Network Working Group
Request for Comments: 4093
Category: Informational

F. Adrangi, Ed.
Intel
H. Levkowetz, Ed.
Ericsson

August 2005

Problem Statement: Mobile IPv4 Traversal of Virtual Private Network (VPN) Gateways

Status of This Memo

Copyright Notice

Abstract

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

Table of Contents

1. Introduction

Mobile IP [RFC3344] 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 with wireless WAN data networks such as General Packet Radio Service (GPRS), the need is increasing 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. 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 document is to:

- Identify and describe practical deployment scenarios for Mobile IP and VPN in enterprise and operator environments.

- Identify example usage scenarios for remote users roaming outside the "home" network protected by a VPN gateway.

- Articulate the problems resulting from Mobile IP and VPN coexistence.

- Specify a set of framework guidelines to evaluate proposed solutions for supporting multi-vendor seamless IPv4 mobility across IPsec-based VPN gateways.

1.1. Overview of the Problem

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 Intranet (also referred to as a VPN domain), roaming between the Intranet (i.e., trusted domain) and the Internet (i.e., untrusted domain) 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 (using VPN for remote access) who desire to add mobility support because of a need to have continuous access to Intranet resources while roaming outside the Intranet from one subnet to another, or between the VPN domain and the Internet.

From the beginning, one explicitly stated restriction was that it was assumed that installed firewalls and VPN gateways had to be kept unchanged, rather than replaced or upgraded, because they have much wider deployments than MIP at the time of writing. Therefore, any solutions would need to minimize the impact on existing VPN and firewall deployments, related standards, and "de facto" standards.

1.2. Specification of Requirements

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.3. Terminology

MIPv4	Mobile IP for IPv4 [RFC3344]
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.

MN Refers to a mobile node that runs both MIP and IPsec-based VPN client software.

MIPv4 inside IPsec-ESP tunnel
MIPv4 packets are encapsulated in an IPsec-ESP tunnel established between the Mobile Node and the VPN gateway.

IPsec-ESP inside MIPv4 tunnel
IPsec-ESP packets are encapsulated in a 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 wherein 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.

The network topology in the following diagrams consists of an Intranet connected to the public network (i.e., the Internet). Here, the word "Intranet" refers to a private network (where private addresses [RFC1918] are typically used) protected by a VPN gateway and perhaps by a layer-3 transparent or non-transparent firewall. When private addresses are used, the non-transparent firewall also functions as a Network Address Translator (NAT) or Network Address Port Translator (NAPT) bridging between the two address realms (i.e., the Intranet and the Internet).

Firewalls may be placed on either side of the VPN gateway; these are referred to as inner and outer firewalls. The inner and outer firewalls typically inspect outbound traffic (i.e., from the Intranet to the Internet) and inbound traffic (i.e., from the Internet to the

Intranet), respectively. When a firewall is present, it MUST be configured to allow Mobile IP traffic (both control and tunneled data packets) to go through. As our focus here is the relationship between MIP and VPN, we have purposely omitted firewalls from the following scenarios in order to keep things simple.

It is assumed 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, and 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.

In order to give a reasonably complete survey of MIPv4 and VPN co-existence scenarios, those in Sections 2.3 and 2.5 are included even though, as covered in more detail below, there are no co-existence problems to be solved in these two cases.

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.

