NSIS Working Group
Internet-Draft
Expires: April 19, 2004

C. Aoun
Nortel Networks
M. Brunner
M. Stiemerling
M. Martin
NEC
H. Tschofenig
Siemens

October 20, 2003

NATFirewall NSLP migration and intra-realm communication considerations
draft-aoun-nsis-nslp-natfw-migration-00

Status of this Memo

Copyright Notice

Abstract

This document discusses NAT/FW migration to support the NSIS NAT/FW NSLP as well as intra-realm communications considerations. The document will serve as input to the NSIS NATFW NSLP document.

Table of Contents

1. Introduction

Whether interim NAT and Firewall traversal mechanisms will already be deployed in networks or not, it is important to understand NSIS NATFW NSLP co-existence with non-NSIS aware NATs and Firewalls until their migration to NSIS would have been accomplished. The NSIS NATFW NSLP responder provides the NSIS NATFW NSLP Initiator with its address, in some cases the NSIS responder may have several addresses: a (or several) global scoped address and its locally scoped address(es). Which address does it provide to the NSIS initiator? This document will cover both the above issues and serve as input to the NSIS NATFW NSLP main document.

2. NSIS Responder address selection for intra realm communications optimization

An NSIS Element hosting an NSIS NATFW NSLP may have more than one address, its local scoped address(es) and global scoped address(es). A default address selection that priorities arbitrarily a global scoped address over a local scoped address may imply none optimal routing for communications between NSIS elements hosted within the same addressing realm.

In Figure 1 the arbitrary selection of the global scoped address, works properly to receive NSIS signaling from Host B. However in deployment scenario shown in Figure 2, host A and host C end-up communicating through their local MB1 middlebox resulting in a non optimal data path with all its implications (additional delay, additional bandwidth in certain parts of the network infrastructure, etc ...).

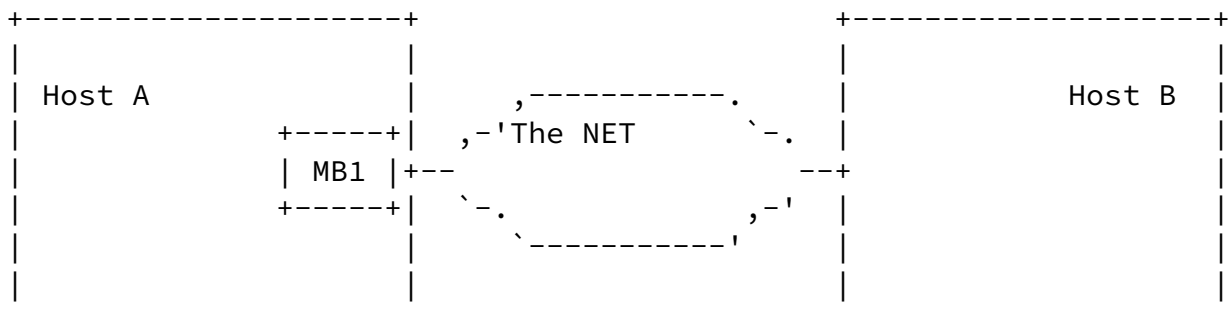

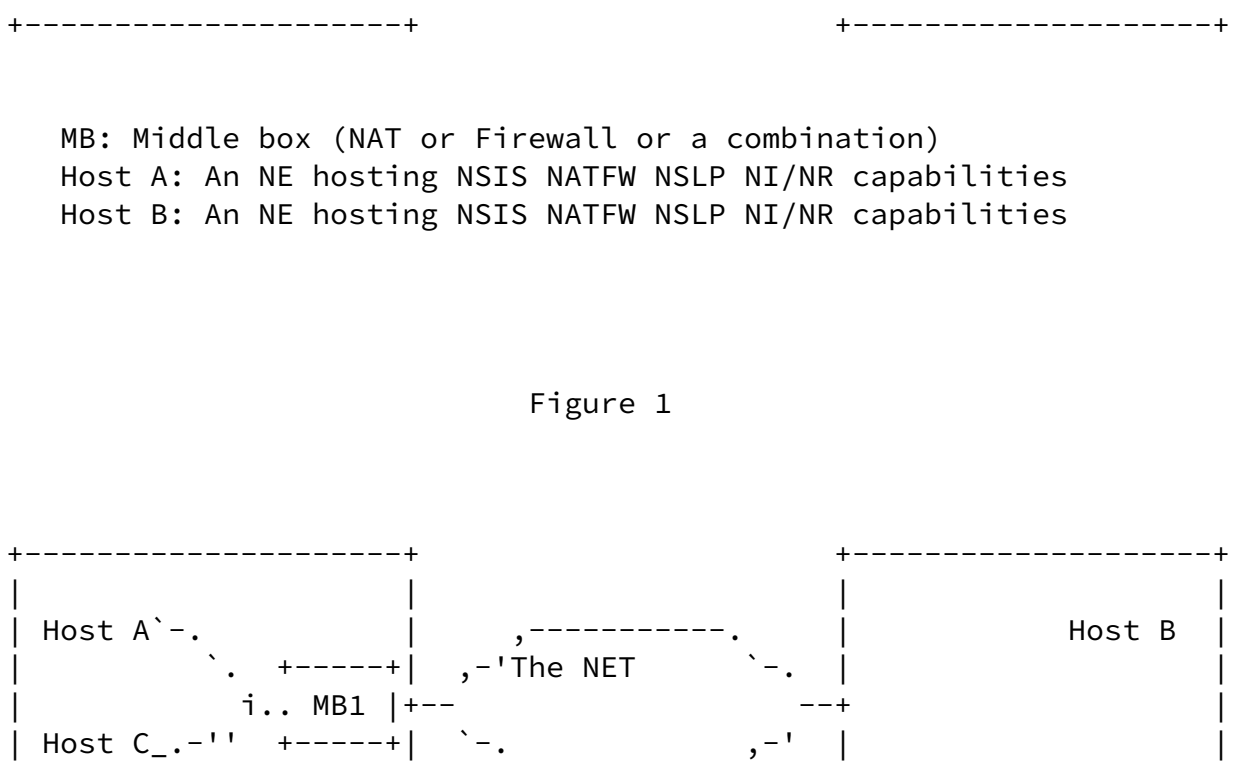

Figure 1

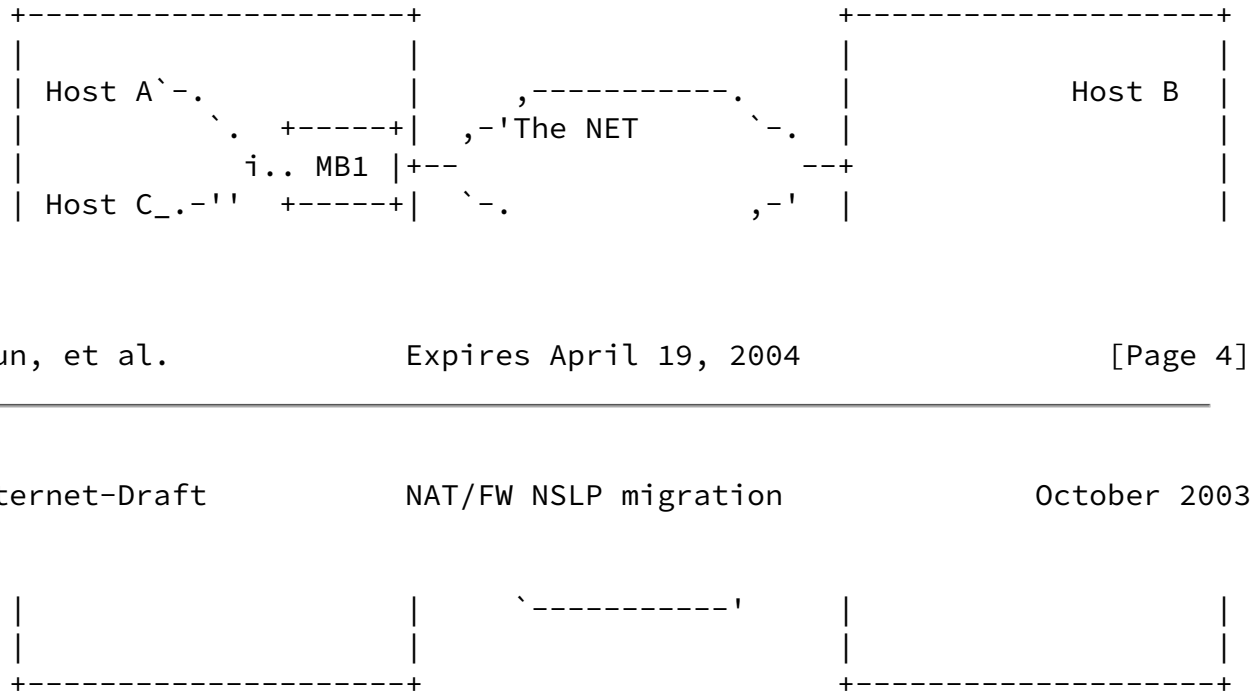

Figure 2

2.1 Potential solutions to the problem

There are two ways to deal with this non-optimal routing of packets for intra-realm communications between NSIS entities. The NSIS Responder should either:

1. Decide based on local decision, which address to provide to the NI to signal it or,

2. Let the far end NSIS Initiator decide which address to select based on the NSIS Responder's provided addresses

(1) lets the NSIS Responder decide on its own, which address to provide based on certain hints, which may not be the most optimal decision since the NSIS Responder may not have sufficient knowledge about the NSIS Initiator at the time of delivering its address to the responder. In most cases none arbitrary local decision for address selection requires to know about the far end host, which is not always practical.

An example of such local based address selection is the IPv6 default address selection mechanism ([2]) where an IPv6 (or IPv6/v4) node has to select which source and destination address to pick in order to communicate with a far end node. The mechanism relies on receiving input from the local resolver (DNS client or local hosts file) about the far end node . In the context of multimedia applications (and probably others as well), the NSIS Initiator is provided the NSIS Responder contact information (includes IP address and transport port [1] via the user application protocols (example: SIP, H.323 etc ...). Within that specific context, the NSIS Responder does not have sufficient information about the NSIS Initiator to provide it a valid address.

Hence it can not decide successfully in all cases which address to provide to the NSIS Initiator. Hence (2) is more efficient as it requires the NSIS Responder to provide all its addresses (local scoped and global scoped ones) to the NSIS Initiator. As a result, the NSIS Initiator will need to decide on which address to signal the NSIS Responder among all the provided ones. One possible way for the NSIS Initiator to decide from which address it will send NSIS signaling to the NSIS Responder and which NSIS Responder address to use is to check on which address the NSIS responder will reply back. An existing proposal [3] discusses the usage of (2) for SIP User Agents (SIP UA), the SIP UA will probe the far end SIP UA to see from which address it will response back. In [3], the STUN protocol [7] is used to check the responsiveness of the far end host. In [6], the reverse routability used to check which address to respond to

counters RTP DOS attacks. The same required reverse routability could be achieved by the NSIS NATFW NSLP.

****Include message sequences in the next iteration of the discussion*******

In the context of NSIS, the NSIS protocol itself should be used as the probing mechanism.
Consequently the NSIS Initiator will send simultaneously several NSIS messages towards the NSIS Responder, one message per provided address.
The following occur during the process:

- In case the NSIS Responder and Initiator are located within the same addressing realm, the NSIS Responder should receive a response from all the addresses to which it has sent NSIS messages. The NSIS Initiator will then choose the local scoped address as the destination address for messages destined to the NSIS Responder.

- In case the NSIS Responder and Initiator are not located within the same addressing realm, the NSIS Initiator would only receive a response from the valid global scope address. In event that there is another NE within the network that has the same local scoped address as the originally targeted NSIS Responder, the wrongly targeted NSIS Responder should send back an error or discard the message (the later would be preferable).

2.1.1 Intra realm solution protocol operation impacts

As opposed to the normal NSIS mode of operation, an NI that has already a security association with the neighboring NE on the path to a specific NR would need to verify that the target local scoped NR

address is the same as one already cached in its NSIS neighbor table cache (association of Next hop NE with the target NR table).This would be required to avoid any confusion with an NSIS node that could be hosted within the same addressing realm and that would have the same local scoped address.
There is a potential threat where an malicious NE with the same local scoped address as the target NR respond back positively to the NSIS

message and consequently hijack the data flow. This threat could be counter-measured by proper cryptographic authentication of the NSIS Responder response by using keying material provided by the application signaling (assuming that it was secured).

The procedure requiring an NSIS Initiator to send NSIS messages to several NR addresses, requires that the NI sends his NSIS messages at the same time to avoid impacting real-time sensitive applications. In case the response to the message sent to the global scoped is received first, it could potentially be used as a hint that no response will be received from the NR's local scoped address. Hence there is no point to continue to delay the address selection process and proceed with the NSIS protocol operations. In case the first response is received from a local scoped address and has been authenticated with the keying material provided by the application signaling protocol then the address selection process ends, and the NSIS protocol operations continue.

2.1.2 Intra realm solution application protocols impacts

The proposed intra realm path optimization proposal requiring an NR to provide several data recipient addresses within the application protocol, has obviously a certain impact on the application protocol. [5] discusses the impact to SDP [8] and should be used by all the application protocols using SDP. A similar approach should be followed by other protocols, not using SDP, including H.323 [9] and others requiring usage of NSIS with multiple addresses for the NR.

3. NATFW NSLP migration analysis

This section goes through a detailed analysis of the NATFW NSLP transition challenges and its possible solutions. The conception of the NSIS NATFW NSLP MUST ensure that upon its deployments in networks, it should not disrupt applications using interim NAT and Firewall traversal mechanisms.
In addition:

- The NATFW NSLP should ensure that an NR would not always be required to have minimal service, particularly when the NI has a simple network configuration without asymmetric route issues. In the early phases of the NSIS NATFW NSLP migration, this situation will be quite frequent and hence this scenario must be supported.

- The NSIS protocol should be designed to traverse non-NSIS aware NATs and Firewalls, to allow usage of non NAT and Firewall related NSLP (Qos NSLP for example). As the reader will notice in the subsequent sections NRs behind

- It is advised for an NSIS aware NAT and Firewall implementation to keep its existing currently used stateful behavior (depending on its applicability) until the transition to NSIS has ended (in order to have a smoother transition).

Several deployment scenario will be considered within this analysis, for simplicity the discussed scenario cover at most two Middleboxes on the path between the NI and NR. In Figure 3, a total of 144 deployment scenario could be possible only the ones having an NI or NR (only when there is a NAT) are considered. Implied but not shown scenario within Figure 5 are ones in the NI column or NR row. Obviously not all the scenario will be covered in this section, only the most interesting scenario are discussed in the next sections. A check list will be added later on in the annex of the next iteration of the analysis.

