NSIS
Internet-Draft
Expires: December 30, 2004

A. Fessi
M. Stiemerling
NEC
S. Thiruvengadam
H. Tschofenig
Siemens
C. Aoun
Nortel Networks
July 2004

Security Threats for the NATFW NSLP
draft-fessi-nsis-natfw-threats-02

Status of this Memo

Copyright Notice

Abstract

Opening a firewall pinhole or creating a NAT binding is a very security sensitive issue. This memo identifies security threats and security requirements that need to be addressed for the NATFW NSLP. Generic security threats to the NSIS protocols have been already

discussed in the NSIS Working Group.

Table of Contents

1. Introduction

This document provides an analysis of the security threats that are specific for the NATFW NSLP. The NATFW NSLP is used to install the required policy rules (firewall pinhole and/or NAT binding) on middleboxes along the path to allow the traversal of a data flow.

Opening a pinhole in the firewall or creating a NAT binding is a very security sensitive issue. Thus, we need to examine carefully who is allowed to install these policy rules and what security threats need to be addressed. In this document we will analyze different types of possible attacks to networks running NSIS for middlebox configuration.

2. Terminology

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 [5].

Furthermore, we use the same terminology as in [1], [3] and [4].

3. Attacks related to authentication and authorization

As described in [1] the NSIS message which installs policy rules at a middlebox is the CREATE message. The CREATE message travels from the Data Sender (DS) toward the Data Receiver (DR). The packet filter or NAT binding is marked as pending by the middleboxes along the path.

If it is confirmed with a success RESPONSE message from the DR the requested policy rules on the middleboxes are installed to allow the traversal of a data flow.

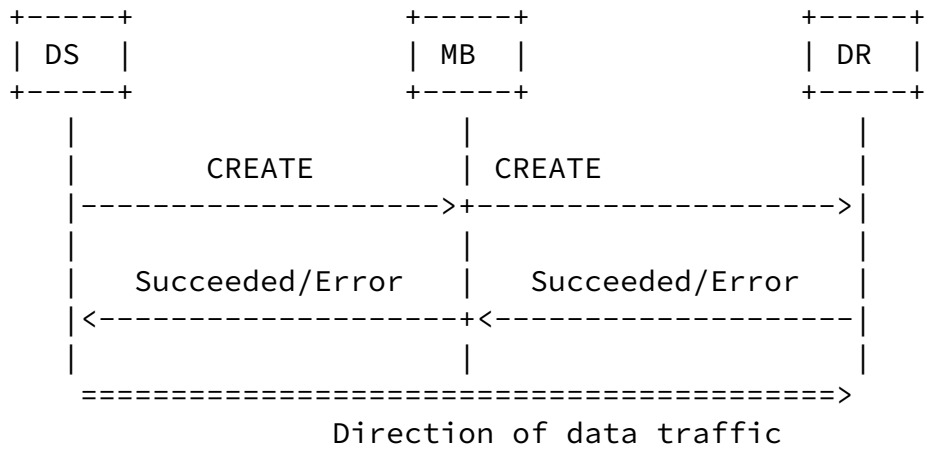

Figure 1: CREATE Mode

In this section we will consider some simple scenarios for middlebox configuration:

- Data Sender (DS) behind a firewall
- Data Sender (DS) behind a NAT
- Data Receiver (DR) behind a firewall
- Data Receiver (DR) behind a NAT

A real scenario could include a combination of one or more cases together, i.e., DS and/or DR is behind a chain of NATs and firewalls. Figure 2 shows such a possible scenario:

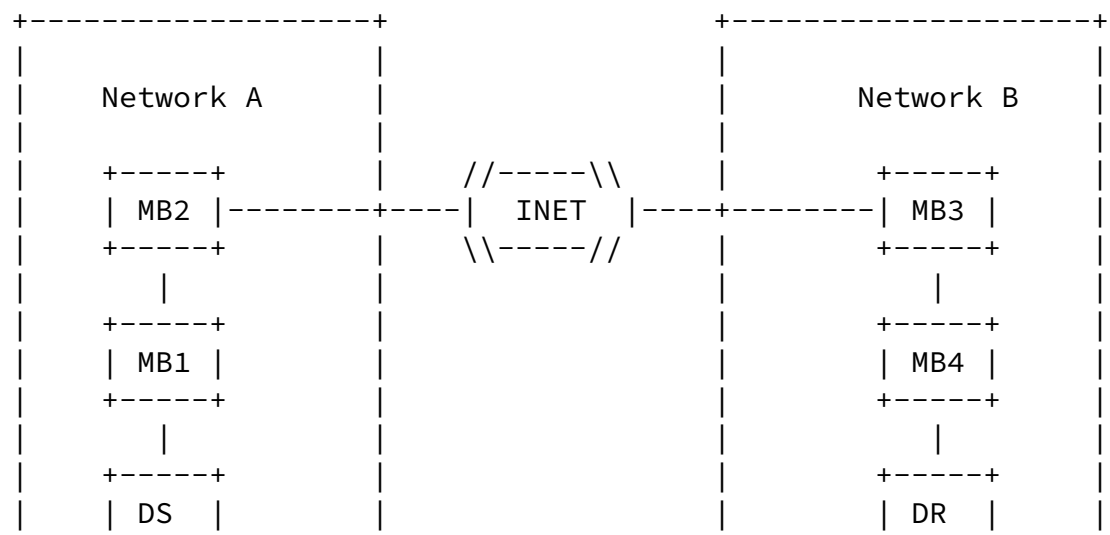

MB: Middle box (NAT or Firewall or a combination)
DS: Data Sender
DR: Data Receiver

Figure 2: Several middleboxes per network

3.1 Data Sender (DS) behind a firewall

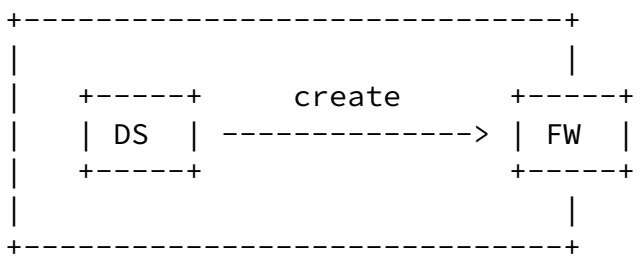

Figure 3: DS behind a firewall

DS sends a CREATE message to request the traversal of a data flow.