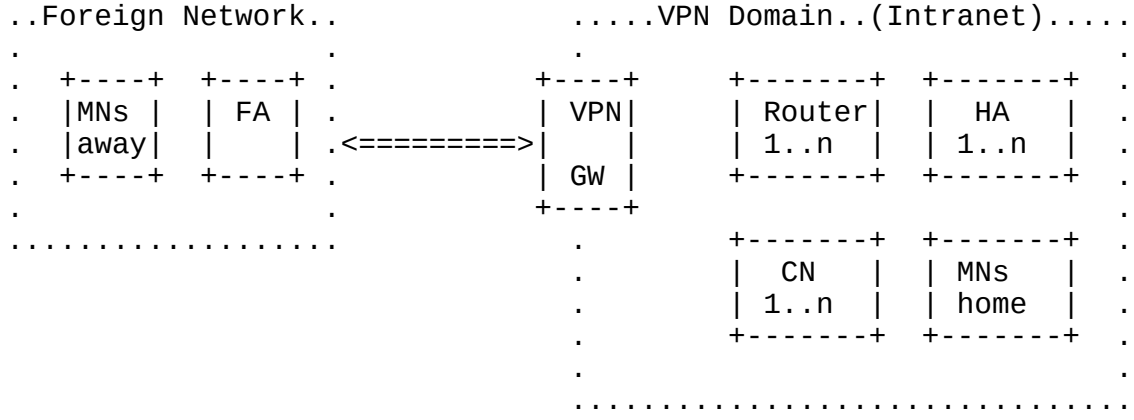

Figure 1

Direct application of MIPv4 standards [RFC3344] 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 (e.g., as in CDMA 2000), MIPv4 registration becomes impossible. This is because the MIPv4 traffic between MN and VPN gateway is encrypted, and the FA (which is likely in a different administrative domain) cannot inspect the MIPv4 headers needed for relaying the MIPv4 packets. Please see Section 4.2 for more details.

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 is 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) on the VPN Domain Border

A MIPv4 HA is deployed on the VPN domain border (e.g., in the DMZ) together with the VPN gateway, and it is directly reachable by MNs inside or outside the Intranet.

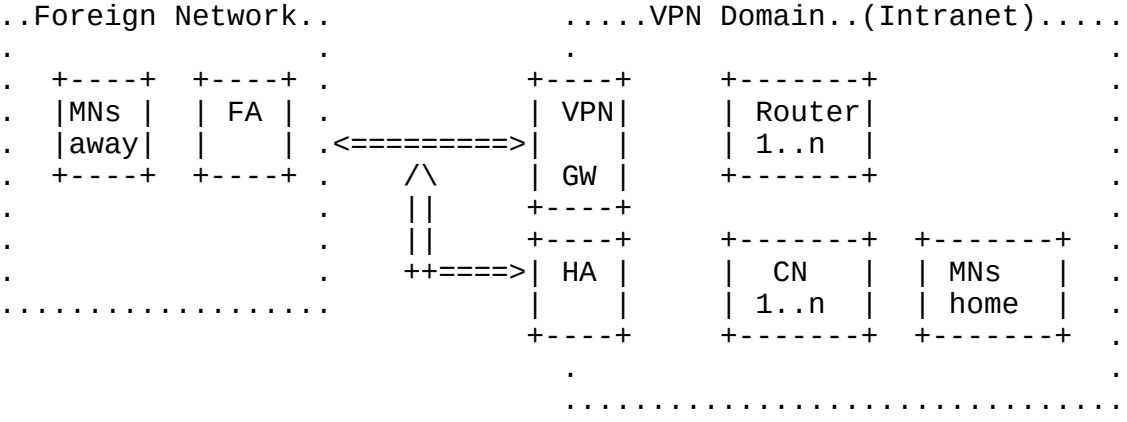

Figure 2

Please note that in deployments where the security policy prohibits direct communication between the MN (roaming outside the Intranet) and outside machines, the HA can be configured to forward only encrypted traffic from/to the MN.

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 thus 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 to 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 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.) The "IPsec-ESP inside MIPv4 tunnel" does not have the 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 the deployment scenario described in Section 2.2, with the exception that the VPN gateway and MIPv4 HA are running on the same physical machine.

