NSIS Working Group
Internet-Draft
Expires: January 17, 2005

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

July 19, 2004

NAT/Firewall NSLP Intra-Realm Considerations
draft-aoun-nsis-nslp-natfw-intrarealm-01

Status of this Memo

Copyright Notice

Abstract

This document discusses NAT/FW NSLP issues and solutions for cases where NATFW NSLP NEs within the same addressing realm communicate with each other. The document will serve as input to the NSIS NAT/FW NSLP document to meet these intra-realm communications' requirements.

Table of Contents

1. Introduction

The NSIS NATFW NSLP responder provides the NSIS NATFW NSLP Initiator with its address, in some cases the NSIS responder may have several addresses: one (or several) global scoped address and its locally scoped address(es).

The following question arises: Which address does it provide to the NSIS initiator?

This document will cover the NSIS Responder address selection as well as the impact on applications and the NSIS protocol suite. The document will serve as input to the NSIS WG documents.

2. Terminology

The terminology used in this document is defined in [1].

3. Problem statement

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 a global scoped address over a local scoped address arbitrarily may imply non-optimal routing for communication between NSIS elements hosted within the same addressing realm. For certain NAT/Firewall implementations it is not only a matter of path optimization: if a global scoped address is used to reach an internal host, the NSIS messages may get dropped due to a conflict with the anti-spoofing rules on the NAT/Firewall NE, hence no communication could be established.

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 (for example, additional delay or additional bandwidth consumption in certain parts of the network).

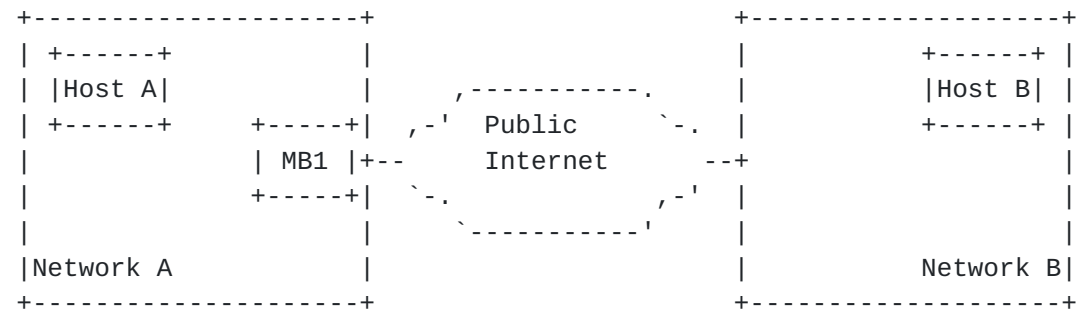

MB: Middle box (NAT or Firewall or a combination)
Host A: An NE hosting NSIS NATFW NSLP NI/NR capabilities
Host B: An NE hosting NSIS NATFW NSLP NI/NR capabilities

Figure 1: Intra-Realm Signaling Scenario

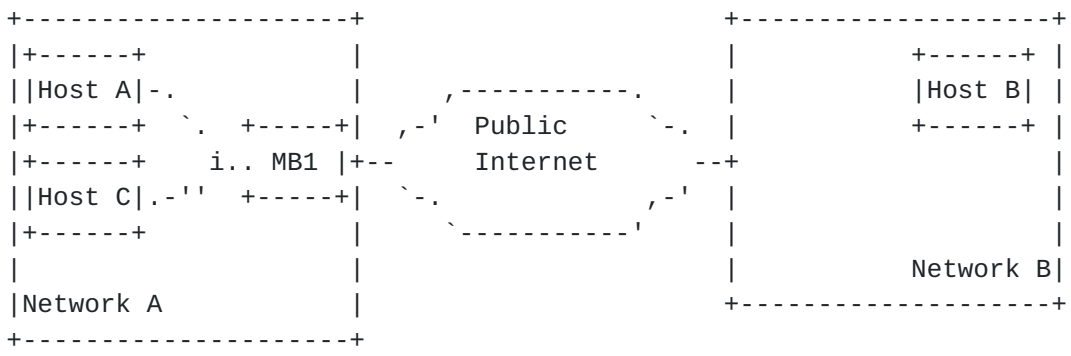

MB: Middle box (NAT or Firewall or a combination)
Host A: An NE hosting NSIS NATFW NSLP NI/NR capabilities
Host B: An NE hosting NSIS NATFW NSLP NI/NR capabilities
Host C: An NE hosting NSIS NATFW NSLP NI/NR capabilities

