPN Domain

    An intranet protected by a VPN gateway.

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

Home Network

    A network, possibly virtual, having a network prefix matching that of a mobile node's home address.

Home Agent

    A router on a mobile node's home network which tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information for the mobile node.

MIPv4 inside IPsec-ESP tunnel

    MIPv4 packet is encapsulated in an IPsec-ESP tunnel established between the Mobile Node and the VPN gateway.



### IPsec-ESP inside MIPv4 tunnel

IPsec-ESP packet is encapsulated in an MIPv4 tunnel established between the Mobile Node and the home agent.

## 2. MIP and VPN Deployment Scenarios

This section describes a set of deployment scenarios where MIP agents and VPN gateways have to coexist to provide mobility and security. The intention is to identify practical deployment scenarios for MIP and VPNs where MIP technology might be extended to solve problems resulting from the desire for co-existence.

In all scenarios, "MN" refers to a mobile node that runs both MIP and IPsec-based VPN client software. The foreign network might or might not employ a foreign agent. And, the term "Intranet" refers to a private network protected by a VPN gateway and perhaps a layer-3 transparent or non-transparent firewall. Please note that firewalls are purposely omitted from the following scenarios, because they may be installed in a number of different ways, and the fact that this draft's focus is the relationship between MIP and VPN.

Finally, the scenarios assume that encryption is not enforced inside the VPN domain because 1) the VPN domain (Intranet) is viewed as a trusted network, and users allowed inside the Intranet are also trusted 2) it is a common VPN deployment practice where the VPN is used to guard the Intranet resources from unauthorized users attached to an untrusted network, and to provide a secure communication channel for authorized users to access resources inside the Intranet from outside.

The following sub-sections introduce five representative combinations of MIPv4 HA and VPN gateway placement.

### 2.1 MIPv4 HA(s) Inside the Intranet behind a VPN Gateway

MIPv4 HAs are deployed inside the Intranet protected by a VPN gateway, and are not directly reachable by the MNs outside the Intranet.

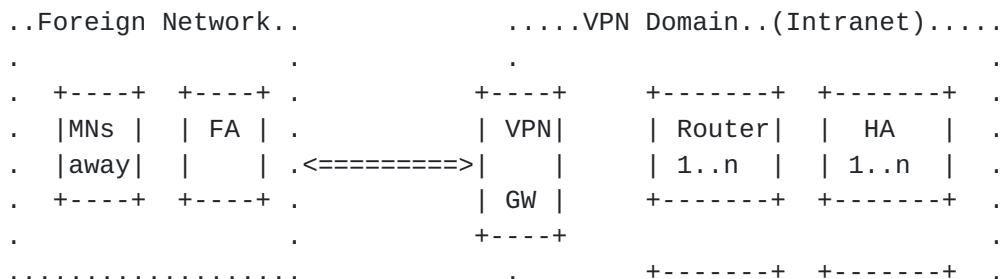




Figure 1

Direct application of MIPv4 standards [1] is successfully used to provide mobility for users inside the Intranet. However, mobile users outside the Intranet can only access the intranet resources (e.g., MIP agents) through the VPN gateway, which will allow only authenticated IPsec traffic inside. This implies that the MIPv4 traffic has to run inside IPsec, which leads to two distinct problems:

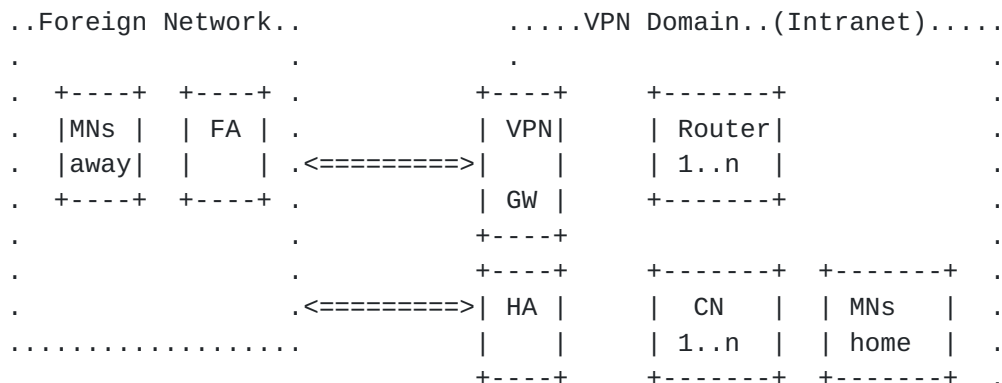
1. When the foreign network has an FA deployed (as in e.g. CDMA 2000), MIPv4 registration becomes impossible because the traffic between MN and VPN gateway, which is what the FA sees, is encrypted and the FA is not set up to decrypt it.
2. In co-located mode, successful registration is possible but the VPN tunnel has to be re-negotiated every time the MN changes its point of network attachment.