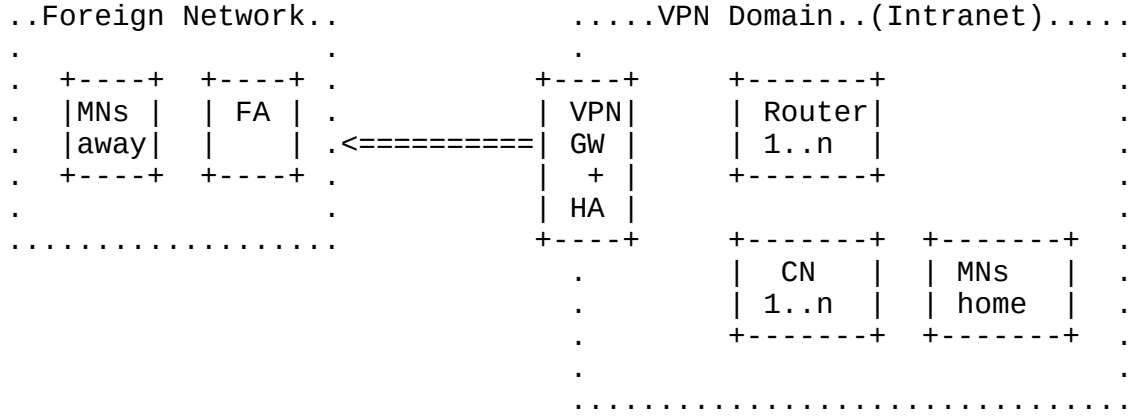

Figure 3

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.

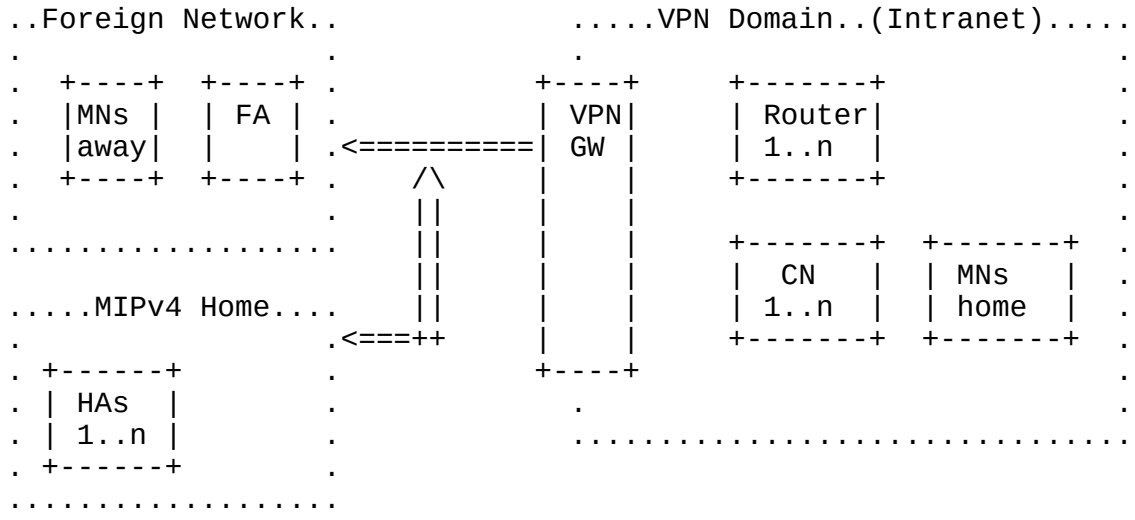

Figure 4

The IPsec tunnel endpoints will be the MN and the VPN gateway. The 'home network' will most likely be a virtual home network, located at the HA, through which authorized remote users (i.e., those that have successfully established a connection to the corporate VPN) can reach the Corporate Intranet and maintain their transport session connectivity while roaming outside the Intranet from one subnet to another. Please note that this deployment scenario does not support mobility inside the Intranet.

In this case, it is most practical to run IPsec-ESP inside a MIPv4 tunnel (i.e., the MIPv4 tunnel endpoints are the MN and the HA; the IPsec-ESP packet from the MN and to the VPN gateway is encapsulated in the MIPv4 tunnel). This is because the MNs can register with the HA without establishing an IPsec tunnel to 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 case, the VPN gateway and HA would most naturally be co-located in the same box, although this is in no way a requirement.

The VPN/HA is assumed to be reachable from the external network; i.e., it is assumed to have a public IP address, and the firewall is assumed to be configured to allow direct access to the VPN/HA from the external network.