It is up to network operators to decide how far they can trust users inside their networks. However, there are several reasons why they should not.

The following attacks are possible:

- DS could open a firewall pinhole with a source address different from its own host.
- DS could open firewall pinholes for incoming data flows that are not supposed to enter the network.

- DS could request installation of any policy rules and allow all traffic go through.

SECURITY REQUIREMENT: As already mentioned in [1] Section (3.2), the middlebox MUST authenticate and authorize the neighboring NAT/FW NSLP node which requests an action. Authentication and authorization of the initiator SHOULD be provided to NATs and Firewalls along the path

as motivated with Section 2.2.3 of [1].

3.2 Data Sender (DS) behind a NAT

The case 'DS behind a NAT' is analogous to the case 'DS behind a firewall'.

It is worth mentioning that authentication based on IP address is not possible if NATs are deployed. Figure 4 illustrates such a scenario:

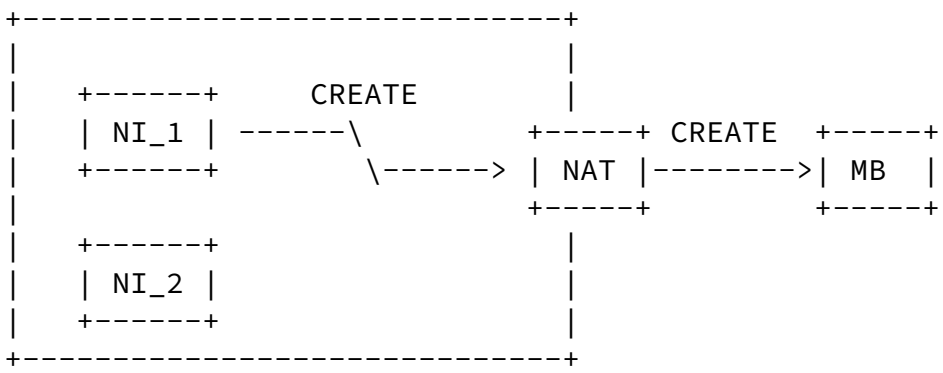

Figure 4: Several NIs behind a NAT

In this case the middlebox MB does not know who is the NSIS Initiator since both NI_1 and NI_2 are behind a NAT. Authentication needs to be provided by other means such as the NSLP or the application layer.

SECURITY REQUIREMENT: The middlebox MUST authenticate and ensure that the neighboring NAT/FW NSLP node is authorized to request an action. Authentication and authorization of the initiator (which is the DR in this scenario) MAY be provided to the middleboxes.

3.3 Data Receiver (DR) behind a firewall

In this case a CREATE message comes from an entity DS outside the network towards the DR inside the network.

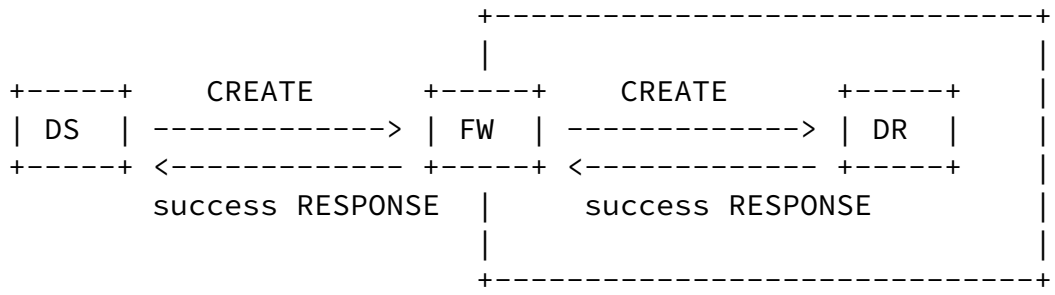

Figure 5: DR behind a firewall

According to [1] (Section 3.3) "Policy rules at middleboxes MUST be only installed upon receiving a successful response of type success RESPONSE".

This means that the middlebox waits until the Data Receiver DR confirms the request of the Data Sender DS with a success RESPONSE message. This is, however, only necessary

- if the policy rule creation/deletion/update at a firewall along the path cannot be authorized and
- if the middlebox is still forwarding the signaling message towards the end host (without state creation/deletion/modification).

This confirmation implies that the data receiver is expecting the data flow.

At this point we differentiate 2 cases:

1. DR knows the IP address of the DS (for instance because of some previous application layer signaling) and is expecting the data flow.
2. DR might be expecting the data flow (for instance because of some previous application layer signaling) but does not know the IP address of the Data Sender DS.

For the second case, Figure 6 illustrates a possible attack: an adversary Mallory M could be sniffing the application layer signaling and thus knows the address and port number where DR is expecting the data flow. Thus it could pretend to be DS and send a CREATE message towards DR with the data flow description (M -> DR). Since DR does not know the IP address of DS, it is not able to recognize that the request is coming from the "wrong guy". It will send a success RESPONSE message back and the middlebox will install policy rules that will allow Mallory M to inject its data into the network.