These problems are articulated in [Section 4](#).

This deployment scenario may not be common yet, but it is practical and becoming important as there is an increasing need for providing corporate remote users with continuous access to the Intranet resources.

## 2.2 VPN Gateway and MIPv4 HA(s) in parallel

A MIPv4 HA is deployed in parallel with the VPN gateway, and it is directly reachable by MNs inside or outside the Intranet.





```

.
.
.....

```

Figure 2

The MIPv4 HA has a public interface connected to the Internet, and a private interface attached to the Intranet. Mobile users will most likely have a virtual home network associated with the MIPv4 HA's private interface, so that the mobile users are always away from home and hence registered with the MIPv4 HA. Furthermore, in deployments where the VPN gateway and the HA are placed in a corporate DMZ, this implies that MIPv4 traffic will always be routed through the DMZ (regardless of whether MNs are located outside or inside the Intranet), which may not be acceptable by IT departments in large corporations.

This deployment can be used with two different configurations: "MIPv4 inside IPsec-ESP tunnel" and "IPsec-ESP inside MIPv4 tunnel". The "MIPv4 inside IPsec-ESP tunnel" has the same problems as the scenario of [Section 2.1](#) (namely, MIPv4 registration becomes impossible when the registration is to be done via an FA and furthermore in co-located mode, the VPN tunnel has to be re-negotiated every time the MN changes its point of attachment). The "IPsec-ESP inside MIPv4 tunnel" does not have problems described in [Section 2.1](#), however it will require some modifications to the routing logic of the MIPv4 HA or the VPN gateway.

### [2.3 Combined VPN Gateway and MIPv4 HA](#)

This is similar to deployment scenario described in [Section 2.2](#), with the exception that the VPN gateway and MIPv4 HA are running on the same physical machine.

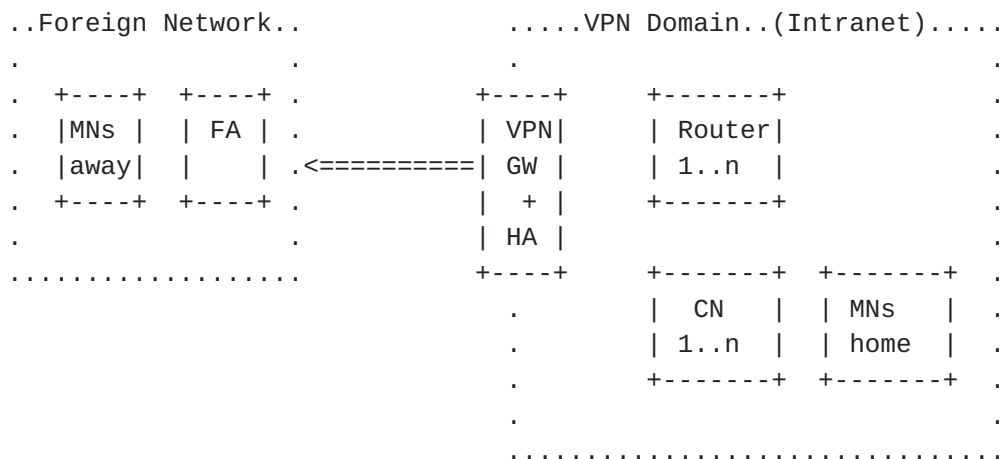


Figure 3



Running MIPv4 HA and VPN on the same machine resolves routing related issues that exist in [Section 2.2](#) when a "IPsec-ESP inside MIPv4 tunnel" configuration is used. However, it does not promote multi-vendor interoperability in environments where MIPv4 HA and VPN technologies must be acquired from different vendors.