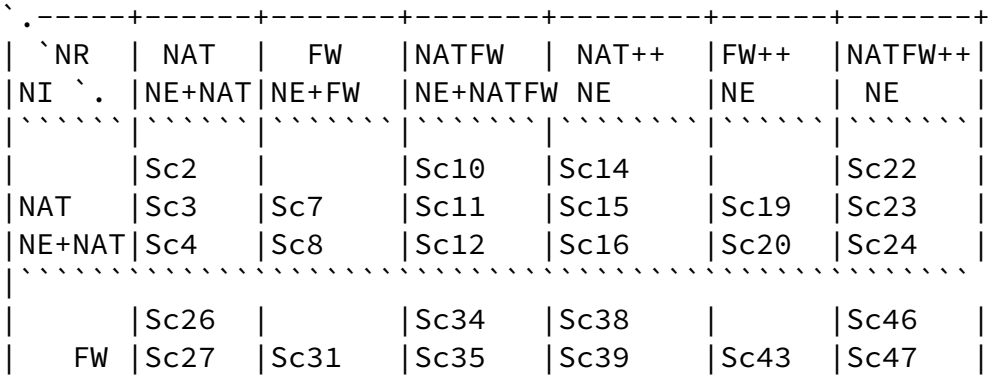

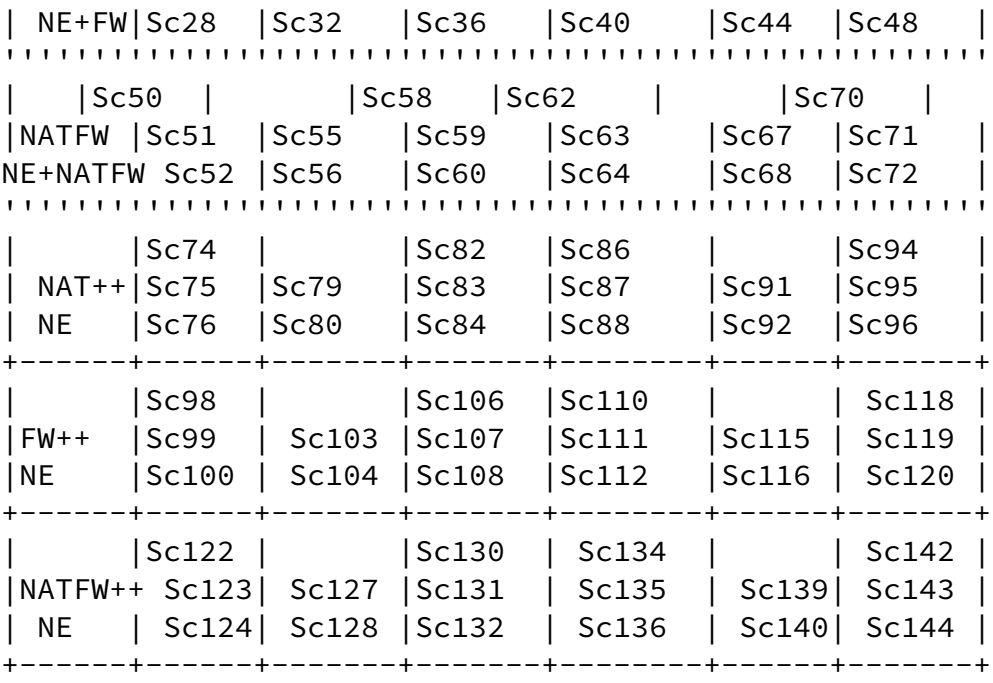

Figure 3

Scxyz: Scenario number xyz
NAT : NSIS Unaware NAT
NE : NSIS Entity with NI and NR functions
NE+NAT: NE hosted within a network connected to the NSIS unaware NAT FW : NSIS Unaware Firewall
NE+FW: NE hosted within a network protected by an NSIS Unaware FW
NATFW: NSIS Unaware NAT and Firewall hosted within the same Middlebox
NE+NATFW: NE hosted within a network protected by an NSIS Unaware NATFW
NAT++: NSIS Aware NAT
NAT++NE: NE hosted within a network connected to an NAT++
FW++ : NSIS Aware FW
FW++NE: NE hosted within a network connected to a FW++
NATFW++: NSIS Aware NATFW
NATFW++NE: NE hosted within a network connected to a NATFW++
Every cell in Figure 3 is the combination of the NI's network and the NR's network. Deployments scenario, where there are no NI and NR, no NI (without an NR with a NAT), are not shown. For example:

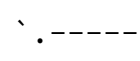

Figure 4

Scenario 1: NAT x NAT is not considered as there is no NI and no NR with a NAT
Scenario 2: NAT x NE+NAT is considered as there is an NR with a NAT (even if it is not NSIS aware)
Scenario 3: NE+NAT x NAT, is considered since there is an NI as part of the NE functions
Scenario 4: NE+NAT x NE+NAT, is considered since there is an NI as part of the NE functions
The same logic is applicable to all the cells in Figure 3
For simplicity only network deployments are shown with a maximum of 2 MBs on the path between the NI and the NR. Handled scenario within the NI column or NR raw are the ones having an NI or NR. In the next sections, we shall go through the various issues that are specific to the scenario mentioned in Figure 3.

3.1 Global scoped address determination with NSIS unaware NATs

This section discusses the potentional role that an NE with a NATFW NSLP could still have to determine a global scoped address translated by a none NSIS aware NAT. Upon detection that the NE is attached to a network hosting an NSIS unaware NAT, it could have the two alternatives, either:

1. The NSIS API could invoke the services of a STUN client [7] as shown is Figure 7, this would allow applications using UDP transport to work (only applicable for cone NAT [7]).

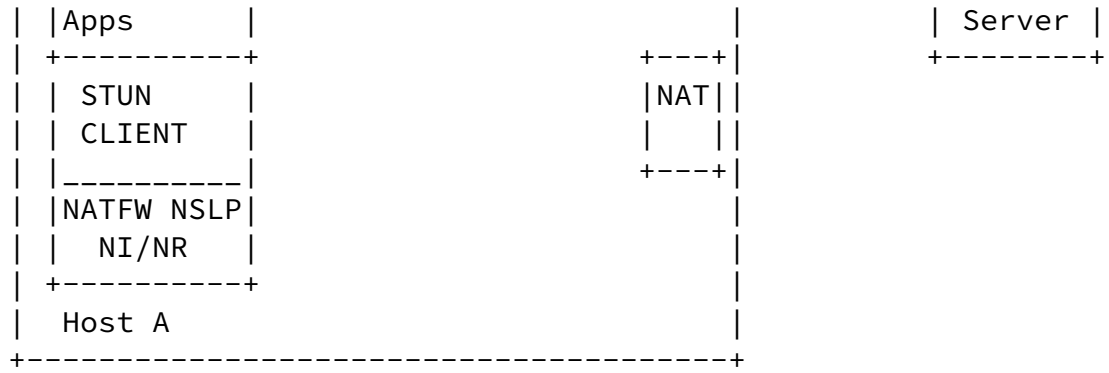

Figure 5

2. The NE could send, through the NSIS unaware NAT, a bind request message towards a network entity hosting a simplified NR function responding to the message with the translated address. The translated address and port would correspond to the translated source address and port of the received NSIS message by the NE. This translated address and port information would only be useful for UDP transported data, and would imply that the NSIS protocol would be sent from the same address and port as the data flows. This behavior implies a change to the protocol operations, since the NTLP does not normally require transport protocol changes for a given NSLP (at least for now); the other implied modification is to support a minimum set of operations on the responding NE hosting the NATFW NSLP. The minimum set of operations would consist of the ability to authenticate the NE, providing the translated address and port, and support of UDP transport (as well as TCP if certificates need to be sent first). This mechanism would only be applicable to cone NATs [7].

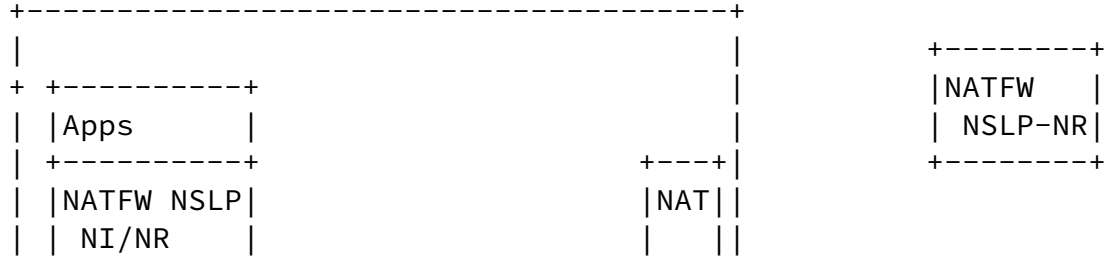

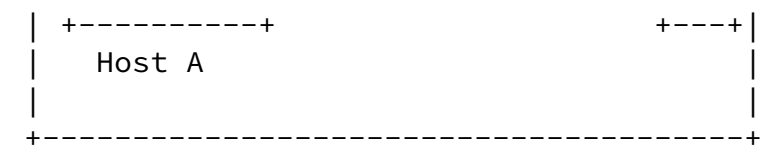
Host A

Figure 6

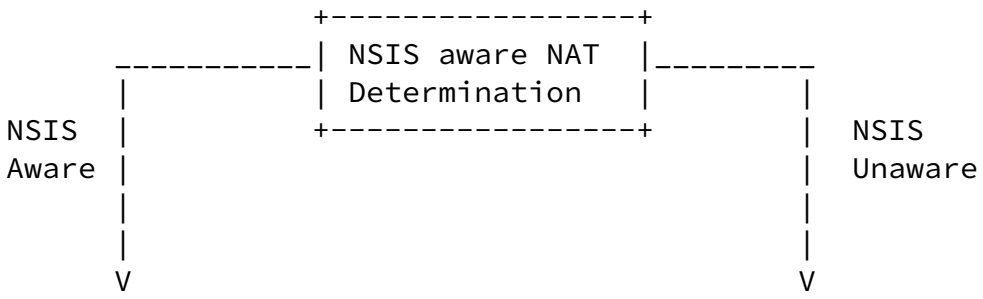
NSIS aware NAT
Determination
NSIS
Aware
NSIS
Unaware

