

An NSIS-based Approach for Firewall Traversal in Mobile IPv6 Networks

Niklas Steinleitner,
Xiaoming Fu,
Dieter Hogrefe
University of Göttingen
Göttingen, Germany
{steinleitner, fu, hogrefe}
@cs.uni-goettingen.de

Thomas Schreck
University of Applied Sciences
Landshut
Landshut, Germany
thomas.schreck
@fh-landshut.de

Hannes Tschofenig
Nokia Siemens Networks and
University of Göttingen
Munich, Germany
hannes.tschofenig@nsn.com

ABSTRACT

Firewalls have been successfully deployed in today's network infrastructure in various environments and will also be used in IPv6 networks. However, most of the current firewalls do not support Mobile IPv6, the best known standardized solution for mobility support in IPv6. As a result, Mobile IPv6 traffic will be most likely dropped when used without an appropriate firewall traversal solution.

This paper describes the problems and impacts of having firewalls in Mobile IPv6 environments and presents a firewall traversal solution based on the IETF's Next Steps In Signaling framework to address these issues. Compared with other candidates such as STUN, TURN, ICE, ALG, MID-COM and COPS, this approach does not rely on specific firewall placements and can be applied in various operational modes without additional introducing entities. In this paper we also explore security aspects since they are typically difficult to handle.

1. INTRODUCTION

Middleboxes, such as firewalls, are an important aspect for a majority of IP networks today. Current IP networks are predominantly based on IPv4 technology, and hence various firewalls (as well as Network Address Translators(NATs)) have been originally designed for these networks. Deployment of IPv6 networks is currently work in progress. However, some firewall products for IPv6 networks are already available. It is foreseen that firewalls will become an indispensable means for protecting against unwanted traffic in operational IPv6 networks, especially in enterprise environments.

Given the fact that Mobile IPv6 [1] is a recent standard, most firewalls available for IPv6 networks still do not support Mobile IPv6. Unless firewalls are aware of Mobile IPv6 protocol details, they will have to either block Mobile IPv6

communication traffic, or carefully deal with the traffic by per-user or per-connection, or allow this traffic in general through manual pre-configuration. This could be a major impediment to the successful deployment of Mobile IPv6.

Some existing firewall traversal solutions, such as STUN [2], TURN [3], ICE [4], Application Layer Gateways, Middlebox Communication [5], COPS [6] or policy-based solutions potentially can be extended for performing firewall and middlebox traversal even in mobile networks. However, some of them require prior knowledge of the existence of firewalls and most do not address the issue of discovering firewalls. Furthermore, they do not support the node mobility case and thus may require significant efforts to be extended for use in Mobile IPv6 networks.

A recent initiative within the IETF, Next Steps in Signaling (NSIS) [7], has developed a signaling protocol for firewall and NAT traversal, the NAT/Firewall NSLP (NAT/FW NSLP) [8]. NSIS utilizes a two-layer signaling paradigm, which defines a lower layer for general extensible IP signaling and a layer for various signaling applications such as signaling for NAT/Firewall traversal. Since its initial design, NSIS has been considering node mobility as its potential use scenarios. However, how the NSIS firewall/NAT traversal signaling protocol supports IPv6 mobility is not specified.

This paper will give an overview of the problems when firewalls are placed in Mobile IPv6 networks, identify potential approaches and present how to use NSIS to address the Mobile IPv6 firewall traversal issues.

The paper is structured as follows. In Section 2 we shortly describe the problems and impacts of having firewalls in Mobile IPv6 environments, as described in RFC 4487 [9], and identify potential state-of-the-art solutions. In Section 3 we present a middlebox traversal solution based on the NSIS signaling layer protocol for NAT/firewall traversal [8] and show how it can be used for firewall traversal in Mobile IPv6. Section 4 provides an analysis of potential authorization solutions and Section 5 discusses open issues and further work. Section 6 summarizes this paper.

2. PROBLEM STATEMENT

To study how firewall traversal can be achieved in Mobile IPv6 environments, it is necessary to understand the problems and impacts of having firewalls in such environments. Mobile IPv6 [1, 10] introduces several new types of messages, which can be categorized into registration messages

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON 2007, October 22-24, 2007, Austin, Texas, USA.
Copyright 2007 ACM 987-963-9799-04-2/07/10 ...\$5.00.

(Binding Update(BU), Binding Acknowledgements(BA)), Home/Care-of-testing messages (Home-of-Test-Init (HoTI), Home-of-Test (HoT), Care-of-Test-Init (CoTI), Care-of-Test (CoT)) and data traffic. A new mobility header is introduced in all this new messages, and all messages between the mobile node (MN) and the home agent (HA) are IPsec ESP [10] encapsulated.

When a user moves to a visited network, a firewall – no matter whether it is located in the home network, the visited network or the access network of the corresponding node – will affect the Mobile IPv6 signaling and data messages. For instance, route optimization, an integral part of Mobile IPv6 specification, does not work with the state-of-the-art firewalls that utilize stateful packet filtering (SPF). This set of extensions is a fundamental part of the protocol, enabling optimized routing of packets between a mobile node and its correspondent node, thus providing optimized communication performance. However, firewall technologies do not support Mobile IPv6 or are not even aware of IPv6 mobility extension headers. Since most networks in the current business environment deploy firewalls, this may prevent future large-scale deployment of Mobile IPv6. Secondly, another mode of communication in Mobile IPv6, namely bi-directional tunneling, does not work under some scenarios, e.g., when a firewall is placed in the access network or the home network. In addition, it is difficult for the Mobile IPv6 binding update packets (encapsulated using IPsec ESP) to traverse firewalls. In summary, these deployment issues with firewalls occur due to the nature that the commonly used firewalls possess [9]:

- do not understand Mobile IPv6 mobility header,
- do not allow IPsec – which is used for Mobile IPv6 registration messages between MN and HA – traffic to traverse,
- do not understand data packets encapsulated in Mobile IPv6 and likely drop them.