Figure 2: Intra-Realm Signaling Scenario

Figure 3 and Figure 4, show clearly the impact when the global scoped address is used to signal an NE within the same addressing realm. Figure 3 show how host A will use the Reserve External Address message (defined in [1] to get its global scoped address (i.e. the external address), same happens to host B. Figure 4 shows how the CREATE/RESPONSE messages (defined in [1]) are exchanged to host B/A global addresses and the impact on the data path.

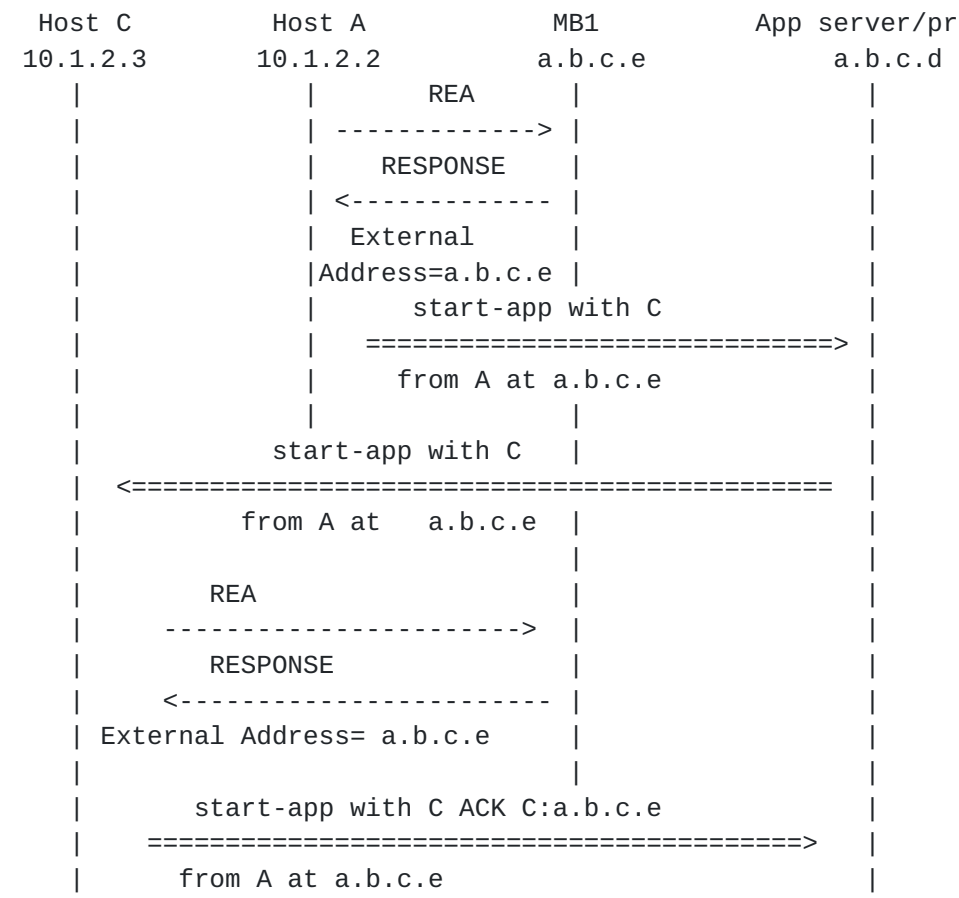

--- NATFW NSLP signaling

=== User Application signaling (SIP, H323, MGCP, MEGACO etc)