#### [2.4 MIPv4 HA\(s\) Outside the VPN domain](#)

In this scenario, MIPv4 HAs are deployed outside the Intranet (e.g., in an operator network), as depicted in Figure 4 below.

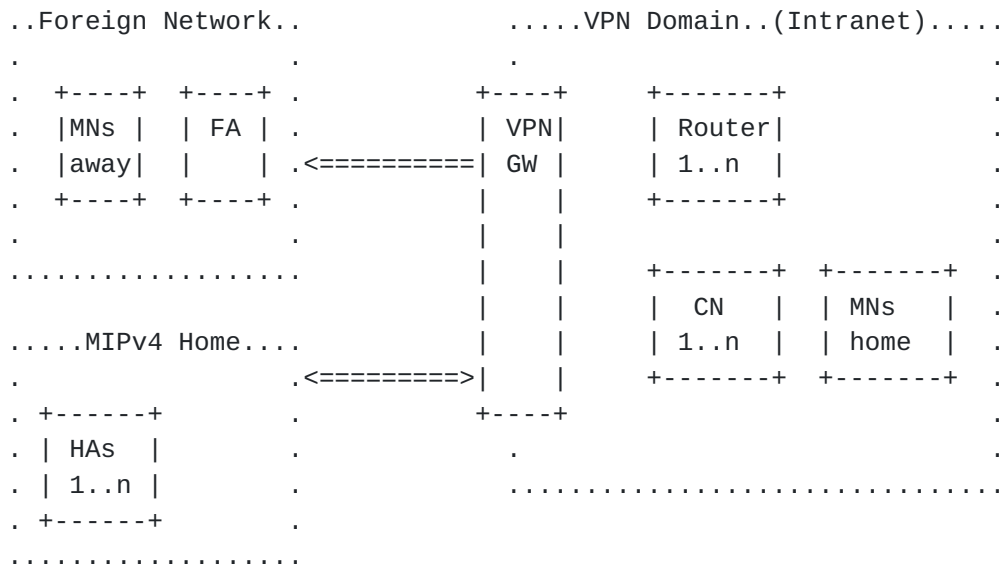


Figure 4

In this deployment scenario the goal is to provide remote users with continuous access to the Intranet resources while they are roaming outside the Intranet only (i.e., mobility is not supported inside the Intranet). In this case it is most practical to run IPsec-ESP inside a MIPv4 tunnel (i.e., MIPv4 tunnel end-points are the MN and the HA; the IPsec-ESP packet from the MN and to the VPN gateway is encapsulated in the MIPv4 tunnel) as the MNs can register with the HA without establishing an IPsec tunnel to the VPN gateway. This should work without any technical problems. The IPsec tunnel end-points will be the MN and the VPN gateway. The 'home network' will be a virtual home network, located at the HA, from which it is possible to reach the Corporate intranet through the VPN gateway.

#### [2.5 Combined VPN Gateway and MIPv4 HA\(s\) on the Local Link](#)

This is similar to the deployment scenario described in [Section 2.3](#), with the difference that the VPN gateway/HA is sitting on the local link. In this the VPN gateway and HA would most naturally be



co-located in the same box, although this is in no way a requirement.