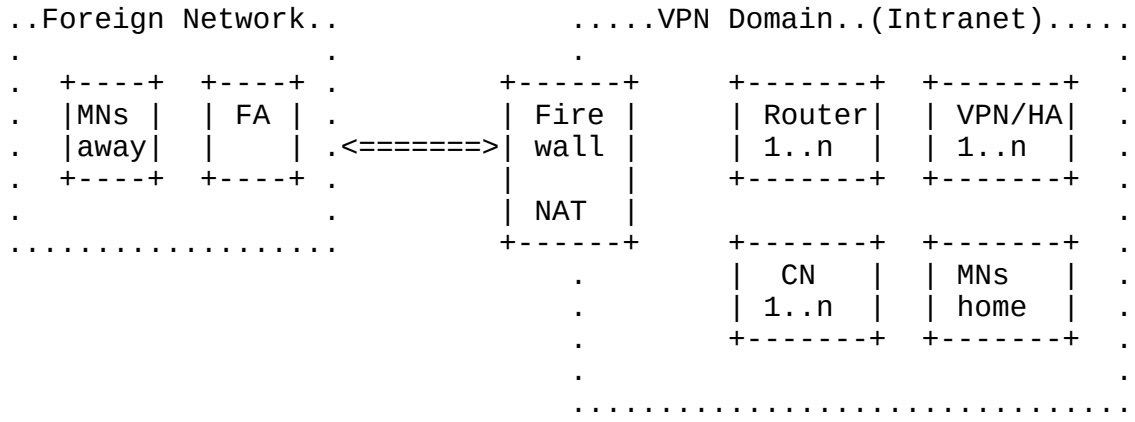

Figure 5

This deployment works today without any technical problems with IPsec-ESP running inside a MIPv4 tunnel. If you were to run MIPv inside the IPsec-ESP tunnel, it would have the same problems as in Section 2.1, so it is deployed with the IPsec-ESP running inside the MIPv4 tunnel. This deployment is not 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 those 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?

The authors believe that the scenario in Section 2.1 is the most important and practical one because of a rising need for providing corporate remote users with continuous access to their Intranet resources. After analyzing each scenario, one realizes that problems

occurring in scenarios in Sections 2.2 and 2.4 are either the same as those in the scenario in Section 2.1 or a subset of them. Therefore, solving the scenario in Section 2.1 will also solve the scenarios in Sections 2.2 and 2.4. The scenarios in Sections 2.3 and 2.5 do not introduce functional problems resulting from MIPv4 and VPN co-existence, and thus 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. In subsequent sections, 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.

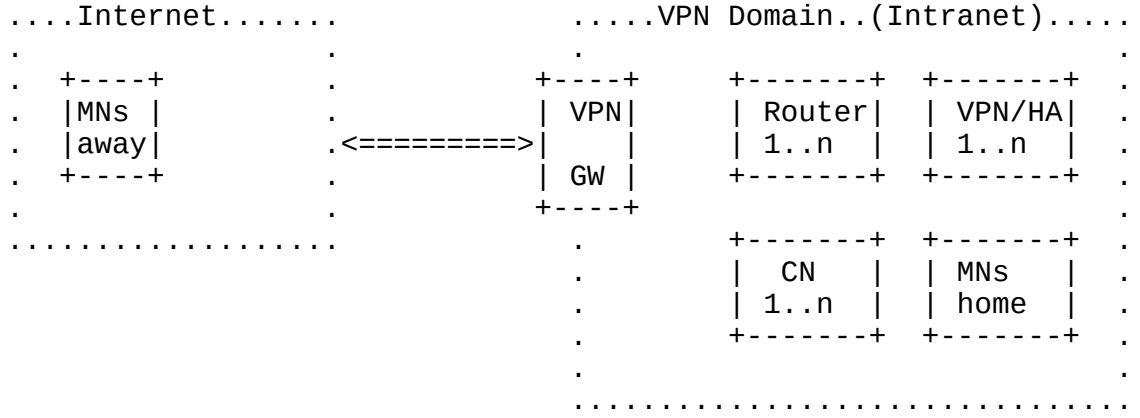

Figure 6: Intranet protected by a VPN gateway

The Intranet, as 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.