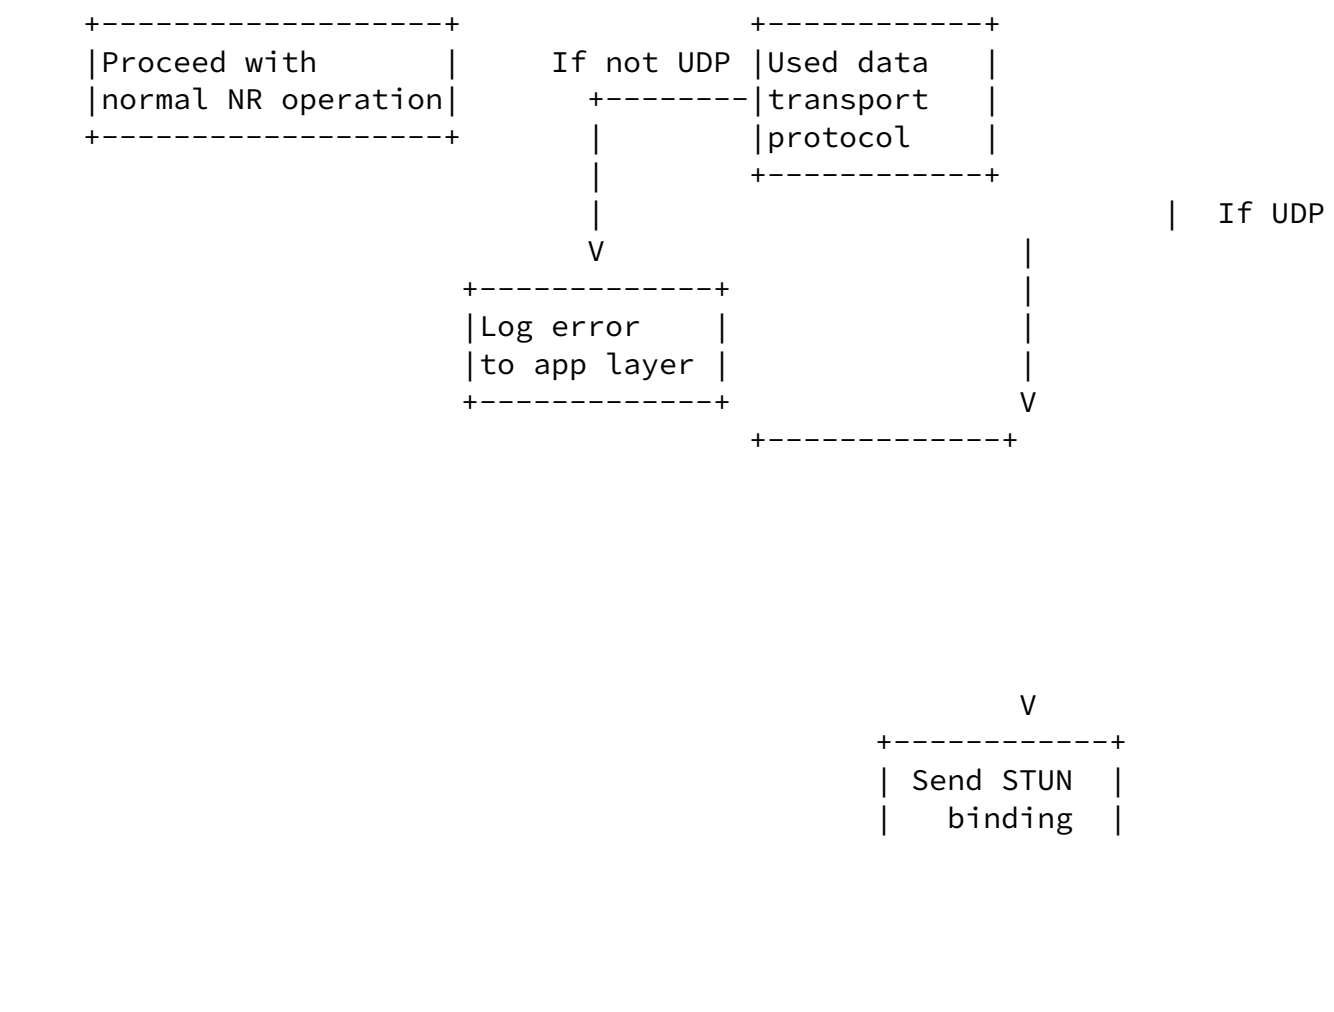
Proceed with
normal NR operation
If not UDP
Used data
transport
protocol
Log error
to app layer
If UDP
Send STUN
binding

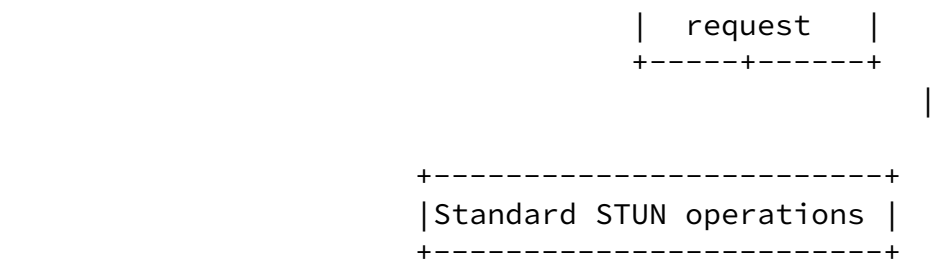
request
Standard STUN operations

Figure 7

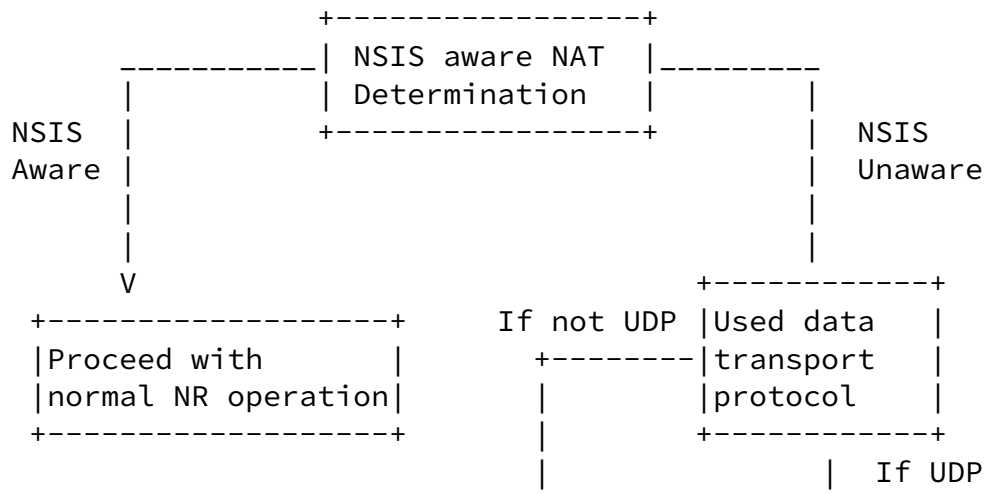
NSIS aware NAT
Determination
NSIS
Aware
NSIS
Unaware
Proceed with
normal NR operation
If not UDP
Used data
transport
protocol
If UDP

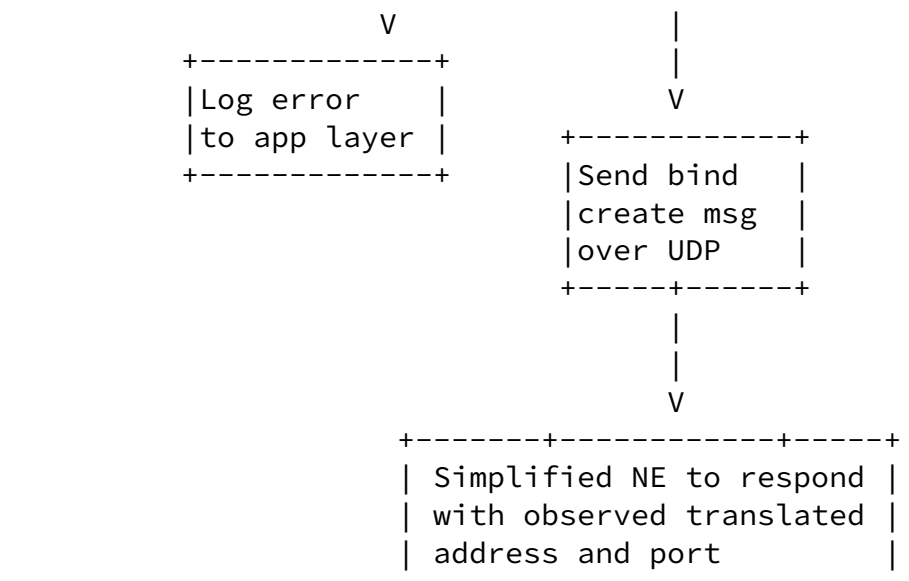
Log error
to app layer
Send bind
create msg
over UDP
Simplified NE to respond
with observed translated
address and port

