

Network Working Group
Request for Comments: 4487
Category: Informational

F. Le
CMU
S. Faccin
B. Patil
Nokia
H. Tschofenig
Siemens
May 2006

Mobile IPv6 and Firewalls: Problem Statement

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document captures the issues that may arise in the deployment of IPv6 networks when they support Mobile IPv6 and firewalls. The issues are not only applicable to firewalls protecting enterprise networks, but are also applicable in 3G mobile networks such as General Packet Radio Service / Universal Mobile Telecommunications System (GPRS/UMTS) and CDMA2000 networks.

The goal of this document is to highlight the issues with firewalls and Mobile IPv6 and act as an enabler for further discussion. Issues identified here can be solved by developing appropriate solutions.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Abbreviations	4
4. Overview of Firewalls	4
5. Analysis of Various Scenarios Involving MIPv6 Nodes and Firewalls	6
5.1. Scenario Where the Mobile Node Is in a Network Protected by Firewall(s)	7
5.2. Scenario Where the Correspondent Node Is in a Network Protected by Firewall(s)	9
5.3. Scenario Where the HA Is in a Network Protected by Firewall(s)	12
5.4. Scenario Where the MN Moves to a Network Protected by Firewall(s)	12
6. Conclusions	13
7. Security Considerations	14
8. Acknowledgements	14
9. References	14
9.1. Normative References	14
9.2. Informative References	14
Appendix A. Applicability to 3G Networks	15

1. Introduction

Network elements such as firewalls are an integral aspect of a majority of IP networks today, given the state of security in the Internet, threats, and vulnerabilities to data networks. Current IP networks are predominantly based on IPv4 technology, and hence firewalls have been designed for these networks. Deployment of IPv6 networks is currently progressing, albeit at a slower pace. Firewalls for IPv6 networks are still maturing and in development.

Mobility support for IPv6 has been standardized as specified in RFC 3775. Given the fact that Mobile IPv6 is a recent standard, most firewalls available for IPv6 networks do not support Mobile IPv6.

Unless firewalls are aware of Mobile IPv6 protocol details, these security devices will interfere with the smooth operation of the protocol and can be a detriment to deployment.

Mobile IPv6 enables IP mobility for IPv6 nodes. It allows a mobile IPv6 node to be reachable via its home IPv6 address irrespective of any link that the mobile attaches to. This is possible as a result of the extensions to IPv6 defined in the Mobile IPv6 specification [1].

Mobile IPv6 protocol design also incorporates a feature termed Route Optimization. This set of extensions is a fundamental part of the protocol that enables optimized routing of packets between a mobile node and its correspondent node and therefore optimized performance of the communication.

In most cases, current firewall technologies, however, do not support Mobile IPv6 or are not even aware of Mobile IPv6 headers and extensions. Since most networks in the current business environment deploy firewalls, this may prevent future large-scale deployment of the Mobile IPv6 protocol.

This document presents in detail some of the issues that firewalls present for Mobile IPv6 deployment, as well as the impact of each issue.

2. Terminology

Return Routability Test (RRT): The Return Routability Test is a procedure defined in RFC 3775 [1]. It is performed prior to the Route Optimization (RO), where a mobile node (MN) instructs a correspondent node (CN) to direct the mobile node's data traffic to its claimed care-of address (CoA). The Return Routability procedure provides some security assurance and prevents the misuse of Mobile IPv6 signaling to maliciously redirect the traffic or to launch other attacks.

3. Abbreviations

This document uses the following abbreviations:

- o CN: Correspondent Node
- o CoA: Care of Address
- o CoTI: Care of Test Init
- o HA: Home Agent
- o HoA: Home Address
- o HoTI: Home Test Init
- o HoT: Home Test
- o MN: Mobile Node
- o RO: Route Optimization
- o RRT: Return Routability Test

4. Overview of Firewalls

The following section provides a brief overview of firewalls. It is intended as background information so that issues with the Mobile IPv6 protocol can then be presented in detail in the following sections.

There are different types of firewalls, and state can be created in these firewalls through different methods. Independent of the adopted method, firewalls typically look at five parameters of the traffic arriving at the firewalls:

- o Source IP address
- o Destination IP address
- o Protocol type
- o Source port number
- o Destination port number

Based on these parameters, firewalls usually decide whether to allow the traffic or to drop the packets. Some firewalls may filter only incoming traffic, while others may also filter outgoing traffic.

According to Section 3.29 of RFC 2647 [2], stateful packet filtering refers to the process of forwarding or rejecting traffic based on the contents of a state table maintained by a firewall. These types of firewalls are commonly deployed to protect networks from different threats, such as blocking unsolicited incoming traffic from the external networks. The following briefly describes how these firewalls work since they can create additional problems with the Mobile IPv6 protocol as described in the subsequent sections.