Figure 3: Message flow for routing via the Middlebox

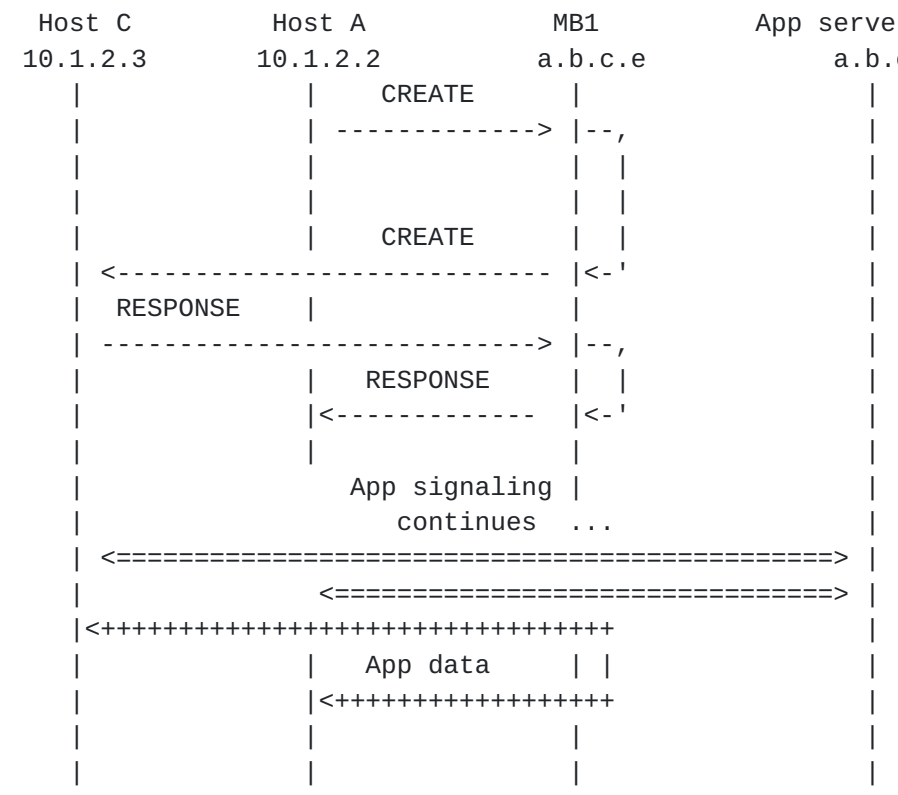

---- NATFW NSLP signaling

==== User Application signaling (SIP, H323, MGCP, MEGACO etc)

Figure 4: NSIS signaling for routing via the Middlebox

4. Solution Approaches

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

1. Use a local policy to decide which address to provide to the NSIS Initiator
2. Make all addresses available to the NSIS Initiator to let it decide which address to use

Approach (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 via user applications. In most cases local decision for address selection requires knowledge about the far end host. This might not always be practical.

An example of such local based address selection is the IPv6 default address selection mechanism (see [6]) 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), this address selection mechanism will not be sufficient since the far end application device is not necessarily known (the resolver client may return the address(es) of the application signaling and not the addresses of the actual application data flow recipient). Hence it can not decide successfully in all cases which address to provide to the NSIS Initiator.

Approach (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. The NSIS Initiator needs to decide which address to use. One possible way for the NSIS Initiator to decide which NSIS Responder address to use is to check which address the NSIS responder will reply back. This security mechanism is typically referred as return-routability check. ICE [7] discusses the usage of approach (2) for SIP User Agents (SIP UA) whereby the SIP UA will probe the far end SIP UA to see from which address it will send a response back to the probe. ICE [7] uses the STUN protocol [14] for the return-routability check. In [9] the reverse routability, provided by the STUN response, is used to check which address to respond to counter RTP Denial of Service attacks. The same reverse routability check could be achieved by the NSIS NATFW NSLP.

In the context of NSIS, the NSIS protocol itself should be used as

the probing mechanism. Consequently, the NSIS Initiator will simultaneously send several NSIS messages towards the NSIS Responder, one message per provided address. Figure 5 and Figure 6 show approach (2) graphically.

