

Internet Engineering Task Force (IETF)
Request for Comments: 5973
Category: Experimental
ISSN: 2070-1721

M. Stiemerling
NEC
H. Tschofenig
Nokia Siemens Networks
C. Aoun
Consultant
E. Davies
Folly Consulting
October 2010

NAT/Firewall NSIS Signaling Layer Protocol (NSLP)

Abstract

This memo defines the NSIS Signaling Layer Protocol (NSLP) for Network Address Translators (NATs) and firewalls. This NSLP allows hosts to signal on the data path for NATs and firewalls to be configured according to the needs of the application data flows. For instance, it enables hosts behind NATs to obtain a publicly reachable address and hosts behind firewalls to receive data traffic. The overall architecture is given by the framework and requirements defined by the Next Steps in Signaling (NSIS) working group. The network scenarios, the protocol itself, and examples for path-coupled signaling are given in this memo.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5973>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	5
1.1.	Scope and Background	5
1.2.	Terminology and Abbreviations	8
1.3.	Notes on the Experimental Status	10
1.4.	Middleboxes	10
1.5.	General Scenario for NATFW Traversal	11
2.	Network Deployment Scenarios Using the NATFW NSLP	13
2.1.	Firewall Traversal	13
2.2.	NAT with Two Private Networks	14
2.3.	NAT with Private Network on Sender Side	15
2.4.	NAT with Private Network on Receiver Side Scenario	15
2.5.	Both End Hosts behind Twice-NATs	16
2.6.	Both End Hosts behind Same NAT	17
2.7.	Multihomed Network with NAT	18
2.8.	Multihomed Network with Firewall	18
3.	Protocol Description	19
3.1.	Policy Rules	19
3.2.	Basic Protocol Overview	20
3.2.1.	Signaling for Outbound Traffic	20
3.2.2.	Signaling for Inbound Traffic	22
3.2.3.	Signaling for Proxy Mode	23
3.2.4.	Blocking Traffic	24
3.2.5.	State and Error Maintenance	24
3.2.6.	Message Types	25
3.2.7.	Classification of RESPONSE Messages	25
3.2.8.	NATFW NSLP Signaling Sessions	26
3.3.	Basic Message Processing	27
3.4.	Calculation of Signaling Session Lifetime	27
3.5.	Message Sequencing	31
3.6.	Authentication, Authorization, and Policy Decisions	32
3.7.	Protocol Operations	32
3.7.1.	Creating Signaling Sessions	32
3.7.2.	Reserving External Addresses	35
3.7.3.	NATFW NSLP Signaling Session Refresh	43
3.7.4.	Deleting Signaling Sessions	45
3.7.5.	Reporting Asynchronous Events	46
3.7.6.	Proxy Mode of Operation	48
3.8.	Demultiplexing at NATs	53
3.9.	Reacting to Route Changes	54
3.10.	Updating Policy Rules	55
4.	NATFW NSLP Message Components	55
4.1.	NSLP Header	56
4.2.	NSLP Objects	57
4.2.1.	Signaling Session Lifetime Object	58
4.2.2.	External Address Object	58
4.2.3.	External Binding Address Object	59

- 4.2.4. Extended Flow Information Object 59
- 4.2.5. Information Code Object 60
- 4.2.6. Nonce Object 64
- 4.2.7. Message Sequence Number Object 64
- 4.2.8. Data Terminal Information Object 64
- 4.2.9. ICMP Types Object 66
- 4.3. Message Formats 67
 - 4.3.1. CREATE 67
 - 4.3.2. EXTERNAL 68
 - 4.3.3. RESPONSE 68
 - 4.3.4. NOTIFY 69
- 5. Security Considerations 69
 - 5.1. Authorization Framework 70
 - 5.1.1. Peer-to-Peer Relationship 70
 - 5.1.2. Intra-Domain Relationship 71
 - 5.1.3. End-to-Middle Relationship 72
 - 5.2. Security Framework for the NAT/Firewall NSLP 73
 - 5.2.1. Security Protection between Neighboring NATFW NSLP
Nodes 73
 - 5.2.2. Security Protection between Non-Neighboring NATFW
NSLP Nodes 74
 - 5.3. Implementation of NATFW NSLP Security 75
- 6. IAB Considerations on UNSAF 76
- 7. IANA Considerations 77
 - 7.1. NATFW NSLP Message Type Registry 77
 - 7.2. NATFW NSLP Header Flag Registry 77
 - 7.3. NSLP Message Object Registry 78
 - 7.4. NSLP Response Code Registry 78
 - 7.5. NSLP IDs and Router Alert Option Values 78
- 8. Acknowledgments 78
- 9. References 79
 - 9.1. Normative References 79
 - 9.2. Informative References 79
- Appendix A. Selecting Signaling Destination Addresses for
EXTERNAL 81
- Appendix B. Usage of External Binding Addresses 82
- Appendix C. Applicability Statement on Data Receivers behind
Firewalls 83
- Appendix D. Firewall and NAT Resources 84
 - D.1. Wildcarding of Policy Rules 84
 - D.2. Mapping to Firewall Rules 84
 - D.3. Mapping to NAT Bindings 85
 - D.4. NSLP Handling of Twice-NAT 85
- Appendix E. Example for Receiver Proxy Case 86

1. Introduction

1.1. Scope and Background

Firewalls and Network Address Translators (NATs) have both been used throughout the Internet for many years, and they will remain present for the foreseeable future. Firewalls are used to protect networks against certain types of attacks from internal networks and the Internet, whereas NATs provide a virtual extension of the IP address space. Both types of devices may be obstacles to some applications, since they only allow traffic created by a limited set of applications to traverse them, typically those that use protocols with relatively predetermined and static properties (e.g., most HTTP traffic, and other client/server applications). Other applications, such as IP telephony and most other peer-to-peer applications, which have more dynamic properties, create traffic that is unable to traverse NATs and firewalls without assistance. In practice, the traffic of many applications cannot traverse autonomous firewalls or NATs, even when they have additional functionality that attempts to restore the transparency of the network.

Several solutions to enable applications to traverse such entities have been proposed and are currently in use. Typically, application-level gateways (ALGs) have been integrated with the firewall or NAT to configure the firewall or NAT dynamically. Another approach is middlebox communication (MIDCOM). In this approach, ALGs external to the firewall or NAT configure the corresponding entity via the MIDCOM protocol [RFC3303]. Several other work-around solutions are available, such as Session Traversal Utilities for NAT (STUN) [RFC5389]. However, all of these approaches introduce other problems that are generally hard to solve, such as dependencies on the type of NAT implementation (full-cone, symmetric, etc.), or dependencies on certain network topologies.

NAT and firewall (NATFW) signaling shares a property with Quality-of-Service (QoS) signaling -- each must reach any device that is on the data path and is involved in (respectively) NATFW or QoS treatment of data packets. This means that for both NATFW and QoS it is convenient if signaling travels path-coupled, i.e., the signaling messages follow exactly the same path that the data packets take. The Resource Reservation Protocol (RSVP) [RFC2205] is an example of a current QoS signaling protocol that is path-coupled. [rsvp-firewall] proposes the use of RSVP as a firewall signaling protocol but does not include NATs.

This memo defines a path-coupled signaling protocol for NAT and firewall configuration within the framework of NSIS, called the NATFW NSIS Signaling Layer Protocol (NSLP). The general requirements for

NSIS are defined in [RFC3726] and the general framework of NSIS is outlined in [RFC4080]. It introduces the split between an NSIS transport layer and an NSIS signaling layer. The transport of NSLP messages is handled by an NSIS Network Transport Layer Protocol (NTLP, with General Internet Signaling Transport (GIST) [RFC5971] being the implementation of the abstract NTLP). The signaling logic for QoS and NATFW signaling is implemented in the different NSLPs. The QoS NSLP is defined in [RFC5974].

The NATFW NSLP is designed to request the dynamic configuration of NATs and/or firewalls along the data path. Dynamic configuration includes enabling data flows to traverse these devices without being obstructed, as well as blocking of particular data flows at inbound firewalls. Enabling data flows requires the loading of firewall rules with an action that allows the data flow packets to be forwarded and NAT bindings to be created. The blocking of data flows requires the loading of firewall rules with an action that will deny forwarding of the data flow packets. A simplified example for enabling data flows: a source host sends a NATFW NSLP signaling message towards its data destination. This message follows the data path. Every NATFW NSLP-enabled NAT/firewall along the data path intercepts this message, processes it, and configures itself accordingly. Thereafter, the actual data flow can traverse all these configured firewalls/NATs.

It is necessary to distinguish between two different basic scenarios when operating the NATFW NSLP, independent of the type of the middleboxes to be configured.

1. Both the data sender and data receiver are NSIS NATFW NSLP aware. This includes the cases in which the data sender is logically decomposed from the initiator of the NSIS signaling (the so-called NSIS initiator) or the data receiver logically decomposed from the receiver of the NSIS signaling (the so-called NSIS receiver), but both sides support NSIS. This scenario assumes deployment of NSIS all over the Internet, or at least at all NATs and firewalls. This scenario is used as a base assumption, if not otherwise noted.
2. Only one end host or region of the network is NSIS NATFW NSLP aware, either the data receiver or data sender. This scenario is referred to as proxy mode.

The NATFW NSLP has two basic signaling messages that are sufficient to cope with the various possible scenarios likely to be encountered before and after widespread deployment of NSIS:

CREATE message: Sent by the data sender for configuring a path outbound from a data sender to a data receiver.

EXTERNAL message: Used by a data receiver to locate inbound NATs/firewalls and prime them to expect inbound signaling and used at NATs to pre-allocate a public address. This is used for data receivers behind these devices to enable their reachability.

CREATE and EXTERNAL messages are sent by the NSIS initiator (NI) towards the NSIS responder (NR). Both types of message are acknowledged by a subsequent RESPONSE message. This RESPONSE message is generated by the NR if the requested configuration can be established; otherwise, the NR or any of the NSLP forwarders (NFs) can also generate such a message if an error occurs. NFs and the NR can also generate asynchronous messages to notify the NI, the so-called NOTIFY messages.

If the data receiver resides in a private addressing realm or behind a firewall, and it needs to preconfigure the edge-NAT/edge-firewall to provide a (publicly) reachable address for use by the data sender, a combination of EXTERNAL and CREATE messages is used.

During the introduction of NSIS, it is likely that one or the other of the data sender and receiver will not be NSIS aware. In these cases, the NATFW NSLP can utilize NSIS-aware middleboxes on the path between the data sender and data receiver to provide proxy NATFW NSLP services (i.e., the proxy mode). Typically, these boxes will be at the boundaries of the realms in which the end hosts are located.

The CREATE and EXTERNAL messages create NATFW NSLP and NTLF state in NSIS entities. NTLF state allows signaling messages to travel in the forward (outbound) and the reverse (inbound) direction along the path between a NAT/firewall NSLP sender and a corresponding receiver. This state is managed using a soft-state mechanism, i.e., it expires unless it is refreshed from time to time. The NAT bindings and firewall rules being installed during the state setup are bound to the particular signaling session. However, the exact local implementation of the NAT bindings and firewall rules are NAT/firewall specific and it is out of the scope of this memo.

This memo is structured as follows. Section 2 describes the network environment for NATFW NSLP signaling. Section 3 defines the NATFW signaling protocol and Section 4 defines the message components and the overall messages used in the protocol. The remaining parts of the main body of the document cover security considerations Section 5, IAB considerations on UNilateral Self-Address Fixing

(UNSAF) [RFC3424] in Section 6, and IANA considerations in Section 7. Please note that readers familiar with firewalls and NATs and their possible location within networks can safely skip Section 2.

1.2. Terminology and Abbreviations

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].

This document uses a number of terms defined in [RFC3726] and [RFC4080]. The following additional terms are used:

- o Policy rule: A policy rule is "a basic building block of a policy-based system. It is the binding of a set of actions to a set of conditions - where the conditions are evaluated to determine whether the actions are performed" [RFC3198]. In the context of NSIS NATFW NSLP, the conditions are the specification of a set of packets to which the rule is applied. The set of actions always contains just a single element per rule, and is limited to either action "deny" or action "allow".
- o Reserved policy rule: A policy rule stored at NATs or firewalls for activation by a later, different signaling exchange. This type of policy rule is kept in the NATFW NSLP and is not loaded into the firewall or NAT engine, i.e., it does not affect the data flow handling.
- o Installed policy rule: A policy rule in operation at NATs or firewalls. This type of rule is kept in the NATFW NSLP and is loaded into the firewall or NAT engine, i.e., it is affecting the data flow.
- o Remembered policy rule: A policy rule stored at NATs and firewalls for immediate use, as soon as the signaling exchange is successfully completed.
- o Firewall: A packet filtering device that matches packets against a set of policy rules and applies the actions.
- o Network Address Translator: Network Address Translation is a method by which IP addresses are mapped from one IP address realm to another, in an attempt to provide transparent routing between hosts (see [RFC2663]). Network Address Translators are devices that perform this work by modifying packets passing through them.
- o Data Receiver (DR): The node in the network that is receiving the data packets of a flow.

- o Data Sender (DS): The node in the network that is sending the data packets of a flow.
- o NATFW NSLP peer (or simply "peer"): An NSIS NATFW NSLP node with which an NTLF adjacency has been created as defined in [RFC5971].
- o NATFW NSLP signaling session (or simply "signaling session"): A signaling session defines an association between the NI, NFs, and the NR related to a data flow. All the NATFW NSLP peers on the path, including the NI and the NR, use the same identifier to refer to the state stored for the association. The same NI and NR may have more than one signaling session active at any time. The state for the NATFW NSLP consists of NSLP state and associated policy rules at a middlebox.
- o Edge-NAT: An edge-NAT is a NAT device with a globally routable IP address that is reachable from the public Internet.
- o Edge-firewall: An edge-firewall is a firewall device that is located on the borderline of an administrative domain.
- o Public Network: "A Global or Public Network is an address realm with unique network addresses assigned by Internet Assigned Numbers Authority (IANA) or an equivalent address registry. This network is also referred as external network during NAT discussions" [RFC2663].
- o Private/Local Network: "A private network is an address realm independent of external network addresses. Private network may also be referred alternately as Local Network. Transparent routing between hosts in private realm and external realm is facilitated by a NAT router" [RFC2663].
- o Public/Global IP address: An IP address located in the public network according to Section 2.7 of [RFC2663].
- o Private/Local IP address: An IP address located in the private network according to Section 2.8 of [RFC2663].
- o Signaling Destination Address (SDA): An IP address generally taken from the public/global IP address range, although, the SDA may, in certain circumstances, be part of the private/local IP address range. This address is used in EXTERNAL signaling message exchanges, if the data receiver's IP address is unknown.

1.3. Notes on the Experimental Status

The same deployment issues and extensibility considerations described in [RFC5971] and [RFC5978] also apply to this document.

1.4. Middleboxes

The term "middlebox" covers a range of devices and is well-defined in [RFC3234]: "A middlebox is defined as any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and a destination host". As such, middleboxes fall into a number of categories with a wide range of functionality, not all of which is pertinent to the NATFW NSLP. Middlebox categories in the scope of this memo are firewalls that filter data packets against a set of filter rules, and NATs that translate packet addresses from one address realm to another address realm. Other categories of middleboxes, such as QoS traffic shapers, are out of the scope of this memo.

The term "NAT" used in this document is a placeholder for a range of different NAT flavors. We consider the following types of NATs:

- o Traditional NAT (basic NAT and NAT)
- o Bi-directional NAT
- o Twice-NAT
- o Multihomed NAT

For definitions and a detailed discussion about the characteristics of each NAT type, please see [RFC2663].

All types of middleboxes under consideration here use policy rules to make a decision on data packet treatment. Policy rules consist of a flow identifier that selects the packets to which the policy applies and an associated action; data packets matching the flow identifier are subjected to the policy rule action. A typical flow identifier is the 5-tuple selector that matches the following fields of a packet to configured values:

- o Source and destination IP addresses
- o Transport protocol number
- o Transport source and destination port numbers

Actions for firewalls are usually one or more of:

- o Allow: forward data packet
- o Deny: block data packet and discard it
- o Other actions such as logging, diverting, duplicating, etc.

Actions for NATs include (amongst many others):

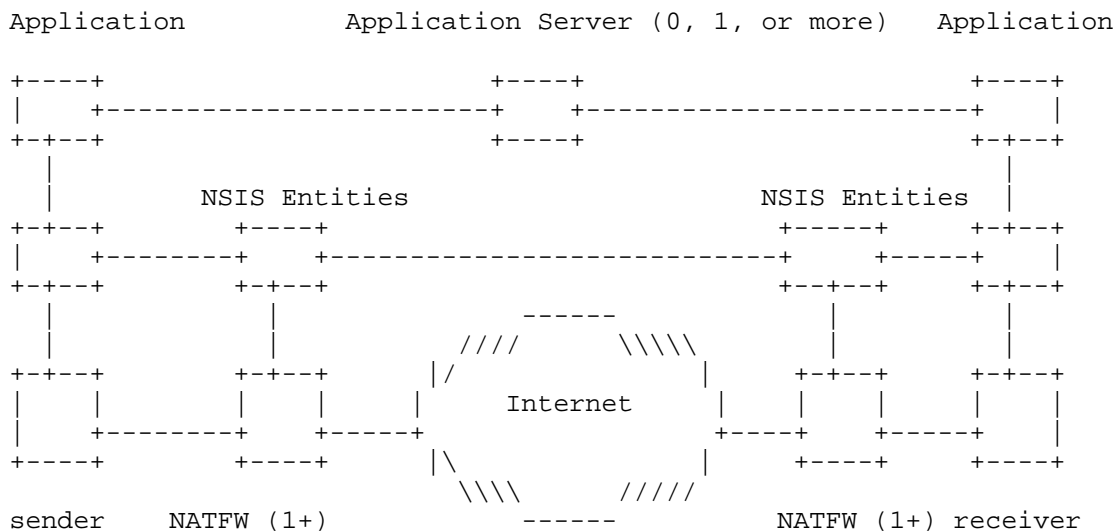
- o Change source IP address and transport port number to a globally routable IP address and associated port number.
- o Change destination IP address and transport port number to a private IP address and associated port number.

It should be noted that a middlebox may contain two logical representations of the policy rule. The policy rule has a representation within the NATFW NSLP, comprising the message routing information (MRI) of the NTLF and NSLP information (such as the rule action). The other representation is the implementation of the NATFW NSLP policy rule within the NAT and firewall engine of the particular device. Refer to Appendix D for further details.

1.5. General Scenario for NATFW Traversal

The purpose of NSIS NATFW signaling is to enable communication between endpoints across networks, even in the presence of NAT and firewall middleboxes that have not been specially engineered to facilitate communication with the application protocols used. This removes the need to create and maintain application layer gateways for specific protocols that have been commonly used to provide transparency in previous generations of NAT and firewall middleboxes. It is assumed that these middleboxes will be statically configured in such a way that NSIS NATFW signaling messages themselves are allowed to reach the locally installed NATFW NSLP daemon. NSIS NATFW NSLP signaling is used to dynamically install additional policy rules in all NATFW middleboxes along the data path that will allow transmission of the application data flow(s). Firewalls are configured to forward data packets matching the policy rule provided by the NSLP signaling. NATs are configured to translate data packets matching the policy rule provided by the NSLP signaling. An additional capability, that is an exception to the primary goal of NSIS NATFW signaling, is that the NATFW nodes can request blocking of particular data flows instead of enabling these flows at inbound firewalls.

The basic high-level picture of NSIS usage is that end hosts are located behind middleboxes, meaning that there is at least one middlebox on the data path from the end host in a private network to the external network (NATFW in Figure 1). Applications located at these end hosts try to establish communication with corresponding applications on other such end hosts. This communication establishment may require that the applications contact an application server that serves as a rendezvous point between both parties to exchange their IP address and port(s). The local applications trigger the NSIS entity at the local host to control provisioning for middlebox traversal along the prospective data path (e.g., via an API call). The NSIS entity, in turn, uses NSIS NATFW NSLP signaling to establish policy rules along the data path, allowing the data to travel from the sender to the receiver without obstruction.



Note that 1+ refers to one or more NATFW nodes.

Figure 1: Generic View of NSIS with NATs and/or Firewalls

For end-to-end NATFW signaling, it is necessary that each firewall and each NAT along the path between the data sender and the data receiver implements the NSIS NATFW NSLP. There might be several NATs and FWs in various possible combinations on a path between two hosts. Section 2 presents a number of likely scenarios with different combinations of NATs and firewalls. However, the scenarios given in the following sections are only examples and should not be treated as limiting the scope of the NATFW NSLP.