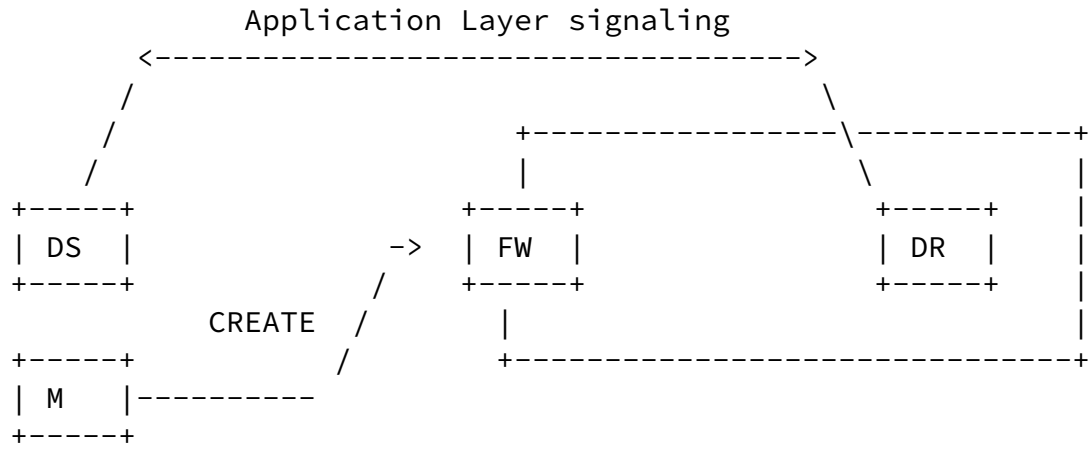

Figure 6: DR behind a firewall with an adversary

Network administrators will probably not rely on a DR to check the IP address of the DS. Thus we have to assume the worst case with an attack such as in Figure 6. Many operators might not allow NSIS signaling message to traverse the firewall in Figure 6 without proper authorization. In this case the threat is not applicable.

SECURITY REQUIREMENT: A binding between the application layer and the NSIS signaling SHOULD be provided.

3.4 Data Receiver (DR) behind a NAT

We will describe briefly the NSIS message flow that takes place to install the necessary rules for the traversal of a data flow from DS towards DR. For detailed description please refer to [1] Section 3.3.

DR sends a RESERVE-EXTERNAL-ADDRESS (REA) message to get a public reachable address that can be used by potential DSs. The NAT reserves an external address and port number and sends them back to DR. The NAT adds an address mapping entry in its reservation list which links the public and private addresses as follows:

(DR_ext <=> DR_int) (*).

The NAT sends a RESPONSE message with 'return external address' object back to the DR with the address DR_ext. DR informs DS about the public address that it has recently received, for instance, by

means of application layer signaling.

Now DS sends the CREATE message towards DR_ext. When the 'create session' message arrives at the NAT, the NAT looks up its reservation list and finds the entry (*).

Now the NAT knows the address of DS and stores it as a part of the policy rule to be loaded. It forwards the message towards DR and waits for the confirmation with the success RESPONSE message.

At the arrival of the success RESPONSE message from DR, the NAT installs the policy rule to forward the data flow correctly from DS to DR.

Possible attack:

We assume that the adversary obtains the external address allocated at the NAT (possibly by eavesdropping on the application layer signaling) and triggers the CREATE message before the NAT binding expires. The CREATE message is assumed to travel from DS to DR through NAT. An attacker Mallory M could send a CREATE message to install a NAT binding to forward the data flow from M to DR instead of from DS to DR. This kind of attack is equivalent to the attack described in Section 3.3 above.

In order for this attack to work the following pre-requisities need to hold:

- The adversary needs to be authorized to create a NAT binding at the NAT.
- The adversary needs to know when a DR creates a NAT binding at the DR. A certain timing is required and some specific information, such as the message routing identifier and session identifier must be known

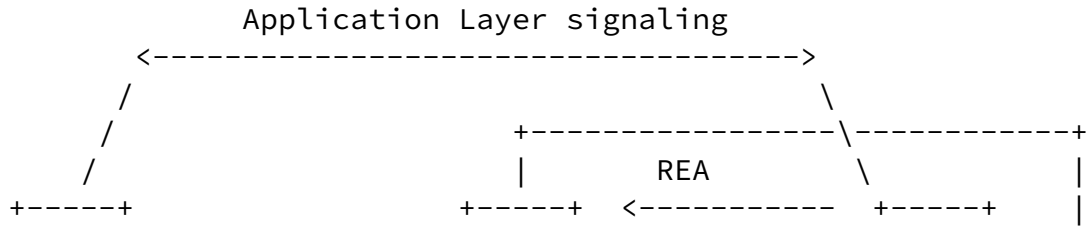

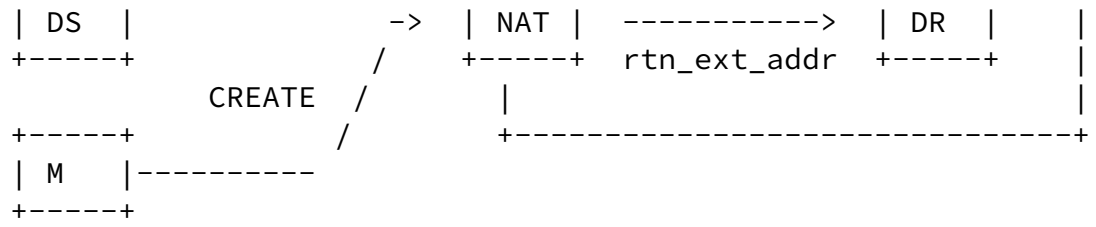

Figure 7: DR behind a NAT with an adversary

SECURITY REQUIREMENT: TBD

3.5 NSLP message injection

Malicious Hosts, locate either off-path or on-path, could inject arbitrary NATFW NSLP messages into the signaling path, causing several problems. These problems apply when no proper authorization and authentication scheme is available.

By injecting a bogus CREATE message with lifetime set to zero, a malicious host could try to teardown NATFW NSLP session state partially or completely on a data path, causing a service interruption.

By injecting a bogus responses message, for instance, timeout, a malicious host could try to teardown NATFW NSLP session state as well. This could affect the data path partially or complete, causing a service interruption.

Other messages, such as TRIGGER, can be misused by malicious hosts, causing a service interruption. Following versions of this document will investigate the impact of these messages as well.

4. Denial-of-Service Attacks

In this section we describe several ways how an adversary could launch a Denial of service (DoS) attack on networks running NSIS for middlebox configuration to exhaust their resources.

4.1 Flooding with CREATE messages from outside

4.1.1 Attacks due to NSLP state

A CREATE message requests the NSLP to store some state information such as Session-ID and flow identifier.

The policy rules requested in the CREATE message will be installed at the arrival of a confirmation from the Data Receiver with a success RESPONSE message. The success RESPONSE message includes the session ID. So the NSLP looks up the NSIS session and installs the requested policy rules.

An adversary from outside could launch a DoS attack with arbitrary CREATE messages. For each of these messages the middlebox needs to store state information such as the policy rules to be loaded, i.e., the middlebox could run out of memory. This kind of attack is also mentioned in [2] Section 4.8.