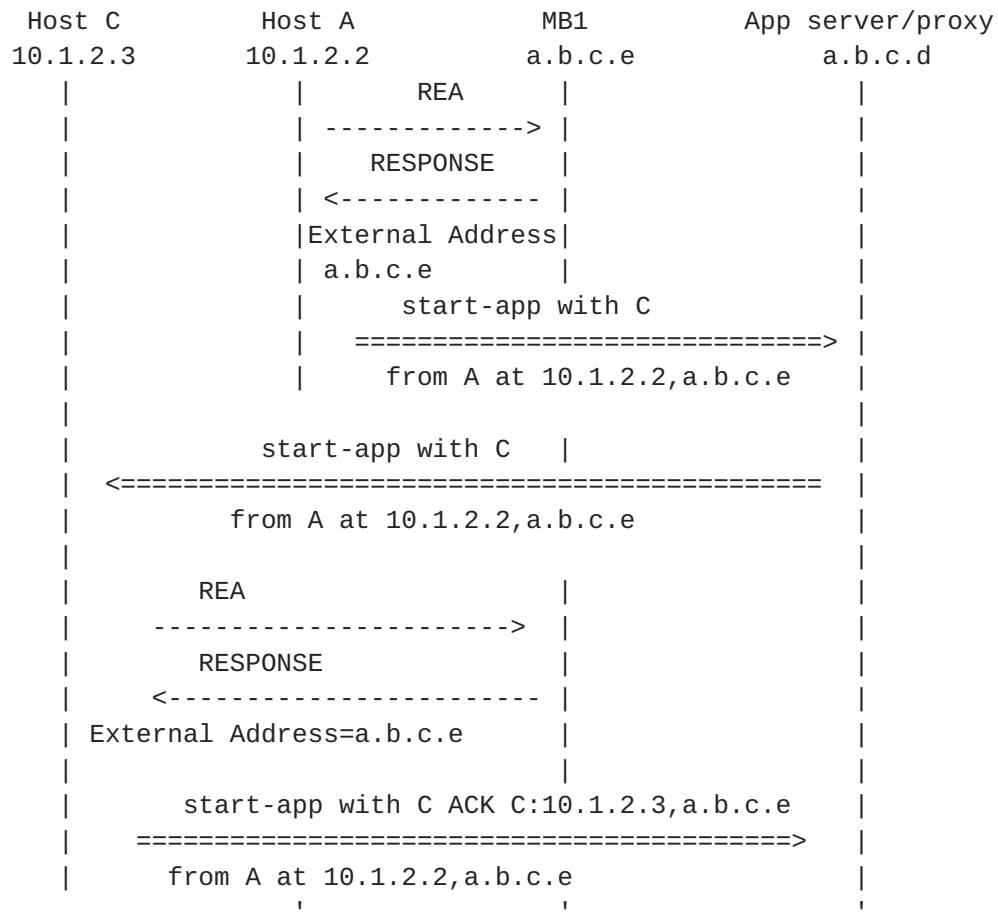

---- NATFW NSLP signaling

==== User Application signaling (SIP, H323, MGCP, MEGACO etc)

Figure 5: Message flow for optimal routing

In Figure 5 Host A (same one as in Figure 2) uses a Reserve External Address (REA) message which is intercepted by the edge NAT (in this case MB1). The edge NAT responds with a RESPONSE message providing the External Address (the global scoped address) to host A. Host A, collects all available addresses where it could receive application data, and then informs its application server that it wishes to

communicate with Host C while receiving data on the global and local scoped address (and ports), 10.1.2.2 and a.b.c.e. The same will happen with Host C, and Host C will be able to respond with the application protocol to Host A about its data recipient addresses (local and global scoped ones). Figure 6 shows how the CREATE messages are simultaneously sent by A to all the returned addresses for C. Host A receives both CREATE ACKs, the local scoped address will therefore be considered as the address to use to send NSIS NATFW messages to C. The problem is that an NSIS host might not be able to distinguish a global scoped address from a local scoped address. To cope with that problem, the following method is proposed:

The NATFW NSLP will have a NAT counter object, inserted in CREATE messages and echoed back within CREATE responses. Every time an NSIS aware NAT is traversed the NAT counter object is incremented. When the NI host receives several RESPONSE messages, it will compare the received NAT counter objects, the NR that responded with the smallest NAT count will be the NR with which the NI will continue communicating with.

In Figure 6, since no NSIS aware NAT (one returned NAT counter was 0) and no Firewalls (NTLP layer [2] confirmed that there was no NATFW NSLP intermediaries) are on the path between host A and host C (when using the local scoped addresses), host A would either send a DELETE message or let the NSIS state expire on its own; the NATFW NSLP protocol is sufficiently robust to handle both. The behavior is left to implementors and network administrators, since it has performance implications.