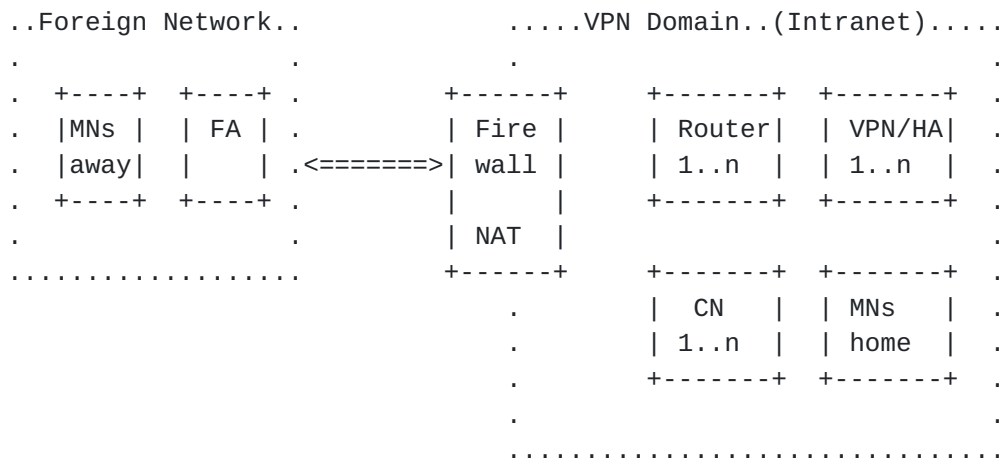


Figure 5

This deployment works today without any technical problems with IPsec-ESP running inside a MIPv4 tunnel. If however MIPv4 is run inside the IPsec-ESP tunnel it has the same problems as in [Section 2.1](#) (namely, MIPv4 registration becomes impossible when the registration is to be done via an FA and furthermore in co-located mode, the VPN tunnel has to be re-negotiated every time the MN changes its point of attachment). This deployment is not common or practical for large deployments (on the order of thousands of users) because of the large and distributed security perimeter.

### 3. Deployment Scenarios Selection

The deployment scenarios described in [Section 2](#) were evaluated to identify the ones most in need of solving. The evaluation was done based on two main criteria: 1) Is the deployment scenario common and practical? and 2) Does the deployment scenario reveal any problems resulting from MIPv4 and VPN coexistence?

There was a consensus about importance and practicality of the scenario in [Section 2.1](#) because of rising needs to provide corporate remote users with continuous access to their Intranet resources. After analyzing each scenario one realizes that problems occurring in scenarios in [Section 2.2](#) and [Section 2.4](#) are either the same as or a subset of those in the scenario in [Section 2.1](#). Therefore, solving the scenario in [Section 2.1](#) will also solve the scenarios in [Section 2.2](#) and [Section 2.4](#). The scenarios in [Section 2.3](#) and [Section 2.5](#) do not introduce functional problems resulting from MIPv4 and VPN co-existence, hence there is no need to seek a solution. A solution for the deployment scenario in [Section 2.1](#) is therefore seen as essential, and this in turn can also be applied to solve problems in



other scenarios. For the remainder of this draft, we will articulate the roaming scenarios, the problems, and the solution guidelines relevant to the scenario in [Section 2.1](#).

#### 4. Problem statement

This section describes roaming scenarios corresponding to the deployment scenario in [Section 2.1](#) where an MN needs to have continuous access to the Intranet resources regardless of whether it is roaming inside or outside the Intranet, and their associated problems. The scenarios are constructed based on a multi-subnetted, MIPv4-enabled Intranet (hereafter, referred to as Intranet or VPN domain) protected by an IPsec-based VPN gateway as depicted in Figure 6.