SECURITY REQUIREMENT: A NAT/FW NSLP node MUST authorize the 'create-session' message before storing state information.

4.1.2 Attacks due to authentication complexity

This kind of attack is possible if authentication is based on mechanisms that require computing power, for example, digital signatures.

For a more detailed treatment of this kind of attack, the reader is encouraged to see [2].

SECURITY REQUIREMENT: A NAT/FW NSLP node MUST NOT introduce new denial of service attacks based on authentication or key management mechanisms.

4.1.3 Attacks to the endpoints

The NATFW NSLP requires firewalls to forward NSLP messages, a malicious node may keep sending NSLP messages to a target. This may consume the access network resources of the victim, drain the battery of the victim's terminal and may force the victim to pay for the received although undesired data.

This threat may be more particularly be relevant in networks where

access link is a limited resource, for instance in cellular networks, and where the terminal capacities are limited.

SECURITY REQUIREMENT: A NATFW NSLP aware firewall or NAT MUST be able to block unauthorized signaling message, if this threat is a concern.

4.1.4 Attacks to the NTLP

Flooding a middlebox with CREATE messages affects also the NTLP.

The success RESPONSE message needs to take the same route as the previous CREATE message. Thus the NTLP needs to store routing information for each CREATE message. This kind of attack is also described in [2] Section 4.8.

SECURITY REQUIREMENT: A NAT/FW NSLP node MUST NOT introduce new denial of service attacks based on authentication or key management mechanisms.

4.2 Flooding with REA messages from inside

Although we are more concerned with possible attacks from outside the network, we need also to consider possible attacks from inside the network.

An adversary inside the network could send arbitrary

RESERVE-EXTERNAL-ADDRESS messages. At a certain point the NAT will run out of port numbers and the access for other users to the outside will be disabled.

SECURITY REQUIREMENT: The NAT/FW NSLP node MUST authorize state creation for the RESERVE-EXTERNAL-ADDRESS message. Furthermore, the NAT/FW NSLP implementation MUST prevent denial of service attacks involving the allocation of an arbitrary number of NAT bindings or the installation of a large number of packet filters.

5. Man-in-the-Middle Attacks

Figure 8 illustrates a possible man-in-the-middle attack using the 'reserve external address' (REA) message. This message travels from DR towards the public Internet. The message might not be intercepted

by any NAT either because there are no NATs or because there are NSIS unaware.

Mallory M returns a faked RESPONSE message with an IP address of its choosing. This IP address is meant to be used by the DR as the public external IP address. Malory might insert it own IP address in the response, the IP address of a third party or the address of a black hole. In the first case, the DR thinks that the address of Mallory M is its public address and will inform the DS about it. As a consequence, the DS will send the data traffic to Mallory M.

The data traffic from the DS to the DR will re-directed to Mallory M. Mallory M will be able to read, modify or block the data traffic. Eavesdropping and modification is only possible if the data traffic is itself unprotected.

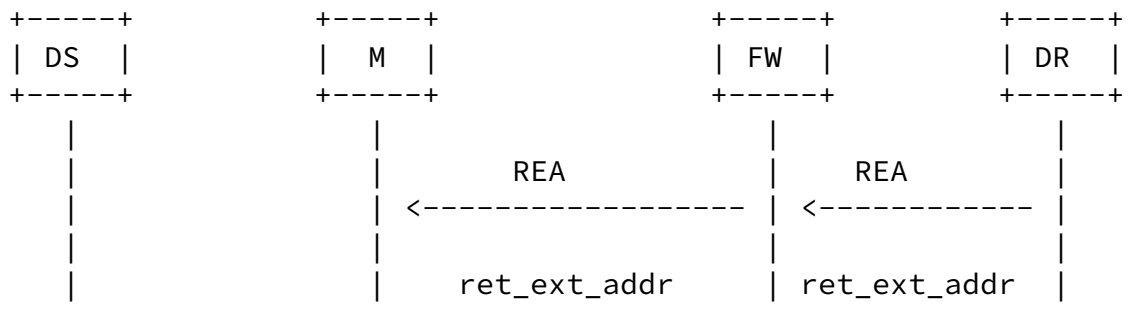

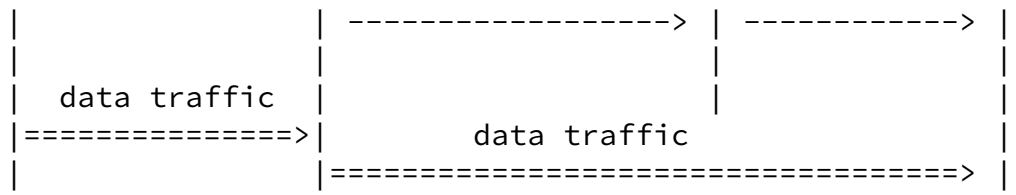

Figure 8: MITM attack using the RESERVE-EXTERNAL-ADDRESS message

Please note that the NSIS aware firewall in Figure 8 might not be present when DR communicates directly with the adversary.

SECURITY REQUIREMENT: Mutual authentication between neighboring NATFW NSLP MUST be provided. To ensure that only legitimate nodes along the path act as NSIS entities the initiator MUST authorize the responder. In the example in Figure 8 the firewall FW must perform an authorization with the neighboring entities.

6. Message Modification

Any on-path subverted node en route to the destination could easily modify, inject or just drop an NSIS message. It could also hijack or disrupt the communication.

SECURITY REQUIREMENT: Message integrity, replay protection and data origin authentication between neighboring NAT/FW NSLPs MUST be provided.

Message modification by a subverted NSIS node could create arbitrary pinholes or NAT bindigs. For example:

- NATs need to modify the source/destination of the data flow in the 'create session' message.
- Each middlebox along the path may change the requested lifetime in the CREATE message to fit their needs and/or local policy (see also section 3.2.7 of [1] with regard to calculation of refresh interval).

SECURITY REQUIREMENT: None. Malicous NSIS NATs and Firewalls will not be addressed.