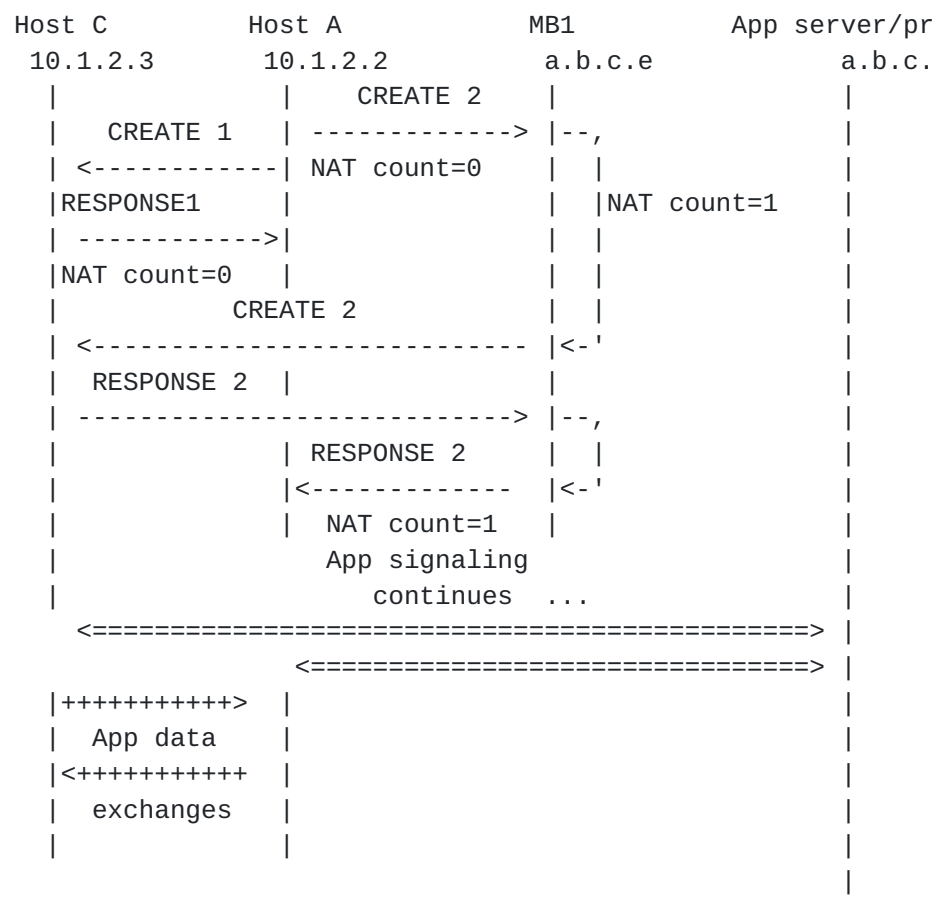

---- NATFW NSLP signaling

==== User Application signaling (SIP, H323, MGCP, MEGACO etc)

Figure 6: NSIS signaling for optimal (intra realm) routing

With regard to the message flow in Figure 6 we can distinguish between the following cases:

- In case that the NSIS Responder and the NSIS Initiator are located within the same addressing realm, the NSIS Responder should receive a response from all the addresses to which an NSIS message was sent. The NSIS Initiator will then choose the address from which RESPONSE message was returned with the smallest NAT counter, as the destination address for messages destined to the NSIS Responder.
- In case that the NSIS Responder and the NSIS Initiator are not located within the same addressing realm, the NSIS Initiator would only receive a response from the valid global scope address. In case 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), Section 5 discusses the related security implications for this behavior. This requires that the NR knows if it is the intended recipient, this knowledge could be provided by the user application which is aware of the application context requiring the establishment of an NSIS session. However this assumption is no longer valid during migration phases where a proxy operation mode is required ([1] and [10]).

5. Security Considerations

Deployments having NRs with local scoped and global scoped address, are subject to the same threats as the ones discussed in [4], [5] and [3] as well as to the potential threat where a 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 requiring the NR to respond back with a challenge response information communicated by the application signaling (assuming that the application was secured). This type of end-to-end security mechanism (irrespectively which degree of security it offers) have an impact on NSIS signaling protocol and on the application layer protocol. If the responding NE is not the application end host then the protocol operation is made more complicated since the NE would have to act on behalf of the application end host. This would be the case when the application end host does not yet support the NATFW NSLP (this is the case of the proxy mode scenario discussed in [1] and [10]).