In TCP, an MN sends a TCP SYN message to connect to another host in the Internet.

Upon receiving that SYN packet, the firewall records the source IP address, the destination IP address, the Protocol type, the source port number, and the destination port number indicated in that packet before transmitting it to the destination.

When an incoming message from the external networks reaches the firewall, it searches the packet's source IP address, destination IP address, Protocol type, source port number, and destination port number in its entries to see if the packet matches the characteristics of a request sent previously. If so, the firewall allows the packet to enter the network. If the packet was not solicited from an internal node, the packet is blocked.

When the TCP close session packets are exchanged or after some configurable period of inactivity, the associated entry in the firewall is deleted. This mechanism prevents entries from remaining when TCP are abruptly terminated.

A similar entry is created when using UDP. The difference with this transport protocol is that UDP is connectionless and does not have packets signaling the initiation or termination of a session. Consequently, the duration of the entries relies solely on timers.

5. Analysis of Various Scenarios Involving MIPv6 Nodes and Firewalls

The following section describes various scenarios involving MIPv6 nodes and firewalls and also presents the issues related to each scenario.

The Mobile IPv6 specifications define three main entities: the mobile node (MN), the correspondent node (CN), and the home agent (HA). Each of these entities may be in a network protected by one or many firewalls:

- o Section 5.1 analyzes the issues when the MN is in a network protected by firewall(s)
- o Section 5.2 analyzes the issues when the CN is in a network protected by firewall(s)
- o Section 5.3 analyzes the issues when the HA is in a network protected by firewall(s)

The MN may also be moving from an external network, to a network protected by firewall(s). The issues of this case are described in Section 5.4.

Some of the described issues (e.g., Sections 5.1 and 5.2) may require modifications to the protocols or to the firewalls, and others (e.g., Section 5.3) may require only that appropriate rules and configuration be in place.

5.1. Scenario Where the Mobile Node Is in a Network Protected by Firewall(s)

Let's consider MN A, in a network protected by firewall(s).




Figure 1: Issues between MIPv6 and firewalls when MN is in a network protected by firewalls

A number of issues need to be considered:

Issue 1: When MN A connects to the network, it should acquire a local IP address (CoA) and send a Binding Update (BU) to its Home Agent to update the HA with its current point of attachment. The Binding Updates and Acknowledgements should be protected by IPsec ESP according to the MIPv6 specifications [1]. However, as a default rule, many firewalls drop IPsec ESP packets because they cannot determine whether inbound ESP packets are legitimate. It is difficult or impossible to create useful state by observing the outbound ESP packets. This may cause the Binding Updates and Acknowledgements between the mobile nodes and their home agent to be dropped.

Issue 2: Let's now consider a node in the external network, B, trying to establish a communication with MN A.

- * B sends a packet to the mobile node's home address.
- * The packet is intercepted by the MN's home agent, which tunnels it to the MN's CoA [1].
- * When arriving at the firewall(s) protecting MN A, the packet may be dropped since the incoming packet may not match any existing state. As described in Section 4, stateful inspection packet filters (for example) typically drop unsolicited incoming traffic.

- * B will thus not be able to contact MN A and establish a communication.

Even though the HA is updated with the location of an MN, firewalls may prevent correspondent nodes from establishing communications when the MN is in a network protected by firewall(s).

Issue 3: Let's assume a communication between MN A and an external node B. MN A may want to use Route Optimization (RO) so that packets can be directly exchanged between the MN and the CN without passing through the HA. However, the firewalls protecting the MN might present issues with the Return Routability procedure that needs to be performed prior to using RO.

According to the MIPv6 specifications, the Home Test message of the RRT must be protected by IPsec in tunnel mode. However, firewalls might drop any packet protected by ESP, since the firewalls cannot analyze the packets encrypted by ESP (e.g., port numbers). The firewalls may thus drop the Home Test messages and prevent the completion of the RRT procedure.

Issue 4: Let's assume that MN A successfully sends a Binding Update to its home agent (resp. correspondent nodes) -- which solves issue 1 (resp. issue 3) -- and that the subsequent traffic is sent from the HA (resp. CN) to the MN's CoA. However there may not be any corresponding state in the firewalls. The firewalls protecting A may thus drop the incoming packets.

The appropriate states for the traffic to the MN's CoA need to be created in the firewall(s).

Issue 5: When MN A moves, it may move to a link that is served by a different firewall. MN A might be sending a BU to its CN; however, incoming packets may be dropped at the firewall, since the firewall on the new link that the MN attaches to does not have any state that is associated with the MN.

The issues described above result from the fact that the MN is behind the firewall. Consequently, the MN's communication capability with other nodes is affected by the firewall rules.