7. Session Modification/Deletion

The Session ID is included in signaling messages as a reference to the established state. If an adversary is able to obtain the Session Identifier for example by eavesdropping on signaling messages, it would be able to add the same Session Identifier to a new a signaling message and effect some modifications.

Consider the scenario described in Figure 9. Here an adversary pretends to be 'DS in mobility'. The signaling messages start from the DS and go through a series of routers towards the DR. We assume that an off-path adversary is connected to one of the routers along the old path (here Router 3). We also assume that the adversary knows the Session ID of the NSIS session initiated by the DS. Knowing the Session ID, the adversary now sends signalling messages towards the DR. When the signaling message reaches Router3 then existing state information can be modified or even deleted. The adversary can modify or delete the established reservation causing unexpected behavior for the legitimate user. The source of the problem is that the Router 3 (cross-over router) is unable to decide whether the new signaling message was initiated from the owner of the session. In this scenario, the adversary need not even be located in the DS-DR path. This problem and the solution approaches are described in more detail in [6].

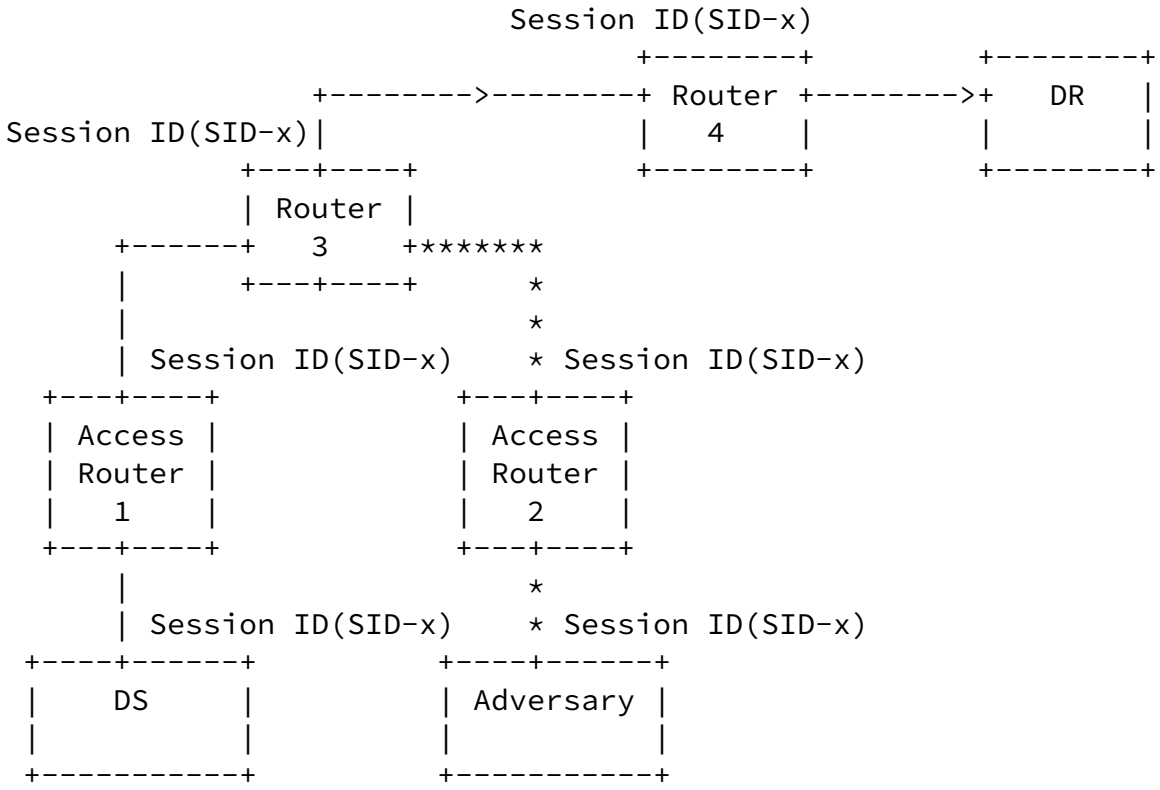

Figure 9: State Modification by off-path adversary

Summary: Off-path adversary's knowledge of Session-ID could cause session modification/deletion.

SECURITY REQUIREMENT: TBD: This is not a NAT/FW NSLP specific problem but a GIMPS problem. The initiator MUST be able to demonstrate ownership of the session it wishes to modify.

7.1 Misuse of mobility in NAT handling

Another kind of session modification is related to mobility scenarios. NSIS allows end hosts to be mobile it is possible that an NSIS node behind a NAT needs to update its NAT binding in case of address change. Whenever a host behind a NAT initiates a data transfer, it is assigned an external IP and port number. In typical mobility scenarios, the DR might also obtain a new address according to the topology and it should convey the NAT binding updates. The NAT is assumed to modify these NAT bindings based on the new IP address conveyed by the endhost.