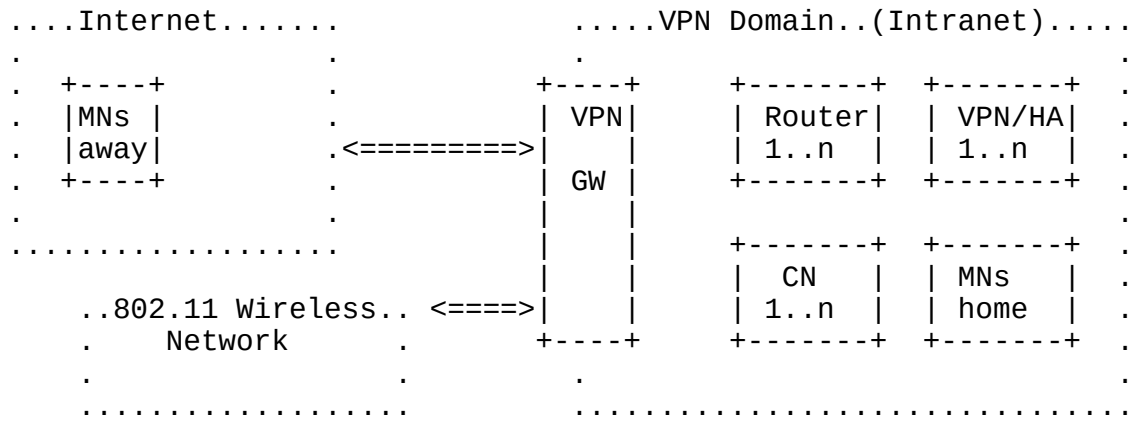

Figure 7: IEEE 802.11 Wireless deployment outside the home network

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.

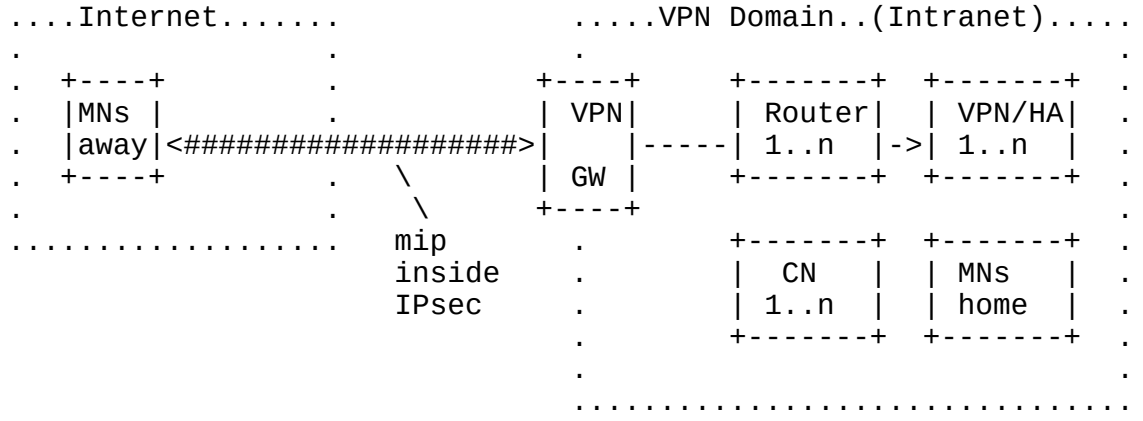

Figure 8

The MN obtains an address at its point of attachment (via DHCP [RFC2131] or some other means), and then sets up an IPsec tunnel to the VPN gateway, after which it can successfully register with its HA through the IPsec tunnel. The IPsec tunnel SA (Security Association) is identified by a triplet consisting of SPI (Security Parameter Index), MN's IP destination address (i.e., the address obtained at the point of attachment), and Security Protocol (AH or ESP) Identifier as described in [RFC2401]. This means that as the MN's IP

destination address changes on each IP subnet handoff, the IPsec tunnel needs to be re-established. This could have noticeable performance implications on real-time applications and in resource-constrained wireless networks. 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 DHCP allocated address and co-located mode is possible, we run into severe trouble. This is illustrated in Figure 9 and explained below:

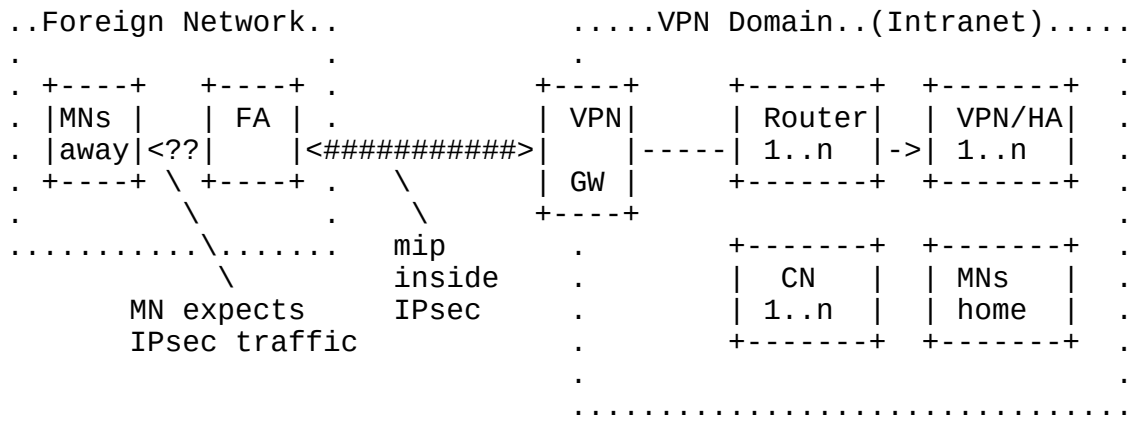

Figure 9

When arriving at the visited network on the left in this figure, the MN has to reach the FA with registration requests in order to have the FA send them on to the HA. However, the MN in all likelihood cannot register with the FA because the registration requests will be sent encrypted, and the FA will not be able to decrypt them. If the MN would have a policy that allowed split tunneling so that it could reach the FA with clear text messages, then the FA would still not be able to get through the VPN gateway unless the HA is reachable from outside and the Intranet security policy allows MIP registration packets to bypass the VPN gateway.

Even if the HA is reachable and the MIP registration succeeds, the FA (which is likely in a different administrative domain) will not be able to relay packets between the MN and the VPN gateway. Packets from the MN will be encapsulated by the FA with IP-in-IP [RFC2003], which the VPN gateway will drop, and packets from the VPN gateway will have ESP payloads (with IP-in-IP inside), which the FA will drop (as it expects IP-in-IP-encapsulated traffic to the MN).

The use of a 'trusted FA' has also been suggested in this scenario, meaning an FA that is actually a combined VPN GW and FA. The scenario will work fine in this case, as the tunnel end-points are at the FA and the VPN gateway as shown in Figure 10 below. However, we cannot expect that the FA in access networks (e.g., wireless hot-spots or CDMA 2000 networks) will have security associations with any given corporate network, so this is not particularly realistic in the general mobility case.

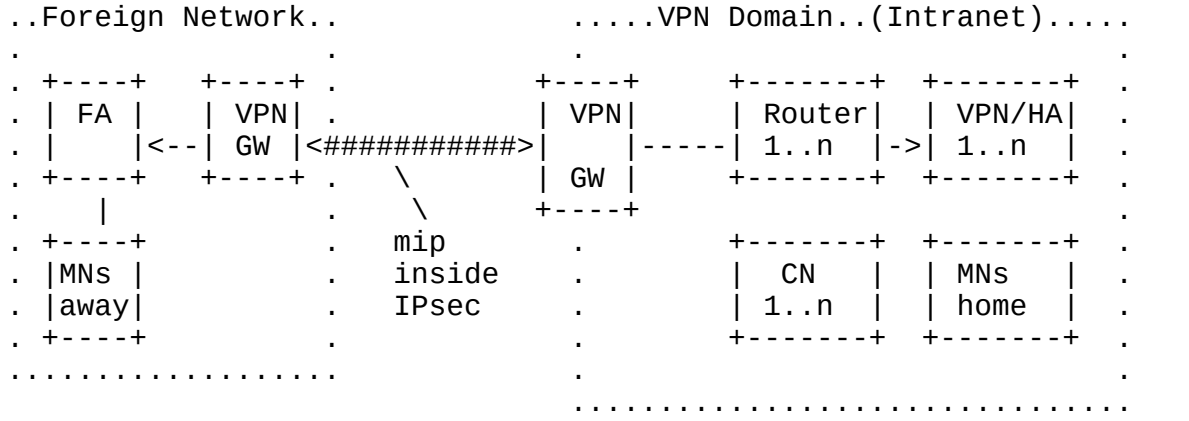

Figure 10

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 thus that it 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.