5.2. Scenario Where the Correspondent Node Is in a Network Protected by Firewall(s)

Let's consider an MN in a network, communicating with a Correspondent Node C in a network protected by firewall(s). There are no issues with the presence of a firewall in the scenario where the MN is sending packets to the CN via a reverse tunnel that is set up between the MN and HA. However, firewalls may present different issues to Route Optimization.




Figure 2: Issues between MIPv6 and firewalls when a CN is in a network protected by firewalls

The following issues need to be considered:

Issue 1: The MN (MN B) should use its Home Address (HoA B) when establishing the communication with the CN (CN C), if MN B wants to take advantage of the mobility support provided by the Mobile IPv6 protocol for its communication with CN C. The state created by the firewall protecting CN C is therefore created based on the IP address of C (IP C) and the home address of Node B (IP HoA B). The states may be created via different means, and the protocol type as well as the port numbers depend on the connection setup.

Uplink packet filters (1)

Source IP address: IP C

Destination IP address: HoA B

Protocol Type: TCP/UDP

Source Port Number: #1

Destination Port Number: #2

Downlink packet filters (2)

Source IP address: HoA B

Destination IP address: IP C

Protocol Type: TCP/UDP

Source Port Number: #2

Destination Port Number: #1

Nodes C and B might be topologically close to each other, while B's home agent may be far away, resulting in a trombone effect that can create delay and degrade the performance. MN B may decide to initiate the route optimization procedure with Node C. Route optimization requires MN B to send a Binding Update to Node C in order to create an entry in its binding cache that maps the MN's home address to its current care-of-address. However, prior to sending the binding update, the mobile node must first execute a Return Routability Test:

- * Mobile Node B has to send a Home Test Init (HoTI) message via its home agent and
- * a Care of Test Init (COTI) message directly to its Correspondent Node C.

The Care of Test Init message is sent using the CoA of B as the source address. Such a packet does not match any entry in the protecting firewall (2). The CoTi message will thus be dropped by the firewall.

The HoTI is a Mobility Header packet, and as the protocol type differs from the established state in the firewall (see (2)), the HoTI packet will also be dropped.

As a consequence, the RRT cannot be completed, and route optimization cannot be applied. Every packet has to go through Node B's home agent and tunneled between B's home agent and B.




Figure 3: Issues with Return Routability Test

Issue 2: Let's assume that the Binding Update to the CN is successful; the firewall(s) might still drop packets that are:

1. coming from the CoA, since these incoming packets are sent from the CoA and do not match the Downlink Packet filter (2).
2. sent from the CN to the CoA if uplink packet filters are implemented. The uplink packets are sent to the MN's CoA and do not match the uplink packet filter (1).

The packet filters for the traffic sent to (resp. from) the CoA need to be created in the firewall(s).

Requiring the firewalls to update the connection state upon detecting Binding Update messages from a node outside the network protected by the firewall does not appear feasible or desirable, since currently the firewall does not have any means to verify the validity of Binding Update messages and therefore to modify the state information securely. Changing the firewall states without verifying the validity of the Binding Update messages could lead to denial of service attacks. Malicious nodes may send fake binding updates, forcing the firewall to change its state information, and therefore leading the firewall to drop packets from the connections that use the legitimate addresses. An adversary might also use an address update to enable its own traffic to pass through the firewall and enter the network.

Issue 3: Let's assume that the Binding Update to the CN is successful. The CN may be protected by different firewalls, and as a result of the MN's change of IP address, incoming and outgoing traffic may pass through a different firewall. The new

firewall may not have any state associated with the CN, and incoming packets (and potentially outgoing traffic as well) may be dropped at the firewall.

Firewall technology allows clusters of firewalls to share state [3]. This, for example, allows the support of routing asymmetry. However, if the previous and the new firewalls, through which the packets are routed after the Binding Update has been sent, do not share state, this may result in packets being dropped at the new firewall. As the new firewall does not have any state associated with the CN, incoming packets (and potentially outgoing traffic as well) may be dropped at the new firewall.

5.3. Scenario Where the HA Is in a Network Protected by Firewall(s)

In the scenarios where the home agent is in a network protected by firewall(s), the following issues may exist:

Issue 1: If the firewall(s) protecting the home agent block ESP traffic, much of the MIPv6 signaling (e.g., Binding Update, HoT) may be dropped at the firewall(s), preventing MN(s) from updating their binding cache and performing Route Optimization, since Binding Update, HoT, and other MIPv6 signaling must be protected by IPsec ESP.

Issue 2: If the firewall(s) protecting the home agent block unsolicited incoming traffic (e.g., as stateful inspection packet filters do), the firewall(s) may drop connection setup requests from CNs, and packets from MNs.