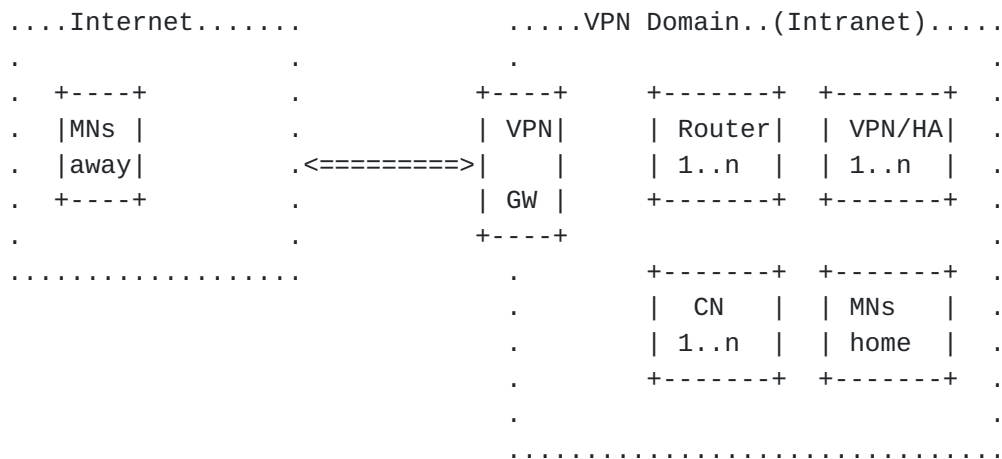


Figure 6: Intranet protected by a VPN Gateway

The Intranet, depicted in Figure 6, may include both wired (IEEE 802.3) and IEEE 802.11 wireless LAN deployments. However, it is also possible to see IEEE 802.11 deployments outside the Intranet due to the perceived lack of current 802.11 security, as depicted in Figure 7.

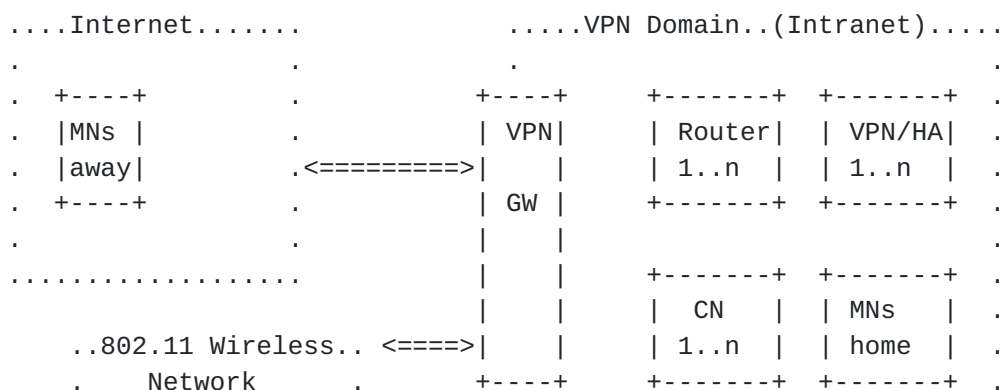






Figure 7: IEEE 802.11 Wireless deployment outside the home network

#### 4.1 Registering in co-located mode

In co-located mode, the IPsec tunnel endpoints would be at the MN and the VPN gateway, which (supposing we have the scenario described in [Section 2.1](#)) results in the mobile-ip tunnel from MN to HA being encapsulated inside the IPsec tunnel. See Figure 8 below. This scenario is still possible, but has some major drawbacks.

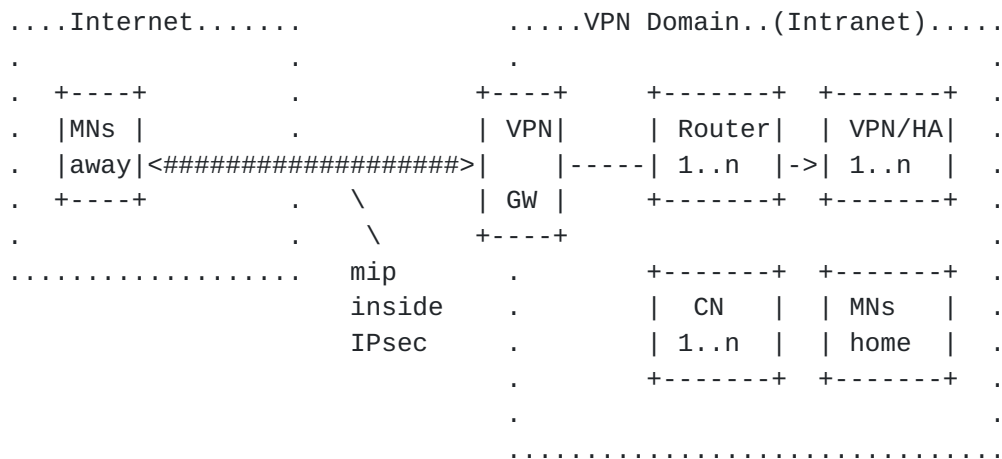


Figure 8

The MN obtains an address at its point of attachment (via DHCP[7] or some other means), and then first sets up an IPsec tunnel to the VPN gateway, after which it can successfully register with its HA through the IPsec tunnel. The problem is that in an end-to-end security model, an IPsec tunnel that terminates at the VPN gateway must protect the IP traffic originating at the MN. As the MN's IPsec tunnel address is the address obtained at the point of attachment, it will change during movement, and the VPN tunnel security association must be refreshed after each IP subnet handoff. This could have noticeable performance implications on real-time applications. In effect, we don't have mobility support for the tunnel endpoint changes associated with MN movements.