In co-located mode, the IPsec tunnel needs to be re-established on each IP subnet handoff, which will have performance implications on real-time applications and resource-constrained wireless networks.

In non-co-located mode (i.e., using an FA care-of address), the problem becomes severe, as the MN may be unable to register with its HA through the FA because the FA cannot understand MIPv4 registration

requests if they are encrypted in the IPsec tunnel (i.e., split tunneling is not supported). Even if the MN could reach the FA with non-encrypted registration requests (i.e., split tunneling is supported), and the requests going from the FA to the HA can pass through the VPN gateway, there would still be a problem with routing of data packets between the Intranet and the internet. This is because the VPN will not allow IP-in-IP-encapsulated packets from the FA to go through. And furthermore, ESP-encapsulated packets from the VPN gateway to the MN will be dropped by the FA, as it expects IP-in-IP-encapsulated traffic to the MN.

5. Solution Guidelines

This section describes guidelines for a solution to MIPv4 traversal across VPN gateways.

5.1. Preservation of Existing VPN Infrastructure

- The solution MUST work with currently deployed VPN gateways. This is the whole raison d'etre of this investigation: Finding a way to deploy Mobile-IP in cases where a VPN solution is already in place.

5.2. Software Upgrades to Existing VPN Client and Gateways

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

5.3. IPsec Protocol

- The solution SHOULD NOT require any changes to existing IPsec or key-exchange standard protocols implemented by VPN gateways.

- 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

- The solution MUST provide multi-vendor interoperability, whereby 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 that purchase and deploy best-of-breed multi-vendor solutions for IP routing, VPNs, firewalls, etc.

5.5. MIPv4 Protocol

- The solution MUST adhere to MIPv4 protocol [RFC3344]. That is, the solution MUST NOT impose any changes that violate MIPv4 protocol.

- The solution MAY introduce new extensions to MIPv4 nodes per guidelines specified in the MIPv4 protocol [RFC3344]. However, in order to overcome barriers to deployment, it is highly desirable to avoid any changes to MIPv4 mobility agents such as the FA and HA.

- The solution MAY require more than one instance of MIPv4 running in parallel (multiple encapsulation).

5.6. Handoff Overhead

- It is imperative to keep the key management overhead down to a minimum, in order to support fast handoffs across IP subnets. Therefore, 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

- The solution complexity MUST increase at most linearly with the number of MNs registered and accessing resources inside the Intranet.

- The solution MAY introduce additional header or tunneling overhead if needed.

5.8. Functional Entities

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

5.9. Implications of Intervening NAT Gateways

- The solution MUST be able to work with the existing MIPv4 and IPsec NAT traversal solutions [RFC3519] [RFC3715] [RFC3947].

5.10. Security Requirements

- o The solution MUST provide security that 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.

6. Security Considerations

This document describes an existing problem and proposes guidelines for possible solutions; as such, its security implications are indirect, through the guidelines it proposes for the solutions. Section 5.10 gives the relevant security requirements.

7. Acknowledgements

The authors who contributed text to this document were, in no particular order: Farid Adrangi, Milind Kulkarni, Gopal Dommety, 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, and Brijesh Kumar, for their feedback and help in improving this document.

8. References

8.1. Normative References

[RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.

8.2. Informative References

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC3519] Levkowetz, H. and S. Vaarala, "Mobile IP Traversal of Network Address Translation (NAT) Devices", RFC 3519, May 2003.

[RFC3715] Aboba, B. and W. Dixon, "IPsec-Network Address Translation (NAT) Compatibility Requirements", RFC 3715, March 2004.

[RFC3947] Kivinen, T., Swander, B., Huttunen, A., and V. Volpe, "Negotiation of NAT-Traversal in the IKE", RFC 3947, January 2005.

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
Ericsson Research
Torshamsgatan 23
SE-164 80 Stockholm
SWEDEN

Phone: +46 7 08 32 16 08
EMail: henrik@levkowetz.com

Full Copyright Statement

Intellectual Property

Acknowledgement

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