Issue 3: If the home agent is in a network protected by several firewalls, an MN/CN's change of IP address may result in the passage of traffic to and from the home agent through a different firewall that may not have the states corresponding to the flows. As a consequence, packets may be dropped at the firewall.

5.4. Scenario Where the MN Moves to a Network Protected by Firewall(s)

Let's consider an HA in a network protected by firewall(s). The following issues need to be investigated:

Issue 1: Similarly to issue 1 described in Section 5.1, the MN will send a Binding Update to its home agent after acquiring a local IP address (CoA). The Binding Updates and Acknowledgements should be protected by IPsec ESP according to the MIPv6 specifications [1]. However, as a default rule, many firewalls drop ESP packets. This may cause the Binding Updates and Acknowledgements between the mobile nodes and their home agent to be dropped.

Issue 2: The MN may be in a communication with a CN, or a CN may be attempting to establish a connection with the MN. In both cases, packets sent from the CN will be forwarded by the MN's HA to the MN's CoA. However, when the packets arrive at the firewall(s), the incoming traffic may not match any existing state, and the firewall(s) may therefore drop it.

Issue 3: If the MN is in a communication with a CN, the MN may attempt to execute an RRT for packets to be route optimized. Similarly to issue 3, Section 5.1, the Home Test message that should be protected by ESP may be dropped by firewall(s) protecting the MN. Firewall(s) may as a default rule drop any ESP traffic. As a consequence, the RRT cannot be completed.

Issue 4: If the MN is in a communication with a CN, and assuming that the MN successfully sent a Binding Update to its CN to use Route Optimization, packets will then be sent from the CN to the MN's CoA and from the MN's CoA to the CN.

Packets sent from the CN to the MN's CoA may, however, not match any existing entry in the firewall(s) protecting the MN, and therefore be dropped by the firewall(s).

If packet filtering is applied to uplink traffic (i.e., traffic sent by the MN), packets sent from the MN's CoA to the CN may not match any entry in the firewall(s) either and may be dropped as well.

6. Conclusions

Current firewalls may not only prevent route optimization but may also prevent regular TCP and UDP sessions from being established in some cases. This document describes some of the issues between the Mobile IPv6 protocol and current firewall technologies.

This document captures the various issues involved in the deployment of Mobile IPv6 in networks that would invariably include firewalls. A number of different scenarios are described, which include configurations where the mobile node, correspondent node, and home agent exist across various boundaries delimited by the firewalls. This enables a better understanding of the issues when deploying Mobile IPv6 as well as the issues for firewall design and policies to be installed therein.

7. Security Considerations

This document describes several issues that exist between the Mobile IPv6 protocol and firewalls.

Firewalls may prevent Mobile IP6 signaling in addition to dropping incoming/outgoing traffic.

If the firewall configuration is modified in order to support the Mobile IPv6 protocol but not properly configured, many attacks may be possible as outlined above: malicious nodes may be able to launch different types of denial of service attacks.

8. Acknowledgements

We would like to thank James Kempf, Samita Chakrabarti, Giaretta Gerardo, Steve Bellovin, Henrik Levkowetz, and Spencer Dawkins for their valuable comments. Their suggestions have helped improve both the presentation and the content of the document.

9. References

9.1. Normative References

- [1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

9.2. Informative References

- [2] Newman, D., "Benchmarking Terminology for Firewall Performance", RFC 2647, August 1999.
- [3] Noble, J., Doug, D., Hourihan, K., Hourihan, K., Stephens, R., Stiefel, B., Amon, A., and C. Tobkin, "Check Point NG VPN-1/Firewall-1 Advanced Configuration and Troubleshooting", Syngress Publishing Inc., 2003.
- [4] Chen, X., Rinne, J., Wiljakka, J., and M. Watson, "Problem Statement for MIPv6 Interactions with GPRS/UMTS Packet Filtering", Work in Progress, January 2006.

Appendix A. Applicability to 3G Networks

In 3G networks, different packet filtering functionalities may be implemented to prevent malicious nodes from flooding or launching other attacks against the 3G subscribers. The packet filtering functionality of 3G networks is further described in [4]. Packet filters are set up and applied to both uplink and downlink traffic: outgoing and incoming data not matching the packet filters is dropped. The issues described in this document also apply to 3G networks.

Authors' Addresses

Franck Le
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
USA

Email: franckle@cmu.edu

Stefano Faccin
Nokia Research Center
6000 Connection Drive
Irving, TX 75039
USA

Email: sfaccinstd@gmail.com

Basavaraj Patil
Nokia
6000 Connection Drive
Irving, TX 75039
USA

Email: Basavaraj.Patil@nokia.com

Hannes Tschofenig
Siemens
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

Email: Hannes.Tschofenig@siemens.com
URI: <http://www.tschofenig.com>

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 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.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