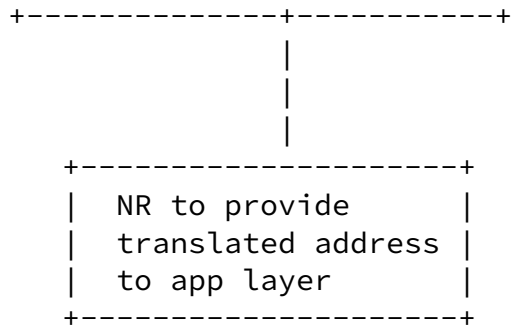

Figure 8

3.2 Analysis of unilateral NSIS signaling

When NSIS NAT/FW signaling will start to be deployed, it is quite possible that an NI sends an NSIS message without having an NR to respond to it. The NATFW NSLP should have the ability to have a mechanism that would allow it to handle this type of deployments. NSIS NATFW NSLP signaling for NAT binds is already local within the trust domain, however this is not the case with firewall signaling that should be end to end.

There are three interesting cases to be analyzed:

1. The local trust domain (from an NI perspective) has at least one NSIS aware Firewall, there is no NR on the far end as well as no NSIS aware NAT or Firewall.

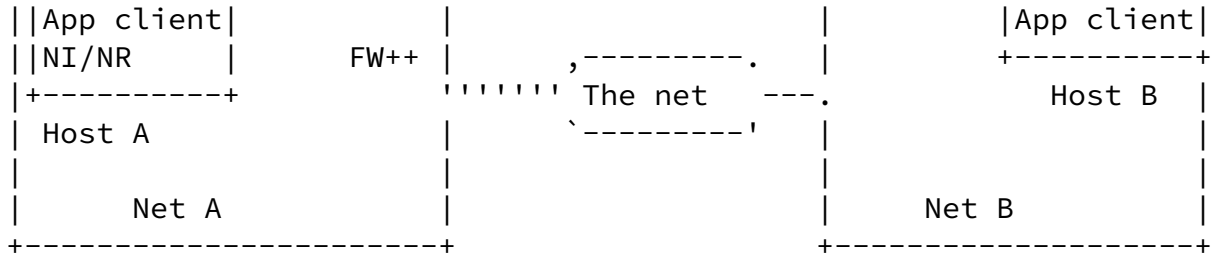

Figure 9

2. The local trust domain has no NSIS aware Firewall, there is no NR at the far end but there is at least an NSIS aware firewall with which the local NI has no direct trust relation (which implies an authorization issue and possibly authentication issues).

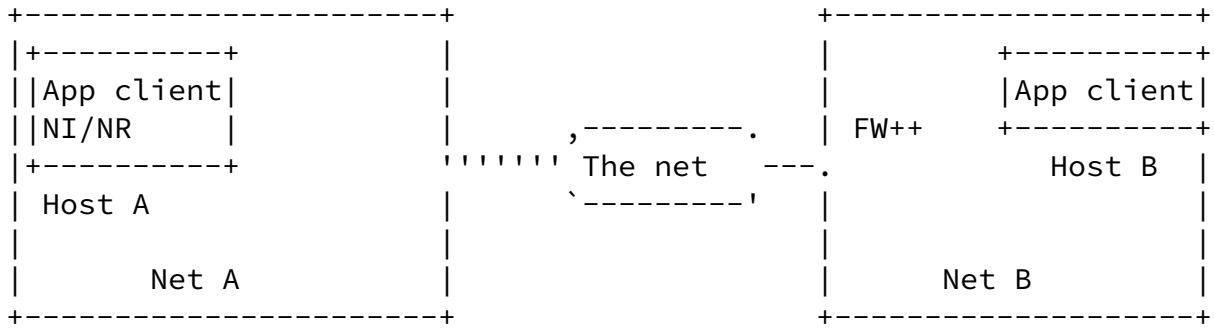

Figure 10

3. The local trust domain (from an NI perspective) has at least one Firewall, there is no NR on the far end but there is at least one NSIS aware Firewall with which the local NI has no direct trust relation (which implies an authorization issue and possibly authentication issues).

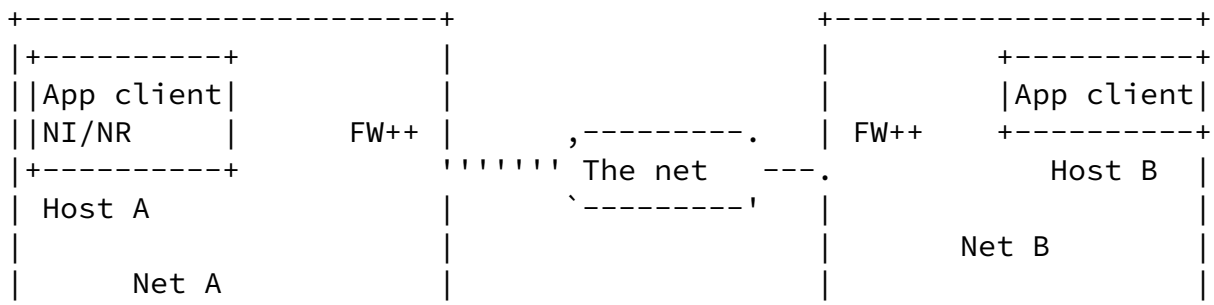

Figure 11

In 1), the NI sends its firewall policy rule creation message, it traverses the first NF (its own firewall) but there is no NR to respond back. If we consider to have a response timer on the last NF being traversed by an NATFW NSLP message then if no response is received to the NSIS message, the last NF will respond back to the NI with a notification of no far end NR response. This will imply that the signaling will be scoped to the last NF on the path that responded back. Using the network deployment shown in Figure 9, the following mode of operation would apply:

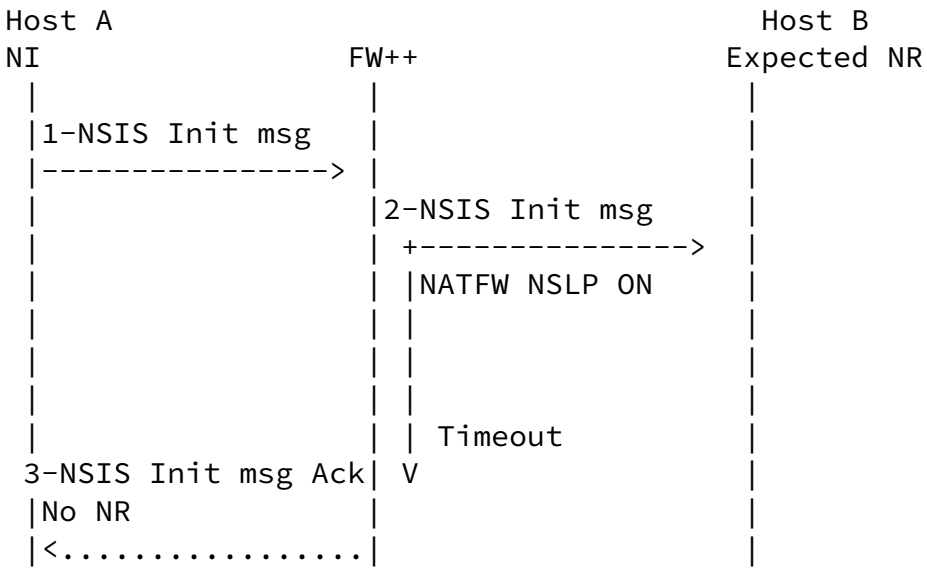

Figure 12

When more than one NSIS aware NAT or Firewall is deployed within the same trust domain, upon determination of a previous NSIS hop, an NSIS aware node will notify the previous NSIS hop of its existence to avoid launching the timer that triggers the sending of an NSIS message back to the NI. The last NF on the path will launch the timer since no valid NSIS neighbor was sent to it.