#### 4.2 Registering via an FA

In the case where a mobile node is in a network where mobility support is provided through the use of an FA, and no dhcpd allocated address and co-located mode is possible, we run into severe trouble.



Figure 9 below illustrates this:

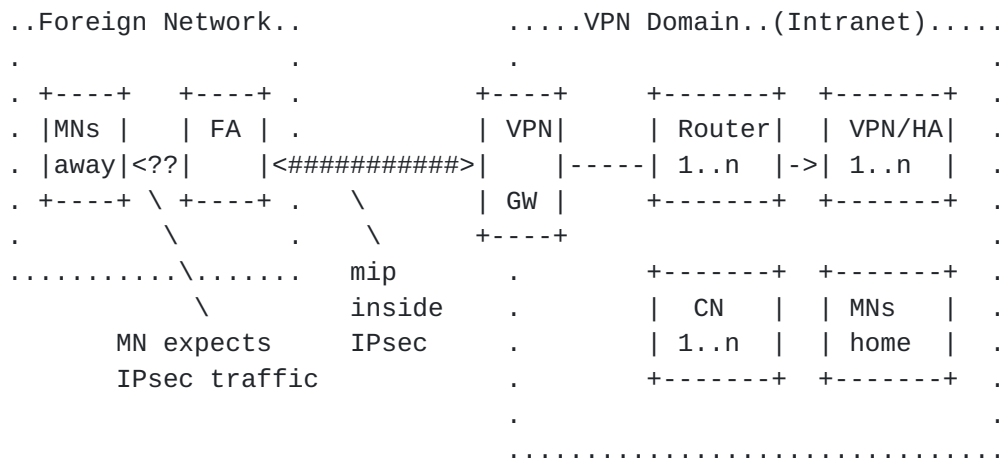


Figure 9

The mobile node, when arriving at this network, may have a IPsec session going with its VPN gateway. This session will not be passed through the FA as long as the MN has not registered and a mip tunnel has been set up. But the MN, which is secure inside the IPsec based VPN, will not even hear the FA advertisements. And any IPsec traffic from the intranet (via the VPN gateway and IPsec tunnel) will not be understood by the FA. Simply put, you could say that the FA needs to see the mip tunnel outermost, while the VPN-GW needs to see the IPsec tunnel outermost. Or in more details:

Firstly, the MN must have a IPsec tunnel established with the VPN-GW in order to reach the HA, which places the IPsec tunnel outside the mip traffic between MN and HA. The FA (which is likely in a different administrative domain) cannot decrypt MIPv4 packets between the MN and the VPN gateway, and will consequently be not able to relay the MIPv4 packets. This is because the MIPv4 headers (which the FA should be able to interpret) will be encrypted and protected by IPsec.

Secondly, when the MN is communicating with the VPN-GW, an explicit bypass policy for MIP packets is required, so that the MN can hear FA advertisements and send and receive MIP registration packets. Although not a problem in principle, there may be practical problems when VPN and MIP clients from different vendors are used.

The use of a 'trusted FA' has been suggested in this scenario; meaning an FA which is actually a combined VPN GW and FA. The scenario will work fine in this case; effectively we are then operating within the VPN established between the two VPN gateways, and the case is analogous to deploying mobile-ip within a corporate



intranet which is not physically disjoint. See Figure 10 below. However, we cannot expect that e.g. wireless hot-spots or CDMA 2000 FAs will have VPN gateways with security associations with any given corporate network, so this is not particularly realistic in the general mobility case.