In the following subsections, we first explore these problems in detail from both operational and technical aspects regarding some relevant scenarios.

2.1 Scenarios and issues

Without loss of generality, let us consider a typical roaming scenario, where a mobile user with a PDA (MN) is roaming outside of his or her company (hereafter, the so-called “Mobile Service Provider”, or MSP) into a visited network (“Access Service Provider”, or ASP) which is also a corporate network. The MN wants to communicate with his home network or its HA (in order to register its new location) and additionally with another node, the corresponding node (CN), for data communication. The visited network could be protected by a firewall, thus parts of the traffic to the MN may be blocked. Besides, both the home network and the network of the CN may deploy firewalls. These three possible firewall placements introduce several problems, which could prevent Mobile IPv6 from operating successfully in the presence of firewalls. In all cases, pinholes have to be open on the firewalls for enabling successful communication. These problems can be differentiated under three basic scenarios.

- Firewall located at the edge of the MN’s ASP,
- Firewall located at the edge of the CN’s ASP,
- Firewall located at the edge of the MN’s MSP.

In the following sections we investigate these three basic scenarios individually, and show how a firewall might prevent

Mobile IPv6 from a successful operation.

2.1.1 Firewall located at the edge of MN’s ASP

The first scenario assumes that the MN roaming to another network (i.e., ASP, which deploys a firewall (ASP-FW)) wants to enjoy communication with his company or ISP (MSA/MSP/ASA). Therefore, the MN needs to traverse the ASP-FW. Figure 1 depicts how the components are placed in this scenario. Several issues need to be considered:




Figure 1: Firewall located at the edge of MN’s ASP

- Binding Updates and Binding Acknowledgements, should be protected by IPsec ESP, but many firewalls drop IPsec ESP packets because they cannot determine whether inbound ESP packets are authorized. A possible solution might be to manually pre-configure the ASP-FW so that MIPv6 traffic is allowed to traverse it. However, not every administrator would permit IPsec traffic in general, so it must also be possible to dynamically install these firewall rules.
- The ASP-FW may drop the Home Test messages and may prevent the completion of the Return Routability Test (RRT) procedure, as the Home Test messages of the RRT are protected by IPsec ESP in the tunnel mode. Therefore, either manual pre-configuration or dynamic on-demand configuration of rules on the ASP-FW is a possible solution for this type of messages.
- If the MN successfully sends a Binding Update to its HA and the subsequent traffic is sent from HA to MN (in bi-directional tunneling), there is also no corresponding state on the firewalls, and the firewalls drops the incoming packets. Hence, it is necessary to dynamically configure the ASP-FW to let this data traffic traverse.
- The ASP-FW may prevent correspondent nodes from establishing communications (e.g. route optimization traffic) because incoming packets are dropped since the packets do not match any existing state.
- If the MN roams and moves to another access network protected by a different firewall, all new incoming packets are dropped as they do not match any existing “allow” state.

2.1.2 Firewall located at the edge of CN’s ASP

Here, a MN wants to communicate with a CN that deploys a firewall. Therefore, the traffic from the MN to the CN needs to traverse the CN’s ASP-FW. Figure 2 depicts how the components are placed in the second scenario. Several issues need to be considered:




Figure 2: Firewall located at the edge of CN's ASP

- The Care-of-Test-Init message is sent using the Care-of-Address (CoA) of the MN as the source address. Such a packet does not match any entry in the protecting firewall, as the states in the firewall are bounded to the old address of the MN. The CoTI message will thus be dropped by that firewall. As a consequence, the RRT cannot be completed, and route optimization cannot be performed. Every packet has to be tunneled through the HA.
- If the BU to the CN is successful, the firewall still drops packets that are coming from the CoA, because these incoming packets are sent from the CoA and do not match any existing firewall state.

2.1.3 Firewall located at the edge of MN's MSP

In this scenario, the MN roaming to another company/ISP (i.e., ASP) wants to enjoy communicating with a CN and his own company (MSP), and the MSP deploys a firewall at its network border. The MN needs to traverse the MSP-FW to run Mobile IPv6.

Figure 3 depicts how the components are placed in third scenario. Several issues need to be considered:




Figure 3: Firewall located at the edge of MN's MSP

- If the firewall protects the home agent by blocking ESP traffic, some of the MIPv6 signaling (e.g., Binding Update, HoTI) may be dropped at the firewall. This prevents MNs from updating their binding cache and performing Route Optimization, since the messages must be protected by IPsec ESP. Manual pre-configuration is a solution, but also has some problems as mentioned before.
- If the firewall is a stateful packet filter and protects the home agent from unsolicited incoming traffic, the firewall may drop connection setup requests from CNs, and packets from MNs.

2.2 Requirements and Solution Alternatives

To get Mobile IPv6 working in these scenarios it is necessary to allow all this messages to traverse the firewall. This requires the usage of a middlebox configuration solution. In general we can distinguish between two types of middlebox configuration; the implicit and the explicit approaches. The implicit middlebox configuration is triggered by data traffic. A stateful packet filtering firewall or a NAT establish state information based on the header information in the data traffic itself. To the category of implicit approaches also belong STUN and TURN since these signaling protocols do not interact with the firewall itself but rather implement a hole-punching behavior as middleboxes treats these messages as ordinary data traffic. In contrast, with an explicit approach the intention is to interact with the middlebox and therefore the middlebox has to implement additional protocols. Application Layer Gateways, MIDCOM or the NAT/Firewall NSLP are examples of this approach.

The main differences between the two approaches is flexibility regarding the pinhole creation vs. the need to enhance existing middleboxes to understand additional protocols. Implicit approaches are less flexible regarding the creation of pinholes that often leads to the need to tunnel one protocol on top of another one. Additionally, since there is no interaction with the middlebox and the end host it is unclear how long established state is kept alive at the middlebox. As a consequence, more frequent refresh messages have to be transmitted to ensure that state information is not discarded. Finally, explicit approaches provide better security properties. However, the major disadvantage is the slow adoption of new protocols at middleboxes.