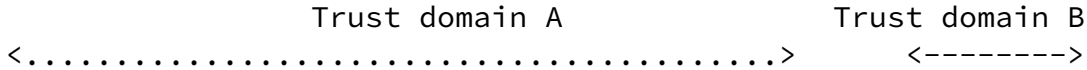

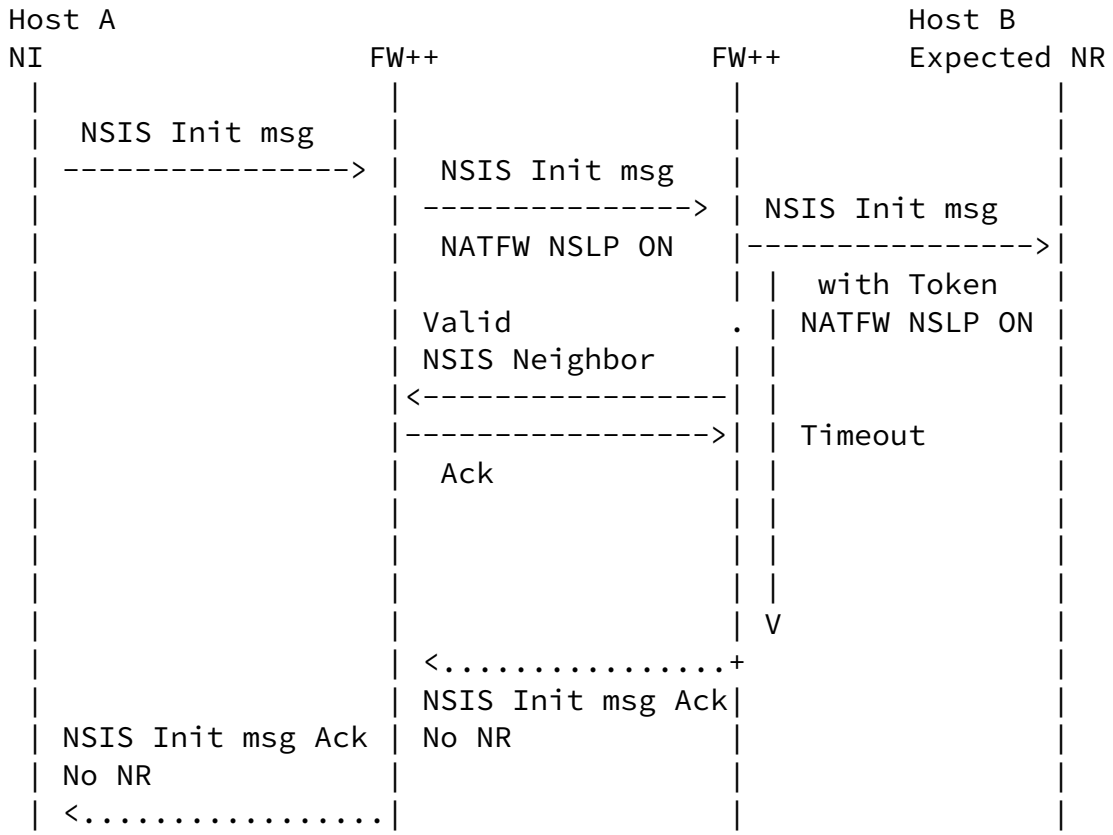

Figure 13

In 2), the NI sends its firewall policy rule creation message, it traverses the FW hosted in Host B's network, but host A is not authorized to install a policy rule unless the policy rule creation is approved by a trusted entity within Net B. Unfortunately Host B was not yet upgraded to support the NATFW NSLP, another entity needs to authorize the policy rule installation.
Potentially a trusted third party already aware of the application session held between Host A and Host B could provide an authorization token to Host A [13], the token would be encapsulated within the NATFW NSLP message and would allow the NSIS aware Firewall in Net B to authorize Host A's requested policy rule to be installed. This approach would obviously require to put in place a mechanism to provide the authorization token to Host A. The token could be requested by the NI and included in the NSLP signaling by default or after receiving an error message from the far end NSIS aware Firewall indicating that authorization data is required. The authorization token would need to be associated with the identity of the NI, associating the authorization token with an IP address is not sufficient, a proper mechanism should be put in place to allow proper authentication of the legal token user. The next revision of the

discussion will cover in more details this aspect.