To avoid the leakage of network topology information, when the CREATE message is leaving the network the network Edge NAT or Edge Firewall supporting the NATFW NSLP will need to remove the NAT count object.

6. Application Signaling Protocol Impacts

The proposed intra realm path optimization proposal requires that an NR provides several data recipient addresses within the application protocol, has obviously a certain impact on the application protocol. In addition the application signaling needs to provide a challenge response allowing the NR to respond back properly. This information either need to be added to the application protocols or existing protocol fields need to be used (preferred way).

Certain applications already provide the ability to advertise several recipient addresses, [8] discusses the impact to SDP [11] 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 [12] and others requiring usage of NSIS with multiple addresses for the NR.In addition [7] proposes the usage of a challenge response parameter within SDP in the context of applications using SIP, a similar approach should be used for other applications.

7. NATFW NSLP required extensions

As shown in Section 4, to solve the intra-realm problem a new object is required for the NATFW NSLP, a NAT count object. It is suggested that this parameter be used in all CREATE messages ([1]).

Another required change for the NATFW NSLP are more of behavioral type:

None Edge-NAT NATs traversed by REA messages: when the Edge NAT responds to REA messages, these NATs will add to the RESPONSE message an External Address object. This new behavior will allow the support of several level of NATs, as shown in Figure 7, in the network while supporting the intra-realm solution (approach 2) discussed in Section 4.

When the Edge NAT responds to REA messages, these NATs will add to the RESPONSE message an External Address object. This new behavior will allow the support of several level of NATs in the network, as shown in Figure 7, Figure 8 and Figure 9, while supporting the intra-realm solution (approach 2) discussed in Section 4.