Since Mobile IPv6 networks yet have to be deployed on a wide scale there is a chance to enhance middlebox with additional protocols and hence we have chosen an explicit signaling approach.

Application Layer Gateways

Application Layer Gateways rely on the installation of an enhanced Firewall/NAT, called an ALG. This ALG understands the application layer protocol semantic. The ALG processes the signaling and data traffic and can modify the signaling to match the public IP addresses and ports that are used by the signaling and media traffic. The ALG is often transparent to end hosts. The ALG might be co-located with the middlebox itself or it interacts with it to setup state information, such as packet filters, or even modifies application specific payloads.

ALGs typically destroy the end-to-end semantic of a protocol and harm end-to-end security since they often modify bypassing payloads. These middleboxes make it more difficult to deploy new extensions since they often drop unknown extensions. Finally, there are sometimes performance problems caused by the deep packet inspection nature of the devices. A Session Border Controller is an example of an ALG in the context of SIP.

MIDCOM

One possible alternative is to use MIDCOM [5]. The main idea of MIDCOM is to move application logic from the middlebox into a trusted third entity. There are three main entities in the MIDCOM framework: middleboxes, MIDCOM agent, and MIDCOM Policy Decision Point

(PDP). MIDCOM agent is an entity performing ALG functions, which reside outside of the middlebox. It interacts with a middlebox to set up states, access control filters, extract middlebox state information, modify application specific payload, or perform other tasks necessary to enable middlebox traversal. The MIDCOM PDP acts as a policy repository, holding MIDCOM related policy profiles in order to make authorization decisions.

The decomposition in MIDCOM provides a number of advantages, including improved performance, lower software development and maintenance costs, and easier deployment of new applications. Nevertheless some disadvantages still exist. MIDCOM assumes to have knowledge about the network topology and the middlebox has to be contacted for starting data transmission. However, for a complex topologies, the task of middlebox discovery becomes a problem.

ICE/M-ICE

Another possible framework that could be used is Mobile IP Interactive Connectivity Establishment (M-ICE) [11] that builds on top of the Interactive Connectivity Establishment (ICE) [4] methodology. ICE is based on STUN [2] and TURN [3]. With M-ICE the ICE framework is applied to Mobile IPv6. M-ICE uses STUN for connectivity checks, a modified return routability procedure and UDP encapsulation of the signaling traffic as described in [12].

First M-ICE will gather the MN's candidates and afterwards will signal them to the CN by including them in the CoTI-ICE message. When the CN receives this message it will also start gathering its candidates and provides them in the CoT-ICE message. At that time both nodes could pair this candidates up and start connectivity checks by using STUN. After they finished this they both have a prioritized list of working candidate pairs.

ICE works reliably, as it is widely used for VoIP. For interacting with middleboxes an extension to STUN has been proposed that enables STUN-aware middleboxes to participate in the signaling exchange, see [13]. Authorization functionality has been proposed with [14].

3. MOBILE IPV6 FIREWALL TRAVERSAL BASED ON NSIS

This section describes how an extended NSIS [7] NAT/Firewall NSLP [8] could be utilized to compose the Mobile IPv6 firewall pinhole creation. This approach has the advantage of being a modular IETF standard protocol able to configure stateful packet filters. One particular advantage is that the NSIS NAT/FW NSLP framework relies on a soft-state approach. Therefore, established sessions will be automatically torn down after a specified timeout. This is very useful in a mobile scenario as it is not necessary to delete a session after roaming to another network. The University of Göttingen has developed an open source implementation of NSIS protocol stack [15], including a NAT/FW NSLP implementation, which allows customized extensions for development. The following section gives an overview of the NSIS framework and the NAT/Firewall NSLP framework, developed by the IETF NSIS Working Group. It also describes how NSIS and the NAT/FW NSLP is applicable for Mobile IPv6 firewall traversal.

3.1 NSIS Introduction

The NSIS framework [7] has been developed with the goal of supporting various signaling applications, which install and manipulate certain control states in the network. Such states are meaningful for data flows and are installed and manipulated on network nodes supporting NSIS (NSIS Entities, NEs) along the data path. Not every node has to be such an NE, for instance, in the the NAT/FW NSLP case only NAT/Firewall boxes need to be the NEs along the data path of a data flow besides the end hosts. The basic protocol concept does not depend on any signaling application. This section describes the fundamental entities involved in NSIS signaling and their basic interactions. Two NSIS entities that communicate directly are said to be in a "peer relationship". Thereby, either or both NEs can store state information about the other NE, but it is not mandatory to establish a long-term signaling connection between them.




Figure 4: Simple Signaling and Data Flow Example

Figure 4 shows one of the simplest possible signaling configurations. A data flow is flowing from the sender via different routers to the receiver. The two end hosts and two of the routers contain NEs that exchange signaling messages about the flow. R3 does not contain an NE and forwards only the data. The signaling messages exchange is possible in both directions. Before a data flow is sent, an NSIS signaling procedure will take place along the NEs in the data path, including discovering their existence and signaling the application-specific states (e.g., firewall configurations for corresponding data traversal).

3.2 NSIS Layered Model Overview

In order to meet the modular requirements for NSIS, the NSIS protocol is structured in two layers:

- The NSIS Transport Layer Protocol (NTLP), which is responsible for moving signaling messages around and nevertheless independent from the underlying signaling application. The NTLP is implemented by GIST [15].
- The NSIS Signaling Layer Protocol (NSLP), which allows application based functionalities, such as message formats and sequences. Figure 5 illustrates this modular NSIS approach and the mutual influence between the NTLP and the NSLP.

Functionality within the NTLP is restricted only to transport and lower-layer operations. Other operations are relocated to the signaling application layer. A short introduction of the NTLP can be described as follows.