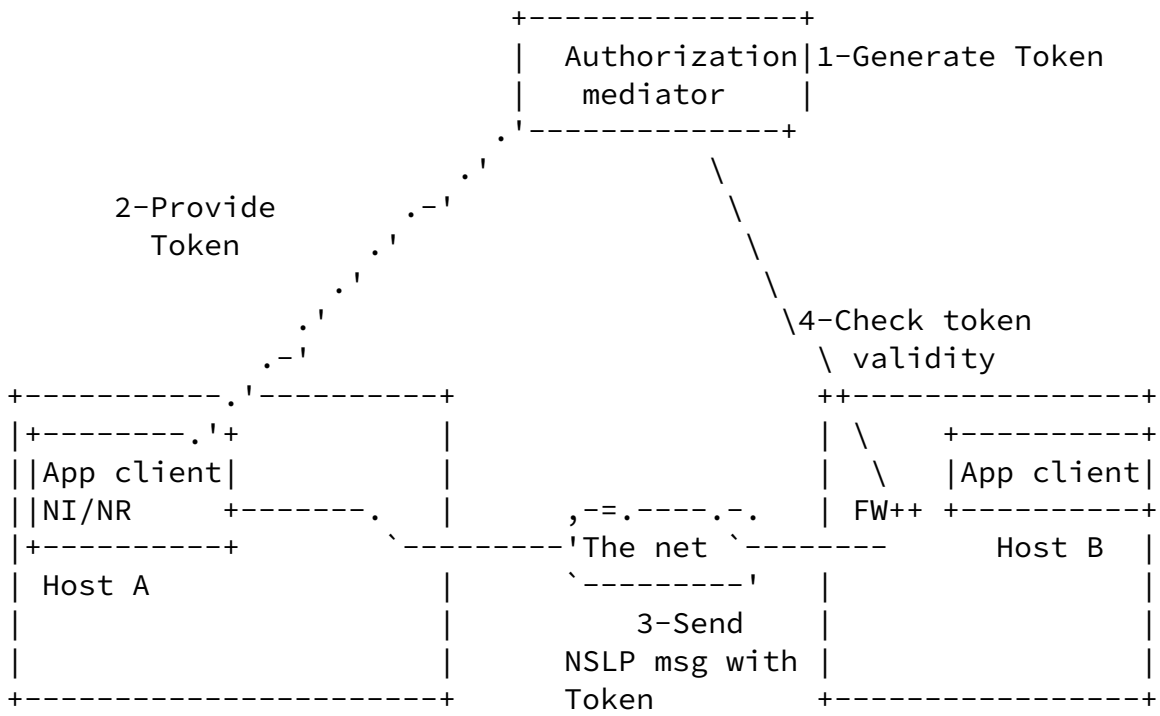

Figure 14

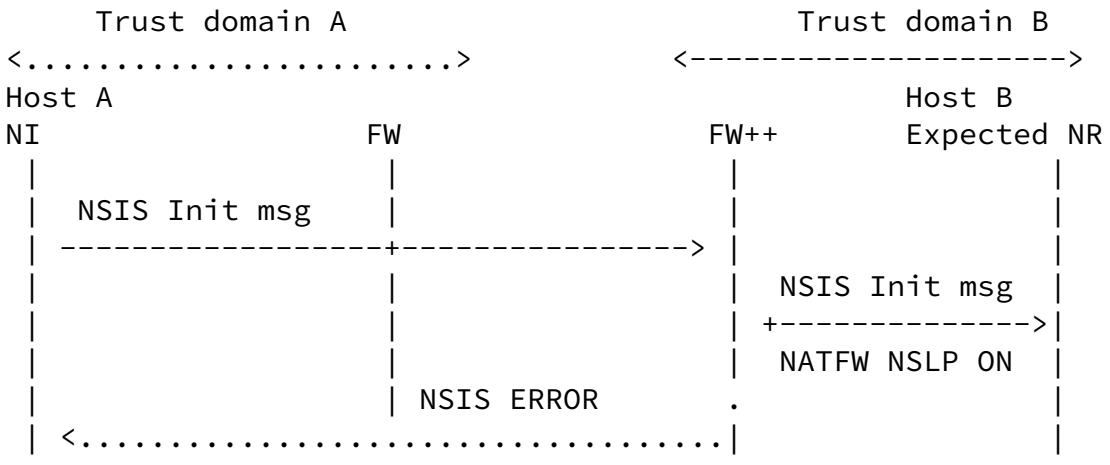

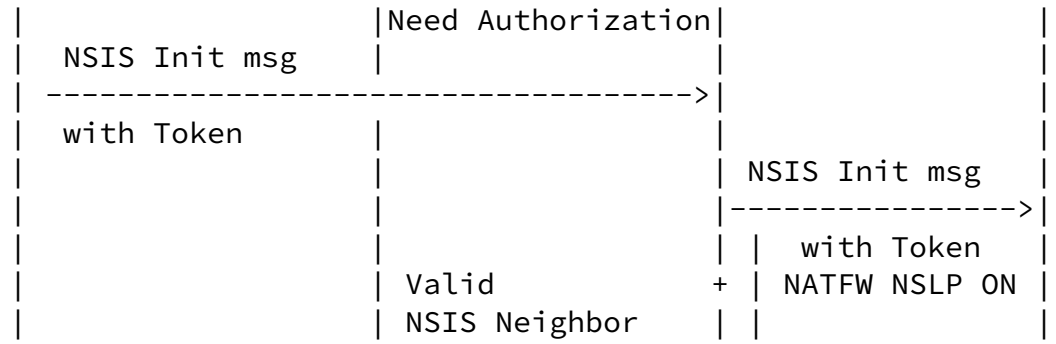

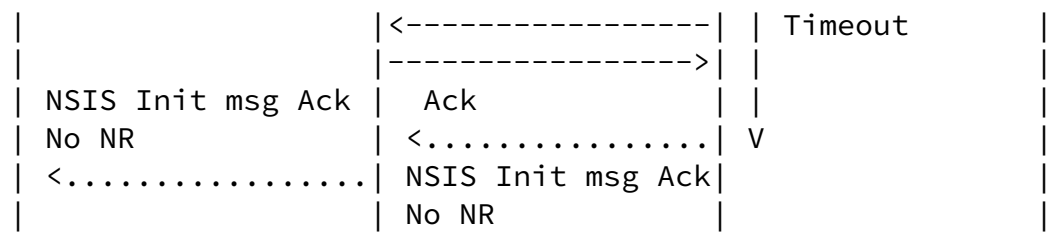

Figure 15

In 3),the NI sends its message to the non-existing NR at host B, it traverses the first NSIS aware Firewall, the policy rule installation succeeds; the message continues to be forwarded until it reaches the 2nd NATFW NSLP aware firewall.
In case no authorization material is provided in the NSLP message, the Firewall will send an error message notifying the NI to send authorization data. If the NI can't send any authorization data, then it will decide to scope the NSIS signaling message to the last NF on which the state installation succeeded.

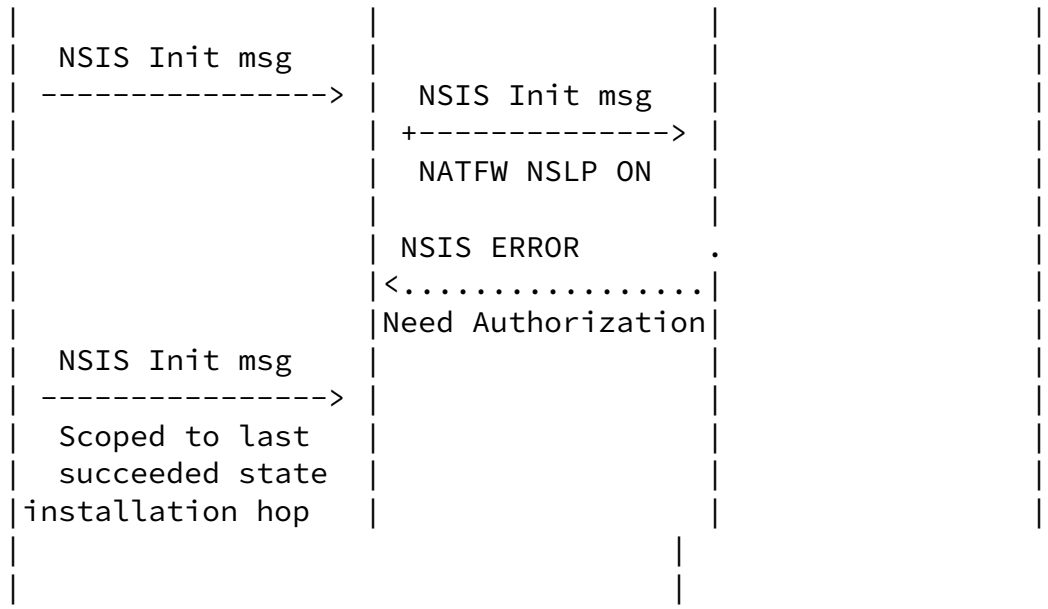

Figure 16

Since the signaling is unilateral (no NI available to do the installation for the other direction), the installed policy rules should be bi-directional. Although bi-directional policy rules could be problematic as discussed in [1], it is the only solution available when no remote NI would be available.

3.3 Co-existence with existing NAT traversal mechanisms

Section 3.1 discussed how a NATFW NE could be used when an NSIS un-aware NAT is deployed within the network infrastructure. This section discusses how the NATFW NSLP could co-exist with interim NAT traversal mechanisms [10]. In Figure 17, a STUN client (Host A) [7], an NE (host B), a host using a Media Proxy [10] and host using a TURN client [11] co-exist in the same network with a NATFW NSLP aware NAT. There are no reasons for the existing mechanisms to be mutually exclusive every host could continue using the existing interim

solutions, meanwhile the unilateral NSIS signaling would be used until both ends support the NSIS NATFW NSLP.

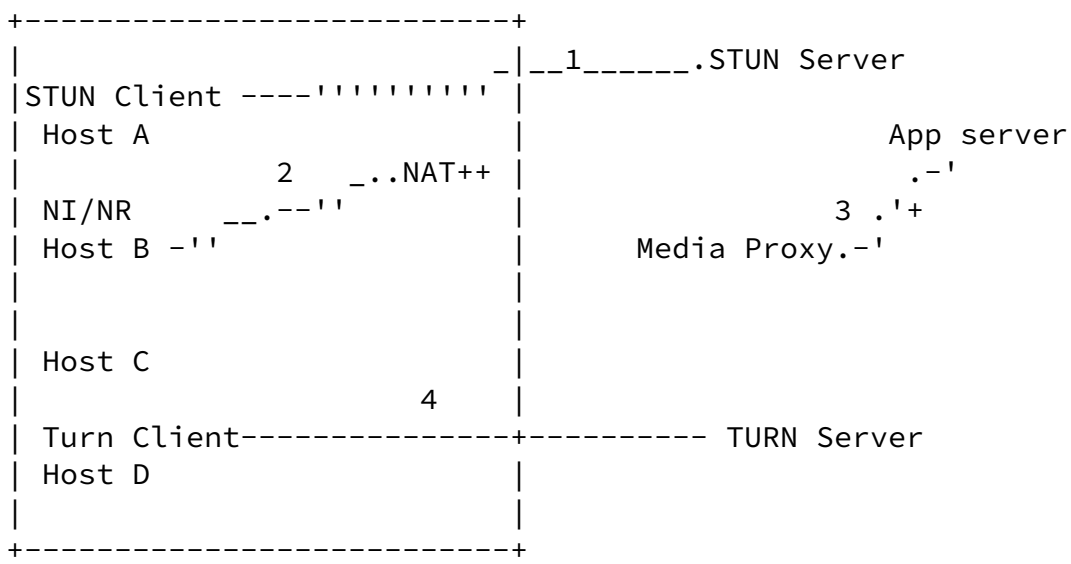

Figure 17

3.4 NSIS protocol traversal of NSIS Unaware Firewalls and NATs

3.4.1 NSIS protocol traversal of NSIS Unaware Firewalls

In case an NSIS unaware firewall is traversed by NSIS messages, NSIS messages should be allowed to go through it, as well as the exchanged data flows between the user application clients. This is not necessarily an obvious task to perform in case the NSIS messages can't be identified by the NSIS unaware firewall. Same applies for the user application data flows.