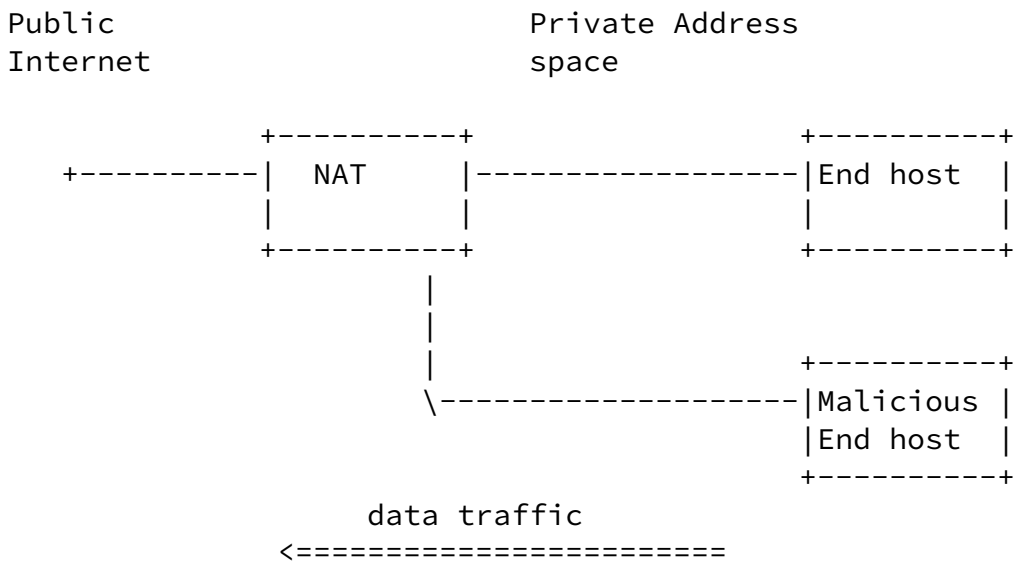

Figure 10: Misuse of mobility in NAT binding

For this description , we assume that a NAT binding state can be changed with the help of NSIS signalling. When the DR is receiving data traffic, if it happens to move to a new location, it sends an NSIS signalling message to modify the NAT binding. It would use the Session-ID and the new flow-id to update the state. The NAT updates the binding and the DR continues to receive the data traffic. Consider the scenario in Figure 10 where an the endhost(DR) and the adversary are behind a NAT. The adversary pretending that it is the end host could generate a spurious signaling message to update the state at the NAT. This could be done for these purposes:

1. Connection hijacking by redirecting packets to the attacker as in Figure 11

2. Third party flooding by redirecting packets to arbitrary hosts

3. Service disruption by redirecting to non-existing hosts

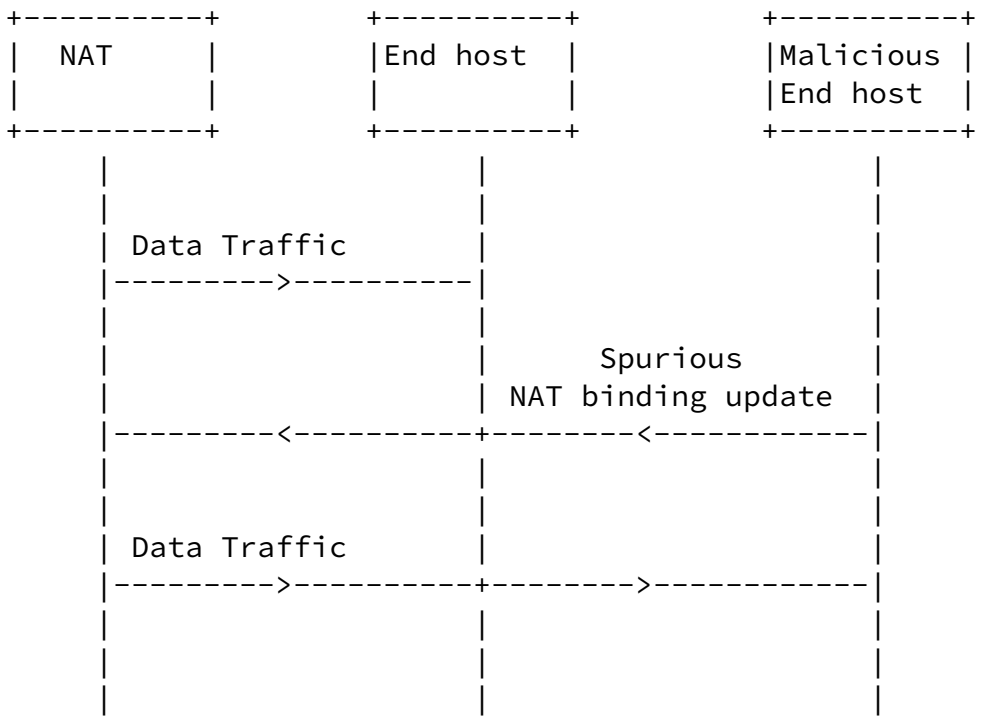

Figure 11: Connection Hijacking

SECURITY REQUIREMENT: A NAT/FW signaling message MUST be authenticated, authorized, integrity protected and replay protected between neighboring NAT/FW NSLP nodes.

8. Misuse of unreleased sessions

Assume that DS (N1) initiates NSIS session with DR (N2) through a series of middleboxes as in Figure 12. When the DS is sending data to DR, it might happen that the DR disconnects from the network

(crashes or moves out of the network in mobility scenarios). In such cases, it is possible that another node N3 (which recently entered the network protected by the same firewall) is assigned the same IP address that was previously allocated to N2. The DS could take advantage of the firewall policies installed already, if the refresh interval time is very high. The DS can flood the node (N3), which will consume the access network resources of the victim forcing it to pay for unwanted traffic as shown in Figure 13.