When an NSLP signaling message needs to be sent, the NSLP gives it over to the NTLP together with the information to which flow it belongs (so-called flow identifier). The NTLP has to care about how the message is sent to the next NE along the path. The NTLP does not need to




Figure 5: The NSIS Protocol Components

have any knowledge about addresses, capabilities, or status of other NEs along the path, only for the NEs that it directly peers with.

Upon receipt of an NSIS message, each intermediate NTLP either directly forwards it or - if the signaling application runs locally - passes the message to the NSLP for further processing. After processing, the NSLP can use the original message or generate another message and hands it over to the NTLP. With this procedure end-to-end NSIS message delivery can be achieved. This restriction of the NTLP to peer-relationship scope simplifies the management and the complexity of the NTLP, at the cost of an increased functionality, complexity of the NSLPs and deployment complexity, as some components (e.g., middleboxes) on the path need to run NSIS.

3.3 The NAT/FW NSLP Protocol

The IETF NSIS working group is currently finalizing the NAT/Firewall NSIS Signaling Layer protocol (NAT/FW NSLP) specification [8], which describes scenarios, problems and solutions for path-coupled network address translator and firewall signaling. The NAT/FW NSLP is one of the two NSLPs that the working group has been developing. Our previous work [16] has shown that NSIS and the NAT/FW NSLP framework is able to support firewall signaling for up to tens of thousands of flows in parallel even in a low-end environment; and the overall performance bottleneck was the utilized firewall implementation, not on the signaling implementation.

The main goal of NSIS NAT/FW signaling is to enable communications between two endpoints across different networks in case of the existence of NATs and firewall middleboxes. Firstly, it is assumed that these middleboxes will be configured in such a way that NSIS NAT/FW signaling messages can traverse them. Then, the NSIS NAT/FW NSLP protocol is used to dynamically install additional policy rules in all NAT/FW NSLP-aware middleboxes along the path. Firewalls will be configured to forward desired data packets according to the policy rules which are established by the NAT/FW NSLP signaling.

The signaling traffic of an application behind a middlebox has to traverse all middleboxes along the data path to establish communication with a corresponding application on the other end host. To achieve middlebox traversal, the application triggers the local NSIS entity to signal along the data path. If the local NSIS entity supports NAT/FW NSLP signaling, the knowledge of these application is used to establish policy rules and NAT bindings in all middleboxes along the path, which allows the data to travel from

the sender to the receiver. Clearly, it is necessary for intermediate middleboxes to support NAT/FW NSLP, but not necessary for other intermediate nodes to support NAT/FW NSLP or even NSIS.

Figure 6 shows a common topology for the use of NAT/FW NSLP. This network is separated into two distinct administrative domains, namely “Domain A” and “Domain B”.




Figure 6: A Firewall Traversal Scenario

The NSLP Initiator (NI) sends NSIS NAT/FW NSLP signaling messages along the data path to the NSLP Responder (NR). It is assumed that NI, NR and every intermediate middlebox implements the NAT/FW NSLP. Thereby, no knowledge about the next middlebox along the path is required; this is done by on-path next-hop discovery. The signaling messages reach different intermediate NSIS nodes (i.e., NSLP Forwarder or NF) and every NAT/FW NSLP node processes the signaling messages and, if necessary, installs additional rules for the following data packets. The NAT/FW NSLP supports several types of signaling messages, most notably the CREATE and the EXT messages:

- The CREATE message is sent from the source address to the destination address and processed by every middlebox and forwarded to the destination.
- The EXT message is sent from the source address to an external address (e.g. the HA’s address or the CN’s address) and is intercepted by the edge firewall and not forwarded to the destination address. This allows signaling pinholes at the edge-firewall without introducing long end-to-end signaling delays.
- The RESPONSE message is used as a response to CREATE and EXT request messages.

Policy rules for firewalls are represented by a common 5-tuple, namely the source and destination addresses, the transport protocol and the source and destination port, in addition to the rule action with the value “allow” or “deny”. Such a policy rule in NAT/FW NSLP is bounded to a specified session. Different from other signaling applications where policy rules are carried in one object, the policy rules in NAT/FW NSLP are divided into an action (allow/deny), the flow identifier and further information. The message routing information (MRI) in the NTLP carries the filter specification, the additional information such as lifetime, session ID, message sequence number, authorization objects and the specified action are carried in NSLP’s objects.

3.4 NSIS for Mobile IPv6 Firewall Traversal

As described in Section 2, the standard Mobile IPv6 does not work with the existence of firewalls. To tackle these issues, one approach is to utilize a signaling protocol to install some firewall rules to allow these Mobile IPv6 messages to pass through. The NSIS NAT/FW NSLP, as described in [8], allows an end system to establish, maintain and delete middlebox state (i.e., firewall rules), and as well as allows packets to traverse these boxes. This protocol thus provides

a possible way to address the aforementioned problems [17]. The following subsections introduce how we could extend the NSIS NAT/FW NSLP to solve the problems.

3.4.1 Firewall located at the edge of MN's ASP

In Figure 1, the MN is protected by a firewall that employs stateful packet filtering. The external CN and the HA are also shown in the figure. The MN is located in a visited network and is expecting to communicate with the CN. If the MN initiated normal data traffic there is no problem with the SPF firewall, as the communication is initiated from internal. The following subsections explain how this approach manages the MIPv6 signaling traffic problems as described in Section 2.

Binding updates

IPsec protected binding updates cause problems in some deployment environments, as described in RFC4487 [9]. As a solution, NAT/FW NSLP can be used to dynamically configure the firewall(s) to allow the IPsec packets and associated traffic like IKE/IKEv2 packets to traverse, before sending the binding updates. Therefore, IP Protocol ID 50 should be allowed in the filter policies in order to allow IPsec ESP and IP Protocol ID 51 to allow IPsec AH. The firewall should also allow IKE packets (to UDP port 500) to bypass, which can also be signaled before.




Figure 7: Signaling for BU and BA

Figure 7 shows the message flow for this signaling. As the firewall is a SPF, the subsequent binding acknowledgement from the HA to the CoA can pass the firewall, as it matches an existing state in the table.