NSIS message identification should be supported by existing firewalls.
Currently firewalls support flow identification by using the 5 tuple or a sub-set of it. We can not assume that the firewall will support the router alert option [14], hence it should not be the only element of the used identification filter.

User application data flow identification, should be deterministic at a specific address and port range level. This means that the application clients uses a combination of an address and specific transport port range. This combination should be configured on the firewall.

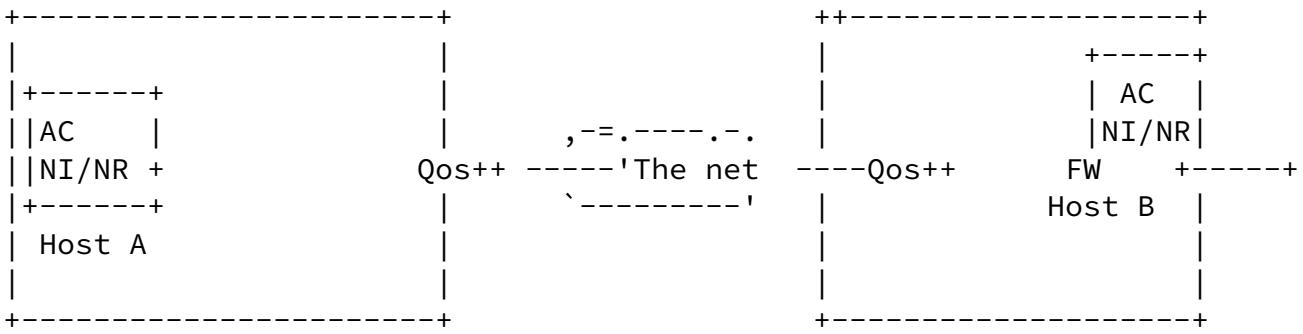

FW: NSIS unaware FW
Qos++: NE with Qos NSLP

Figure 18

3.4.1.1 NSIS protocol traversal of a mix of NSIS Unaware Firewalls and NSIS aware NATs

In case a NAT is deployed on the path and it is NSIS-NATFW, the assigned bind should be consistent with policy rules configured with the NSIS unaware firewall.

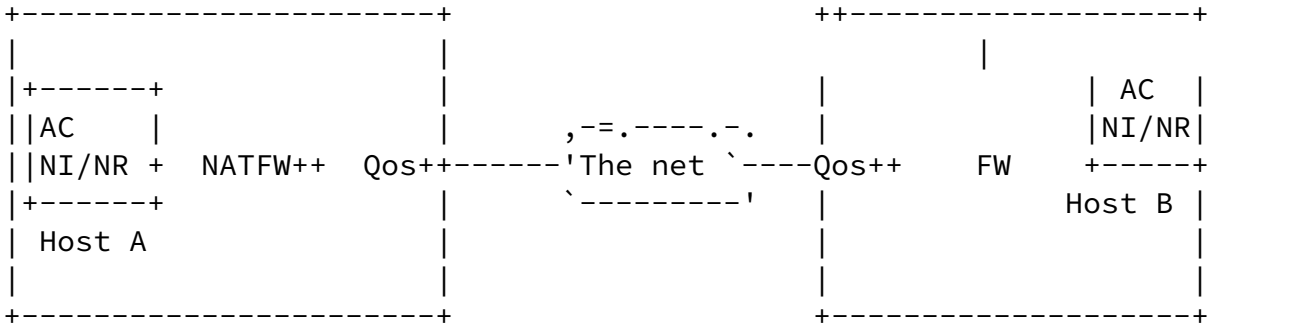

FW: NSIS unaware FW
Qos++: NE with Qos NSLP

NATFW++: NSIS aware NATFW
Policy rules configured on the NSIS unaware FW to allow
specific filters for NSIS signaling and user Application data flows

Figure 19

Even though the deployed FW is not NSIS aware, the application data would still be forwarded if existing interim solutions were used such as a mix of stateless policy rules and flow based states with initial packets sent in the outbound direction (inside to outside a trust domain).

3.4.1.2 NSIS protocol traversal of NSIS Unaware NATs

NATs create an address bind state for flows having well known patterns part of a predefined filter matching expression. In most cases the patterns consist of the protocol number within the IP header and transport port numbers. When a packet flow has different paterns within in its filter matching expression, not all NATs will be able to forward the packets. When several NEs are deployed behind NATs, a mandatory demultiplexing field is required for the NSIS protocol in order to match a source or destination to another source or destination. Prior to the NSIS protocol, NATs had to work with IPSEC before IPSEC UDP encapsulation was used ([4]), the SPI parameter was used as demultiplexing field, but this capability was not native in all NATs. Hence IPSEC had to wait for UDP encapsulation to be be forwarded through consumer market NATs. The learned lesson is that the best approach for the NSIS protocol to be backward compatible with existing NATs, would be to be transported over existing transport protocols and not to be sent as raw IP payload.

4. NATFW NSLP NTLP requirements

The NATFW NSLP transition requires the NTLP to change transport protocol to UDP when the data is transported over UDP, as discussed

in Section 3.1.

If the valid next neighbor determination described in Section 3.2, is applicable to other NSLPs it would potentially make sense to have a part of it incorporated in the NTLP. Further investigation would be required to define what should be done in the NTLP (NSLP independent) and what should be done within the NSLP.

The NATFW NSLP does not have any next NSIS hop failure detection mechanism, the NSLP relies on the the NTLP layer for this capability.

NTLP security requirements: TBD

5. Security Considerations

Section Section 3.2 and [1] discuss the security considerations for the NSIS NATFW NSLP.

6. IANA Considerations

There are no IANA considerations defined in this document.

7. Open Issues

Need to close on: the intra-realm security issues, the editorial issues, linking the authorization token with the NI cryptographic authentification mechanisms, NTLP required NAT handling capabilities.

Normative References

[1] Brunner, M., Stiemerling, M., Martin, M., Tschofenig, H., Schulzrinne, H. and C. Aoun, "A NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", DRAFT draft-ietf-nsis-nslp-natfw-00.txt, October 2003.

Informative References

[2] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[3] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for the Session Initiation Protocol (SIP)", DRAFT draft-rosenberg-sipping-ice-01.txt, June 2003.

[4] A. Huttunen et all, A., "UDP Encapsulation of IPsec Packets", DRAFT draft-camarillo-mmusic-alt-01.txt, Jan 2003.

[5] Camarillo, J. and J. Rosenberg, "The Alternative Semantics for the Session Description Protocol Grouping Framework", DRAFT draft-camarillo-mmusic-alt-01.txt, June 2003.

[6] Rosenberg, J., "The Real Time Transport Protocol (RTP) Denial of Service (Dos) Attack and its Prevention", DRAFT

draft-camarillo-mmusic-alt-01.txt, June 2003.

[7] Rosenberg, J., Weinberger, J., Huitema, C. and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

[8] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

[9] ITU-T SG16, "Packet-based multimedia communications systems", ITU-T H.323, November 2000.

[10] Rosenberg, J., "NAT and Firewall Scenarios and Solutions for SIP", draft-rosenberg-sipping-nat-scenarios-00 (work in progress), November 2001.

[11] Rosenberg, J., "Traversal Using Relay NAT (TURN)", draft-rosenberg-midcom-turn-01 (work in progress), March 2003.

[12] Swale, R., Mart, P., Sijben, P., Brim, S. and M. Shore, "Middlebox Communications (midcom) Protocol Requirements", RFC 3304, August 2002, <reference.RFC.3304.xml>.

[13] Hamer, L-N., Gage, B. and H. Shieh, "Framework for Session Set-up with Media Authorization", RFC 3521, April 2003, <reference.RFC.3521.xml>.

[14] Katz, D., "IP Router Alert Option", RFC 2113, February 1997, <reference.RFC.2113.xml>.

Authors' Addresses

Cedric Aoun
Nortel Networks

France

EMail: cedric.aoun@nortelnetworks.com

Marcus Brunner
Network Laboratories, NEC Europe Ltd.

Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 905 11 29
EMail: brunner@ccrle.nec.de
URI: http://www.brubers.org/marcus

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

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

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

Phone: +49 (0) 6221 905 11 16
EMail: miquel.martin@ccrle.nec.de
URI:

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
Munich 81739
Germany

Phone:
EMail: Hannes.Tschofenig@siemens.com
URI:

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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