2. Network Deployment Scenarios Using the NATFW NSLP

This section introduces several scenarios for middlebox placement within IP networks. Middleboxes are typically found at various different locations, including at enterprise network borders, within enterprise networks, as mobile phone network gateways, etc. Usually, middleboxes are placed more towards the edge of networks than in network cores. Firewalls and NATs may be found at these locations either alone or combined; other categories of middleboxes may also be found at such locations, possibly combined with the NATs and/or firewalls.

NSIS initiators (NI) send NSIS NATFW NSLP signaling messages via the regular data path to the NSIS responder (NR). On the data path, NATFW NSLP signaling messages reach different NSIS nodes that implement the NATFW NSLP. Each NATFW NSLP node processes the signaling messages according to Section 3 and, if necessary, installs policy rules for subsequent data packets.

Each of the following sub-sections introduces a different scenario for a different set of middleboxes and their ordering within the topology. It is assumed that each middlebox implements the NSIS NATFW NSLP signaling protocol.

2.1. Firewall Traversal

This section describes a scenario with firewalls only; NATs are not involved. Each end host is behind a firewall. The firewalls are connected via the public Internet. Figure 2 shows the topology. The part labeled "public" is the Internet connecting both firewalls.

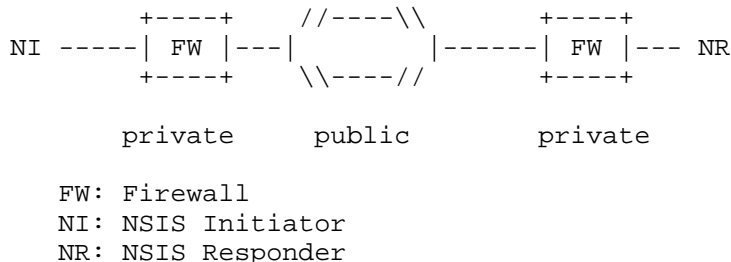


Figure 2: Firewall Traversal Scenario

Each firewall on the data path must provide traversal service for NATFW NSLP in order to permit the NSIS message to reach the other end host. All firewalls process NSIS signaling and establish appropriate policy rules, so that the required data packet flow can traverse them.

There are several very different ways to place firewalls in a network topology. To distinguish firewalls located at network borders, such as administrative domains, from others located internally, the term edge-firewall is used. A similar distinction can be made for NATs, with an edge-NAT fulfilling the equivalent role.

2.2. NAT with Two Private Networks

Figure 3 shows a scenario with NATs at both ends of the network. Therefore, each application instance, the NSIS initiator and the NSIS responder, are behind NATs. The outermost NAT, known as the edge-NAT (MB2 and MB3), at each side is connected to the public Internet. The NATs are generically labeled as MBX (for middlebox No. X), since those devices certainly implement NAT functionality, but can implement firewall functionality as well.

Only two middleboxes (MBs) are shown in Figure 3 at each side, but in general, any number of MBs on each side must be considered.

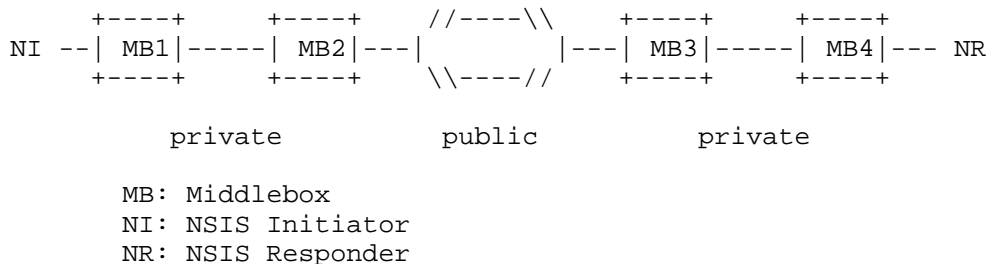


Figure 3: NAT with two Private Networks Scenario

Signaling traffic from the NI to the NR has to traverse all the middleboxes on the path (MB1 to MB4, in this order), and all the middleboxes must be configured properly to allow NSIS signaling to traverse them. The NATFW signaling must configure all middleboxes and consider any address translation that will result from this configuration in further signaling. The sender (NI) has to know the IP address of the receiver (NR) in advance, otherwise it will not be possible to send any NSIS signaling messages towards the responder. Note that this IP address is not the private IP address of the responder but the NAT's public IP address (here MB3's IP address). Instead, a NAT binding (including a public IP address) has to be previously installed on the NAT MB3. This NAT binding subsequently allows packets reaching the NAT to be forwarded to the receiver within the private address realm. The receiver might have a number of ways to learn its public IP address and port number (including the NATFW NSLP) and might need to signal this information to the sender using an application-level signaling protocol.

2.3. NAT with Private Network on Sender Side

This scenario shows an application instance at the sending node that is behind one or more NATs (shown as generic MB, see discussion in Section 2.2). The receiver is located in the public Internet.

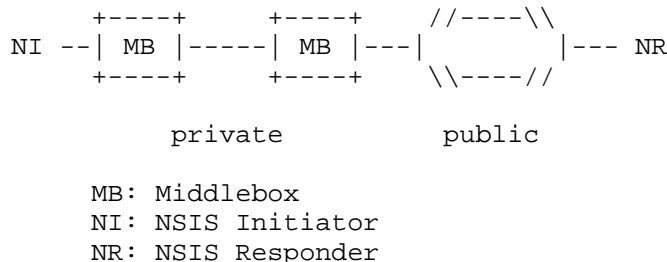


Figure 4: NAT with Private Network on Sender Side

The traffic from NI to NR has to traverse middleboxes only on the sender's side. The receiver has a public IP address. The NI sends its signaling message directly to the address of the NSIS responder. Middleboxes along the path intercept the signaling messages and configure accordingly.

The data sender does not necessarily know whether or not the receiver is behind a NAT; hence, it is the receiving side that has to detect whether or not it is behind a NAT.

2.4. NAT with Private Network on Receiver Side Scenario

The application instance receiving data is behind one or more NATs shown as MB (see discussion in Section 2.2).

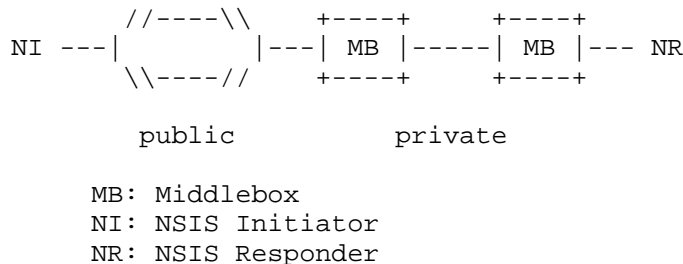


Figure 5: NAT with Private Network on Receiver Scenario

Initially, the NSIS responder must determine its publicly reachable IP address at the external middlebox and notify the NSIS initiator about this address. One possibility is that an application-level

protocol is used, meaning that the public IP address is signaled via this protocol to the NI. Afterwards, the NI can start its signaling towards the NR and therefore establish the path via the middleboxes in the receiver side private network.

This scenario describes the use case for the EXTERNAL message of the NATFW NSLP.

2.5. Both End Hosts behind Twice-NATs

This is a special case, where the main problem arises from the need to detect that both end hosts are logically within the same address space, but are also in two partitions of the address realm on either side of a twice-NAT (see [RFC2663] for a discussion of twice-NAT functionality).

Sender and receiver are both within a single private address realm, but the two partitions potentially have overlapping IP address ranges. Figure 6 shows the arrangement of NATs.

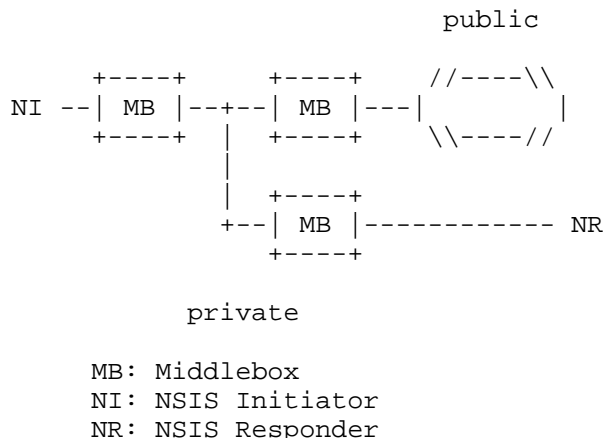


Figure 6: NAT to Public, Sender and Receiver on Either Side of a Twice-NAT Scenario


The middleboxes shown in Figure 6 are twice-NATs, i.e., they map IP addresses and port numbers on both sides, meaning the mapping of source and destination IP addresses at the private and public interfaces.

This scenario requires the assistance of application-level entities, such as a DNS server. The application-level entities must handle requests that are based on symbolic names and configure the middleboxes so that data packets are correctly forwarded from NI to

NR. The configuration of those middleboxes may require other middlebox communication protocols, such as MIDCOM [RFC3303]. NSIS signaling is not required in the twice-NAT only case, since middleboxes of the twice-NAT type are normally configured by other means. Nevertheless, NSIS signaling might be useful when there are also firewalls on the path. In this case, NSIS will not configure any policy rule at twice-NATs, but will configure policy rules at the firewalls on the path. The NSIS signaling protocol must be at least robust enough to survive this scenario. This requires that twice-NATs must implement the NATFW NSLP also and participate in NATFW signaling sessions, but they do not change the configuration of the NAT, i.e., they only read the address mapping information out of the NAT and translate the Message Routing Information (MRI, [RFC5971]) within the NSLP and NTLF accordingly. For more information, see Appendix D.4.

2.6. Both End Hosts behind Same NAT

When the NSIS initiator and NSIS responder are behind the same NAT (thus, being in the same address realm, see Figure 7), they are most likely not aware of this fact. As in Section 2.4, the NSIS responder must determine its public IP address in advance and transfer it to the NSIS initiator. Afterwards, the NSIS initiator can start sending the signaling messages to the responder's public IP address. During this process, a public IP address will be allocated for the NSIS initiator at the same middlebox as for the responder. Now, the NSIS signaling and the subsequent data packets will traverse the NAT twice: from initiator to public IP address of responder (first time) and from public IP address of responder to responder (second time).



MB: Middlebox
NI: NSIS Initiator
NR: NSIS Responder

Figure 7: NAT to Public, Both Hosts behind Same NAT

2.7. Multihomed Network with NAT

The previous sub-sections sketched network topologies where several NATs and/or firewalls are ordered sequentially on the path. This section describes a multihomed scenario with two NATs placed on alternative paths to the public network.




Figure 8: Multihomed Network with Two NATs

Depending on the destination, either one or the other middlebox is used for the data flow. Which middlebox is used, depends on local policy or routing decisions. NATFW NSLP must be able to handle this situation properly, see Section 3.7.2 for an extended discussion of this topic with respect to NATs.

2.8. Multihomed Network with Firewall

This section describes a multihomed scenario with two firewalls placed on alternative paths to the public network (Figure 9). The routing in the private and public networks decides which firewall is being taken for data flows. Depending on the data flow's direction, either outbound or inbound, a different firewall could be traversed. This is a challenge for the EXTERNAL message of the NATFW NSLP where the NSIS responder is located behind these firewalls within the private network. The EXTERNAL message is used to block a particular data flow on an inbound firewall. NSIS must route the EXTERNAL message inbound from NR to NI probably without knowing which path the data traffic will take from NI to NR (see also Appendix C).




Figure 9: Multihomed Network with Two Firewalls

3. Protocol Description

This section defines messages, objects, and protocol semantics for the NATFW NSLP.

3.1. Policy Rules

Policy rules, bound to a NATFW NSLP signaling session, are the building blocks of middlebox devices considered in the NATFW NSLP. For firewalls, the policy rule usually consists of a 5-tuple and an action such as allow or deny. The information contained in the tuple includes source/destination IP addresses, transport protocol, and source/destination port numbers. For NATs, the policy rule consists of the action 'translate this address' and further mapping information, that might be, in the simplest case, internal IP address and external IP address.

The NATFW NSLP carries, in conjunction with the NTLF's Message Routing Information (MRI), the policy rules to be installed at NATFW peers. This policy rule is an abstraction with respect to the real policy rule to be installed at the respective firewall or NAT. It conveys the initiator's request and must be mapped to the possible configuration on the particular used NAT and/or firewall in use. For pure firewalls, one or more filter rules must be created, and for pure NATs, one or more NAT bindings must be created. In mixed firewall and NAT boxes, the policy rule must be mapped to filter rules and bindings observing the ordering of the firewall and NAT engine. Depending on the ordering, NAT before firewall or vice

versa, the firewall rules must carry public or private IP addresses. However, the exact mapping depends on the implementation of the firewall or NAT that is possibly different for each implementation.

The policy rule at the NATFW NSLP level comprises the message routing information (MRI) part, carried in the NTLP, and the information available in the NATFW NSLP. The information provided by the NSLP is stored in the 'extend flow information' (NATFW_EFI) and 'data terminal information' (NATFW_DTINFO) objects, and the message type. Additional information, such as the external IP address and port number, stored in the NAT or firewall, will be used as well. The MRI carries the filter part of the NAT/firewall-level policy rule that is to be installed.


The NATFW NSLP specifies two actions for the policy rules: deny and allow. A policy rule with action set to deny will result in all packets matching this rule to be dropped. A policy rule with action set to allow will result in all packets matching this rule to be forwarded.

3.2. Basic Protocol Overview

The NSIS NATFW NSLP is carried over the General Internet Signaling Transport (GIST, the implementation of the NTLP) defined in [RFC5971]. NATFW NSLP messages are initiated by the NSIS initiator (NI), handled by NSLP forwarders (NFs) and received by the NSIS responder (NR). It is required that at least NI and NR implement this NSLP, intermediate NFs only implement this NSLP when they provide relevant middlebox functions. NSLP forwarders that do not have any NATFW NSLP functions just forward these packets as they have no interest in them.

3.2.1. Signaling for Outbound Traffic

A data sender (DS), intending to send data to a data receiver (DR), has to start NATFW NSLP signaling. This causes the NI associated with the DS to launch NSLP signaling towards the address of the DR (see Figure 10). Although it is expected that the DS and the NATFW NSLP NI will usually reside on the same host, this specification does not rule out scenarios where the DS and NI reside on different hosts, the so-called proxy mode (see Section 3.7.6).



=====>
 Data Traffic Direction (outbound)

```

----> : NATFW NSLP request signaling
~~~>  : NATFW NSLP response signaling
DS/NI : Data sender and NSIS initiator
DR/NR : Data receiver and NSIS responder
MB1   : Middlebox 1 and NSLP forwarder 1
MB2   : Middlebox 2 and NSLP forwarder 2
    