Route optimization

Immediately after moving into a new network, the MN acquires a new CoA, performs the pinhole creation as described before and runs the Binding Update to the HA. The HoTI message from the MN is IPsec encapsulated in tunnel mode and as it does not belong to the session initiated by the MN or match a previously installed rule, it will be dropped by the firewall. Using CREATE, the MN initiates NSIS signaling to the firewall and open pinholes for the HoTI message. The message flow is comparable to the flow in Figure 7, whereas the CREATE message install different pinholes. The HoT message can re-use this pinhole and is able to reach the MN. The CoTI message and the CoT message can traverse the MN's ASP-firewall, as the CoTI message is not IPsec encapsulated and the CoT message correspond to the state previously installed by the CoTI message.

Once the RRT is successful, the binding update message is sent to the CN. If the MN wants to continue sending data

traffic, no NSIS signaling is needed at all for this scenario. However, if the CN wants to send data traffic and the rules installed before matching again the addresses, the ports and the IPsec encapsulation, the relevant packet filter rules have to be installed at the firewall. If the rules installed before only matching again source and destination address, the data traffic exchanged with the CN in RO-case can also traverse the firewall with no need of installing additional rules. However, that would allow all kind of traffic from the CN and is rejected. Hence, the MN has to initiate sending data traffic to the CN but this happens after the RRT.

Bi-directional tunnelling

Consider the scenario where the MN is protected by a SPF. Even though the MN had earlier initiated a connection for the purpose of binding update, new filter rules have to be installed to allow the tunnelled data traffic as the rules before installed rules match again the addresses, the ports and the IPsec ESP encapsulation. The message flow is shown in Figure 8. If the MN is the data sender, no signaling is necessary at all. Otherwise, the MN opens pinholes using EXT to let the data messages traverse.




Figure 8: Signaling for data traffic

3.4.2 Firewall located at the edge of CN's ASP

Route Optimization

In Figure 2, the CN is protected by a firewall that employs the stateful packet filtering. The external MN and its associated HA are also shown in the figure. The MN communicates with the CN. If the CN initiated normal data traffic there is no problem with the SPF, as the communication is initiated from internal. The following subsections explain how this approach manages the MIPv6 signaling traffic problems as described in Section 2.

The MN moves out of its home network and has to perform the return routability test before sending the binding update to the CN. It sends a HoTI message through the HA to the CN and expects a HoT message from the CN along the same path. It also sends a CoTI message directly to the CN and expects CoT message in the same path from the CN. The SPF will only allow packets that belong to an existing session and hence both the packets (HoTI, CoTI) will be dropped as these packets are Mobile IPv6 packets and these packets have a different header structure. The existing rules at the firewall might have been installed for some kind of data traffic. As the RRT procedure can not be executed, the firewall rules have to be modified to allow these MIPv6 messages to go through. The MN initiates the NSIS session by sending a CREATE message to the CN to install rules for the CoTI message. The NSIS signaling to allow the CoTI message is shown in Figure 9. However, such

an approach where an external node is able to install filter rules in an ASP-FW clearly requires a strong authentication framework. Section 4 discusses this in more detail and presents several potential candidates.




Figure 9: Signaling for CoTI and CoT

If the MN signal as described in the previous section, the HoTI is able to reach the HA. Nevertheless, the HoTI message from the HA to the CN is not able to traverse, as it does not match any state at the CN's ASP-FW. Therefore, either the HA or the CN has to signal install rules to let the HoTI traverse. When the MN receives both CoT and HoT messages, it performs binding update to the CN which is possible, as the BU can re-use the previously installed rules. Note that the aforementioned signaling was only to allow the Mobile IPv6 messages.

If the CN wants to continue sending data traffic (CN is the data sender(DS)) to the new CoA, it can do so without any additional signaling. This is because the SPF will allow the traffic initiated by the nodes that it protects. But if the MN wants to continue sending data traffic (MN is the DS), it has to install filter rules for data traffic. The approach of combined signaling (for control and data traffic) could be useful, but currently the NSIS NAT/FW protocol does not support installing multiple rules at the same time. This will be discussed in Section 5 in detail.

This solution works under the assumption that the firewalls will allow NSIS messages from external network to bypass, by applying a delayed packet filter state establishment and authorization from the CN. However, operators might be reluctant to allow NSIS message from external network as this might lead to Denial of Service (DoS) attacks. The CN might therefore be required to authorize the traversal of NSIS signaling message implicitly to reduce unwanted traffic. To avoid this complexity, it is also possible to ask the CN to open pinholes in the firewall on behalf of the MN. However, this solution may not work in some scenarios due to routing asymmetry as explained in [8].

Bi-directional Tunnelling

If the CN is protected by a SPF firewall, there is no need for any signaling if the CN starts sending data traffic. The CN sends the data traffic and hence the SPF will store relevant state information and accepts packets from the reverse direction.

If the HA is the DS, then either the CN has to initiate the signaling using EXT or the HA using CREATE, in order to configure the firewall to allow the data traffic traverse from the HA to CN. To support that function, Mobile IPv6 module at the HA or CN will need to be changed so that it

triggers the local MIP6-firewall-traversal-application in the event of receiving a CoTI message from the MN. The local MIP6-firewall-traversal-application is then able to trigger the pinhole creation process. The message flow if the CN should signal for this pinhole is shown in Figure 10.




Figure 10: Signaling for data traffic

3.4.3 Firewall located at the edge of the MN's MSP

Route Optimization

In Figure 3, the Mobile Node's MSP is protected by a firewall that employs the stateful packet filtering. The MN and the CN are also shown in the figure. The MN, after entering a new network, sends a Binding Update to the HA. But as it is initiated by the MN, it first has to install some filter rules in the firewall before sending the Binding Update.

