NSIS Working Group
Internet-Draft
Expires: August 15, 2004

M. Martin
M. Brunner
M. Stiemerling
J. Girao
NEC
C. Aoun
Nortel Networks

February 15, 2004

A NSIS NAT/Firewall NSLP Security Infrastructure
draft-martin-nsis-nslp-security-01

Status of this Memo

Copyright Notice

Abstract

This document proposes a security infrastructure for the NAT/FW traversal NSLP of the NSIS protocol. We begin with a description of the problem, followed by the proposed solution, based on public key infrastructure. The document finally deals with examples that clarify the proposed methods.

Table of Contents

1. Introduction 3

2. Problem Description 4
2.1 Message changes along the data path 4
2.2 Points of change along the data path 4

3. Solution requirements 7

4. Proposed solution 8
4.1 Assumptions 8
4.2 Proposed signature scheme 8
4.3 Trust relationship establishment 9

5. Security Considerations 11

6. Conclusions 12

7. Contributors 13

Normative References 14

Informative References 15

Authors' Addresses 15

A. Appendix A: generic pros and cons for digital signatures . . . 17
A.1 Public key availability 17
A.2 Computational cost 17
A.3 Incurred overhead 18

Intellectual Property and Copyright Statements 19

1. Introduction

The NAT/Firewall traversal NSLP for the NSIS protocol is a highly sensitive service. Because of its functionality, it can potentially open paths to networks otherwise protected by NAT/FW infrastructures. It also has the potential to render NAT/FW infrastructures inoperative, closing paths or exhausting the resources of the involved boxes.

For this reason, a tight security scheme has to be devised, to allow both fine and coarse access control. This draft aims at solving this problem by using cryptographic digital signatures to authenticate the peers. Decisions on whether to allow access or not are based on the authenticity of the requesting peer and the security policy configured in the box.

The text or part of it is intended to be integrated into the full NSIS NAT/FW NSLP, but several issues concerning the NTLP and NAT/FW NSLP need to be resolved first.

2. Problem Description

The NAT/FW traversal NSLP provides the following security sensitive services:

1. Firewall traversal: override specifically set access rules, and allow data transfer through firewall devices, both from the inside out and the outside in.

2. NAT traversal: reach machines in a private network, which were not meant to be accessible from the outside without specific setups.

3. Resource allocation: Install packet filters and NAT bindings on a machine that can only allocate a certain number of them. For instance due to the rule engine capacity, the number of policy rules (firewall filter specifications) are limited

Misuse of these services can compromise the network and even render it inoperable. The NSIS-Threats document [1] shows a number of ways in which such services can be exploited. The following sections detail the specific problems that a security mechanism for this NSLP must face.

2.1 Message changes along the data path

The NSIS Create and Reserve messages transport a tuple that specifies what pinhole or NAT binding should be installed on the Middleboxes discovered on the path. This tuple defines the flow to be allowed, based on its source and destination addresses and ports. Still, such data will change upon traversing a NAT. This means also that the tuple transported in the NSIS packet must also change, in order to remain current with the way the flow looks at that point on the data path.

When using some signatures or other security means, this means that any security of the message will be potentially broken, if a NAT is found on the data path. The solution described in Section 4.2 addresses this problem and proposes a possible solution.

2.2 Points of change along the data path

This section will cover the issues related to messages changing as they travel on the data path. Specifically, we will analyze what information changes and at what point, as a prior step to proposing a methodology that keeps as much authentication information as possible.

The scenario in Figure 1 shows a data path between DS and DR. The path is obstructed by two NATs. According to the NAT solution of the NSLP, DR must first signal on the reverse path, towards DS, to allocate a public address in NAT2. Once this is done, DS sends a Create message pointed to that address at NAT2, and NAT2 forwards the packet inside to DR.

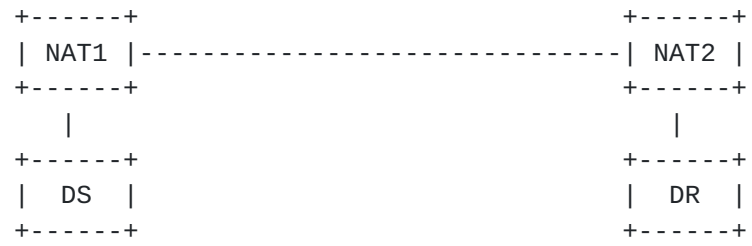

DS: Data Sender
DR: Data Receiver

Figure 1: Message alterations when traversing NATs

The NATs on the way change the source or destination of the data flow, depending on whether they are traversed outwards or inwards, respectively. In the case of twice NATs, both source and destination are changed. In Figure 1, NAT1 will change the data source and NAT2 will change the data destination. A twice NAT would behave as the immediate concatenation of NAT1 and NAT2.