```

Figure 10: General NSIS Signaling

The following list shows the normal sequence of NSLP events without detailing the interaction with the NTLF and the interactions on the NTLF level.

- o NSIS initiators generate request messages (which are either CREATE or EXTERNAL messages) and send these towards the NSIS responder. This request message is the initial message that creates a new NATFW NSLP signaling session. The NI and the NR will most likely already share an application session before they start the NATFW NSLP signaling session. Note well the difference between both sessions.
- o NSLP request messages are processed each time an NF with NATFW NSLP support is traversed. Each NF that is intercepting a request message and is accepting it for further treatment is joining the particular NATFW NSLP signaling session. These nodes process the message, check local policies for authorization and authentication, possibly create policy rules, and forward the signaling message to the next NSIS node. The request message is forwarded until it reaches the NSIS responder.
- o NSIS responders will check received messages and process them if applicable. NSIS responders generate RESPONSE messages and send them hop-by-hop back to the NI via the same chain of NFs (traversal of the same NF chain is guaranteed through the established reverse message routing state in the NTLF). The NR is also joining the NATFW NSLP signaling session if the request message is accepted.

- o The RESPONSE message is processed at each NF that has been included in the prior NATFW NSLP signaling session setup.
- o If the NI has received a successful RESPONSE message and if the signaling NATFW NSLP session started with a CREATE message, the data sender can start sending its data flow to the data receiver. If the NI has received a successful RESPONSE message and if the signaling NATFW NSLP session started with an EXTERNAL message, the data receiver is ready to receive further CREATE messages.

Because NATFW NSLP signaling follows the data path from DS to DR, this immediately enables communication between both hosts for scenarios with only firewalls on the data path or NATs on the sender side. For scenarios with NATs on the receiver side, certain problems arise, as described in Section 2.4.

3.2.2. Signaling for Inbound Traffic

When the NR and the NI are located in different address realms and the NR is located behind a NAT, the NI cannot signal to the NR address directly. The DR/NR is not reachable from other NIs using the private address of the NR and thus NATFW signaling messages cannot be sent to the NR/DR's address. Therefore, the NR must first obtain a NAT binding that provides an address that is reachable for the NI. Once the NR has acquired a public IP address, it forwards this information to the DS via a separate protocol. This application-layer signaling, which is out of the scope of the NATFW NSLP, may involve third parties that assist in exchanging these messages.

The same holds partially true for NRs located behind firewalls that block all traffic by default. In this case, NR must tell its inbound firewalls of inbound NATFW NSLP signaling and corresponding data traffic. Once the NR has informed the inbound firewalls, it can start its application-level signaling to initiate communication with the NI. This mechanism can be used by machines hosting services behind firewalls as well. In this case, the NR informs the inbound firewalls as described, but does not need to communicate this to the NIs.

NATFW NSLP signaling supports this scenario by using the EXTERNAL message.

1. The DR acquires a public address by signaling on the reverse path (DR towards DS) and thus making itself available to other hosts. This process of acquiring public addresses is called reservation. During this process the DR reserves publicly reachable addresses and ports suitable for further usage in application-level

signaling and the publicly reachable address for further NATFW NSLP signaling. However, the data traffic will not be allowed to use this address/port initially (see next point). In the process of reservation, the DR becomes the NI for the messages necessary to obtain the publicly reachable IP address, i.e., the NI for this specific NATFW NSLP signaling session.

2. Now on the side of the DS, the NI creates a new NATFW NSLP signaling session and signals directly to the public IP address of the DR. This public IP address is used as NR's address, as the NI would do if there is no NAT in between, and creates policy rules at middleboxes. Note, that the reservation will only allow forwarding of signaling messages, but not data flow packets. Policy rules allowing forwarding of data flow packets set up by the prior EXTERNAL message signaling will be activated when the signaling from NI towards NR is confirmed with a positive RESPONSE message. The EXTERNAL message is described in Section 3.7.2.

3.2.3. Signaling for Proxy Mode

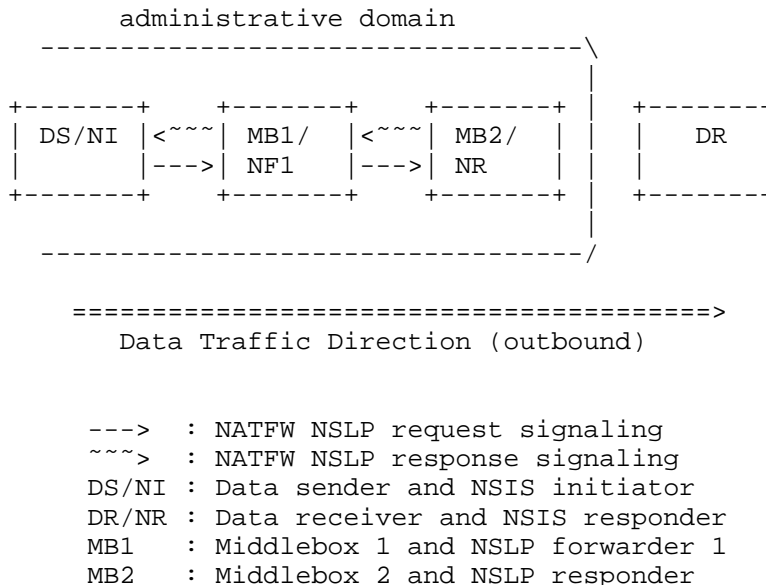


Figure 11: Proxy Mode Signaling for Data Sender

The above usage assumes that both ends of a communication support NSIS, but fails when NSIS is only deployed at one end of the path. In this case, only one of the sending side (see Figure 11) or receiving side (see Figure 12) is NSIS aware and not both at the same

time. NATFW NSLP supports both scenarios (i.e., either the DS or DR does not support NSIS) by using a proxy mode, as described in Section 3.7.6.

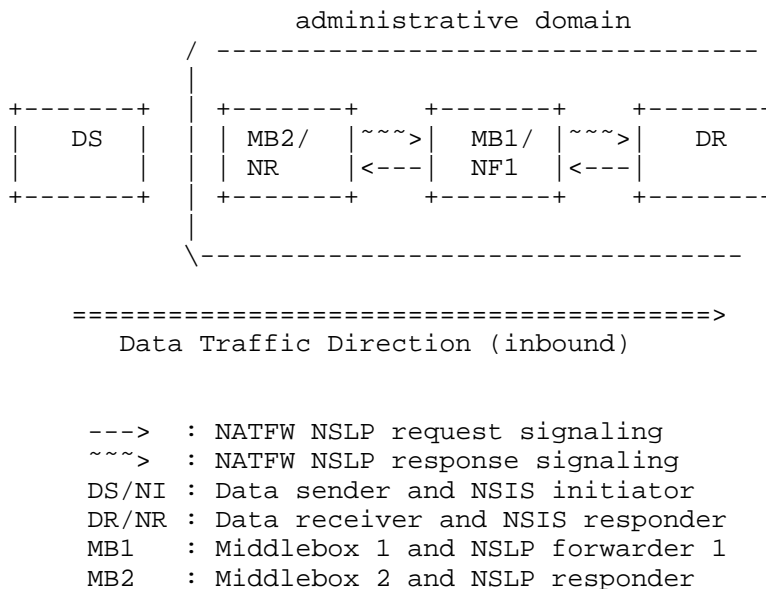


Figure 12: Proxy Mode Signaling for Data Receiver

3.2.4. Blocking Traffic

The basic functionality of the NATFW NSLP provides for opening firewall pin holes and creating NAT bindings to enable data flows to traverse these devices. Firewalls are normally expected to work on a "deny-all" policy, meaning that traffic not explicitly matching any firewall filter rule will be blocked. Similarly, the normal behavior of NATs is to block all traffic that does not match any already configured/installed binding or NATFW NSLP session. However, some scenarios require support of firewalls having "allow-all" policies, allowing data traffic to traverse the firewall unless it is blocked explicitly. Data receivers can utilize NATFW NSLP's EXTERNAL message with action set to "deny" to install policy rules at inbound firewalls to block unwanted traffic.

3.2.5. State and Error Maintenance

The protocol works on a soft-state basis, meaning that whatever state is installed or reserved on a middlebox will expire, and thus be uninstalled or forgotten after a certain period of time. To prevent premature removal of state that is needed for ongoing communication,

the NATFW NI involved will have to specifically request a NATFW NSLP signaling session extension. An explicit NATFW NSLP state deletion capability is also provided by the protocol.

If the actions requested by a NATFW NSLP message cannot be carried out, NFs and the NR must return a failure, such that appropriate actions can be taken. They can do this either during the request message handling (synchronously) by sending an error RESPONSE message or at any time (asynchronously) by sending a NOTIFY notification message.

The next sections define the NATFW NSLP message types and formats, protocol operations, and policy rule operations.

3.2.6. Message Types

The protocol uses four messages types:

- o CREATE: a request message used for creating, changing, refreshing, and deleting NATFW NSLP signaling sessions, i.e., open the data path from DS to DR.
- o EXTERNAL: a request message used for reserving, changing, refreshing, and deleting EXTERNAL NATFW NSLP signaling sessions. EXTERNAL messages are forwarded to the edge-NAT or edge-firewall and allow inbound CREATE messages to be forwarded to the NR. Additionally, EXTERNAL messages reserve an external address and, if applicable, port number at an edge-NAT.
- o NOTIFY: an asynchronous message used by NATFW peers to alert other NATFW peers about specific events (especially failures).
- o RESPONSE: used as a response to CREATE and EXTERNAL request messages.

3.2.7. Classification of RESPONSE Messages

RESPONSE messages will be generated synchronously to CREATE and EXTERNAL messages by NSLP forwarders and responders to report success or failure of operations or some information relating to the NATFW NSLP signaling session or a node. RESPONSE messages MUST NOT be generated for any other message, such as NOTIFY and RESPONSE.

All RESPONSE messages MUST carry a NATFW_INFO object that contains an error class code and a response code (see Section 4.2.5). This section defines terms for groups of RESPONSE messages depending on the error class.

- o Successful RESPONSE: Messages carrying NATFW_INFO with error class 'Success' (2).
- o Informational RESPONSE: Messages carrying NATFW_INFO with error class 'Informational' (1) (only used with NOTIFY messages).
- o Error RESPONSE: Messages carrying NATFW_INFO with error class other than 'Success' or 'Informational'.

3.2.8. NATFW NSLP Signaling Sessions

A NATFW NSLP signaling session defines an association between the NI, NFs, and the NR related to a data flow. This association is created when the initial CREATE or EXTERNAL message is successfully received at the NFs or the NR. There is signaling NATFW NSLP session state stored at the NTLF layer and at the NATFW NSLP level. The NATFW NSLP signaling session state for the NATFW NSLP comprises NSLP state and the associated policy rules at a middlebox.

The NATFW NSLP signaling session is identified by the session ID (plus other information at the NTLF level). The session ID is generated by the NI before the initial CREATE or EXTERNAL message is sent. The value of the session ID MUST be generated as a cryptographically random number (see [RFC4086]) by the NI, i.e., the output MUST NOT be easily guessable by third parties. The session ID is not stored in any NATFW NSLP message but passed on to the NTLF.

A NATFW NSLP signaling session has several conceptual states that describe in what state a signaling session is at a given time. The signaling session can have these states at a node:

- o Pending: The NATFW NSLP signaling session has been created and the node is waiting for a RESPONSE message to the CREATE or EXTERNAL message. A NATFW NSLP signaling session in state 'Pending' MUST be marked as 'Dead' if no corresponding RESPONSE message has been received within the time of the locally granted NATFW NSLP signaling session lifetime of the forwarded CREATE or EXTERNAL message (as described in Section 3.4).
- o Established: The NATFW NSLP signaling session is established, i.e., the signaling has been successfully performed and the lifetime of NATFW NSLP signaling session is counted from now on. A NATFW NSLP signaling session in state 'Established' MUST be marked as 'Dead' if no refresh message has been received within the time of the locally granted NATFW NSLP signaling session lifetime of the RESPONSE message (as described in Section 3.4).

- o Dead: Either the NATFW NSLP signaling session is timed out or the node has received an error RESPONSE message for the NATFW NSLP signaling session and the NATFW NSLP signaling session can be deleted.
- o Transitory: The node has received an asynchronous message, i.e., a NOTIFY, and can delete the NATFW NSLP signaling session if needed after some time. When a node has received a NOTIFY message, it marks the signaling session as 'Transitory'. This signaling session SHOULD NOT be deleted before a minimum hold time of 30 seconds, i.e., it can be removed after 30 seconds or more. This hold time ensures that the existing signaling session can be reused by the NI, e.g., a part of a signaling session that is not affected by the route change can be reused once the updating request message is received.

3.3. Basic Message Processing

All NATFW messages are subject to some basic message processing when received at a node, independent of the message type. Initially, the syntax of the NSLP message is checked and a RESPONSE message with an appropriate error of class 'Protocol error' (3) code is generated if a non-recoverable syntax error is detected. A recoverable error is, for instance, when a node receives a message with reserved flags set to values other than zero. This also refers to unknown NSLP objects and their handling, according to Section 4.2. If a message is delivered to the NATFW NSLP, this implies that the NTLP layer has been able to correlate it with the session ID (SID) and MRI entries in its database. There is therefore enough information to identify the source of the message and routing information to route the message back to the NI through an established chain of NTLP messaging associations. The message is not further forwarded if any error in the syntax is detected. The specific response codes stemming from the processing of objects are described in the respective object definition section (see Section 4). After passing this check, the NATFW NSLP node performs authentication- and authorization-related checks, described in Section 3.6. Further processing is executed only if these tests have been successfully passed; otherwise, the processing stops and an error RESPONSE is returned.

Further message processing stops whenever an error RESPONSE message is generated, and the EXTERNAL or CREATE message is discarded.

3.4. Calculation of Signaling Session Lifetime

NATFW NSLP signaling sessions, and the corresponding policy rules that may have been installed, are maintained via a soft-state mechanism. Each signaling session is assigned a signaling session

lifetime and the signaling session is kept alive as long as the lifetime is valid. After the expiration of the signaling session lifetime, signaling sessions and policy rules MUST be removed automatically and resources bound to them MUST be freed as well. Signaling session lifetime is handled at every NATFW NSLP node. The NSLP forwarders and NSLP responder MUST NOT trigger signaling session lifetime extension refresh messages (see Section 3.7.3): this is the task of the NSIS initiator.

The NSIS initiator MUST choose a NATFW NSLP signaling session lifetime value (expressed in seconds) before sending any message, including the initial message that creates the NATFW NSLP signaling session, to other NSLP nodes. It is RECOMMENDED that the NATFW NSLP signaling session lifetime value is calculated based on:

- o the number of lost refresh messages with which NFs should cope;
- o the end-to-end delay between the NI and NR;
- o network vulnerability due to NATFW NSLP signaling session hijacking ([RFC4081]), NATFW NSLP signaling session hijacking is made easier when the NI does not explicitly remove the NATFW NSLP signaling session;
- o the user application's data exchange duration, in terms of time and networking needs. This duration is modeled as R , with R the message refresh period (in seconds);
- o the load on the signaling plane. Short lifetimes imply more frequent signaling messages;
- o the acceptable time for a NATFW NSLP signaling session to be present after it is no longer actually needed. For example, if the existence of the NATFW NSLP signaling session implies a monetary cost and teardown cannot be guaranteed, shorter lifetimes would be preferable;
- o the lease time of the NI's IP address. The lease time of the IP address must be longer than the chosen NATFW NSLP signaling session lifetime; otherwise, the IP address can be re-assigned to a different node. This node may receive unwanted traffic, although it never has requested a NAT/firewall configuration, which might be an issue in environments with mobile hosts.

The RSVP specification [RFC2205] provides an appropriate algorithm for calculating the NATFW NSLP signaling session lifetime as well as a means to avoid refresh message synchronization between NATFW NSLP signaling sessions. [RFC2205] recommends:

1. The refresh message timer to be randomly set to a value in the range $[0.5R, 1.5R]$.
2. To avoid premature loss of state, lt (with lt being the NATFW NSLP signaling session lifetime) must satisfy $lt \geq (K + 0.5) * 1.5 * R$, where K is a small integer. Then, in the worst case, $K-1$ successive messages may be lost without state being deleted. Currently, $K = 3$ is suggested as the default. However, it may be necessary to set a larger K value for hops with high loss rate. Other algorithms could be used to define the relation between the NATFW NSLP signaling session lifetime and the refresh message period; the algorithm provided is only given as an example.

It is RECOMMENDED to use a refresh timer of 300 s (5 minutes), unless the NI or the requesting application at the NI has other requirements (e.g., flows lasting a very short time).

This requested NATFW NSLP signaling session lifetime value lt is stored in the NATFW_LT object of the NSLP message.

NSLP forwarders and the NSLP responder can execute the following behavior with respect to the requested lifetime handling:

Requested signaling session lifetime acceptable:

No changes to the NATFW NSLP signaling session lifetime values are needed. The CREATE or EXTERNAL message is forwarded, if applicable.

Signaling session lifetime can be lowered:

An NSLP forwarded or the NSLP responder MAY also lower the requested NATFW NSLP signaling session lifetime to an acceptable value (based on its local policies). If an NF changes the NATFW NSLP signaling session lifetime value, it MUST store the new value in the NATFW_LT object. The CREATE or EXTERNAL message is forwarded.

Requested signaling session lifetime is too big:

An NSLP forwarded or the NSLP responder MAY reject the requested NATFW NSLP signaling session lifetime value as being too big and MUST generate an error RESPONSE message of class 'Signaling session failure' (7) with response code 'Requested lifetime is too big' (0x02) upon rejection. Lowering the lifetime is preferred instead of generating an error message.

Requested signaling session lifetime is too small:


An NSLP forwarder or the NSLP responder MAY reject the requested NATFW NSLP signaling session lifetime value as being too small and MUST generate an error RESPONSE message of class 'Signaling session failure' (7) with response code 'Requested lifetime is too small' (0x10) upon rejection.

NFs or the NR MUST NOT increase the NATFW NSLP signaling session lifetime value. Messages can be rejected on the basis of the NATFW NSLP signaling session lifetime being too long when a NATFW NSLP signaling session is first created and also on refreshes.

The NSLP responder generates a successful RESPONSE for the received CREATE or EXTERNAL message, sets the NATFW NSLP signaling session lifetime value in the NATFW_LT object to the above granted lifetime and sends the message back towards NSLP initiator.

Each NSLP forwarder processes the RESPONSE message and reads and stores the granted NATFW NSLP signaling session lifetime value. The forwarders MUST accept the granted NATFW NSLP signaling session lifetime, if the lifetime value is within the acceptable range. The acceptable value refers to the value accepted by the NSLP forwarder when processing the CREATE or EXTERNAL message. For received values greater than the acceptable value, NSLP forwarders MUST generate a RESPONSE message of class 'Signaling session failure' (7) with response code 'Modified lifetime is too big' (0x11), including a Signaling Session Lifetime object that carries the maximum acceptable signaling session lifetime for this node. For received values lower than the values acceptable by the node local policy, NSLP forwarders MUST generate a RESPONSE message of class 'Signaling session failure' (7) with response code 'Modified lifetime is too small' (0x12), including a Signaling Session Lifetime object that carries the minimum acceptable signaling session lifetime for this node. In both cases, either 'Modified lifetime is too big' (0x11) or 'Modified lifetime is too small' (0x12), the NF MUST generate a NOTIFY message and send it outbound with the error class set to 'Informational' (1) and with the response code set to 'NATFW signaling session terminated' (0x05).