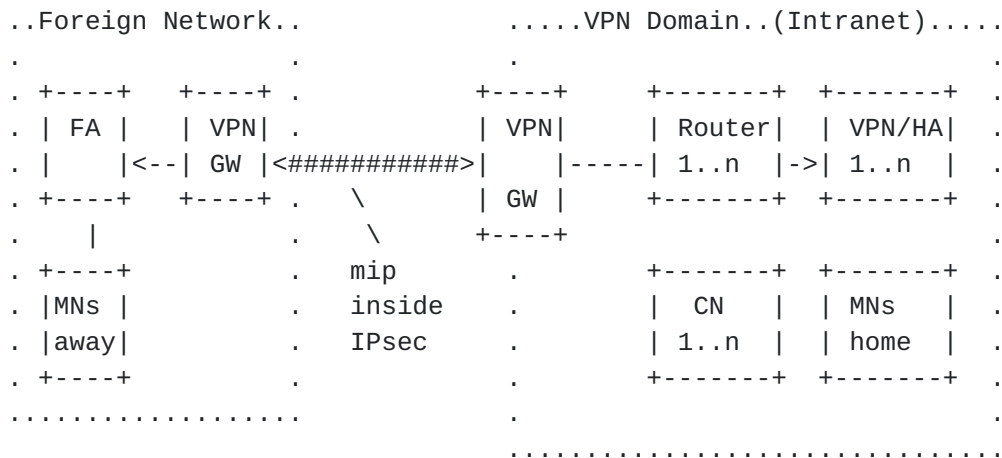


Figure 10

Furthermore, this solution would leave the traffic between FA and MN unprotected, and as this link in particular may be a wireless link, this is clearly undesirable.

#### 4.3 Summary: MIP Incompatibilities with IPsec-based VPN Gateways

An MN roaming outside the Intranet has to establish an IPsec tunnel to its home VPN gateway first, in order to be able to register with its home agent. This is because the MN cannot reach its HA (inside the private protected network) directly from the outside. This implies that the MIPv4 traffic from the MN to a node inside the Intranet is forced to run inside an IPsec tunnel, and hence will not be in the clear. This in turn leads to two distinct problems depending on whether the MN uses co-located or non co-located modes to register with its HA.

### 5. Solution Guidelines

This section describes guidelines for a solution to MIPv4 traversal across VPN gateways. The subsections discuss the guidelines in a decreasing order of importance.

#### 5.1 Preservation of Existing VPN Infrastructure

- o The solution MUST preserve the investment in existing VPN gateways.



- o The solution MUST provide security which is not inferior to what is already provided to existing "nomadic computing" remote access users, i.e. for confidentiality, authentication, message integrity, protection against replay attacks and related security services.

## **5.2 Software Upgrades to Existing VPN Client and Gateways**

- o The solution SHOULD minimize changes to existing VPN client/gateway software.

## **5.3 IPsec Protocol**

- o The solution SHOULD NOT require any changes to existing IPsec or key exchange standard protocols implemented by VPN gateways.
- o The solution SHOULD NOT require that the VPN gateway or the VPN client implement any new protocols in addition to the existing standard protocols.

## **5.4 Multi-Vendor Interoperability**

- o The solution MUST provide multi-vendor interoperability, where MIPv4 mobility agents, mobility clients (MN), VPN server, and VPN client solutions may come from four different vendors. This is typical for medium and large enterprises which purchase and deploy best-of-breed multi-vendor solutions for IP routing, VPNs, firewalls etc.

## **5.5 MIPv4 Protocol**

- o The solution MUST adhere to MIPv4 protocol [[1](#)]. That is, the solution MUST NOT impose any changes that violates MIPv4 protocol.
- o The solution MAY introduce new extensions to MIPv4 nodes per guidelines specified in the MIPv4 protocol [[1](#)]. However, it is highly desirable to avoid any changes to MIPv4 mobility agents such as the FA and HA in order to overcome barriers to deployment.
- o The solution MAY require more than one instance of MIPv4 running in parallel (multiple encapsulation).