The current Create message defined in the NAT/FW NSLP carries three pieces of information:

- Session ID (ID): unique number that identifies the states installed along the data path for a given flow.
- Source (S): the source address-port pair. This is where the data being signaled by this message will come from.
- Destination (D): Same as Source, but for destination addresses.

With this defined, we can specify the points of change:

1. NAT1: As the create message traverses NAT1, or any other NAT in the outwards direction (private to public) the NSIS part the NAT must change S to the address it allocates on the public side.

2. NAT2: When traversing NAT2, or any other NAT in the inwards direction (public to private), the NSIS part of the NAT must change D to the address listed in its reservations list.

Firewalls never change the messages, and twice NATs will act as the added behavior of NAT1 and NAT2

3. Solution requirements

Given the problems and its exploit possibilities, a solution for security management for the NSIS NAT/FW NSLP will have to cover the following requirements:

1. Middleboxes must be able to verify the authenticity of the requests including its parameters.

2. Message integrity must be guaranteed.

3. Nodes must have the last word in deciding whether they accept a session or not. A node must enable a local decision process of whether to accept a certain action or not. The decision must be based on the identity of the requester, the level of trust into the request, and the parameters of the actions to be taken.

4. The security mechanism must be flexible in the entity it trusts. Trust relationship might exist with NI an/or NR, or somebody else.

5. Trust established end-to-end by other means, for instance, application level signalling, should be re-used on the NAT/FW signalling.

The requirements listed can be met by use of cryptography, namely, digital signatures. Through its use, nodes could be sure of who is making a request, and what request it is, and match it with an access list, or any other arbitrary decision mechanism.

This approach provides great flexibility, as the decision process is entirely based on the configured policies, operating on trustworthy data.

4. Proposed solution

4.1 Assumptions

Unless otherwise noted, the proposed solution assumes:

- The end hosts, that is, the Data Sender (DS) and the Data Receiver (DR) have established a trust relationship before the steps considered here ever happen. For instance, they have exchanged some key material, in a safe manner and they trust these keys to be genuine. How this has happen is out of scope of that document, but could have been done by a application level signalling protocol.

- Either the NTLP of the NSLP packet contains some randomly chosen information to prevent replay attacks. This information is included in all the signatures that will be described in the following sections, even if not explicitly mentioned. The nature of this information, its choice or generation is not covered in the current version of this draft.

4.2 Proposed signature scheme

Out of the three message parts we have in a Create or Reserve Message, S can be changed on outgoing NATs, D can change on incoming NATs, and ID remains always constant.

Since we assume DS and DR know each other, it is necessary to keep the signature of ID by DS, which can safely travel end to end. The first signature we add to the packet is thus: sigDS(ID) which reads "the signature, by DS, of ID".

S only changes when going out a NAT. Only it can be the first signer of the data, since it's it that makes the change. Also, this information should be linked to the sessionID to prevent mixing S, D and ID with each other. for this reason, the second signature, is that of the last one to change S. At the beginning this will be DS, and later, the latest NAT traversed in the private to public direction. This signature is: SigX(ID, S) which reads "the signature, by X, of ID and S", where X is either DS, or the last NAT the packet traversed.

Finally, we are left with the changes of D. Once a packet starts going inwards into NATs, it is highly unlikely that it will go outwards again, so S should become immutable . If we were to follow the approach of the last sender signature, we should sign ID and D, but given this special condition, it makes sense to include ID, D and

S in the signature.

This would, in effect, be the signature of the whole packet, and arises a new question: should we only sign the whole packet on the inwards NATs? In fact, since only NATs are concerned with address changes, and we assume no outwards NATs appear after the first inwards NAT, signing on every NAT is in fact, the same as signing on every inwards NAT. At least, after the first inward NAT is encountered.

Signing on every NAT provides a last hop integrity check on the whole packet. Also, neighboring NSIS nodes are more likely to know (and trust) each other. For these reasons, it seems reasonable to sign the whole message not only on every NAT, but in every Middlebox, even if they don't change the data and break any signature. The third signature becomes thus: SigX(ID, S, D) and reads "signature by X of ID, S and D" where X is either the previous NSIS aware Middlebox that the message traversed, or DS.

Summarizing, the signature scheme proposes 3 signatures:

1. Session signature: sigDS(ID),provides authentication of the signaling initiator for this session ID. It is not changed on path. (Might be a function already implemented in the NTLP).