Figure 13 shows the procedure with an example, where an initiator requests 60 seconds lifetime in the CREATE message and the lifetime is shortened along the path by the forwarder to 20 seconds and by the responder to 15 seconds. When the NSLP forwarder receives the RESPONSE message with a NATFW NSLP signaling session lifetime value of 15 seconds it checks whether this value is lower or equal to the acceptable value.



lt = lifetime

Figure 13: Signaling Session Lifetime Setting Example

3.5. Message Sequencing

NATFW NSLP messages need to carry an identifier so that all nodes along the path can distinguish messages sent at different points in time. Messages can be lost along the path or duplicated. So, all NATFW NSLP nodes should be able to identify messages that have been received before (duplicated) or lost before (loss). For message replay protection, it is necessary to keep information about messages that have already been received and requires every NATFW NSLP message to carry a message sequence number (MSN), see also Section 4.2.7.

The MSN MUST be set by the NI and MUST NOT be set or modified by any other node. The initial value for the MSN MUST be generated randomly and MUST be unique only within the NATFW NSLP signaling session for which it is used. The NI MUST increment the MSN by one for every message sent. Once the MSN has reached the maximum value, the next value it takes is zero. All NATFW NSLP nodes MUST use the algorithm defined in [RFC1982] to detect MSN wrap-arounds.

NSLP forwarders and the responder store the MSN from the initial CREATE or EXTERNAL packet that creates the NATFW NSLP signaling session as the start value for the NATFW NSLP signaling session. NFs and NRs MUST include the received MSN value in the corresponding RESPONSE message that they generate.

When receiving a CREATE or EXTERNAL message, a NATFW NSLP node uses the MSN given in the message to determine whether the state being requested is different from the state already installed. The message MUST be discarded if the received MSN value is equal to or lower than the stored MSN value. Such a received MSN value can indicate a duplicated and delayed message or replayed message. If the received MSN value is greater than the already stored MSN value, the NATFW NSLP MUST update its stored state accordingly, if permitted by all security checks (see Section 3.6), and store the updated MSN value accordingly.

3.6. Authentication, Authorization, and Policy Decisions

NATFW NSLP nodes receiving signaling messages MUST first check whether this message is authenticated and authorized to perform the requested action. NATFW NSLP nodes requiring more information than provided MUST generate an error RESPONSE of class 'Permanent failure' (0x5) with response code 'Authentication failed' (0x01) or with response code 'Authorization failed' (0x02).

The NATFW NSLP is expected to run in various environments, such as IP-based telephone systems, enterprise networks, home networks, etc. The requirements on authentication and authorization are quite different between these use cases. While a home gateway, or an Internet cafe, using NSIS may well be happy with a "NATFW signaling coming from inside the network" policy for authorization of signaling, enterprise networks are likely to require more strongly authenticated/authorized signaling. This enterprise scenario may require the use of an infrastructure and administratively assigned identities to operate the NATFW NSLP.

Once the NI is authenticated and authorized, another step is performed. The requested policy rule for the NATFW NSLP signaling session is checked against a set of policy rules, i.e., whether the requesting NI is allowed to request the policy rule to be loaded in the device. If this fails, the NF or NR must send an error RESPONSE of class 'Permanent failure' (5) and with response code 'Authorization failed' (0x02).

3.7. Protocol Operations

This section defines the protocol operations including how to create NATFW NSLP signaling sessions, maintain them, delete them, and how to reserve addresses.

This section requires a good knowledge of the NTLTP [RFC5971] and the message routing method mechanism and the associated message routing information (MRI). The NATFW NSLP uses information from the MRI, e.g., the destination and source ports, and the NATFW NSLP to construct the policy rules used on the NATFW NSLP level. See also Appendix D for further information about this.

3.7.1. Creating Signaling Sessions

Allowing two hosts to exchange data even in the presence of middleboxes is realized in the NATFW NSLP by the use of the CREATE message. The NI (either the data sender or a proxy) generates a CREATE message as defined in Section 4.3.1 and hands it to the NTLTP. The NTLTP forwards the whole message on the basis of the message

routing information (MRI) towards the NR. Each NSLP forwarder along the path that implements NATFW NSLP processes the NSLP message. Forwarding is done hop-by-hop but may pass transparently through NSLP forwarders that do not contain NATFW NSLP functionality and non-NSIS-aware routers between NSLP hop way points. When the message reaches the NR, the NR can accept the request or reject it. The NR generates a response to CREATE and this response is transported hop-by-hop towards the NI. NATFW NSLP forwarders may reject requests at any time. Figure 14 sketches the message flow between the NI (DS in this example), an NF (e.g., NAT), and an NR (DR in this example).

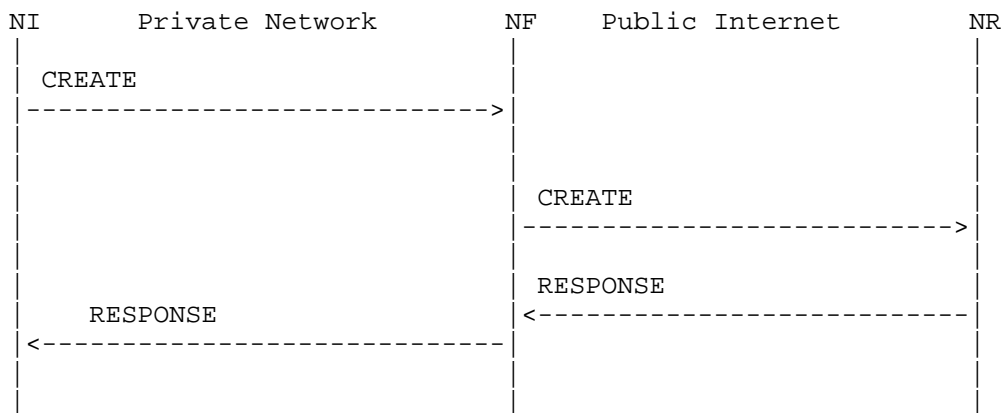


Figure 14: CREATE Message Flow with Success RESPONSE

There are several processing rules for a NATFW peer when generating and receiving CREATE messages, since this message type is used for creating new NATFW NSLP signaling sessions, updating existing ones, and extending the lifetime and deleting NATFW NSLP signaling sessions. The three latter functions operate in the same way for all kinds of CREATE messages, and are therefore described in separate sections:

- o Extending the lifetime of NATFW NSLP signaling sessions is described in Section 3.7.3.
- o Deleting NATFW NSLP signaling sessions is described in Section 3.7.4.
- o Updating policy rules is described in Section 3.10.

For an initial CREATE message creating a new NATFW NSLP signaling session, the processing of CREATE messages is different for every NATFW node type:

- o NSLP initiator: An NI only generates CREATE messages and hands them over to the NTLP. The NI should never receive CREATE messages and MUST discard them.
- o NATFW NSLP forwarder: NFs that are unable to forward the CREATE message to the next hop MUST generate an error RESPONSE of class 'Permanent failure' (5) with response code 'Did not reach the NR' (0x07). This case may occur if the NTLP layer cannot find a NATFW NSLP peer, either another NF or the NR, and returns an error via the GIST API (a timeout error reported by GIST). The NSLP message processing at the NFs depends on the middlebox type:
 - * NAT: When the initial CREATE message is received at the public side of the NAT, it looks for a reservation made in advance, by using an EXTERNAL message (see Section 3.7.2). The matching process considers the received MRI information and the stored MRI information, as described in Section 3.8. If no matching reservation can be found, i.e., no reservation has been made in advance, the NSLP MUST return an error RESPONSE of class 'Signaling session failure' (7) with response code 'No reservation found matching the MRI of the CREATE request' (0x03). If there is a matching reservation, the NSLP stores the data sender's address (and if applicable port number) as part of the source IP address of the policy rule ('the remembered policy rule') to be loaded, and forwards the message with the destination IP address set to the internal (private in most cases) address of the NR. When the initial CREATE message is received at the private side, the NAT binding is allocated, but not activated (see also Appendix D.3). An error RESPONSE message is generated, if the requested policy rule cannot be reserved right away, of class 'Signaling session failure' (7) with response code 'Requested policy rule denied due to policy conflict' (0x4). The MRI information is updated to reflect the address, and if applicable port, translation. The NSLP message is forwarded towards the NR with source IP address set to the NAT's external address from the newly remembered binding.
 - * Firewall: When the initial CREATE message is received, the NSLP just remembers the requested policy rule, but does not install any policy rule. Afterwards, the message is forwarded towards the NR. If the requested policy rule cannot be reserved right away, an error RESPONSE message is generated, of class 'Signaling session failure' (7) with response code 'Requested policy rule denied due to policy conflict' (0x4).
 - * Combined NAT and firewall: Processing at combined firewall and NAT middleboxes is the same as in the NAT case. No policy rules are installed. Implementations MUST take into account

the order of packet processing in the firewall and NAT functions within the device. This will be referred to as "order of functions" and is generally different depending on whether the packet arrives at the external or internal side of the middlebox.

- o NSLP receiver: NRs receiving initial CREATE messages MUST reply with a success RESPONSE of class 'Success' (2) with response code set to 'All successfully processed' (0x01), if they accept the CREATE message. Otherwise, they MUST generate a RESPONSE message with a suitable response code. RESPONSE messages are sent back NSLP hop-by-hop towards the NI, irrespective of the response codes, either success or error.

Remembered policy rules at middleboxes MUST be only installed upon receiving a corresponding successful RESPONSE message with the same SID as the CREATE message that caused them to be remembered. This is a countermeasure to several problems, for example, wastage of resources due to loading policy rules at intermediate NFs when the CREATE message does not reach the final NR for some reason.

Processing of a RESPONSE message is different for every NSIS node type:

- o NSLP initiator: After receiving a successful RESPONSE, the data path is configured and the DS can start sending its data to the DR. After receiving an error RESPONSE message, the NI MAY try to generate the CREATE message again or give up and report the failure to the application, depending on the error condition.
- o NSLP forwarder: NFs install the remembered policy rules, if a successful RESPONSE message with matching SID is received. If an ERROR RESPONSE message with matching SID is received, the NATFW NSLP session is marked as 'Dead', no policy rule is installed and the remembered rule is discarded.
- o NSIS responder: The NR should never receive RESPONSE messages and MUST silently drop any such messages received.

NFs and the NR can also tear down the CREATE session at any time by generating a NOTIFY message with the appropriate response code set.

3.7.2. Reserving External Addresses

NSIS signaling is intended to travel end-to-end, even in the presence of NATs and firewalls on-path. This works well in cases where the data sender is itself behind a NAT or a firewall as described in Section 3.7.1. For scenarios where the data receiver is located

behind a NAT or a firewall and it needs to receive data flows from outside its own network (usually referred to as inbound flows, see Figure 5), the problem is more troublesome.

NSIS signaling, as well as subsequent data flows, are directed to a particular destination IP address that must be known in advance and reachable. Data receivers must tell the local NSIS infrastructure (i.e., the inbound firewalls/NATs) about incoming NATFW NSLP signaling and data flows before they can receive these flows. It is necessary to differentiate between data receivers behind NATs and behind firewalls to understand the further NATFW procedures. Data receivers that are only behind firewalls already have a public IP address and they need only to be reachable for NATFW signaling. Unlike data receivers that are only behind firewalls, data receivers behind NATs do not have public IP addresses; consequently, they are not reachable for NATFW signaling by entities outside their addressing realm.

The preceding discussion addresses the situation where a DR node that wants to be reachable is unreachable because the NAT lacks a suitable rule with the 'allow' action that would forward inbound data. However, in certain scenarios, a node situated behind inbound firewalls that do not block inbound data traffic (firewalls with "default to allow") unless requested might wish to prevent traffic being sent to it from specified addresses. In this case, NSIS NATFW signaling can be used to achieve this by installing a policy rule with its action set to 'deny' using the same mechanisms as for 'allow' rules.

The required result is obtained by sending an EXTERNAL message in the inbound direction of the intended data flow. When using this functionality, the NSIS initiator for the 'Reserve External Address' signaling is typically the node that will become the DR for the eventual data flow. To distinguish this initiator from the usual case where the NI is associated with the DS, the NI is denoted by NI+ and the NSIS responder is similarly denoted by NR+.

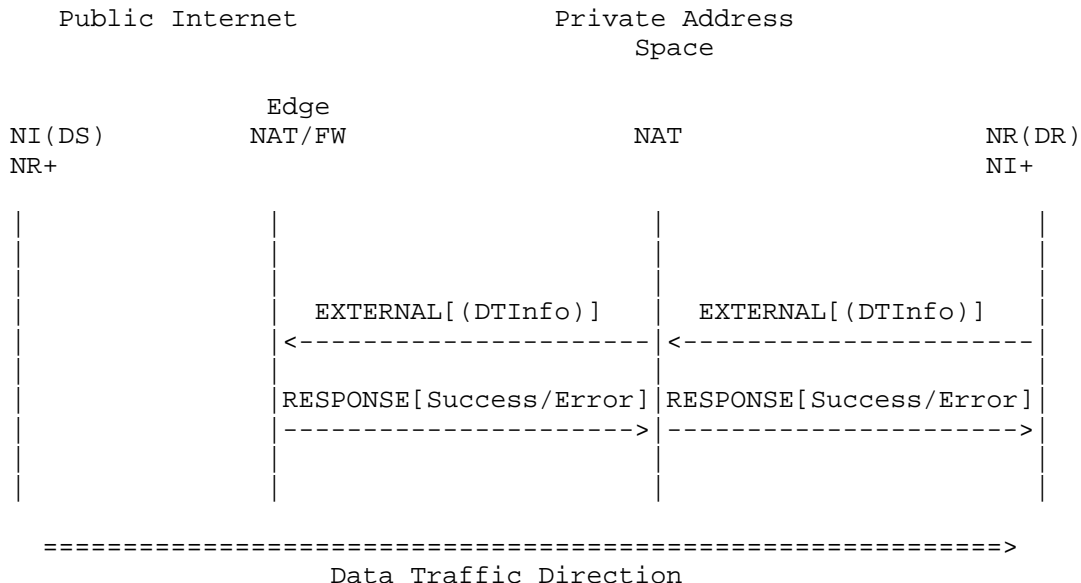


Figure 15: Reservation Message Flow for DR behind NAT or Firewall

Figure 15 shows the EXTERNAL message flow for enabling inbound NATFW NSLP signaling messages. In this case, the roles of the different NSIS entities are:

- o The data receiver (DR) for the anticipated data traffic is the NSIS initiator (NI+) for the EXTERNAL message, but becomes the NSIS responder (NR) for following CREATE messages.
- o The actual data sender (DS) will be the NSIS initiator (NI) for later CREATE messages and may be the NSIS target of the signaling (NR+).
- o It may be necessary to use a signaling destination address (SDA) as the actual target of the EXTERNAL message (NR+) if the DR is located behind a NAT and the address of the DS is unknown. The SDA is an arbitrary address in the outermost address realm on the other side of the NAT from the DR. Typically, this will be a suitable public IP address when the 'outside' realm is the public Internet. This choice of address causes the EXTERNAL message to be routed through the NATs towards the outermost realm and would force interception of the message by the outermost NAT in the network at the boundary between the private address and the public address realm (the edge-NAT). It may also be intercepted by other NATs and firewalls on the path to the edge-NAT.

Basically, there are two different signaling scenarios. Either

1. the DR behind the NAT/firewall knows the IP address of the DS in advance, or
2. the address of the DS is not known in advance.

Case 1 requires the NATFW NSLP to request the path-coupled message routing method (PC-MRM) from the NTLP. The EXTERNAL message MUST be sent with PC-MRM (see Section 5.8.1 in [RFC5971]) with the direction set to 'upstream' (inbound). The handling of case 2 depends on the situation of the DR: if the DR is solely located behind a firewall, the EXTERNAL message MUST be sent with the PC-MRM, direction 'upstream' (inbound), and the data flow source IP address set to 'wildcard'. If the DR is located behind a NAT, the EXTERNAL message MUST be sent with the loose-end message routing method (LE-MRM, see Section 5.8.2 in [RFC5971]), the destination-address set to the signaling destination IP address (SDA, see also Appendix A). For scenarios with the DR behind a firewall, special conditions apply (see applicability statement in Appendix C). The data receiver is challenged to determine whether it is solely located behind firewalls or NATs in order to choose the right message routing method. This decision can depend on a local configuration parameter, possibly given through DHCP, or it could be discovered through other non-NSLP related testing of the network configuration. The use of the PC-MRM with the known data sender's IP address is RECOMMENDED. This gives GIST the best possible handle to route the message 'upstream' (outbound). The use of the LE-MRM, if and only if the data sender's IP address is not known and the data receiver is behind a NAT, is RECOMMENDED.

For case 2 with NAT, the NI+ (which could be on the data receiver DR or on any other host within the private network) sends the EXTERNAL message targeted to the signaling destination IP address. The message routing for the EXTERNAL message is in the reverse direction of the normal message routing used for path-coupled signaling where the signaling is sent outbound (as opposed to inbound in this case). When establishing NAT bindings (and a NATFW NSLP signaling session), the signaling direction does not matter since the data path is modified through route pinning due to the external IP address at the NAT. Subsequent NSIS messages (and also data traffic) will travel through the same NAT boxes. However, this is only valid for the NAT boxes, but not for any intermediate firewall. That is the reason for having a separate CREATE message enabling the reservations made with EXTERNAL at the NATs and either enabling prior reservations or creating new pinholes at the firewalls that are encountered on the outbound path depending on whether the inbound and outbound routes coincide.

The EXTERNAL signaling message creates an NSIS NATFW signaling session at any intermediate NSIS NATFW peer(s) encountered, independent of the message routing method used. Furthermore, it has to be ensured that the edge-NAT or edge-firewall device is discovered as part of this process. The end host cannot be assumed to know this device -- instead the NAT or firewall box itself is assumed to know that it is located at the outer perimeter of the network. Forwarding of the EXTERNAL message beyond this entity is not necessary, and MUST be prohibited as it may provide information on the capabilities of internal hosts. It should be noted, that it is the outermost NAT or firewall that is the edge-device that must be found during this discovery process. For instance, when there are a NAT and (afterwards) a firewall on the outbound path at the network border, the firewall is the edge-firewall. All messages must be forwarded to the topology-wise outermost edge-device to ensure that this device knows about the NATFW NSLP signaling sessions for incoming CREATE messages. However, the NAT is still the edge-NAT because it has a public globally routable IP address on its public side: this is not affected by any firewall between the edge-NAT and the public network.

Possible edge arrangements are:

```

Public Net ----- Private net -----
| Public Net|--|Edge-FW|--|FW|...|FW|--|DR|
| Public Net|--|Edge-FW|--|Edge-NAT|...|NAT or FW|--|DR|
| Public Net|--|Edge-NAT|--|NAT or FW|...|NAT or FW|--|DR|

```

The edge-NAT or edge-firewall device closest to the public realm responds to the EXTERNAL request message with a successful RESPONSE message. An edge-NAT includes a NATFW_EXTERNAL_IP object (see Section 4.2.2), carrying the publicly reachable IP address, and if applicable, a port number.

The NI+ can request each intermediate NAT (i.e., a NAT that is not the edge-NAT) to include the external binding address (and if applicable port number) in the external binding address object. The external binding address object stores the external IP address (and port) at the particular NAT. The NI+ has to include the external binding address (see Section 4.2.3) object in the request message, if it wishes to obtain the information.

There are several processing rules for a NATFW peer when generating and receiving EXTERNAL messages, since this message type is used for creating new reserve NATFW NSLP signaling sessions, updating existing, extending the lifetime, and deleting NATFW NSLP signaling

session. The three latter functions operate in the same way for all kinds of CREATE and EXTERNAL messages, and are therefore described in separate sections:

- o Extending the lifetime of NATFW NSLP signaling sessions is described in Section 3.7.3.
- o Deleting NATFW NSLP signaling sessions is described in Section 3.7.4.
- o Updating policy rules is described in Section 3.10.

The NI+ MUST always include a NATFW_DTINFO object in the EXTERNAL message. Especially, the LE-MRM does not include enough information for some types of NATs (basically, those NATs that also translate port numbers) to perform the address translation. This information is provided in the NATFW_DTINFO (see Section 4.2.8). This information MUST include at least the 'dst port number' and 'protocol' fields, in the NATFW_DTINFO object as these may be required by NATs that are en route, depending on the type of the NAT. All other fields MAY be set by the NI+ to restrict the set of possible NIs. An edge-NAT will use the information provided in the NATFW_DTINFO object to allow only a NATFW CREATE message with a matching MRI to be forwarded. The MRI of the NATFW CREATE message has to use the parameters set in NATFW_DTINFO object ('src IPv4/v6 address', 'src port number', 'protocol') as the source IP address/port of the flow from DS to DR. A NAT requiring information carried in the NATFW_DTINFO can generate a number of error RESPONSE messages of class 'Signaling session failure' (7):

- o 'Requested policy rule denied due to policy conflict' (0x04)
- o 'Unknown policy rule action' (0x05)
- o 'Requested rule action not applicable' (0x06)
- o 'NATFW_DTINFO object is required' (0x07)
- o 'Requested value in sub_ports field in NATFW_EFI not permitted' (0x08)
- o 'Requested IP protocol not supported' (0x09)
- o 'Plain IP policy rules not permitted -- need transport layer information' (0x0A)
- o 'Source IP address range is too large' (0x0C)

- o 'Destination IP address range is too large' (0x0D)
- o 'Source L4-port range is too large' (0x0E)
- o 'Destination L4-port range is too large' (0x0F)

Processing of EXTERNAL messages is specific to the NSIS node type:

- o NSLP initiator: NI+ only generate EXTERNAL messages. When the data sender's address information is known in advance, the NI+ can include a NATFW_DTINFO object in the EXTERNAL message, if not anyway required to do so (see above). When the data sender's IP address is not known, the NI+ MUST NOT include an IP address in the NATFW_DTINFO object. The NI should never receive EXTERNAL messages and MUST silently discard it.
- o NSLP forwarder: The NSLP message processing at NFs depends on the middlebox type:
 - * NAT: NATs check whether the message is received at the external (public in most cases) address or at the internal (private) address. If received at the external address, an NF MUST generate an error RESPONSE of class 'Protocol error' (3) with response code 'Received EXTERNAL request message on external side' (0x0B). If received at the internal (private address) and the NATFW_EFI object contains the action 'deny', an error RESPONSE of class 'Protocol error' (3) with response code 'Requested rule action not applicable' (0x06) MUST be generated. If received at the internal address, an IP address, and if applicable, one or more ports, are reserved. If the NATFW_EXTERNAL_BINDING object is present in the message, any NAT that is not an edge-NAT MUST include the allocated external IP address, and if applicable one or more ports, (the external binding address) in the NATFW_EXTERNAL_BINDING object. If it is an edge-NAT and there is no edge-firewall beyond, the NSLP message is not forwarded any further and a successful RESPONSE message is generated containing a NATFW_EXTERNAL_IP object holding the translated address, and if applicable, port information from the binding reserved as a result of the EXTERNAL message. The edge-NAT MUST copy the NATFW_EXTERNAL_BINDING object to response message, if the object is included in the EXTERNAL message. The RESPONSE message is sent back towards the NI+. If it is not an edge-NAT, the NSLP message is forwarded further using the translated IP address as signaling source IP address in the LE-MRM and translated port in the NATFW_DTINFO object in the field 'DR port number', i.e., the NATFW_DTINFO object is updated to reflect the translated port number. The edge-NAT or any other

NAT MUST reject EXTERNAL messages not carrying a NATFW_DTINFO object or if the address information within this object is invalid or is not compliant with local policies (e.g., the information provided relates to a range of addresses ('wildcarded') but the edge-NAT requires exact information about DS's IP address and port) with the above mentioned response codes.

- * Firewall: Non edge-firewalls remember the requested policy rule, keep NATFW NSLP signaling session state, and forward the message. Edge-firewalls stop forwarding the EXTERNAL message. The policy rule is immediately loaded if the action in the NATFW_EFI object is set to 'deny' and the node is an edge-firewall. The policy rule is remembered, but not activated, if the action in the NATFW_EFI object is set to 'allow'. In both cases, a successful RESPONSE message is generated. If the action is 'allow', and the NATFW_DTINFO object is included, and the MRM is set to LE-MRM in the request, additionally a NATFW_EXTERNAL_IP object is included in the RESPONSE message, holding the translated address, and if applicable port, information. This information is obtained from the NATFW_DTINFO object's 'DR port number' and the source-address of the LE-MRM. The edge-firewall MUST copy the NATFW_EXTERNAL_BINDING object to response message, if the object is included in the EXTERNAL message.
- * Combined NAT and firewall: Processing at combined firewall and NAT middleboxes is the same as in the NAT case.
- o NSLP receiver: This type of message should never be received by any NR+, and it MUST generate an error RESPONSE message of class 'Permanent failure' (5) with response code 'No edge-device here' (0x06).

Processing of a RESPONSE message is different for every NSIS node type:

- o NSLP initiator: Upon receiving a successful RESPONSE message, the NI+ can rely on the requested configuration for future inbound NATFW NSLP signaling sessions. If the response contains a NATFW_EXTERNAL_IP object, the NI can use IP address and port pairs carried for further application signaling. After receiving an error RESPONSE message, the NI+ MAY try to generate the EXTERNAL message again or give up and report the failure to the application, depending on the error condition.

- o NSLP forwarder: NFs simply forward this message as long as they keep state for the requested reservation, if the RESPONSE message contains NATFW_INFO object with class set to 'Success' (2). If the RESPONSE message contains NATFW_INFO object with class set not to 'Success' (2), the NATFW NSLP signaling session is marked as 'Dead'.
- o NSIS responder: This type of message should never be received by any NR+. The NF should never receive response messages and MUST silently discard it.

NFs and the NR can also tear down the EXTERNAL session at any time by generating a NOTIFY message with the appropriate response code set.

Reservations with action 'allow' made with EXTERNAL MUST be enabled by a subsequent CREATE message. A reservation made with EXTERNAL (independent of selected action) is kept alive as long as the NI+ refreshes the particular NATFW NSLP signaling session and it can be reused for multiple, different CREATE messages. An NI+ may decide to tear down a reservation immediately after receiving a CREATE message. This implies that a new NATFW NSLP signaling session must be created for each new CREATE message. The CREATE message does not re-use the NATFW NSLP signaling session created by EXTERNAL.

Without using CREATE (see Section 3.7.1) or EXTERNAL in proxy mode (see Section 3.7.6) no data traffic will be forwarded to the DR beyond the edge-NAT or edge-firewall. The only function of EXTERNAL is to ensure that subsequent CREATE messages traveling towards the NR will be forwarded across the public-private boundary towards the DR. Correlation of incoming CREATE messages to EXTERNAL reservation states is described in Section 3.8.

3.7.3. NATFW NSLP Signaling Session Refresh

NATFW NSLP signaling sessions are maintained on a soft-state basis. After a specified timeout, sessions and corresponding policy rules are removed automatically by the middlebox, if they are not refreshed. Soft-state is created by CREATE and EXTERNAL and the maintenance of this state must be done by these messages. State created by CREATE must be maintained by CREATE, state created by EXTERNAL must be maintained by EXTERNAL. Refresh messages, are messages carrying the same session ID as the initial message and a NATFW_LT lifetime object with a lifetime greater than zero. Messages with the same SID but which carry a different MRI are treated as updates of the policy rules and are processed as defined in Section 3.10. Every refresh CREATE or EXTERNAL message MUST be acknowledged by an appropriate response message generated by the NR. Upon reception by each NSLP forwarder, the state for the given

session ID is extended by the NATFW NSLP signaling session refresh period, a period of time calculated based on a proposed refresh message period. The new (extended) lifetime of a NATFW NSLP signaling session is calculated as current local time plus proposed lifetime value (NATFW NSLP signaling session refresh period). Section 3.4 defines the process of calculating lifetimes in detail.

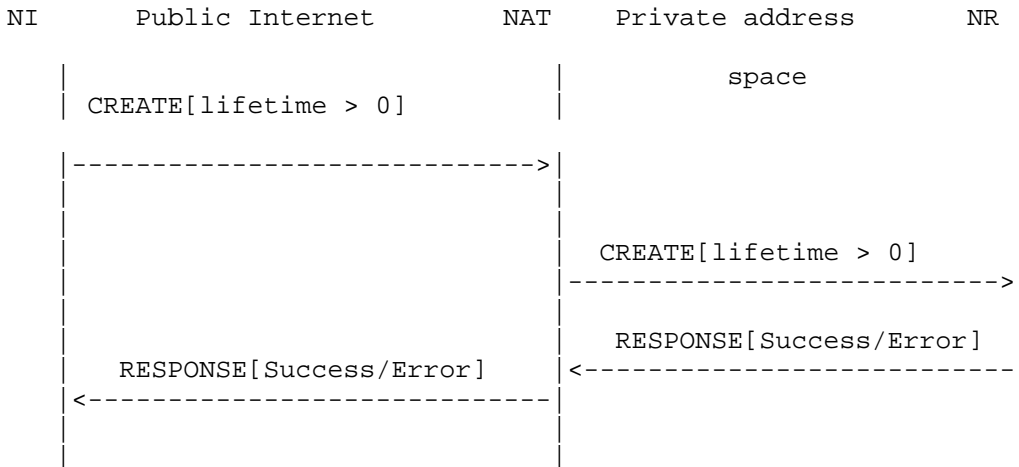


Figure 16: Successful Refresh Message Flow, CREATE as Example

Processing of NATFW NSLP signaling session refresh CREATE and EXTERNAL messages is different for every NSIS node type:

- o NSLP initiator: The NI/NI+ can generate NATFW NSLP signaling session refresh CREATE/EXTERNAL messages before the NATFW NSLP signaling session times out. The rate at which the refresh CREATE/EXTERNAL messages are sent and their relation to the NATFW NSLP signaling session state lifetime is discussed further in Section 3.4.
- o NSLP forwarder: Processing of this message is independent of the middlebox type and is as described in Section 3.4.
- o NSLP responder: NRs accepting a NATFW NSLP signaling session refresh CREATE/EXTERNAL message generate a successful RESPONSE message, including the granted lifetime value of Section 3.4 in a NATFW_LT object.

3.7.4. Deleting Signaling Sessions

NATFW NSLP signaling sessions can be deleted at any time. NSLP initiators can trigger this deletion by using a CREATE or EXTERNAL messages with a lifetime value set to 0, as shown in Figure 17. Whether a CREATE or EXTERNAL message type is use depends on how the NATFW NSLP signaling session was created.

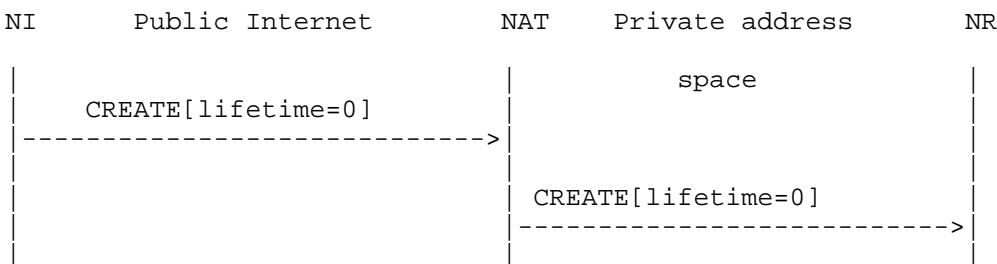


Figure 17: Delete message flow, CREATE as Example

NSLP nodes receiving this message delete the NATFW NSLP signaling session immediately. Policy rules associated with this particular NATFW NSLP signaling session MUST be also deleted immediately. This message is forwarded until it reaches the final NR. The CREATE/EXTERNAL message with a lifetime value of 0, does not generate any response, either positive or negative, since there is no NSIS state left at the nodes along the path.

NSIS initiators can use CREATE/EXTERNAL message with lifetime set to zero in an aggregated way, such that a single CREATE or EXTERNAL message is terminating multiple NATFW NSLP signaling sessions. NIs can follow this procedure if they like to aggregate NATFW NSLP signaling session deletion requests: the NI uses the CREATE or EXTERNAL message with the session ID set to zero and the MRI's source-address set to its used IP address. All other fields of the respective NATFW NSLP signaling sessions to be terminated are set as well; otherwise, these fields are completely wildcarded. The NSLP message is transferred to the NTLF requesting 'explicit routing' as described in Sections 5.2.1 and 7.1.4. in [RFC5971].

The outbound NF receiving such an aggregated CREATE or EXTERNAL message MUST reject it with an error RESPONSE of class 'Permanent failure' (5) with response code 'Authentication failed' (0x01) if the authentication fails and with an error RESPONSE of class 'Permanent failure' (5) with response code 'Authorization failed' (0x02) if the authorization fails. Proof of ownership of NATFW NSLP signaling sessions, as it is defined in this memo (see Section 5.2.1), is not possible when using this aggregation for multiple session

termination. However, the outbound NF can use the relationship between the information of the received CREATE or EXTERNAL message and the GIST messaging association where the request has been received. The outbound NF MUST only accept this aggregated CREATE or EXTERNAL message through already established GIST messaging associations with the NI. The outbound NF MUST NOT propagate this aggregated CREATE or EXTERNAL message but it MAY generate and forward per NATFW NSLP signaling session CREATE or EXTERNAL messages.

3.7.5. Reporting Asynchronous Events

NATFW NSLP forwarders and NATFW NSLP responders must have the ability to report asynchronous events to other NATFW NSLP nodes, especially to allow reporting back to the NATFW NSLP initiator. Such asynchronous events may be premature NATFW NSLP signaling session termination, changes in local policies, route change or any other reason that indicates change of the NATFW NSLP signaling session state.

NFs and NRs may generate NOTIFY messages upon asynchronous events, with a NATFW_INFO object indicating the reason for event. These reasons can be carried in the NATFW_INFO object (class MUST be set to 'Informational' (1)) within the NOTIFY message. This list shows the response codes and the associated actions to take at NFs and the NI:

- o 'Route change: Possible route change on the outbound path' (0x01): Follow instructions in Section 3.9. This MUST be sent inbound and outbound, if the signaling session is any state except 'Transitory'. The NOTIFY message for signaling sessions in state Transitory MUST be discarded, as the signaling session is anyhow Transitory. The outbound NOTIFY message MUST be sent with explicit routing by providing the SII-Handle to the NTLF.
- o 'Re-authentication required' (0x02): The NI should re-send the authentication. This MUST be sent inbound.
- o 'NATFW node is going down soon' (0x03): The NI and other NFs should be prepared for a service interruption at any time. This message MAY be sent inbound and outbound.
- o 'NATFW signaling session lifetime expired' (0x04): The NATFW signaling session has expired and the signaling session is invalid now. NFs MUST mark the signaling session as 'Dead'. This message MAY be sent inbound and outbound.

- o 'NATFW signaling session terminated' (0x05): The NATFW signaling session has been terminated for some reason and the signaling session is invalid now. NFs MUST mark the signaling session as 'Dead'. This message MAY be sent inbound and outbound.

NOTIFY messages are always sent hop-by-hop inbound towards NI until they reach NI or outbound towards the NR as indicated in the list above.

The initial processing when receiving a NOTIFY message is the same for all NATFW nodes: NATFW nodes MUST only accept NOTIFY messages through already established NTLF messaging associations. The further processing is different for each NATFW NSLP node type and depends on the events notified:

- o NSLP initiator: NIs analyze the notified event and behave appropriately based on the event type. NIs MUST NOT generate NOTIFY messages.
- o NSLP forwarder: NFs analyze the notified event and behave based on the above description per response code. NFs SHOULD generate NOTIFY messages upon asynchronous events and forward them inbound towards the NI or outbound towards the NR, depending on the received direction, i.e., inbound messages MUST be forwarded further inbound and outbound messages MUST be forwarded further outbound. NFs MUST silently discard NOTIFY messages that have been received outbound but are only allowed to be sent inbound, e.g., 'Re-authentication required' (0x02).
- o NSLP responder: NRs SHOULD generate NOTIFY messages upon asynchronous events including a response code based on the reported event. The NR MUST silently discard NOTIFY messages that have been received outbound but are only allowed to be sent inbound, e.g., 'Re-authentication required' (0x02).

NATFW NSLP forwarders, keeping multiple NATFW NSLP signaling sessions at the same time, can experience problems when shutting down service suddenly. This sudden shutdown can be as a result of local node failure, for instance, due to a hardware failure. This NF generates NOTIFY messages for each of the NATFW NSLP signaling sessions and tries to send them inbound. Due to the number of NOTIFY messages to be sent, the shutdown of the node may be unnecessarily prolonged, since not all messages can be sent at the same time. This case can be described as a NOTIFY storm, if a multitude of NATFW NSLP signaling sessions is involved.

To avoid the need for generating per NATFW NSLP signaling session NOTIFY messages in such a scenario described or similar cases, NFs SHOULD follow this procedure: the NF uses the NOTIFY message with the session ID in the NTLF set to zero, with the MRI completely wildcarded, using the 'explicit routing' as described in Sections 5.2.1 and 7.1.4 of [RFC5971]. The inbound NF receiving this type of NOTIFY immediately regards all NATFW NSLP signaling sessions from that peer matching the MRI as void. This message will typically result in multiple NOTIFY messages at the inbound NF, i.e., the NF can generate per terminated NATFW NSLP signaling session a NOTIFY message. However, an NF MAY also aggregate the NOTIFY messages as described here.

3.7.6. Proxy Mode of Operation

Some migration scenarios need specialized support to cope with cases where NSIS is only deployed in some areas of the Internet. End-to-end signaling is going to fail without NSIS support at or near both data sender and data receiver terminals. A proxy mode of operation is needed. This proxy mode of operation must terminate the NATFW NSLP signaling topologically-wise as close as possible to the terminal for which it is proxying and proxy all messages. This NATFW NSLP node doing the proxying of the signaling messages becomes either the NI or the NR for the particular NATFW NSLP signaling session, depending on whether it is the DS or DR that does not support NSIS. Typically, the edge-NAT or the edge-firewall would be used to proxy NATFW NSLP messages.

This proxy mode operation does not require any new CREATE or EXTERNAL message type, but relies on extended CREATE and EXTERNAL message types. They are called, respectively, CREATE-PROXY and EXTERNAL-PROXY and are distinguished by setting the P flag in the NSLP header to P=1. This flag instructs edge-NATs and edge-firewalls receiving them to operate in proxy mode for the NATFW NSLP signaling session in question. The semantics of the CREATE and EXTERNAL message types are not changed and the behavior of the various node types is as defined in Sections 3.7.1 and 3.7.2, except for the proxying node. The following paragraphs describe the proxy mode operation for data receivers behind middleboxes and data senders behind middleboxes.

3.7.6.1. Proxying for a Data Sender

The NATFW NSLP gives the NR the ability to install state on the inbound path towards the data sender for outbound data packets, even when only the receiving side is running NSIS (as shown in Figure 18). The goal of the method described is to trigger the edge-NAT/edge-firewall to generate a CREATE message on behalf of the data receiver. In this case, an NR can signal towards the network border

as it is performed in the standard EXTERNAL message handling scenario as in Section 3.7.2. The message is forwarded until the edge-NAT/edge-firewall is reached. A public IP address and port number is reserved at an edge-NAT/edge-firewall. As shown in Figure 18, unlike the standard EXTERNAL message handling case, the edge-NAT/edge-firewall is triggered to send a CREATE message on a new reverse path that traverse several firewalls or NATs. The new reverse path for CREATE is necessary to handle routing asymmetries between the edge-NAT/edge-firewall and the DR. It must be stressed that the semantics of the CREATE and EXTERNAL messages are not changed, i.e., each is processed as described earlier.