The MN-HA Binding Update message is assumed to be IPsec encapsulated. This might cause problems, as some primitive firewalls do not recognize IPsec traffic and hence drop the packets because of the absence of any transport header. One approach is to use UDP encapsulation of IPsec traffic in order to overcome this problem. Another is using NSIS NAT/FW NSLP to signal the firewall to allow such traffic to traverse. The MN initiates the NSIS signaling to create rules that will allow the Binding Update messages to go through the firewall. The MN then sends the Binding Update message to the HA.

By default, the rules previously installed in the firewall will not allow the HoTI message to go through. Hence, the MN has to install a different set of rules for these signaling messages by initiating another NAT/FW NSLP signaling exchange. After that it sends the HoTI message to the HA. The HA installs rules between the HA and the CN and accordingly send the HoTI to the CN. The HoT message from the CN to the HA is also allowed by the SPF as it belongs to the session previously installed by the HA. The HoT message from the HA to the MN is also allowed as it is initiated by the HA. The RRT completes successfully. Detailed message flow between MN and HA is shown in Figure 11.

For the data traffic, there is no additional signaling as the MN sends data directly to CN and none of these networks are protected by firewalls. This is applicable for both cases when either MN or CN is the data senders.

Bi-directional tunnelling

Here, it is necessary that the HA opens pinholes for the data traffic from the CN using EXT. The CN is then allowed to send the data traffic through the firewall. After intercepting a packet, the HA tunnels it to the MN.

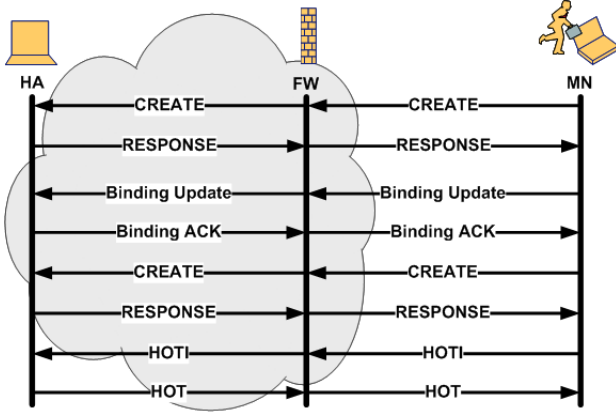


Figure 11: Signaling for BU, BA, HoTI and HoT

4. AUTHENTICATION, AUTHORIZATION AND KEY MANAGEMENT

An important aspect for firewall signaling is how to ensure that only authorized hosts are allowed to perform actions. This leads to the question of how to provide authentication, authorization and key management. Manner et al. [18] specifies how authorization is accomplished for the NSIS QoS and NAT/FW NSLP using an authorization token. That document reuses the authorization token format specified for RSVP and allows information to be exchanged between nodes in order to authorize access to resources. In addition to the already proposed mechanism we discuss three solutions, namely using the Generic Service Authorization Architecture (GSABA) [19], SAML assertions and TLS-EAP [21].

4.1 Generic Service Authorization Architecture

The Generic Service Authorization Architecture (GSABA) [19] is an authentication system with three parties. The goal is to give an end host the ability to access services offered by third parties and to utilize the AAA infrastructure for authentication, authorization and accounting. In this section we will give a short introduction to the GSABA and subsequently discuss a possible integration with the NSIS NAT/FW NSLP for MIPv6 usage.

4.1.1 GSABA Architecture

Figure 12 illustrates the basic architecture elements of GSABA. The Bootstrapping target (BT) is the entity that offers the requested service. In MIPv6 case, the firewall will act as the BT. The Bootstrapping Configuration Agent (BCA) provides necessary bootstrapping information to the MN. The Bootstrapping Authorization Agent (BAA) stores the MN's profile and acts as an Identity Provider. For roaming purposes there will be an additional architectural element, the BAA Proxy. Its function is to forward and, if necessary, to modify policies.

One important interface between the BT and the BCA is the Bootstrapping Target Protocol (TP-p) that provides the mechanism to exchange service related information. RADIUS and Diameter are example protocols for TP-p. The Bootstrapping Protocol (BCA-p) will transmit bootstrap-

ping information to the MN and also informs it about the authorization decision taken by the BAA and BAA Proxy. HTTP is a possible candidate for the BCA-p interface. RADIUS and Diameter is again used for delivering decisions between the BAA and the BCA via the Bootstrapping Agent Protocol (BA-p). The interface between the MN and the BT is the Service Related Protocol (SP) that ideally does not need to be changed to support GSABA when certain minimum requirements are met. The latter aspect is particularly important since it allows a smooth transition for various protocols since no additional standardization work is necessary if basic security mechanisms are already specified.




Figure 12: The GSABA Architecture

4.1.2 GSABA Integration in the NSIS NAT/FW NSLP

The integration of GSABA into the NSIS NAT/FW NSLP requires, in case of firewall traversal for Mobile IPv6, the investigation of the three scenarios described in Section 2.1. In all three scenarios below the firewall acts as the BT. The first scenario additionally details bootstrapping of Mobile IP via the same infrastructure.

Firewall located at the edge of MN's ASP

When the MN wants to install rules at the firewall, it usually uses CREATE or EXT. Therefore, it has to interact with the BAA. When the MN and the BCA/BAA mutually authenticated each other the BAA will send the GSABA Key and the users profile to the GSABA Proxy, which will store this information locally. The MN gets the GSABA Key and is able to request HA information at the GSABA Proxy. The proxy checks whether the MN is authorized and selects a HA. Then, the MN derives the IKEv2 PSK to authenticate against the HA. The HA will fetch the PSK from the GSABA Proxy. After this step, the MN derives the GIST Key and uses it as a PSK in the TLS handshake with the firewall. At this point the firewall fetches the PSK also from the GSABA Proxy.

Now, the MN starts NSIS NAT/FW signaling, for example, by sending a CREATE message through the firewall to the HA. The firewall authorizes the CREATE message. Figure 13 shows an example message flow for the above-described procedure.