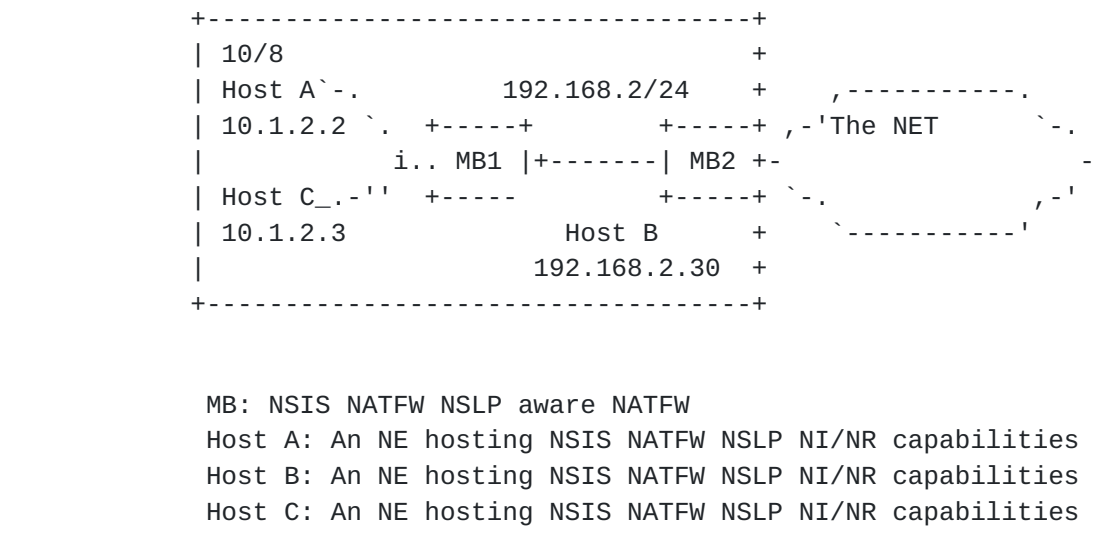

Figure 7: Nested NATs intra-realm communications support

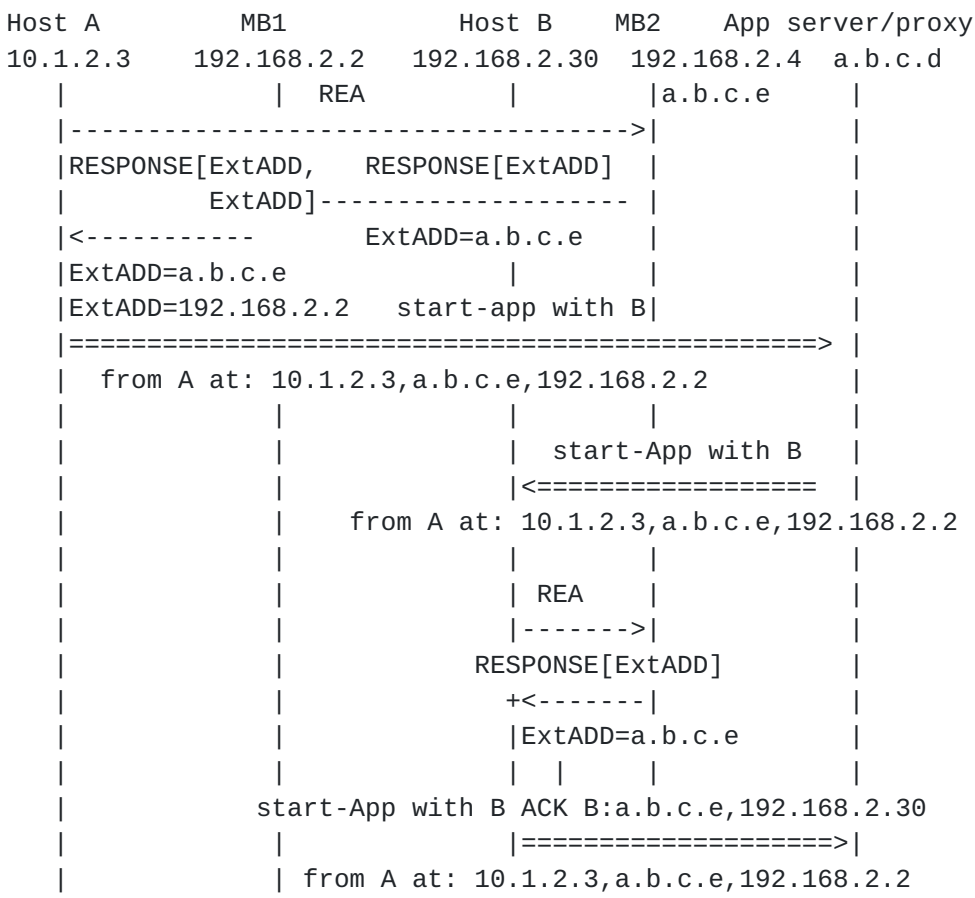

Figure 8: Nested NATs intra-realm message sequences: REA and NR address advertisement