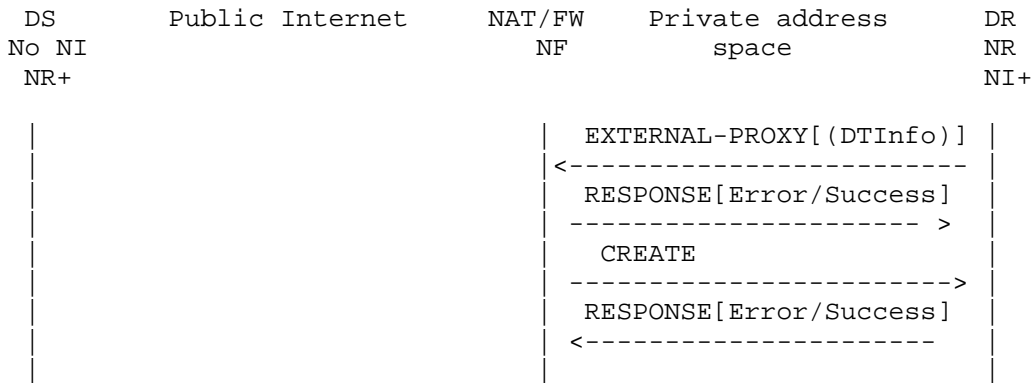


Figure 18: EXTERNAL Triggering Sending of CREATE Message

A NATFW_NONCE object, carried in the EXTERNAL and CREATE message, is used to build the relationship between received CREATEs at the message initiator. An NI+ uses the presence of the NATFW_NONCE object to correlate it to the particular EXTERNAL-PROXY. The absence of a NONCE object indicates a CREATE initiated by the DS and not by the edge-NAT. The two signaling sessions, i.e., the session for EXTERNAL-PROXY and the session for CREATE, are not independent. The primary session is the EXTERNAL-PROXY session. The CREATE session is secondary to the EXTERNAL-PROXY session, i.e., the CREATE session is valid as long as the EXTERNAL-PROXY session is the signaling states 'Established' or 'Transitory'. There is no CREATE session in any other signaling state of the EXTERNAL-PROXY, i.e., 'Pending' or 'Dead'. This ensures fate-sharing between the two signaling sessions.

These processing rules of EXTERNAL-PROXY messages are added to the regular EXTERNAL processing:

- o NSLP initiator (NI+): The NI+ MUST take the session ID (SID) value of the EXTERNAL-PROXY session as the nonce value of the NATFW_NONCE object.
- o NSLP forwarder being either edge-NAT or edge-firewall: When the NF accepts an EXTERNAL-PROXY message, it generates a successful RESPONSE message as if it were the NR, and it generates a CREATE message as defined in Section 3.7.1 and includes a NATFW_NONCE object having the same value as of the received NATFW_NONCE object. The NF MUST NOT generate a CREATE-PROXY message. The NF MUST refresh the CREATE message signaling session only if an EXTERNAL-PROXY refresh message has been received first. This also includes tearing down signaling sessions, i.e., the NF must tear down the CREATE signaling session only if an EXTERNAL-PROXY message with lifetime set to 0 has been received first.

The scenario described in this section challenges the data receiver because it must make a correct assumption about the data sender's ability to use NSIS NATFW NSLP signaling. It is possible for the DR to make the wrong assumption in two different ways:

- a) the DS is NSIS unaware but the DR assumes the DS to be NSIS aware, and
- b) the DS is NSIS aware but the DR assumes the DS to be NSIS unaware.

Case a) will result in middleboxes blocking the data traffic, since the DS will never send the expected CREATE message. Case b) will result in the DR successfully requesting proxy mode support by the edge-NAT or edge-firewall. The edge-NAT/edge-firewall will send CREATE messages and DS will send CREATE messages as well. Both CREATE messages are handled as separated NATFW NSLP signaling sessions and therefore the common rules per NATFW NSLP signaling session apply; the NATFW_NONCE object is used to differentiate CREATE messages generated by the edge-NAT/edge-firewall from the NI-initiated CREATE messages. It is the NR's responsibility to decide whether to tear down the EXTERNAL-PROXY signaling sessions in the case where the data sender's side is NSIS aware, but was incorrectly assumed not to be so by the DR. It is RECOMMENDED that a DR behind NATs use the proxy mode of operation by default, unless the DR knows that the DS is NSIS aware. The DR MAY cache information about data senders that it has found to be NSIS aware in past NATFW NSLP signaling sessions.

There is a possible race condition between the RESPONSE message to the EXTERNAL-PROXY and the CREATE message generated by the edge-NAT. The CREATE message can arrive earlier than the RESPONSE message. An NI+ MUST accept CREATE messages generated by the edge-NAT even if the RESPONSE message to the EXTERNAL-PROXY was not received.

3.7.6.2. Proxying for a Data Receiver

As with data receivers behind middleboxes, data senders behind middleboxes can require proxy mode support. The issue here is that there is no NSIS support at the data receiver's side and, by default, there will be no response to CREATE messages. This scenario requires the last NSIS NATFW NSLP-aware node to terminate the forwarding and to proxy the response to the CREATE message, meaning that this node is generating RESPONSE messages. This last node may be an edge-NAT/edge-firewall, or any other NATFW NSLP peer, that detects that there is no NR available (probably as a result of GIST timeouts but there may be other triggers).

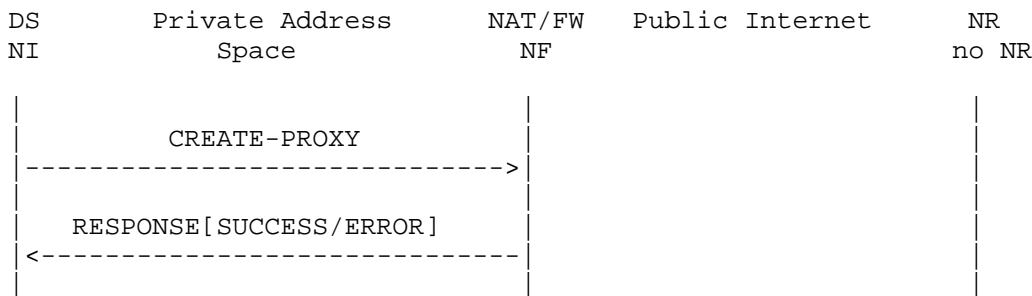


Figure 19: Proxy Mode CREATE Message Flow

The processing of CREATE-PROXY messages and RESPONSE messages is similar to Section 3.7.1, except that forwarding is stopped at the edge-NAT/edge-firewall. The edge-NAT/edge-firewall responds back to NI according to the situation (error/success) and will be the NR for future NATFW NSLP communication.

The NI can choose the proxy mode of operation although the DR is NSIS aware. The CREATE-PROXY mode would not configure all NATs and firewalls along the data path, since it is terminated at the edge-device. Any device beyond this point will never receive any NATFW NSLP signaling for this flow.

3.7.6.3. Incremental Deployment Using the Proxy Mode

The above sections described the proxy mode for cases where the NATFW NSLP is solely deployed at the network edges. However, the NATFW NSLP might be incrementally deployed first in some network edges, but later on also in other parts of the network. Using the proxy mode only would prevent the NI from determining whether the other parts of the network have also been upgraded to use the NATFW NSLP. One way of determining whether the path from the NI to the NR is NATFW-NSLP-capable is to use the regular CREATE message and to wait for a successful response or an error response. This will lead to extra messages being sent, as a CREATE message, in addition to the CREATE-PROXY message (which is required anyhow), is sent from the NI.

The NATFW NSLP allows the usage of the proxy-mode and a further probing of the path by the edge-NAT or edge-firewall. The NI can request proxy-mode handling as described, and can set the E flag (see Figure 20) to request the edge-NAT or edge-firewall to probe the further path for NATFW NSLP enabled NFs or an NR.

The edge-NAT or edge-firewall MUST continue to send the CREATE-PROXY or EXTERNAL-proxy towards the NR, if the received proxy-mode message has the E flag set, in addition to the regular proxy mode handling. The edge-NAT or edge-firewall relies on NTLP measures to determine whether or not there is another NATFW NSLP reachable towards the NR. A failed attempt to forward the request message to the NR will be silently discarded. A successful attempt of forwarding the request message to the NR will be acknowledged by the NR with a successful response message, which is subject to the regular behavior described in the proxy-mode sections.

3.7.6.4. Deployment Considerations for Edge-Devices

The proxy mode assumes that the edge-NAT or edge-firewall are properly configured by network operator, i.e., the edge-device is really the final NAT or firewall of that particular network. There is currently no known way of letting the NATFW NSLP automatically detect which of the NAT or firewalls are the actual edge of a network. Therefore, it is important for the network operator to configure the edge-NAT or edge-firewall and also to re-configure these devices if they are not at the edge anymore. For instance, an edge-NAT is located within an ISP and the ISP chooses to place another NAT in front of this edge-NAT. In this case, the ISP needs to reconfigure the old edge-NAT to be a regular NATFW NSLP NAT and to configure the newly installed NAT to be the edge-NAT.

3.8. Demultiplexing at NATs

Section 3.7.2 describes how NSIS nodes behind NATs can obtain a publicly reachable IP address and port number at a NAT and how the resulting mapping rule can be activated by using CREATE messages (see Section 3.7.1). The information about the public IP address/port number can be transmitted via an application-level signaling protocol and/or third party to the communication partner that would like to send data toward the host behind the NAT. However, NSIS signaling flows are sent towards the address of the NAT at which this particular IP address and port number is allocated and not directly to the allocated IP address and port number. The NATFW NSLP forwarder at this NAT needs to know how the incoming NSLP CREATE messages are related to reserved addresses, meaning how to demultiplex incoming NSIS CREATE messages.

The demultiplexing method uses information stored at the local NATFW NSLP node and in the policy rule. The policy rule uses the LE-MRM MRI source-address (see [RFC5971]) as the flow destination IP address and the network-layer-version (IP-ver) as IP version. The external IP address at the NAT is stored as the external flow destination IP address. All other parameters of the policy rule other than the flow destination IP address are wildcarded if no NATFW_DTINFO object is included in the EXTERNAL message. The LE-MRM MRI destination-address MUST NOT be used in the policy rule, since it is solely a signaling destination address.

If the NATFW_DTINFO object is included in the EXTERNAL message, the policy rule is filled with further information. The 'dst port number' field of the NATFW_DTINFO is stored as the flow destination port number. The 'protocol' field is stored as the flow protocol. The 'src port number' field is stored as the flow source port number. The 'data sender's IPv4 address' is stored as the flow source IP address. Note that some of these fields can contain wildcards.

When receiving a CREATE message at the NATFW NSLP, the NATFW NSLP uses the flow information stored in the MRI to do the matching process. This table shows the parameters to be compared against each other. Note that not all parameters need be present in an MRI at the same time.

Flow parameter (Policy Rule)	MRI parameter (CREATE message)
IP version	network-layer-version
Protocol	IP-protocol
source IP address (w)	source-address (w)
external IP address	destination-address
destination IP address (n/u)	N/A
source port number (w)	L4-source-port (w)
external port number (w)	L4-destination-port (w)
destination port number (n/u)	N/A
IPsec-SPI	ipsec-SPI

Table entries marked with (w) can be wildcarded and entries marked with (n/u) are not used for the matching.

Table 1

It should be noted that the Protocol/IP-protocol entries in Table 1 refer to the 'Protocol' field in the IPv4 header or the 'next header' entry in the IPv6 header.

3.9. Reacting to Route Changes

The NATFW NSLP needs to react to route changes in the data path. This assumes the capability to detect route changes, to perform NAT and firewall configuration on the new path and possibly to tear down NATFW NSLP signaling session state on the old path. The detection of route changes is described in Section 7 of [RFC5971], and the NATFW NSLP relies on notifications about route changes by the NTLP. This notification will be conveyed by the API between NTLP and NSLP, which is out of the scope of this memo.

A NATFW NSLP node other than the NI or NI+ detecting a route change, by means described in the NTLP specification or others, generates a NOTIFY message indicating this change and sends this inbound towards NI and outbound towards the NR (see also Section 3.7.5).

Intermediate NFs on the way to the NI can use this information to decide later if their NATFW NSLP signaling session can be deleted locally, if they do not receive an update within a certain time period, as described in Section 3.2.8. It is important to consider the transport limitations of NOTIFY messages as mandated in Section 3.7.5.

The NI receiving this NOTIFY message MAY generate a new CREATE or EXTERNAL message and send it towards the NATFW NSLP signaling session's NI as for the initial message using the same session ID.

All the remaining processing and message forwarding, such as NSLP next-hop discovery, is subject to regular NSLP processing as described in the particular sections. Normal routing will guide the new CREATE or EXTERNAL message to the correct NFs along the changed route. NFs that were on the original path receiving these new CREATE or EXTERNAL messages (see also Section 3.10), can use the session ID to update the existing NATFW NSLP signaling session; whereas NFs that were not on the original path will create new state for this NATFW NSLP signaling session. The next section describes how policy rules are updated.

3.10. Updating Policy Rules

NSIS initiators can request an update of the installed/reserved policy rules at any time within a NATFW NSLP signaling session. Updates to policy rules can be required due to node mobility (NI is moving from one IP address to another), route changes (this can result in a different NAT mapping at a different NAT device), or the wish of the NI to simply change the rule. NIs can update policy rules in existing NATFW NSLP signaling sessions by sending an appropriate CREATE or EXTERNAL message (similar to Section 3.4) with modified message routing information (MRI) as compared with that installed previously, but using the existing session ID to identify the intended target of the update. With respect to authorization and authentication, this update CREATE or EXTERNAL message is treated in exactly the same way as any initial message. Therefore, any node along in the NATFW NSLP signaling session can reject the update with an error RESPONSE message, as defined in the previous sections.

The message processing and forwarding is executed as defined in the particular sections. An NF or the NR receiving an update simply replaces the installed policy rules installed in the firewall/NAT. The local procedures on how to update the MRI in the firewall/NAT is out of the scope of this memo.

4. NATFW NSLP Message Components

A NATFW NSLP message consists of an NSLP header and one or more objects following the header. The NSLP header is carried in all NATFW NSLP messages and objects are Type-Length-Value (TLV) encoded using big endian (network ordered) binary data representations. Header and objects are aligned to 32-bit boundaries and object lengths that are not multiples of 32 bits must be padded to the next higher 32-bit multiple.

The whole NSLP message is carried as payload of a NTL message.

Note that the notation 0x is used to indicate hexadecimal numbers.

4.1. NSLP Header

All GIST NSLP-Data objects for the NATFW NSLP MUST contain this common header as the first 32 bits of the object (this is not the same as the GIST Common Header). It contains two fields, the NSLP message type and the P Flag, plus two reserved fields. The total length is 32 bits. The layout of the NSLP header is defined by Figure 20.




Figure 20: Common NSLP Header

The reserved field MUST be set to zero in the NATFW NSLP header before sending and MUST be ignored during processing of the header.

The defined messages types are:

- o 0x1: CREATE
- o 0x2: EXTERNAL
- o 0x3: RESPONSE
- o 0x4: NOTIFY

If a message with another type is received, an error RESPONSE of class 'Protocol error' (3) with response code 'Illegal message type' (0x01) MUST be generated.

The P flag indicates the usage of proxy mode. If the proxy mode is used, it MUST be set to 1. Proxy mode MUST only be used in combination with the message types CREATE and EXTERNAL. The P flag MUST be ignored when processing messages with type RESPONSE or NOTIFY.

The E flag indicates, in proxy mode, whether the edge-NAT or edge-firewall MUST continue sending the message to the NR, i.e., end-to-end. The E flag can only be set to 1 if the P flag is set; otherwise, it MUST be ignored. The E flag MUST only be used in combination with the message types CREATE and EXTERNAL. The E flag MUST be ignored when processing messages with type RESPONSE or NOTIFY.

4.2. NSLP Objects

NATFW NSLP objects use a common header format defined by Figure 21. The object header contains these fields: two flags, two reserved bits, the NSLP object type, a reserved field of 4 bits, and the object length. Its total length is 32 bits.




Figure 21: Common NSLP Object Header

The object length field contains the total length of the object without the object header. The unit is a word, consisting of 4 octets. The particular values of type and length for each NSLP object are listed in the subsequent sections that define the NSLP objects. An error RESPONSE of class 'Protocol error' (3) with response code 'Wrong object length' (0x07) MUST be generated if the length given in the object header is inconsistent with the type of object specified or the message is shorter than implied by the object length. The two leading bits of the NSLP object header are used to signal the desired treatment for objects whose treatment has not been defined in this memo (see [RFC5971], Appendix A.2.1), i.e., the Object Type has not been defined. NATFW NSLP uses a subset of the categories defined in GIST:

- o AB=00 ("Mandatory"): If the object is not understood, the entire message containing it MUST be rejected with an error RESPONSE of class 'Protocol error' (3) with response code 'Unknown object present' (0x06).
- o AB=01 ("Optional"): If the object is not understood, it should be deleted and then the rest of the message processed as usual.
- o AB=10 ("Forward"): If the object is not understood, it should be retained unchanged in any message forwarded as a result of message processing, but not stored locally.

The combination AB=11 MUST NOT be used and an error RESPONSE of class 'Protocol error' (3) with response code 'Invalid Flag-Field combination' (0x09) MUST be generated.

The following sections do not repeat the common NSLP object header, they just list the type and the length.

4.2.1. Signaling Session Lifetime Object

The signaling session lifetime object carries the requested or granted lifetime of a NATFW NSLP signaling session measured in seconds.

Type: NATFW_LT (0x00C)

Length: 1




Figure 22: Signaling Session Lifetime Object

4.2.2. External Address Object

The external address object can be included in RESPONSE messages (Section 4.3.3) only. It carries the publicly reachable IP address, and if applicable port number, at an edge-NAT.

Type: NATFW_EXTERNAL_IP (0x00D)

Length: 2




Figure 23: External Address Object for IPv4 Addresses

Please note that the field 'port number' MUST be set to 0 if only an IP address has been reserved, for instance, by a traditional NAT. A port number of 0 MUST be ignored in processing this object.

IP-Ver (4 bits): The IP version number. This field MUST be set to 4.

4.2.3. External Binding Address Object

The external binding address object can be included in RESPONSE messages (Section 4.3.3) and EXTERNAL (Section 3.7.2) messages. It carries one or multiple external binding addresses, and if applicable port number, for multi-level NAT deployments (for an example, see Section 2.5). The utilization of the information carried in this object is described in Appendix B.

Type: NATFW_EXTERNAL_BINDING (0x00E)

Length: 1 + number of IPv4 addresses

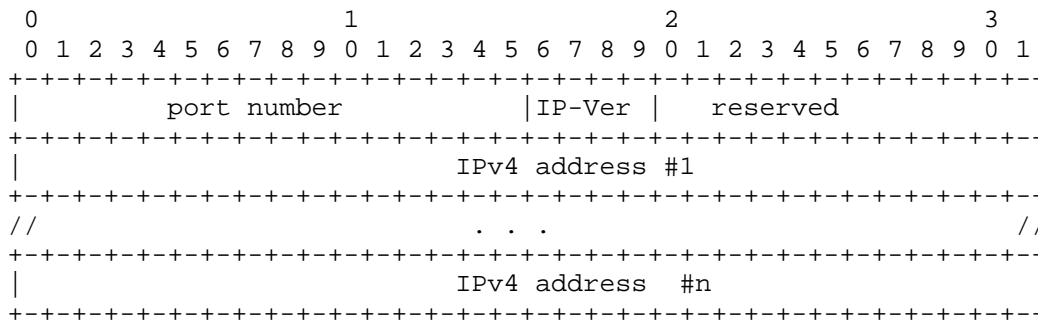


Figure 24: External Binding Address Object

Please note that the field 'port number' MUST be set to 0 if only an IP address has been reserved, for instance, by a traditional NAT. A port number of 0 MUST be ignored in processing this object.

IP-Ver (4 bits): The IP version number. This field MUST be set to 4.

4.2.4. Extended Flow Information Object

In general, flow information is kept in the message routing information (MRI) of the NTLP. Nevertheless, some additional information may be required for NSLP operations. The 'extended flow information' object carries this additional information about the action of the policy rule for firewalls/NATs and a potential contiguous port.

Type: NATFW_EFI (0x00F)

Length: 1




Figure 25: Extended Flow Information

This object has two fields, 'rule action' and 'sub_ports'. The 'rule action' field has these meanings:

- o 0x0001: Allow: A policy rule with this action allows data traffic to traverse the middlebox and the NATFW NSLP MUST allow NSLP signaling to be forwarded.
- o 0x0002: Deny: A policy rule with this action blocks data traffic from traversing the middlebox and the NATFW NSLP MUST NOT allow NSLP signaling to be forwarded.

If the 'rule action' field contains neither 0x0001 nor 0x0002, an error RESPONSE of class 'Signaling session failure' (7) with response code 'Unknown policy rule action' (0x05) MUST be generated.

The 'sub_ports' field contains the number of contiguous transport layer ports to which this rule applies. The default value of this field is 0, i.e., only the port specified in the NTLF's MRM or NATFW_DTINFO object is used for the policy rule. A value of 1 indicates that additionally to the port specified in the NTLF's MRM (port1), a second port (port2) is used. This value of port 2 is calculated as: port2 = port1 + 1. Other values than 0 or 1 MUST NOT be used in this field and an error RESPONSE of class 'Signaling session failure' (7) with response code 'Requested value in sub_ports field in NATFW_EFI not permitted' (0x08) MUST be generated. These two contiguous numbered ports can be used by legacy voice over IP equipment. This legacy equipment assumes two adjacent port numbers for its RTP/RTCP flows, respectively.

4.2.5. Information Code Object

This object carries the response code in RESPONSE messages, which indicates either a successful or failed CREATE or EXTERNAL message depending on the value of the 'response code' field. This object is also carried in a NOTIFY message to specify the reason for the notification.

Type: NATFW_INFO (0x010)

Length: 1




Figure 26: Information Code Object