Firewall located at the edge of CN's ASP

In this scenario, the CN needs to establish a security association between the firewall and itself. When the MN wants to open pinholes at this firewall, it firstly signals this with the CREATE message. As there is no authorization at this point, the firewall responds with an error message including its domain name. The MN now derives a NSLP Key from the GSABA Key and sends the CREATE message again. At this time, it uses the PSK in the TLS handshake with the firewall and the firewall fetches the NSLP Key from the

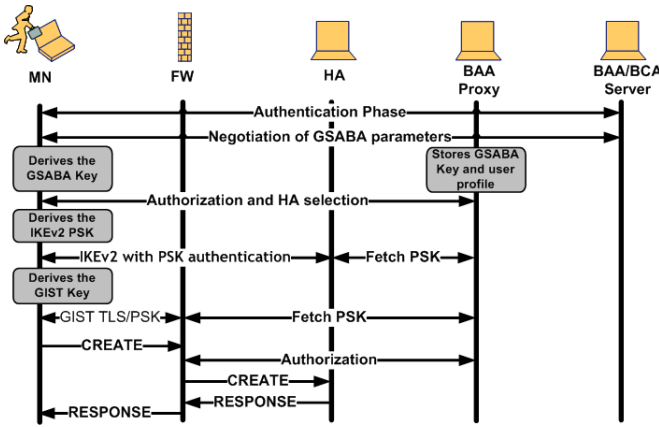


Figure 13: The GSABA message flow, Firewall located at the edge of MN's ASP

GSABA Proxy. Hence, the firewall is able to authorize the message sent by the MN and forwards it to the CN, which replies with a RESPONSE message on the same path. The MN and the CN are now able to send the CoTI/CoT messages for route optimization.

The message flow for the HoTI message is different as the MN tunnels the HoTI message via its HA, which will then trigger a CREATE message for opening pinholes at the firewall on the CN side. The firewall could now authorize the CREATE message. The subsequent BU/BA message exchange between the MN and the CN will be able to traverse the firewall without problems.

Firewall located at the edge of MN's MSP

In this scenario, the MN first needs to be authorized against the GSABA Server to get the GSABA Key. Afterwards the MN derives the GIST Key and uses it as a PSK in the TLS handshake with the firewall. The firewall fetches the PSK from the GSABA Proxy and the MN could send a CREATE message to allow IKEv2 traffic to traverse the firewall. The firewall checks the authorization at the GSABA Server and then decides if the CREATE message can traverse the firewall. The MN derives the IKEv2 PSK to authenticate against the HA. The HA will fetch the PSK from the GSABA Proxy. The GSABA infrastructure may provide additional information to the firewall in order to pre-authorize subsequent NAT/FW messages.

4.2 Security Assertion Markup Language

Security Assertion Markup Language (SAML) is an XML-based framework for creating and exchanging security information. In the course of making, or relying upon such assertions, SAML system entities may use SAML protocols, or other protocols, to communicate an assertion itself, or the subject of an assertion.

Thus one can employ SAML to make and encode statements such as "Alice has these profile attributes and her domain's certificate is available over there, and I'm making this statement that she is allowed to traverse firewalls within this particular domain." Then, an end host can cause such an assertion to be conveyed to some party, for example a firewall, who can then rely on it for computing an authorization

decision, for example using it as input into some local policy evaluation for granting the establishment of a pinhole.

A possible approach of applying SAML for NAT/FW NSLP signaling in Mobile IPv6 environments is as follows. The MN first asks the Identity Provider (IdP) to get such an assertion before starting signaling with the firewall. The IdP will authenticate the user or end host and will return an assertion in case of success. When interacting with a firewall the MN will attach the SAML assertion to the message. After that the firewall verifies whether the assertion is valid and if the MN is authorized to perform the indicated action (e.g., creating a pinhole) for further communication. An error is returned in case the end host is not authorized. Despite the popularity of SAML for identity management there is also a disadvantage, namely the large size of the XML-based assertions when conveyed by value. To overcome this limitation a reference to a SAML assertion can be used instead; the firewall then has to resolve the reference first to obtain the assertion, for example using an HTTP lookup.

4.3 TLS using EAP Authentication

Transport Layer Security (TLS) using EAP Authentication [21] (TLS-EAP) enhances the TLS handshake with support for authorization with the Extensible Authentication Protocol (EAP). NSIS allows TLS to be used the integration with EAP is attractive since the TLS server is able to relay EAP payloads to the existing AAA infrastructure to offload authentication, authorization and accounting tasks. When TLS-EAP is used then the TLS handshake is initiated and EAP messages are exchanged between the TLS client (i.e., NAT/FW client) and the TLS server (i.e., firewall). The TLS server forwards messages to the AAA infrastructure whenever it is not able to handle authentication locally. The TLS server does not need to understand the specific EAP method and acts as a relay until the exchange is completed and the AAA server indicates the success or failure of the EAP exchange to the TLS server. Later, when the NAT/FW client requests the creation of new firewall pinholes the firewall may need to initiate an authorization request towards the AAA server. The AAA server may either grant or deny the request and returns the decision to the firewall. The advantage of using TLS-EAP is the smooth integration with the AAA infrastructure and the simple enhancements needed for EAP integration into TLS due to the extensibility of the NSIS framework, in this particular case GIST. A weakness of using TLS-EAP is the additional message exchanges since EAP is quite chatty.

5. OPEN ISSUES AND FUTURE WORK

The firewall traversal solution based on the IETF NAT/FW NSLP presented in this paper can deal with the problems of having firewalls in Mobile IPv6 environments. However, the approach as described in this paper might not be efficient enough for some environments. As a result, the optimization of the signaling exchange and the reduction of the signaling delay are for further study. Overall, more work on performance optimizations and scalability investigations are necessary.