### **5.6 Handoff Overhead**

- o It is imperative to keep the key management overhead down to a minimum, in order to support fast handoffs across IP subnets. Hence, the solution MUST propose a mechanism to avoid or minimize IPsec tunnel SA renegotiation and IKE renegotiation as the MN changes its current point of network attachment.

### **5.7 Scalability, Availability, Reliability, and Performance**

- o The solution complexity MUST increase at most linearly with the number of MNs registered and accessing resources inside the Intranet.
- o The solution MAY introduce additional header or tunnelling overhead if needed.

### **5.8 Functional Entities**

- o The solution MAY introduce new MIPv4 compliant functional entities.

### **5.9 Implications of Intervening NAT Gateways**

- o The solution MUST be able to leverage the existing MIPv4 and IPsec NAT traversal solutions [[9](#)][[10](#)][[11](#)].

### **5.10 Security Implications**

- o The solution MUST NOT introduce any new vulnerabilities to the MIPv4 or IPsec as specified in related RFCs.

## **6. Acknowledgements**

The authors who contributed text to this document were in no particular order: Farid Adrangi, Milind Kulkarni, Gopal Dommetty, Eli Gelasco, Qiang Zhang, Sami Vaarala, Dorothy Gellert, Nitsan Baider and Henrik Levkowetz.

The authors would like to thank other contributors, especially Prakash Iyer, Mike Andrews, Ranjit Narjala, Joe Lau, Kent Leung, Alpesh Patel, Phil Roberts, Hans Sjostrand, Serge Tessier, Antti Nuopponen, Alan O'Neill, Gaetan Feige, Brijesh Kumar for their



continuous feedback and helping us improve this draft.

#### Normative References

- [1] Perkins, C., "IP Mobility Support for IPv4", [RFC 3344](#), August 2002.

#### Informative References

- [2] Hanks, S., Li, T., Farinacci, D. and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 1701](#), October 1994.
- [3] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [4] Perkins, C., "Minimal Encapsulation within IP", [RFC 2004](#), October 1996.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [6] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [7] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [8] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", [RFC 3024](#), January 2001.
- [9] Vaarala, S. and O. Levkowetz, "Mobile IP NAT/NAPT Traversal using UDP Tunnelling", [draft-ietf-mobileip-nat-traversal-07](#) (work in progress), November 2002.
- [10] Aboba, B. and W. Dixon, "IPsec-NAT Compatibility Requirements", [draft-ietf-ipsec-nat-reqts-04](#) (work in progress), March 2003.
- [11] Kivinen, T., "Negotiation of NAT-Traversal in the IKE", [draft-ietf-ipsec-nat-t-ike-06](#) (work in progress), May 2003.



Authors' Addresses

Farid Adrangi  
Intel Corporation  
2111 N.E. 25th Avenue  
Hillsboro OR  
USA

Phone: +1 503-712-1791  
EMail: farid.adrangi@intel.com

Henrik Levkowetz  
ipUnplugged AB  
Arenavagen 33  
Stockholm S-121 28  
SWEDEN

Phone: +46 8 725 9513  
EMail: henrik@levkowetz.com

Milind Kulkarni  
Cisco Systems  
170 W. Tasman Drive  
San Jose CA 95134  
USA

Phone: +1 408-527-8382  
EMail: mkulkarn@cisco.com

Gopal Dommety  
Cisco Systems  
170 W. Tasman Drive  
San Jose CA 95134  
USA

EMail: gdommety@cisco.com

Eli Gelasco  
Cisco Systems  
170 W. Tasman Drive  
San Jose CA 95134  
USA

EMail: egelasco@cisco.com



Qiang Zhang  
Liqwid Networks, Inc.  
1000 Wilson Blvd, Suite 900  
Arlington VA 22209  
USA

Phone: +1 703-224-1120 -x 203  
EMail: qzhang@liqwidnet.com

Sami Vaarala  
Netseal  
Niittykatu 6  
Espoo 02201  
FINLAND

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

Dorothy Gellert  
Nokia Corporation  
  
EMail: dorothy.gellert@nokia.com

Nitsan Baider  
Check Point Software Technologies, Inc.  
  
EMail: nitsan@checkpoint.com



## Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

## 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 assignees.

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.

#### Acknowledgment

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