The field 'resv.' is reserved for future extensions and MUST be set to zero when generating such an object and MUST be ignored when receiving. The 'Object Type' field contains the type of the object causing the error. The value of 'Object Type' is set to 0, if no object is concerned. The leading fours bits marked with 'r' are always set to zero and ignored. The 4-bit class field contains the error class. The following classes are defined:

- o 0: Reserved
- o 1: Informational (NOTIFY only)
- o 2: Success
- o 3: Protocol error
- o 4: Transient failure
- o 5: Permanent failure
- o 7: Signaling session failure

Within each error class a number of responses codes are defined as follows.

- o Informational:
 - * 0x01: Route change: possible route change on the outbound path.
 - * 0x02: Re-authentication required.
 - * 0x03: NATFW node is going down soon.
 - * 0x04: NATFW signaling session lifetime expired.
 - * 0x05: NATFW signaling session terminated.
- o Success:
 - * 0x01: All successfully processed.

- o Protocol error:
 - * 0x01: Illegal message type: the type given in the Message Type field of the NSLP header is unknown.
 - * 0x02: Wrong message length: the length given for the message in the NSLP header does not match the length of the message data.
 - * 0x03: Bad flags value: an undefined flag or combination of flags was set in the NSLP header.
 - * 0x04: Mandatory object missing: an object required in a message of this type was missing.
 - * 0x05: Illegal object present: an object was present that must not be used in a message of this type.
 - * 0x06: Unknown object present: an object of an unknown type was present in the message.
 - * 0x07: Wrong object length: the length given for the object in the object header did not match the length of the object data present.
 - * 0x08: Unknown object field value: a field in an object had an unknown value.
 - * 0x09: Invalid Flag-Field combination: An object contains an invalid combination of flags and/or fields.
 - * 0x0A: Duplicate object present.
 - * 0x0B: Received EXTERNAL request message on external side.
- o Transient failure:
 - * 0x01: Requested resources temporarily not available.
- o Permanent failure:
 - * 0x01: Authentication failed.
 - * 0x02: Authorization failed.
 - * 0x04: Internal or system error.
 - * 0x06: No edge-device here.

- * 0x07: Did not reach the NR.
- o Signaling session failure:
 - * 0x01: Session terminated asynchronously.
 - * 0x02: Requested lifetime is too big.
 - * 0x03: No reservation found matching the MRI of the CREATE request.
 - * 0x04: Requested policy rule denied due to policy conflict.
 - * 0x05: Unknown policy rule action.
 - * 0x06: Requested rule action not applicable.
 - * 0x07: NATFW_DTINFO object is required.
 - * 0x08: Requested value in sub_ports field in NATFW_EFI not permitted.
 - * 0x09: Requested IP protocol not supported.
 - * 0x0A: Plain IP policy rules not permitted -- need transport layer information.
 - * 0x0B: ICMP type value not permitted.
 - * 0x0C: Source IP address range is too large.
 - * 0x0D: Destination IP address range is too large.
 - * 0x0E: Source L4-port range is too large.
 - * 0x0F: Destination L4-port range is too large.
 - * 0x10: Requested lifetime is too small.
 - * 0x11: Modified lifetime is too big.
 - * 0x12: Modified lifetime is too small.

4.2.6. Nonce Object

This object carries the nonce value as described in Section 3.7.6.

Type: NATFW_NONCE (0x011)

Length: 1




Figure 27: Nonce Object

4.2.7. Message Sequence Number Object

This object carries the MSN value as described in Section 3.5.

Type: NATFW_MSN (0x012)

Length: 1




Figure 28: Message Sequence Number Object

4.2.8. Data Terminal Information Object

The 'data terminal information' object carries additional information that MUST be included the EXTERNAL message. EXTERNAL messages are transported by the NTLP using the Loose-End message routing method (LE-MRM). The LE-MRM contains only the DR's IP address and a signaling destination address (destination IP address). This destination IP address is used for message routing only and is not necessarily reflecting the address of the data sender. This object contains information about (if applicable) DR's port number (the destination port number), the DS's port number (the source port number), the used transport protocol, the prefix length of the IP address, and DS's IP address.

Type: NATFW_DTINFO (0x013)

Length: variable. Maximum 3.

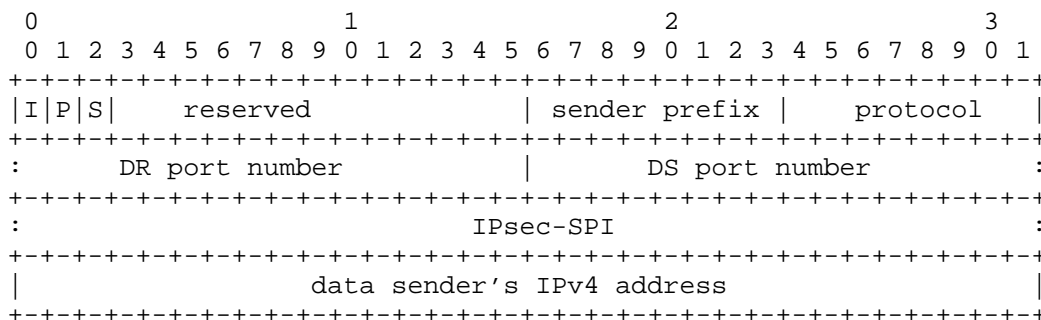


Figure 29: Data Terminal IPv4 Address Object

The flags are:

- o I: I=1 means that 'protocol' should be interpreted.
- o P: P=1 means that 'dst port number' and 'src port number' are present and should be interpreted.
- o S: S=1 means that SPI is present and should be interpreted.

The SPI field is only present if S is set. The port numbers are only present if P is set. The flags P and S MUST NOT be set at the same time. An error RESPONSE of class 'Protocol error' (3) with response code 'Invalid Flag-Field combination' (0x09) MUST be generated if they are both set. If either P or S is set, I MUST be set as well and the protocol field MUST carry the particular protocol. An error RESPONSE of class 'Protocol error' (3) with response code 'Invalid Flag-Field combination' (0x09) MUST be generated if S or P is set but I is not set.

The fields MUST be interpreted according to these rules:

- o (data) sender prefix: This parameter indicates the prefix length of the 'data sender's IP address' in bits. For instance, a full IPv4 address requires 'sender prefix' to be set to 32. A value of 0 indicates an IP address wildcard.
- o protocol: The IP protocol field. This field MUST be interpreted if I=1; otherwise, it MUST be set to 0 and MUST be ignored.

- o DR port number: The port number at the data receiver (DR), i.e., the destination port. A value of 0 indicates a port wildcard, i.e., the destination port number is not known. Any other value indicates the destination port number.
- o DS port number: The port number at the data sender (DS), i.e., the source port. A value of 0 indicates a port wildcard, i.e., the source port number is not known. Any other value indicates the source port number.
- o data sender's IPv4 address: The source IP address of the data sender. This field MUST be set to zero if no IP address is provided, i.e., a complete wildcard is desired (see the dest prefix field above).

4.2.9. ICMP Types Object

The 'ICMP types' object contains additional information needed to configure a NAT of firewall with rules to control ICMP traffic. The object contains a number of values of the ICMP Type field for which a filter action should be set up:

Type: NATFW_ICMP_TYPES (0x014)

Length: Variable = ((Number of Types carried + 1) + 3) DIV 4

Where DIV is an integer division.




Figure 30: ICMP Types Object

The fields MUST be interpreted according to these rules:

count: 8-bit integer specifying the number of 'Type' entries in the object.

type: 8-bit field specifying an ICMP Type value to which this rule applies.

padding: Sufficient 0 bits to pad out the last word so that the total size of the object is an even multiple of words. Ignored on reception.

4.3. Message Formats

This section defines the content of each NATFW NSLP message type. The message types are defined in Section 4.1.

Basically, each message is constructed of an NSLP header and one or more NSLP objects. The order of objects is not defined, meaning that objects may occur in any sequence. Objects are marked either with mandatory (M) or optional (O). Where (M) implies that this particular object MUST be included within the message and where (O) implies that this particular object is OPTIONAL within the message. Objects defined in this memo always carry the flag combination AB=00 in the NSLP object header. An error RESPONSE message of class 'Protocol error' (3) with response code 'Mandatory object missing' (0x04) MUST be generated if a mandatory declared object is missing. An error RESPONSE message of class 'Protocol error' (3) with response code 'Illegal object present' (0x05) MUST be generated if an object was present that must not be used in a message of this type. An error RESPONSE message of class 'Protocol error' (3) with response code 'Duplicate object present' (0x0A) MUST be generated if an object appears more than once in a message.

Each section elaborates the required settings and parameters to be set by the NSLP for the NTLF, for instance, how the message routing information is set.

4.3.1. CREATE

The CREATE message is used to create NATFW NSLP signaling sessions and to create policy rules. Furthermore, CREATE messages are used to refresh NATFW NSLP signaling sessions and to delete them.

The CREATE message carries these objects:

- o Signaling Session Lifetime object (M)
- o Extended flow information object (M)
- o Message sequence number object (M)
- o Nonce object (M) if P flag set to 1 in the NSLP header, otherwise (O)
- o ICMP Types Object (O)

The message routing information in the NTLF MUST be set to DS as source IP address and DR as destination IP address. All other parameters MUST be set according to the required policy rule. CREATE messages MUST be transported by using the path-coupled MRM with the direction set to 'downstream' (outbound).

4.3.2. EXTERNAL

The EXTERNAL message is used to a) reserve an external IP address/port at NATs, b) to notify firewalls about NSIS capable DRs, or c) to block incoming data traffic at inbound firewalls.

The EXTERNAL message carries these objects:

- o Signaling Session Lifetime object (M)
- o Message sequence number object (M)
- o Extended flow information object (M)
- o Data terminal information object (M)
- o Nonce object (M) if P flag set to 1 in the NSLP header, otherwise (O)
- o ICMP Types Object (O)
- o External binding address object (O)

The selected message routing method of the EXTERNAL message depends on a number of considerations. Section 3.7.2 describes exhaustively how to select the correct method. EXTERNAL messages can be transported via the path-coupled message routing method (PC-MRM) or via the loose-end message routing method (LE-MRM). In the case of PC-MRM, the source-address is set to the DS's address and the destination-address is set to the DR's address, the direction is set to inbound. In the case of LE-MRM, the destination-address is set to the DR's address or to the signaling destination IP address. The source-address is set to the DS's address.

4.3.3. RESPONSE

RESPONSE messages are responses to CREATE and EXTERNAL messages. RESPONSE messages MUST NOT be generated for any other message, such as NOTIFY and RESPONSE.

The RESPONSE message for the class 'Success' (2) carries these objects:

- o Signaling Session Lifetime object (M)
- o Message sequence number object (M)
- o Information code object (M)
- o External address object (O)
- o External binding address object (O)

The RESPONSE message for other classes than 'Success' (2) carries these objects:

- o Message sequence number object (M)
- o Information code object (M)
- o Signaling Session Lifetime object (O)

This message is routed towards the NI hop-by-hop, using existing NTLF messaging associations. The MRM used for this message MUST be the same as MRM used by the corresponding CREATE or EXTERNAL message.

4.3.4. NOTIFY

The NOTIFY messages is used to report asynchronous events happening along the signaled path to other NATFW NSLP nodes.

The NOTIFY message carries this object:

- o Information code object (M)

The NOTIFY message is routed towards the next NF, NI, or NR hop-by-hop using the existing inbound or outbound node messaging association entry within the node's Message Routing State table. The MRM used for this message MUST be the same as MRM used by the corresponding CREATE or EXTERNAL message.

5. Security Considerations

Security is of major concern particularly in the case of firewall traversal. This section provides security considerations for the NAT/firewall traversal and is organized as follows.

In Section 5.1, we describe how the participating entities relate to each other from a security point of view. That subsection also motivates a particular authorization model.

Security threats that focus on NSIS in general are described in [RFC4081] and they are applicable to this document as well.

Finally, we illustrate how the security requirements that were created based on the security threats can be fulfilled by specific security mechanisms. These aspects will be elaborated in Section 5.2.

5.1. Authorization Framework

The NATFW NSLP is a protocol that may involve a number of NSIS nodes and is, as such, not a two-party protocol. Figures 1 and 2 of [RFC4081] already depict the possible set of communication patterns. In this section, we will re-evaluate these communication patterns with respect to the NATFW NSLP protocol interaction.

The security solutions for providing authorization have a direct impact on the treatment of different NSLPs. As it can be seen from the QoS NSLP [RFC5974] and the corresponding Diameter QoS work [RFC5866], accounting and charging seems to play an important role for QoS reservations, whereas monetary aspects might only indirectly effect authorization decisions for NAT and firewall signaling. Hence, there are differences in the semantics of authorization handling between QoS and NATFW signaling. A NATFW-aware node will most likely want to authorize the entity (e.g., user or machine) requesting the establishment of pinholes or NAT bindings. The outcome of the authorization decision is either allowed or disallowed, whereas a QoS authorization decision might indicate that a different set of QoS parameters is authorized (see [RFC5866] as an example).

5.1.1. Peer-to-Peer Relationship

Starting with the simplest scenario, it is assumed that neighboring nodes are able to authenticate each other and to establish keying material to protect the signaling message communication. The nodes will have to authorize each other, additionally to the authentication. We use the term 'Security Context' as a placeholder for referring to the entire security procedure, the necessary infrastructure that needs to be in place in order for this to work (e.g., key management) and the established security-related state. The required long-term keys (symmetric or asymmetric keys) used for authentication either are made available using an out-of-band mechanism between the two NSIS NATFW nodes or are dynamically established using mechanisms not further specified in this document. Note that the deployment environment will most likely have an impact on the choice of credentials being used. The choice of these credentials used is also outside the scope of this document.

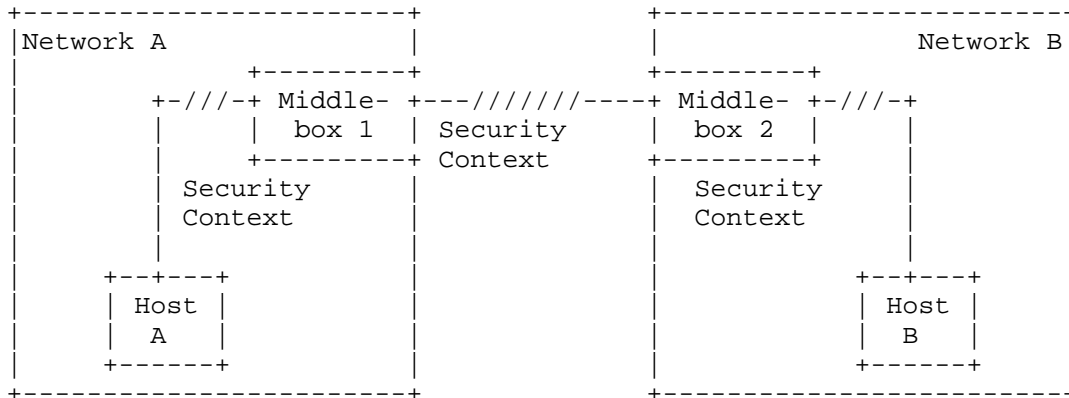


Figure 31: Peer-to-Peer Relationship

Figure 31 shows a possible relationship between participating NSIS-aware nodes. Host A might be, for example, a host in an enterprise network that has keying material established (e.g., a shared secret) with the company's firewall (Middlebox 1). The network administrator of Network A (company network) has created access control lists for Host A (or whatever identifiers a particular company wants to use). Exactly the same procedure might also be used between Host B and Middlebox 2 in Network B. For the communication between Middlebox 1 and Middlebox 2 a security context is also assumed in order to allow authentication, authorization, and signaling message protection to be successful.

5.1.2. Intra-Domain Relationship

In larger corporations, for example, a middlebox is used to protect individual departments. In many cases, the entire enterprise is controlled by a single (or a small number of) security department(s), which give instructions to the department administrators. In such a scenario, the previously discussed peer-to-peer relationship might be prevalent. Sometimes it might be necessary to preserve authentication and authorization information within the network. As a possible solution, a centralized approach could be used, whereby an interaction between the individual middleboxes and a central entity (for example, a policy decision point - PDP) takes place. As an alternative, individual middleboxes exchange the authorization decision with another middlebox within the same trust domain. Individual middleboxes within an administrative domain may exploit their relationship instead of requesting authentication and authorization of the signaling initiator again and again. Figure 32 illustrates a network structure that uses a centralized entity.

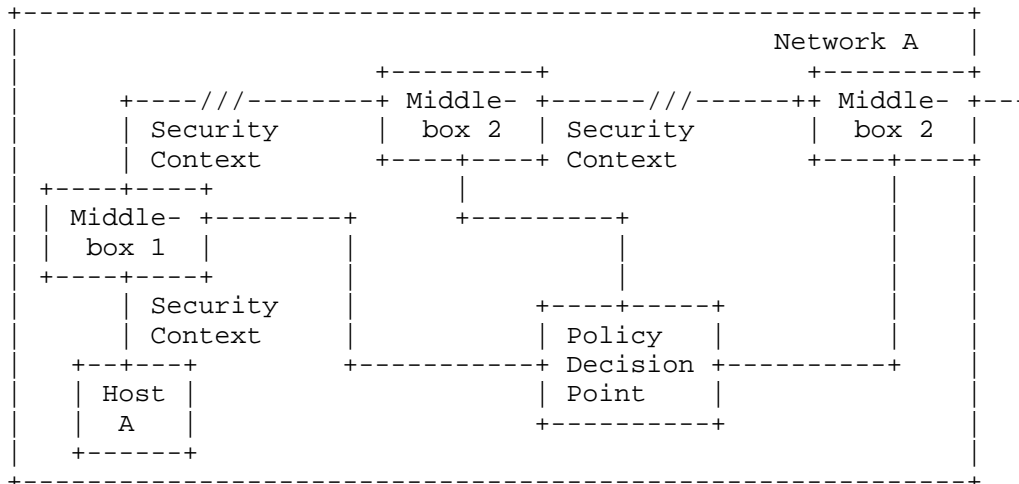


Figure 32: Intra-Domain Relationship

The interaction between individual middleboxes and a policy decision point (or AAA server) is outside the scope of this document.

5.1.3. End-to-Middle Relationship

The peer-to-peer relationship between neighboring NSIS NATFW NSLP nodes might not always be sufficient. Network B might require additional authorization of the signaling message initiator (in addition to the authorization of the neighboring node). If authentication and authorization information is not attached to the initial signaling message then the signaling message arriving at Middlebox 2 would result in an error message being created, which indicates the additional authorization requirement. In many cases, the signaling message initiator might already be aware of the additionally required authorization before the signaling message exchange is executed.

Figure 33 shows this scenario.

by using a mutually authenticated Transport Layer secured connection and by relying on authorization by the neighboring NATFW NSLP entities.

The NATFW NSLP relies on an established security association between neighboring peers to prevent unauthorized nodes from modifying or deleting installed state. Between non-neighboring nodes the session ID (SID) carried in the NTLP is used to show ownership of a NATFW NSLP signaling session. The session ID MUST be generated in a random way and thereby prevents an off-path adversary from mounting targeted attacks. Hence, an adversary would have to learn the randomly generated session ID to perform an attack. In a mobility environment a former on-path node that is now off-path can perform an attack. Messages for a particular NATFW NSLP signaling session are handled by the NTLP to the NATFW NSLP for further processing. Messages carrying a different session ID not associated with any NATFW NSLP are subject to the regular processing for new NATFW NSLP signaling sessions.

5.2.2. Security Protection between Non-Neighboring NATFW NSLP Nodes

Based on the security threats and the listed requirements, it was noted that some threats also demand authentication and authorization of a NATFW signaling entity (including the initiator) towards a non-neighboring node. This mechanism mainly demands entity authentication. The most important information exchanged at the NATFW NSLP is information related to the establishment for firewall pinholes and NAT bindings. This information can, however, not be protected over multiple NSIS NATFW NSLP hops since this information might change depending on the capability of each individual NATFW NSLP node.

Some scenarios might also benefit from the usage of authorization tokens. Their purpose is to associate two different signaling protocols (e.g., SIP and NSIS) and their authorization decision. These tokens are obtained by non-NSIS protocols, such as SIP or as part of network access authentication. When a NAT or firewall along the path receives the token it might be verified locally or passed to the AAA infrastructure. Examples of authorization tokens can be found in RFC 3520 [RFC3520] and RFC 3521 [RFC3521]. Figure 34 shows an example of this protocol interaction.

An authorization token is provided by the SIP proxy, which acts as the assertion generating entity and gets delivered to the end host with proper authentication and authorization. When the NATFW signaling message is transmitted towards the network, the authorization token is attached to the signaling messages to refer to the previous authorization decision. The assertion-verifying entity needs to process the token or it might be necessary to interact with

the assertion-granting entity using HTTP (or other protocols). As a result of a successfully authorization by a NATFW NSLP node, the requested action is executed and later a RESPONSE message is generated.

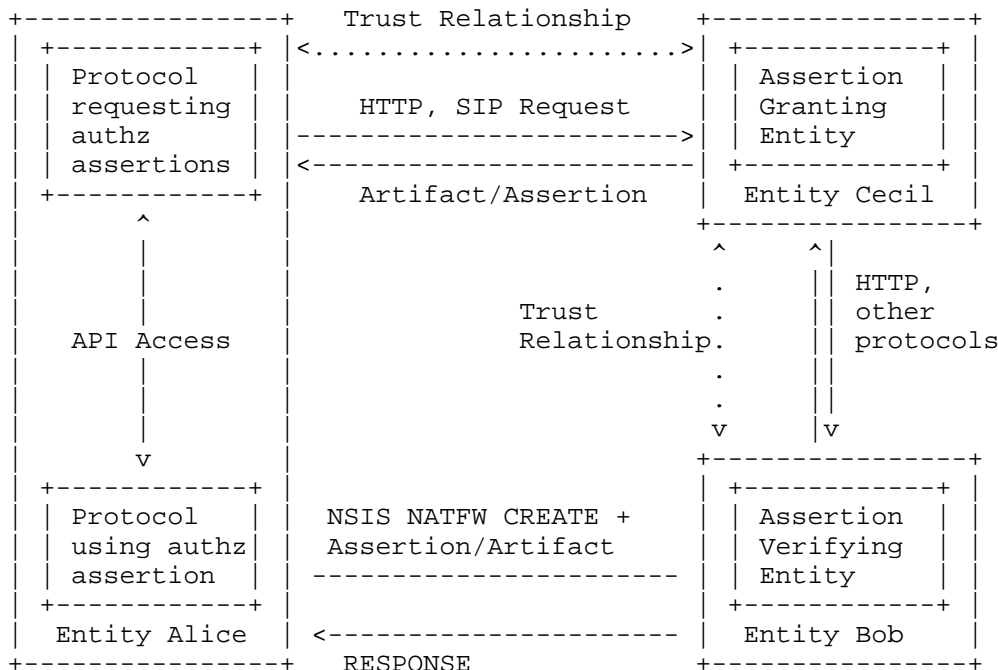


Figure 34: Authorization Token Usage

Threats against the usage of authorization tokens have been mentioned in [RFC4081]. Hence, it is required to provide confidentiality protection to avoid allowing an eavesdropper to learn the token and to use it in another NATFW NSLP signaling session (replay attack). The token itself also needs to be protected against tempering.

5.3. Implementation of NATFW NSLP Security

The prior sections describe how to secure the NATFW NSLP in the presence of established trust between the various players and the particular relationships (e.g., intra-domain, end-to-middle, or peer-to-peer). However, in typical Internet deployments there is no established trust, other than granting access to a network, but not between various sites in the Internet. Furthermore, the NATFW NSLP may be incrementally deployed with a widely varying ability to be able to use authentication and authorization services.

The NATFW NSLP offers a way to keep the authentication and authorization at the "edge" of the network. The local edge network can deploy and use any type of Authentication and Authorization (AA) scheme without the need to have AA technology match with other edges in the Internet (assuming that firewalls and NATs are deployed at the edges of the network and not somewhere in the cores).