2. Last Source signature: sigX(ID, S), provides integrity to the source address, as signed by the last Middlebox, which has changed it.

3. Last Hop signature: sigX(ID, S, D) provides hop by hop integrity on the NSLP. It is not clear at the time of writing this draft, if this functionality will be provided by the NSIS NTLP. An whether authentication is part of it as well.

We assume that not all of these signatures will be needed in all implementations. Depending on the type of security needed and on what the transport layer already provides, only a part of the scheme can suffice.

This draft does not specify what mechanisms should be used to produce and encode the signatures, but the use of a standard, such as CMS [2] is recommended.

4.3 Trust relationship establishment

NSIS provides end to end signaling; This circumstance can be used to provide receiver generated trust. The solution would involve the hosts on the data path remembering the session IDs they distrust, and

awaiting the Path Succeeded message, as defined in the NAT/FW NSLP specification, to include the receiver's signature on the session ID, sigDR(ID).

If the Middlebox trusts DR, and DR confirms the remembered Session ID, the integrity of S and D is no longer doubted, since the signaling message actually reached DR, and it was expecting it.

There are two ways a Middlebox can open a pinhole without directly trusting the signature that covers the information: The first is by direct command and the second by implicit trust.

The first assumes that an entity that the Middlebox trusts (such as DR) can prove its trust on the message and the information in it. By his own signed response to the signalling, he implicitly shows trust on the information and may provide additional weight for the Middlebox to make its decision.

In the latter we assume the existence of certificates issued by a third party. These consist of the signature of the trusted third party over the public key of the node being distrusted. This entity can either be a Certifying Authority that exists globally or a node that the Middlebox trusts and accepts the relayed signed certificate.

5. Security Considerations

This entire memo discusses the security implications of using an NSIS NAT/FW NSLP.

6. Conclusions

The proposed method provides a reasonable way to decide whether to honor or not NSIS NAT/FW NSLP requests. Peer to peer trust is expected within the network, and a certain degree of flexibility is also expected in the pinhole source.

The intention is that this approach be flexible and adequate to the different scenarios where trust between the nodes is rare and, when present, should be exploited as much as possible.

Further field research is required to determine if we are actually covering most of the real life scenarios.

7. Contributors

We would like to acknowledge the excellent contributions of Hannes Tschofenig to this draft.

Normative References

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

[2] Housley, R., "Cryptographic Message Syntax (CMS)", RFC 3369, August 2002.

Informative References

[3] Boneh, D., Lynn, B. and H. Schacham, "Short Signatures from the Weil Pairing", 2001.

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

Authors' Addresses

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

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

Joao Girao
Network Laboratories, NEC Europe Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 905 11 17
EMail: girao@netlab.nec.de
URI:

Cedric Aoun
Nortel Networks

France

EMail: cedric.aoun@nortelnetworks.com

Appendix A. Appendix A: generic pros and cons for digital signatures

Digital signatures provide a secure way of transmitting messages: they proof that the author is who he says he is, and that what he is saying has not been altered. In other words, authentication and integrity, both strictly necessary if a node has to take security sensitive actions.

To meet the requirements stipulated and provide the functionalities needed, we propose the use of a public key cryptographic scheme.

This, of course, comes at a cost; we will face three main problems:

- Signature verification requires the public key of the signer. This key has to be known and trusted before a signature can be validated.
- Public key cryptography has a high computational cost, in comparison to other cryptography algorithms and authentication systems.
- Appending signatures to a message implies a significant overhead.
- Any change in the signed message breaks the signature.

A.1 Public key availability

The NAT/FW NSLP is expected to travel over an unknown data path, connecting nodes that don't necessarily know each other. This collides with the necessity of knowing the public key of the sender to verify it's signatures.

The need for a Public Key Infrastructure is common to other protocols , and it is outside the scope of this document to provide a similar service. Methods such as the use of a Certification Authority or locally initiated trust chains might be able to help solve the problem.

A.2 Computational cost

Although the efficiency of public key cryptography algorithms has kept on improving over the years, it is still slow by comparison to symmetric keys. However, the use of symmetric keys and shared secrets proves a scalability problem which makes these schemes unsuitable for this protocol. The exact impact and cost of their usage is thus left for further analysis and tests; These should particularly look at the possibility of a Denial of Service attack through resource depletion.

A.3 Incurred overhead

Appending signatures to a signaling packet significantly increases its size. This is not a new problem, and has been tackled already in algorithms such as those in [3].

The NSLP Create packet payload proposed in [4] has a size of 23 bytes.

For this reason, we must carefully choose the kind of signature we use, to minimize the introduced overhead.

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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