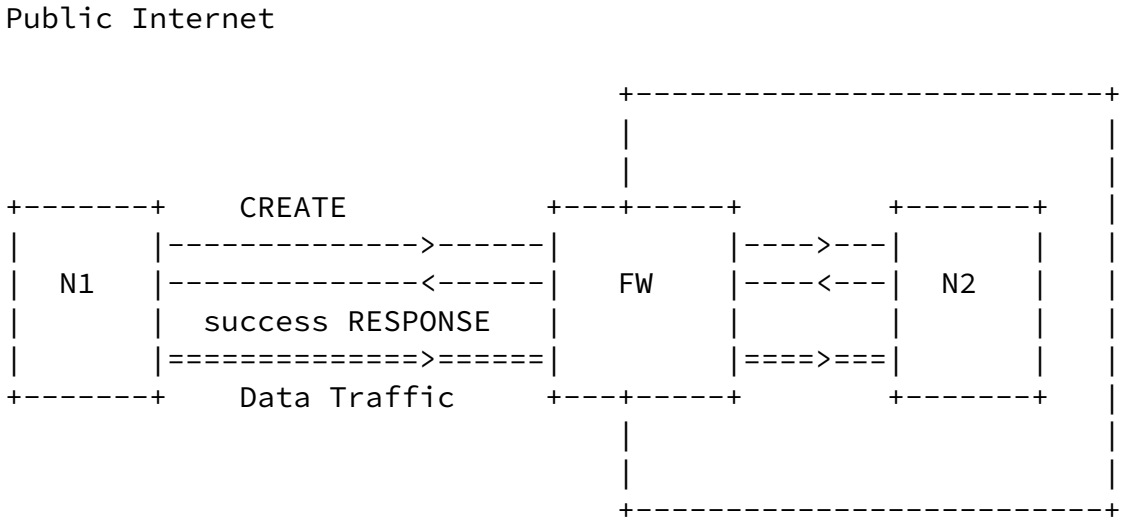

Figure 12: Before mobility

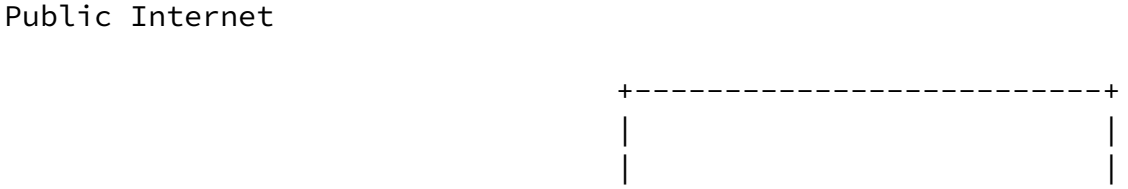

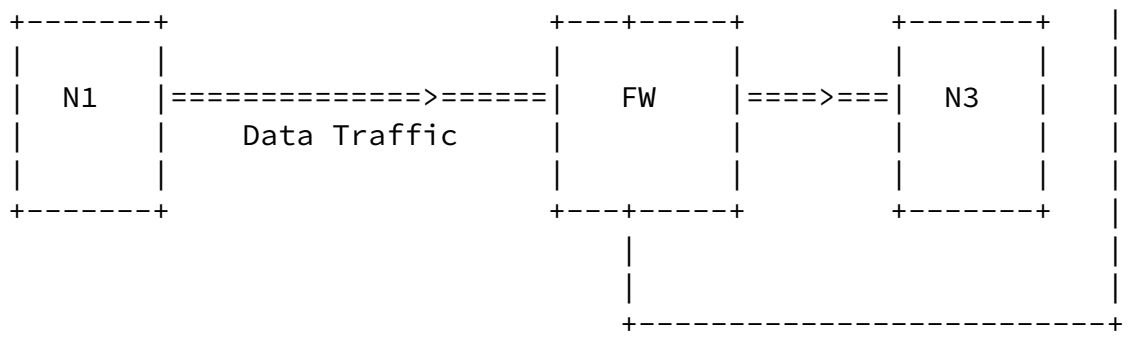

Figure 13: After mobility

Also, this threat is valid for the other direction as well. The DS which is communicating with the DR may disconnect from the network and this IP address may be assigned to a new node that had recently entered the network. This new node could pretend to be the DS and send data traffic to the DR in conformance with the firewall policies and cause service disruption.

SECURITY REQUIREMENT: Data origin authentication is needed to mitigate this threat. However, the described threat is applicable

only for the time until the policy rules are deleted due to NSLP soft state. Awareness for this threat is important especially when the refresh interval time is high. It should be noted, that networks supporting mobility should remove any state at middleboxes when a mobile node is diconnecting, thus leaving a clean state.

9. Data traffic injection

This attack takes place where there exists trust relationship between machines. It is common in corporate networks, where internal machines trust each other and authentication is only based on IP address. Hence by spoofing a connection, an attacker is able to reach the target machines, using the existing firewall rules.

The adversary is able to inject its own data traffic in conformance with the firewall policies simultaneously along with the genuine DS.

SECURITY REQUIREMENT: Since IP spoofing is a general limitation of

non-cryptographic packet filters no security requirement needs to be created for the NAT/FW NSLP. Techniques such as ingress filtering (described below) and data origin authentication (such as provided with IPsec based VPNs) can help mitigate this threat. This issue is, however, outside the scope of this document.

Ingress Filtering: Consider the scenario shown in Figure 14. In this scenario the DS is behind a router (R1) and a malicious node (M) is behind another router (R2). The DS communicates with the DR through a firewall (FW). The DS initiates NSIS signaling and installs firewall policies at FW. But the malicious node is also able to send data traffic using DS's source address. If R2 implements ingress filtering, these spoofed packets will be blocked. But this ingress filtering may not work in all scenarios. If both the DS and the malicious node are behind the same router, then the ingress filter will not be able to detect the spoofed packets as both the DS and the malicious node are in the same address range.