Each network edge that has the NATFW NSLP deployed can use the EXTERNAL request message to allow a secure access to the network. Using the EXTERNAL request message does allow the DR to open the firewall/NAT on the receiver's side. However, the edge-devices should not allow the firewall/NAT to be opened up completely (i.e., should not apply an allow-all policy), but should let DRs reserve very specific policies. For instance, a DR can request reservation of an 'allow' policy rule for an incoming TCP connection for a Jabber file transfer. This reserved policy (see Figure 15) rule must be activated by matching the CREATE request message (see Figure 15). This mechanism allows for the authentication and authorization issues to be managed locally at the particular edge-network. In the reverse direction, the CREATE request message can be handled independently on the DS side with respect to authentication and authorization.

The usage described in the above paragraph is further simplified for an incremental deployment: there is no requirement to activate a reserved policy rule with a CREATE request message. This is completely handled by the EXTERNAL-PROXY request message and the associated CREATE request message. Both of them are handled by the local authentication and authorization scheme.

6. IAB Considerations on UNSAF

Unilateral Self-Address Fixing (UNSAF) is described in [RFC3424] as a process at originating endpoints that attempts to determine or fix the address (and port) by which they are known to another endpoint. UNSAF proposals, such as STUN [RFC5389] are considered as a general class of workarounds for NAT traversal and as solutions for scenarios with no middlebox communication.

This memo specifies a path-coupled middlebox communication protocol, i.e., the NSIS NATFW NSLP. NSIS in general and the NATFW NSLP are not intended as a short-term workaround, but more as a long-term solution for middlebox communication. In NSIS, endpoints are involved in allocating, maintaining, and deleting addresses and ports at the middlebox. However, the full control of addresses and ports at the middlebox is at the NATFW NSLP daemon located at the respective NAT.

Therefore, this document addresses the UNSAF considerations in [RFC3424] by proposing a long-term alternative solution.

7. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to the NATFW NSLP, in accordance with BCP 26, RFC 5226 [RFC5226].

The NATFW NSLP requires IANA to create a number of new registries:

- o NATFW NSLP Message Types
- o NATFW NSLP Header Flags
- o NSLP Response Codes

It also requires registration of new values in a number of registries:

- o NSLP Message Objects
- o NSLP Identifiers (under GIST Parameters)
- o Router Alert Option Values (IPv4 and IPv6)

7.1. NATFW NSLP Message Type Registry

The NATFW NSLP Message Type is an 8-bit value. The allocation of values for new message types requires IETF Review. Updates and deletion of values from the registry are not possible. This specification defines four NATFW NSLP message types, which form the initial contents of this registry. IANA has added these four NATFW NSLP Message Types: CREATE (0x1), EXTERNAL (0x2), RESPONSE (0x3), and NOTIFY (0x4). 0x0 is Reserved. Each registry entry consists of value, description, and reference.

7.2. NATFW NSLP Header Flag Registry

NATFW NSLP messages have a message-specific 8-bit flags/reserved field in their header. The registration of flags is subject to IANA registration. The allocation of values for flag types requires IETF Review. Updates and deletion of values from the registry are not possible. This specification defines only two flags in Section 4.1, the P flag (bit 8) and the E flag (bit 9). Each registry entry consists of value, bit position, description (containing the section number), and reference.

7.3. NSLP Message Object Registry

In Section 4.2 this document defines 9 objects for the NATFW NSLP: NATFW_LT, NATFW_EXTERNAL_IP, NATFW_EXTERNAL_BINDING, NATFW_EFI, NATFW_INFO, NATFW_NONCE, NATFW_MSN, NATFW_DTINFO, NATFW_ICMP_TYPES. IANA has assigned values for them from the NSLP Message Objects registry.

7.4. NSLP Response Code Registry

In addition, this document defines a number of Response Codes for the NATFW NSLP. These can be found in Section 4.2.5 and have been assigned values from the NSLP Response Code registry. The allocation of new values for Response Codes requires IETF Review. IANA has assigned values for them as given in Section 4.2.5 for the error class and also for the number of responses values per error class. Each registry entry consists of response code, value, description, and reference.

7.5. NSLP IDs and Router Alert Option Values

GIST NSLPID

This specification defines an NSLP for use with GIST and thus requires an assigned NSLP identifier. IANA has added one new value (33) to the NSLP Identifiers (NSLPID) registry defined in [RFC5971] for the NATFW NSLP.

IPv4 and IPv6 Router Alert Option (RAO) value

The GIST specification also requires that each NSLP-ID be associated with specific Router Alert Option (RAO) value. For the purposes of the NATFW NSLP, a single IPv4 RAO value (65) and a single IPv6 RAO value (68) have been allocated.

8. Acknowledgments

We would like to thank the following individuals for their contributions to this document at different stages:

- o Marcus Brunner and Henning Schulzrinne for their work on IETF documents that led us to start with this document;
- o Miquel Martin for his large contribution on the initial version of this document and one of the first prototype implementations;
- o Srinath Thiruvengadam and Ali Fessi work for their work on the NAT/firewall threats document;

- o Henning Peters for his comments and suggestions;
- o Ben Campbell as Gen-ART reviewer;
- o and the NSIS working group.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5971] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.
- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, August 1996.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.

9.2. Informative References

- [RFC4080] Hancock, R., Karagiannis, G., Loughney, J., and S. Van den Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.
- [RFC3726] Brunner, M., "Requirements for Signaling Protocols", RFC 3726, April 2004.
- [RFC5974] Manner, J., Karagiannis, G., and A. McDonald, "NSIS Signaling Layer Protocol (NSLP) for Quality-of-Service Signaling", RFC 5974, October 2010.
- [RFC5866] Sun, D., McCann, P., Tschofenig, H., Tsou, T., Doria, A., and G. Zorn, "Diameter Quality-of-Service Application", RFC 5866, May 2010.
- [RFC5978] Manner, J., Bless, R., Loughney, J., and E. Davies, "Using and Extending the NSIS Protocol Family", RFC 5978, October 2010.
- [RFC3303] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A., and A. Rayhan, "Middlebox communication architecture and framework", RFC 3303, August 2002.

- [RFC4081] Tschofenig, H. and D. Kroeselberg, "Security Threats for Next Steps in Signaling (NSIS)", RFC 4081, June 2005.
- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC3198] Westerinen, A., Schnizlein, J., Strassner, J., Scherling, M., Quinn, B., Herzog, S., Huynh, A., Carlson, M., Perry, J., and S. Waldbusser, "Terminology for Policy-Based Management", RFC 3198, November 2001.
- [RFC3520] Hamer, L-N., Gage, B., Kosinski, B., and H. Shieh, "Session Authorization Policy Element", RFC 3520, April 2003.
- [RFC3521] Hamer, L-N., Gage, B., and H. Shieh, "Framework for Session Set-up with Media Authorization", RFC 3521, April 2003.
- [rsvp-firewall] Roedig, U., Goertz, M., Karten, M., and R. Steinmetz, "RSVP as firewall Signalling Protocol", Proceedings of the 6th IEEE Symposium on Computers and Communications, Hammamet, Tunisia, pp. 57 to 62, IEEE Computer Society Press, July 2001.

Appendix A. Selecting Signaling Destination Addresses for EXTERNAL

As with all other message types, EXTERNAL messages need a reachable IP address of the data sender on the GIST level. For the path-coupled MRM, the source-address of GIST is the reachable IP address (i.e., the real IP address of the data sender, or a wildcard). While this is straightforward, it is not necessarily so for the loose-end MRM. Many applications do not provide the IP address of the communication counterpart, i.e., either the data sender or both a data sender and receiver. For the EXTERNAL messages, the case of data sender is of interest only. The rest of this section gives informational guidance about determining a good destination-address of the LE-MRM in GIST for EXTERNAL messages.

This signaling destination address (SDA, the destination-address in GIST) can be the data sender, but for applications that do not provide an address upfront, the destination IP address has to be chosen independently, as it is unknown at the time when the NATFW NSLP signaling has to start. Choosing the 'correct' destination IP address may be difficult and it is possible that there is no 'right answer' for all applications relying on the NATFW NSLP.

Whenever possible, it is RECOMMENDED to chose the data sender's IP address as the SDA. It is necessary to differentiate between the received IP addresses on the data sender. Some application-level signaling protocols (e.g., SIP) have the ability to transfer multiple contact IP addresses of the data sender. For instance, private IP addresses, public IP addresses at a NAT, and public IP addresses at a relay. It is RECOMMENDED to use all non-private IP addresses as SDAs.

A different SDA must be chosen, if the IP address of the data sender is unknown. This can have multiple reasons: the application-level signaling protocol cannot determine any data sender IP address at this point in time or the data receiver is server behind a NAT, i.e., accepting inbound packets from any host. In this case, the NATFW NSLP can be instructed to use the public IP address of an application server or any other node. Choosing the SDA in this case is out of the scope of the NATFW NSLP and depends on the application's choice. The local network can provide a network-SDA, i.e., an SDA that is only meaningful to the local network. This will ensure that GIST packets with destination-address set to this network-SDA are going to be routed to an edge-NAT or edge-firewall.

Appendix B. Usage of External Binding Addresses

The NATFW_EXTERNAL_BINDING object carries information, which has a different utility to the information carried within the NATFW_EXTERNAL_IP object. The NATFW_EXTERNAL_IP object has the public IP address and potentially port numbers that can be used by the application at the NI to be reachable via the public Internet. However, there are cases in which various NIs are located behind the same public NAT, but are subject to a multi-level NAT deployment, as shown in Figure 35. They can use their public IP address port assigned to them to communicate between each other (e.g., NI with NR1 and NR2) but they are forced to send their traffic through the edge-NAT, even though there is a shorter way possible.




Figure 35: Multi-Level NAT Scenario

Figure 35 shows an example that is explored here:

1. NI -> NR1: Both NI and NR1 send EXTERNAL messages towards NAT2 and get an external address+port binding. Then, they exchange that external binding and all traffic gets pinned to NAT2 instead of taking the shortest path by NAT1 to NAT3 directly. However, to do that, NR1 and NI both need to be aware that they also have the address on the external side of NAT1 and NAT3, respectively. If ICE is deployed and there is actually a STUN server in the 10/8 network configured, it is possible to get the shorter path to work. The NATFW NSLP provides all external addresses in the NATFW_EXTERNAL_BINDING towards the public network it could allow for optimizations.
2. For the case NI -> NR2 is even more obvious. Pinning this to NAT2 is an important fallback, but allowing for trying for a direct path between NAT1 and NAT3 might be worth it.

Please note that if there are overlapping address domains between NR and the public Internet, the regular routing will not necessary allow sending the packet to the right domain.

Appendix C. Applicability Statement on Data Receivers behind Firewalls

Section 3.7.2 describes how data receivers behind middleboxes can instruct inbound firewalls/NATs to forward NATFW NSLP signaling towards them. Finding an inbound edge-NAT in an address environment with NAT'ed addresses is quite easy. It is only required to find some edge-NAT, as the data traffic will be route-pinned to the NAT. Locating the appropriate edge-firewall with the PC-MRM sent inbound is difficult. For cases with a single, symmetric route from the Internet to the data receiver, it is quite easy; simply follow the default route in the inbound direction.

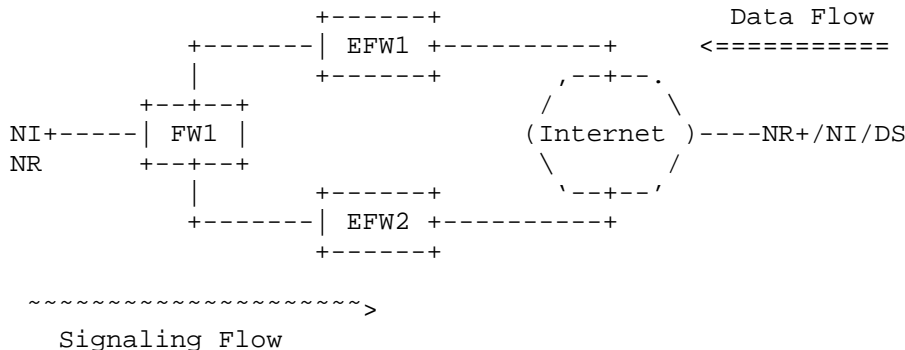


Figure 36: Data Receiver behind Multiple Firewalls Located in Parallel

When a data receiver, and thus NR, is located in a network site that is multihomed with several independently firewalled connections to the public Internet (as shown in Figure 36), the specific firewall through which the data traffic will be routed has to be ascertained. NATFW NSLP signaling messages sent from the NI+/NR during the EXTERNAL message exchange towards the NR+ must be routed by the NTLP to the edge-firewall that will be passed by the data traffic as well. The NTLP would need to be aware about the routing within the Internet to determine the path between the DS and DR. Out of this, the NTLP could determine which of the edge-firewalls, either EFW1 or EFW2, must be selected to forward the NATFW NSLP signaling. Signaling to the wrong edge-firewall, as shown in Figure 36, would install the NATFW NSLP policy rules at the wrong device. This causes either a blocked data flow (when the policy rule is 'allow') or an ongoing attack (when the policy rule is 'deny'). Requiring the NTLP to know all about the routing within the Internet is definitely a tough challenge and usually not possible. In a case as described, the NTLP must basically give up and return an error to the NSLP level, indicating that the next hop discovery is not possible.

Appendix D. Firewall and NAT Resources

This section gives some examples on how NATFW NSLP policy rules could be mapped to real firewall or NAT resources. The firewall rules and NAT bindings are described in a natural way, i.e., in a way that one will find in common implementations.

D.1. Wildcarding of Policy Rules

The policy rule/MRI to be installed can be wildcarded to some degree. Wildcarding applies to IP address, transport layer port numbers, and the IP payload (or next header in IPv6). Processing of wildcarding splits into the NTLP and the NATFW NSLP layer. The processing at the NTLP layer is independent of the NSLP layer processing and per-layer constraints apply. For wildcarding in the NTLP, see Section 5.8 of [RFC5971].

Wildcarding at the NATFW NSLP level is always a node local policy decision. A signaling message carrying a wildcarded MRI (and thus policy rule) arriving at an NSLP node can be rejected if the local policy does not allow the request. For instance, take an MRI with IP addresses set (not wildcarded), transport protocol TCP, and TCP port numbers completely wildcarded. If the local policy allows only requests for TCP with all ports set and not wildcarded, the request is going to be rejected.

D.2. Mapping to Firewall Rules

This section describes how a NSLP policy rule signaled with a CREATE message is mapped to a firewall rule. The MRI is set as follows:

- o network-layer-version=IPv4
- o source-address=192.0.2.100, prefix-length=32
- o destination-address=192.0.50.5, prefix-length=32
- o IP-protocol=UDP
- o L4-source-port=34543, L4-destination-port=23198

The NATFW_EFI object is set to action=allow and sub_ports=0.

The resulting policy rule (firewall rule) to be installed might look like: allow udp from 192.0.2.100 port=34543 to 192.0.50.5 port=23198.

D.3. Mapping to NAT Bindings

This section describes how a NSLP policy rule signaled with an EXTERNAL message is mapped to a NAT binding. It is assumed that the EXTERNAL message is sent by a NI+ located behind a NAT and does contain a NATFW_DTINFO object. The MRI is set following using the signaling destination address, since the IP address of the real data sender is not known:

- o network-layer-version=IPv4
- o source-address= 192.168.5.100
- o destination-address=SDA
- o IP-protocol=UDP

The NATFW_EFI object is set to action=allow and sub_ports=0. The NATFW_DTINFO object contains these parameters:

- o P=1
- o dest prefix=0
- o protocol=UDP
- o dst port number = 20230, src port number=0
- o src IP=0.0.0.0

The edge-NAT allocates the external IP 192.0.2.79 and port 45000.

The resulting policy rule (NAT binding) to be installed could look like: translate udp from any to 192.0.2.79 port=45000 to 192.168.5.100 port=20230.

D.4. NSLP Handling of Twice-NAT

The dynamic configuration of twice-NATs requires application-level support, as stated in Section 2.5. The NATFW NSLP cannot be used for configuring twice-NATs if application-level support is needed. Assuming application-level support performing the configuration of the twice-NAT and the NATFW NSLP being installed at this devices, the NATFW NSLP must be able to traverse it. The NSLP is probably able to traverse the twice-NAT, as is any other data traffic, but the flow information stored in the NTLF's MRI will be invalidated through the translation of source and destination IP addresses. The NATFW NSLP implementation on the twice-NAT MUST intercept NATFW NSLP and NTLF

signaling messages as any other NATFW NSLP node does. For the given signaling flow, the NATFW NSLP node MUST look up the corresponding IP address translation and modify the NTLN/NSLP signaling accordingly. The modification results in an updated MRI with respect to the source and destination IP addresses.

Appendix E. Example for Receiver Proxy Case

This section gives an example on how to use the NATFW NSLP for a receiver behind a NAT, where only the receiving side is NATFW NSLP enabled. We assume FTP as the application to show a working example. An FTP server is located behind a NAT, as shown in Figure 5, and uses the NATFW NSLP to allocate NAT bindings for the control and data channel of the FTP protocol. The information about where to reach the server is communicated by a separate protocol (e.g., email, chat) to the DS side.

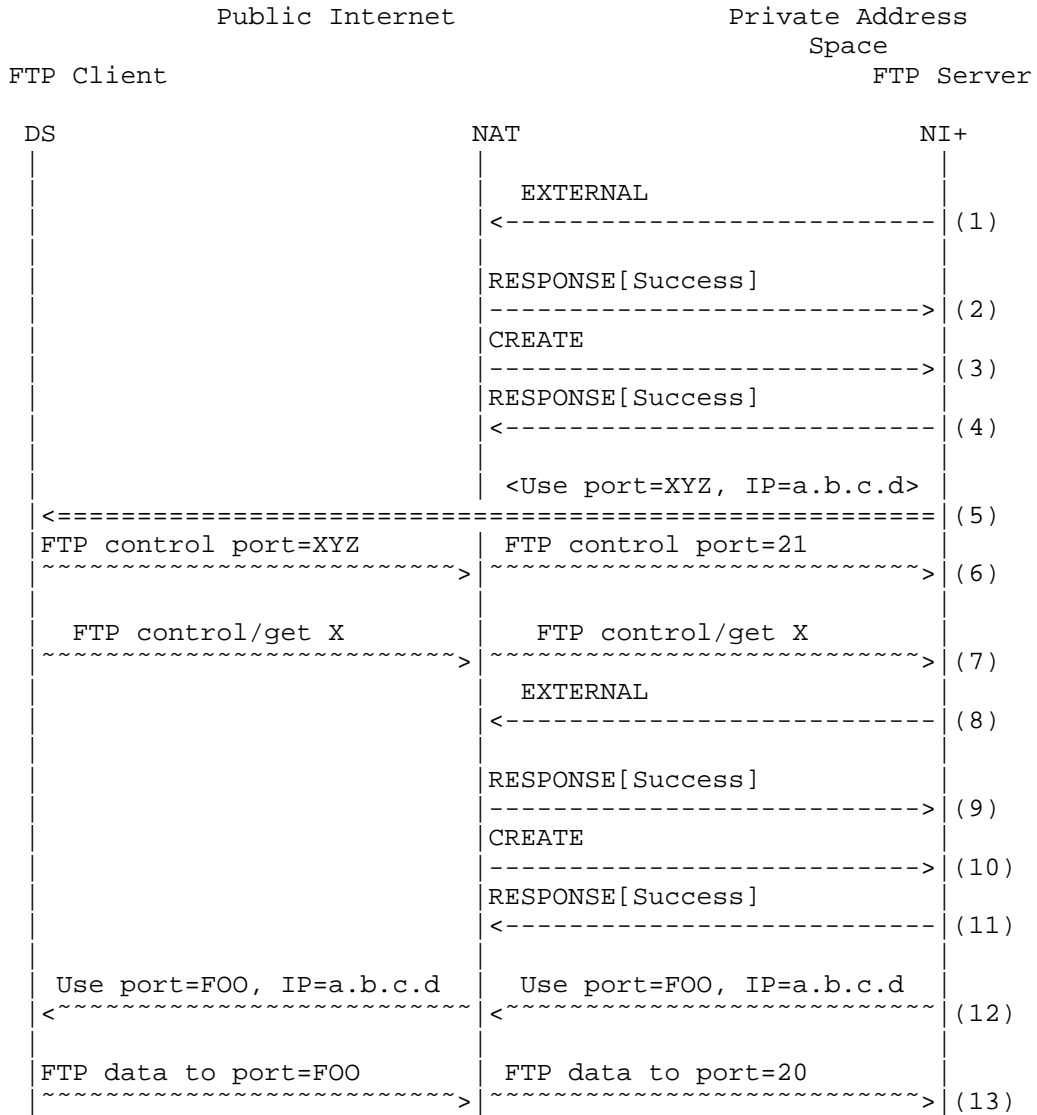


Figure 37: Flow Chart

1. EXTERNAL request message sent to NAT, with these objects: signaling session lifetime, extended flow information object (rule action=allow, sub_ports=0), message sequence number object, nonce object (carrying nonce for CREATE), and the data terminal information object (I/P-flags set, sender prefix=0, protocol=TCP, DR port number = 21, DS's IP address=0); using the LE-MRM. This is used to allocate the external binding for the FTP control channel (TCP, port 21).
2. Successful RESPONSE sent to NI+, with these objects: signaling session lifetime, message sequence number object, information code object ('Success':2), external address object (port=XYZ, IPv4 addr=a.b.c.d).
3. The NAT sends a CREATE towards NI+, with these objects: signaling session lifetime, extended flow information object (rule action=allow, sub_ports=0), message sequence number object, nonce object (with copied value from (1)); using the PC-MRM (src-IP=a.b.c.d, src-port=XYZ, dst-IP=NI+, dst-port=21, downstream).
4. Successful RESPONSE sent to NAT, with these objects: signaling session lifetime, message sequence number object, information code object ('Success':2).
5. The application at NI+ sends external NAT binding information to the other end, i.e., the FTP client at the DS.
6. The FTP client connects the FTP control channel to port=XYZ, IP=a.b.c.d.
7. The FTP client sends a get command for file X.
8. EXTERNAL request message sent to NAT, with these objects: signaling session lifetime, extended flow information object (rule action=allow, sub_ports=0), message sequence number object, nonce object (carrying nonce for CREATE), and the data terminal information object (I/P-flags set, sender prefix=32, protocol=TCP, DR port number = 20, DS's IP address=DS-IP); using the LE-MRM. This is used to allocate the external binding for the FTP data channel (TCP, port 22).
9. Successful RESPONSE sent to NI+, with these objects: signaling session lifetime, message sequence number object, information code object ('Success':2), external address object (port=FOO, IPv4 addr=a.b.c.d).

10. The NAT sends a CREATE towards NI+, with these objects: signaling session lifetime, extended flow information object (rule action=allow, sub_ports=0), message sequence number object, nonce object (with copied value from (1)); using the PC-MRM (src-IP=a.b.c.d, src-port=FOO, dst-IP=NI+, dst-port=20, downstream).
11. Successful RESPONSE sent to NAT, with these objects: signaling session lifetime, message sequence number object, information code object ('Success':2).
12. The FTP server responses with port=FOO and IP=a.b.c.d.
13. The FTP clients connects the data channel to port=FOO and IP=a.b.c.d.

Authors' Addresses

Martin Stiemerling
NEC Europe Ltd. and University of Goettingen
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 4342 113
EMail: Martin.Stiemerling@neclab.eu
URI: <http://www.stiemerling.org>

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
EMail: Hannes.Tschofenig@nsn.com
URI: <http://www.tschofenig.priv.at>

Cedric Aoun
Consultant
Paris, France

EMail: cedaoun@yahoo.fr

Elwyn Davies
Folly Consulting
Soham
UK

Phone: +44 7889 488 335
EMail: elwynd@dial.pipex.com