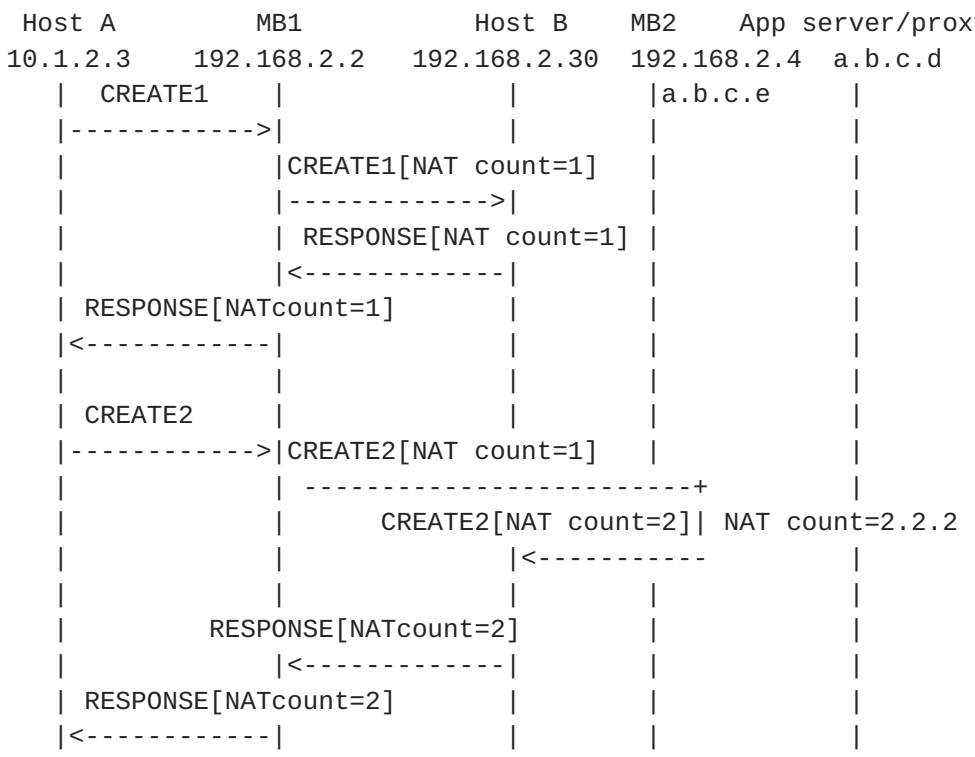

Figure 9: Nested NATs intra-realm message sequences: CREATE and RESPONSE exchanges

8. Perfomance considerations

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 that 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. This assumption may not be always applicable depending on the network topology and network load at the moment of the protocol message exchanges. In case the first response is received from a local scoped address and has succefuly provided the challenge response maerial (Section 5) provided by the application signaling protocol then the address selection process ends, and the NSIS protocol operations continue.

9. IANA Considerations

There are no IANA considerations defined in this document.

10. Open issues

- Get agreement on solving the problem and its associated security issue, is the challenge response sufficient?.
- Get feedback on which parameter is used within certain application protocols (SIP, MEGACO, MGCP, H323) as the challenge response parameter
- Where should the NSIS challenge response be done? within the NTLP or the NSLP? if done in the NSLP several messaging associations would have been established for no reasons, hence it seems more interesting that the challenge response be handled at the NTLP level.

11. Conclusion

Approach (2) provides a reasonable solution to the intra-realm communication problem, while introducing a NATFW NSLP new object to be added within CREATE messages. Although a new threat is added, its mitigation is very similar to other NATFW NSLP threats discussed in [5] hence no additional extensions would be required when this solution is used.

12. Acknowledgments

The idea of using a NAT count object came out of discussions with Georg Kullgren and Kenneth Sundell. Thanks to Elwyn Davies for his comments on the draft.

13. References

13.1 Normative References

[1] Stiemerling, M., Martin, M., Tschofenig, H. and C. Aoun, "A NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", DRAFT draft-ietf-nsis-nslp-natfw-03.txt, July 2004.

[2] Schulzrinne, H. and R. Hancock, "GIMPS: General Internet Messaging Protocol for Signaling", DRAFT draft-ietf-nsis-ntlp-03.txt, November 2004.

[3] Van den Bosch, S., Karagiannis, G. and A. McDonald, "NSLP for Quality-of-Service signaling", DRAFT draft-ietf-nsis-qos-nslp-03.txt, May 2004.

[4] Tschofenig, H. and D. Kroeselberg, "Security Threats for NSIS", DRAFT draft-ietf-nsis-threats-01.txt, January 2003.

[5] Fessi, A., Brunner, M., Stiemerling, M., Thiruvengadam, S., Tschofenig, H. and C. Aoun, "Security Threats for the NAT/Firewall NSLP", DRAFT draft-fessi-nsis-natfw-threats-01.txt, July 2004.

13.2 Informative References

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

[7] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Multimedia Session Establishment Protocols", draft-ietf-mmusic-ice-01 (work in progress), February 2004.

[8] Camarillo, G. and J. Rosenberg, "The Alternative Semantics for the Session Description Protocol Grouping Framework", draft-camarillo-mmusic-alt-02 (work in progress), October 2003.

[9] 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.

[10] Aoun, C., Brunner, M., Stiemerling, M., Martin, M. and H. Tschofenig, "NAT/Firewall NSLP Migration Considerations", DRAFT draft-aoun-nsis-nslp-natfw-migration-01.txt, Februar 2004.

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

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

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

[14] Rosenberg et al, J., "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

Authors' Addresses

Cedric Aoun
Nortel Networks

France

EMail: cedric.aoun@nortelnetworks.com

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

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

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:

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

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:

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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

<