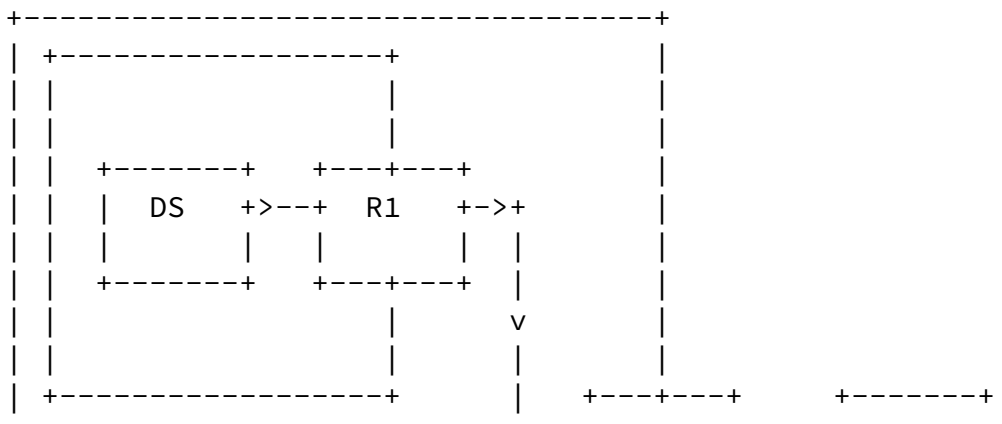

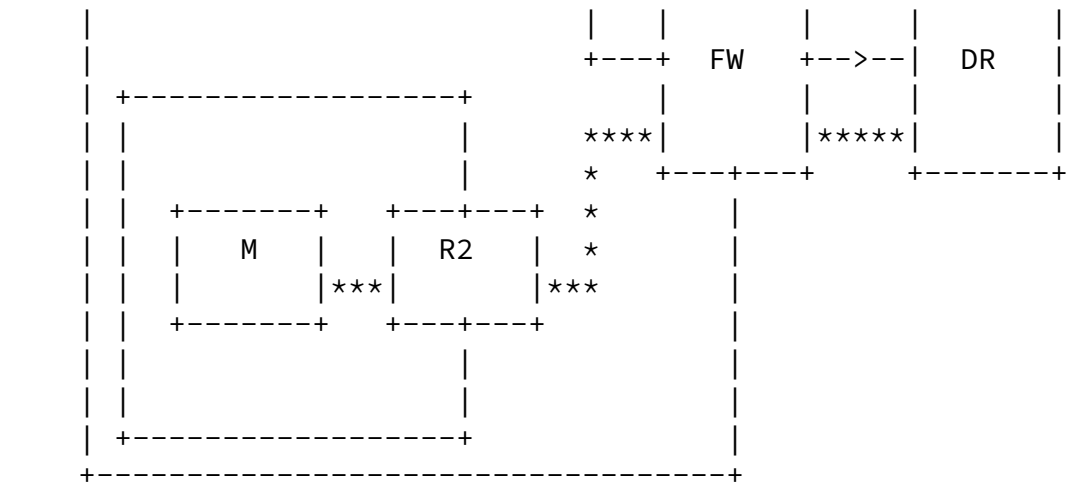

----->---- = genuine data traffic
********* = spoofed data traffic

Figure 14: Ingress filtering

10. Eavesdropping and traffic analysis

By collecting NSLP messages, an adversary is able to learn policy rules for packet filters and knows which ports are open. It can use this to inject its own data traffic due to the IP spoofing capability as already mentioned in Section 9.

An adversary could learn authorization tokens included in CREATE messages and use them to launch reply-attacks or to create a session with its own address as source address. (cut-and-paste attack)

As shown in Section 4.3 of [6] a solution of Section 7 might require confidentiality protection of signaling messages

SECURITY REQUIREMENT: The threat of eavesdropping itself does not

mandate the usage of confidentiality protection since an adversary can also eavesdrop on data traffic. In the context of a particular security solutions (e.g., authorization tokens) it might be necessary to offer confidentiality protection. Confidentiality protection also

needs to be offered to the refresh period.

11. Conclusions

This memo describes security threads that are applicable to the NSIS NATFW NSLP and some related threads inherent to firewalls and NATs. Security requirements are given for the scenarios and some issues to be considered in NTLP design are raised.

The most security threads shown here are related to missing authentication or authorization schemes between all NATFW nodes. Given a proper authentication and authorization scheme, many of these threads can be mitigated. The general problem is the missing identity of the nodes to what authorization and authentication could be bound. IP addresses are in general not suitable, since NATs are involved in any place to imagine and in mobility scenarios they are changed frequently. Other attacks, such as message eavesdropping, can be managed easily between adjacent NSIS nodes if the NTLP or NSLP supports encryption. The flooding, or denial of service, of NSIS nodes can be mitigated not only by authorization and authentication schemes, but also by extensions to NATFW NSLP, where NSIS receivers should be able to communicate upstream which type or from which node, via the flow routing information, signaling traffic is allowed to be forwarded to them.

12. Security Considerations

The entire document highlights security threats that need to be mitigated for the NATFW NSLP. It also addresses security issues related to packet filters. Security requirements have been derived from the relevant threats.

13. Acknowledgments

This document is the result of discussions with many individuals. The authors would like to thank especially: Marcus Brunner, Miquel Martin, Frank Le, Joao Girao, and Elwyn Davis.

14. References

14.1 Normative References

[1] Stiemerling, M., Tschofenig, H. and M. Martin, "A NAT/Firewall NSIS Signaling Layer Protocol (NSLP)",

draft-ietf-nsis-nslp-natfw-03 (work in progress), July 2004, <reference.I-D.ietf-nsis-nslp-natfw.xml>.

[2] Tschofenig, H. and D. Kroeselberg, "Security Threats for NSIS", draft-ietf-nsis-threats-05 (work in progress), June 2004, <reference.I-D.ietf-nsis-threats.xml>.

[3] Schulzrinne, H. and R. Hancock, "GIMPS: General Internet Messaging Protocol for Signaling", draft-draft-ietf-nsis-ntlp-02 (work in progress), May 2004, <reference.I-D.draft-ietf-nsis-ntlp.xml>.

[4] Brunner, M., "Requirements for Signaling Protocols", RFC 3726, April 2004, <reference.I-D.ietf-nsis-.requirements.xml>.

[5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

14.2 Informative References

[6] Tschofenig, H., Schulzrinne, H., Hancock, R., McDonald, A. and X. Fu, "Security Implications of the Session Identifier", June 2003, <reference.I-D.tschofenig-nsis-sid.xml>.

[7] Aoun, C., Brunner, M., Stiemerling, M., Martin, M. and H. Tschofenig, "NAT/Firewall NSLP Migration Considerations", draft-aoun-nsis-nslp-natfw-migration-01 (work in progress), February 2004, <reference.I-D.aoun-nsis-nslp-natfw-migration.xml>.

[8] Bless, R., "Mobility and Internet Signaling Protocols", draft-manyfolks-signaling-protocol-mobility-00 (work in progress), January 2004, <reference.I-D.manyfolks-signaling-protocol-mobility.xml>.

[9] Bosch, S., "NSLP for Quality-of-Service signaling", draft-ietf-nsis-qos-nslp-03 (work in progress), May 2004, <reference.I-D.ietf-nsis-qos-nslp.xml>.

Authors' Addresses

Ali Fessi
Network Laboratories, NEC Europe Ltd.

EMail: alifessi@web.de
URI:

Martin Stiemerling
Network Laboratories, NEC Europe Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 905 11 13
EMail: stiemerling@netlab.nec.de
URI:

Srinath Thiruvengadam
Siemens
Otto-Hahn-Ring 6
Munich, Bayern 81739
Germany

EMail: srinath@mytum.de

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

EMail: Hannes.Tschofenig@siemens.com

Cedric Aoun
Nortel Networks

France

EMail: cedric.aoun@nortelnetworks.com

Intellectual Property Statement

Disclaimer of Validity

WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Acknowledgment

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