Firewall traversal requires strong authentication and authorization. An initial set of security mechanisms are proposed in Section 4 but further work is needed to investigate the details in order to study the security properties, potentially including a formal analysis. For example, currently there

is no SAML “profile” or “binding” defined that describes in detail how SAML assertions are carried within NSIS. Today’s infrastructure mostly supports MIPv4, rarely MIPv6. Therefore, it is necessary to investigate a MIPv6/v4 dual stack solution. We are currently finalizing a prototype implementation to prove the feasibility and usability of such an Mobile IPv6 firewall traversal approach.

6. CONCLUSIONS

This paper shows how the NSIS NAT/FW NSLP can address the issues caused by stateful packet filter firewalls encountered in a Mobile IPv6 network. We described the problems and impacts of having firewalls in Mobile IPv6 environments and presented a firewall traversal solution based on the IETF NSIS framework, which can handle all these issues in the different scenarios. It has to be noted that a real scenario could include a combination of some set of these cases. In contrast to other middlebox configuration solutions, the NSIS solution can offer a solution for all deployment scenarios assuming that the MN, the CN, the HA and the firewalls are NSIS NAT/FW NSLP aware.

In contrast to other explicit middlebox configuration approaches like MIDCOM, NAT/FW NSLP require more signaling but does not require knowledge about the topology and avoids possible performance problems caused by the deep packet inspection of said approaches. M-ICE that builds on STUN is a very recent proposal that is designed based on a different security framework but conceptually similar to the NSIS NAT/FW NSLP.

Finally, this paper also outlines approaches for addressing the security aspects for NAT and firewall traversal. These approaches are based on the recently developed GSABA (a AAA-based bootstrapping framework), TLS-EAP that reuses EAP and the AAA infrastructure and SAML assertions. Further study with respect to the aspects described in Section 5 are necessary.

7. ACKNOWLEDGMENT

We would like to thank Ivano Guardini, Li Cai, Qin Wu, Ingo Juchem, Swen Weiland, Jan Demter and Mehmet Er-sue for their contributions and insightful discussions. This work has been partially supported by the EC FP6 IST research project ENABLE - Enabling Efficient and Operational Mobility in Large Heterogeneous IP Networks.

8. REFERENCES

- [1] D. Johnson, C. Perkins, J. Arkko, “Mobility Support in IPv6”, RFC 3775, June 2004.
- [2] J. Rosenberg, J. Weinberger, C. Huitema, R. Mahy, “STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)”, RFC 3489, March 2003.
- [3] J. Rosenberg, R. Mahy, C. Huitema, “Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)”, I-D (draft-ietf-behave-turn-04), work in progress, July 2007.
- [4] J. Rosenberg, “Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols”, I-D (draft-ietf-mmusic-ice-17), work in progress, July 2007.
- [5] J. Quittek, M. Stiernerling, P. Srisuresh, “Definitions of Managed Objects for Middlebox Communication”, I-D (draft-ietf-midcom-mib-09), work in progress, October 2006.
- [6] D. Durham, J. Boyle, R. Cohen, S. Herzog, R. Rajan, A. Sastry, “The COPS (Common Open Policy Service) Protocol”, RFC 2748, January 2000.
- [7] R. Hancock, G. Karagiannis, J. Loughney, and S. Van den Bosch, “Next Steps in Signaling (NSIS): Framework”, RFC 4080, June 2005.
- [8] M. Stiernerling, H. Tschofenig, and C. Aoun, “A NAT/Firewall NSIS Signaling Layer Protocol (NSLP)”, I-D (draft-ietf-nsis-nslp-natfw-15), work in progress, July 2007.
- [9] F. Le, S. Faccin, B. Patil, and H. Tschofenig, “Mobile IPv6 and Firewalls: Problem Statement”, RFC 4487, May 2006.
- [10] J. Arkko, V. Devarapalli, F. Dupont, “Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents”, RFC 3776, June 2004.
- [11] H. Tschofenig, G. Bajko, “Mobile IP Interactive Connectivity Establishment (M-ICE)”, I-D (draft-tschofenig-mip6-ice-01), work in progress, July 2007.
- [12] H. Tschofenig, G. Bajko, “Firewall friendly Return-Routability Test (RRT) for Mobile IPv6”, I-D (draft-bajko-mip6-rrtfw-02), work in progress, July 2007.
- [13] D. Wing, J. Rosenberg, H. Tschofenig, “Discovering, Querying, and Controlling Firewalls and NATs using STUN”, I-D (draft-wing-behave-nat-control-stun-usage-03), work in progress, July 2007.
- [14] D. Wing, “Media Session Authorization”, I-D (draft-wing-session-auth-00), work in progress, January 2006.
- [15] “An Implementation of the Next Steps in Signaling (NSIS) Protocol Suite at the University of Göttingen”, <http://user.informatik.uni-goettingen.de/~nsis/>.
- [16] N. Steinleitner, H. Peters, H. Tschofenig, X. Fu, “Implementation and Performance Study of a New NAT/Firewall Signaling Protocol”, ADSN2006, in conjunction with ICDCS 2006, Portugal, July 2006.
- [17] S. Thiruvengadam, H. Tschofenig, F. Le, N. Steinleitner, X. Fu, “Mobile IPv6 - NSIS Interaction for Firewall traversal”, I-D (draft-thiruvengadam-nsis-mip6-fw-06), work in progress, March 2007.
- [18] J. Manner, M. Stiernerling, H. Tschofenig, “Authorization for NSIS Signaling Layer Protocols”, I-D (draft-manner-nsis-nslp-auth-03), work in progress, March 2007.
- [19] F. Kohlmayer, H. Tschofenig, R. Falk, R. Lopez, S. Hernandez, P. Segura, A. Skarmeta, “GSABA: A Generic Service Authorization Architecture”, MobiArch’06, San Francisco, December 2006.
- [20] P. Eronen, H. Tschofenig, “Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)”, RFC 4279, December 2005.
- [21] Y. Nir, Y. Sheffer, H. Tschofenig, P. Gutmann, “TLS using EAP Authentication”, I-D (draft-nir-tls-eap-01), work in progress, July 2007.