

**Path-coupled NAT/Firewall Signaling Security Problems
draft-tschofenig-nsis-natfw-security-problems-00.txt**

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 10, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This draft raises some of the open issues in dealing with path-coupled NAT/Firewall signaling and tries to raise awareness of the security issues beyond the NSIS working group. These issues need to be addressed in order to proceed with the security architecture for NAT/Firewall signaling.

Table of Contents

- [1. Introduction](#) [3](#)
- [2. High-level Protocol Overview](#) [6](#)
 - [2.1 GIMPS](#) [6](#)
 - [2.2 NAT/Firewall NSLP](#) [9](#)
- [3. Challenges](#) [12](#)
 - [3.1 Security for NAT vs. Firewall Traversal](#) [12](#)
 - [3.2 Which Security Protection at Which Layer?](#) [13](#)
 - [3.3 Different Requirements for Different Parts of the
Network](#) [13](#)
 - [3.4 Mobility, Sender Invariance, and Authorization Problems](#) [14](#)
 - [3.5 Dependencies among QoS, NAT, and Firewall Signaling](#) [15](#)
 - [3.6 End-to-end security](#) [16](#)
 - [3.7 Asymmetry of Security Protocols](#) [18](#)
- [4. Conclusion](#) [20](#)
- [5. Security Considerations](#) [21](#)
- [6. Contributors](#) [22](#)
- [7. Acknowledgements](#) [23](#)
- [8. References](#) [24](#)
 - [8.1 Normative References](#) [24](#)
 - [8.2 Informative References](#) [24](#)
- [Author's Address](#) [26](#)
- [Intellectual Property and Copyright Statements](#) [27](#)

1. Introduction

The NSIS working group is currently working on three protocols: a lower-layer transport mechanism (NTLP) and two signaling applications (NSLPs). The first signaling application deals with QoS signaling and the other one with NAT/Firewall signaling. The lower-layer transport only carries application-specific payloads between a number of NSIS aware nodes along the path in the forward and the backward direction.

The work on path-coupled QoS signaling is the result of efforts on RSVP. The work on path-coupled NAT/Firewall signaling has its origin in the Midcom working group where NAT and Firewall signaling has to cope with network topology problems. The TIST [[refs.tist](#)] proposal led to a BOF and the work was moved to the NSIS working group. The approach taken by the Midcom working group assumes that the NAT/Firewall is known when the signaling protocol starts and that it can be addressed by the entity controlling the middlebox. A strong trust relationship between the middlebox and the entity controlling the middlebox is assumed. In more complex topologies with multiple NATs and Firewalls the order of these devices need to be considered with respect to the data flow traversing them. Additionally, the entity controlling these devices need to know which device will be hit by which data flow. This often requires some knowledge of the topology.

The NSIS approach is different in the sense that this knowledge about the topology is moved to a discovery mechanism, and it becomes the responsibility of the end host to start signaling.

This document is organized as follows: [Section 1](#) describes what problem the NAT/Firewall NSLP is going to solve. [Section 2](#) presents a basic protocol model and a high-level description of the messages transmitted, as suggested in [[I-D.iab-model](#)]. [Section 3](#) lists challenges and open issues.

The approach of path-coupled signaling has some implications:

- o First, some application logic needs to be added to the end host to control the creation of the NAT binding or Firewall pinholing on a per-flow basis. NSIS allows a proxy in the network to be used to perform this operation on behalf of the end host, but some additional security considerations need to be taken into account.
- o Due to the soft-state principle it is necessary to refresh the state continual. Otherwise, the established state would be deleted.
- o Path-coupled signaling for each direction is required to deal with routing asymmetry.

Tschofenig

Expires January 10, 2005

[Page 3]

- o Changes to the routing path (e.g., due to mobility) require periodic re-discovery. The refresh partially addresses this issue, but the effect on the communication in the NAT/Firewall signaling case is more devastating, since data cannot flow to one of the end hosts if packet filters are not established at Firewalls along the path.
- o The impression exists that Firewalls and NATs are commonly used today and in a way that requires path-coupled signaling. With NATs, a number of protocols deal with creating NAT bindings, but mostly without incorporating security mechanisms between the signaling end points and the NAT(s). With Firewalls the situation is quite different, since deployment heavily depends on the scenario and on the environment. Furthermore, with increasing end-to-end encryption and with protocols heavily overloading HTTP and SIP, it is difficult to estimate the future of traditional, packet-filter-based firewalls (and also for stateful packet filtering firewalls).

Security considerations for NAT and Firewall traversal need to be treated separately.

In the past, mainly two approaches have been used for establishing NAT bindings:

These NAT bindings are typically used to allow data traffic from the outside to be forwarded to a specific host on the inside. Dynamic NAT creation can be categorized into one of the following three categories:

- * Implicit creation by outbound-initiated communications whereby the translated address and port is selected from a configured address and port pool.
- * Explicit creation by the Application Layer Gateway(ALG) either via an API call if the NAT and the ALG are co-located or otherwise via a separate protocol.
- * Separate signaling protocols that requests the creation of a NAT binding

An alternative classification is by the trigger for the creation of a NAT binding. In many cases an outbound data packet itself is used to cause the allocation of a NAT binding. Alternatively, a signaling protocol can be used to accomplish the same goal by directly addressing the NAT itself. The Midcom and the NSIS working groups are trying to develop protocols of the latter category.

There is little doubt that a user needs to have sufficient rights (or be authorized) to create packet filters at a Firewall. The Midcom working group addressed this aspect in a convenient way, since trust between the middlebox and the entity controlling the middlebox is

Tschofenig

Expires January 10, 2005

[Page 4]

assumed. In most scenarios these two entities belong to the same administrative domain. Another common 'firewall' uses cryptographic data protection with IPsec. Protocols for establishing IPsec security associations already exists with IKE [[RFC2409](#)], KINK [[I-D.ietf-kink-kink](#)] and IKEv2 [[I-D.ietf-ipsec-ikev2](#)], and hence there is little motivation to focus on these cases.

2. High-level Protocol Overview

NSIS decided to use a two-layer architecture with one lower-layer transport (NTLP) and multiple upper-layer application signaling protocols (NSLPs). The NTLP itself is meaningless if it is not used in conjunction with an upper-layer NSLP. An upper layer protocol, such as the NAT/Firewall NSLP, cannot work without the lower layer. The layering provides a functional split and has to ensure that future applications can be easily integrated without modifying other parts of the protocol.

This two-layer architecture is explained and the relationship between the GIMPS and the NTLP is described in [[I-D.ietf-nsis-fw](#)]. For this document the difference between the GIMPS and the NTLP is not too important.

This section addresses the protocol functionality of the NAT/Firewall NSLP and also the NTLP, since the former depends on the latter.

2.1 GIMPS

GIMPS (see [[I-D.draft-ietf-nsis-ntlp](#)]) establishes installed NTLP "routing" state, which allows signaling messages to be routed backwards along the same path. This is not possible without installed state (or similar mechanisms such as record route) due to routing asymmetry. This state is different from application-specific state (such as QoS reservations).

GIMPS provides two ways to send signaling messages:

- o The first is an RSVP-like signaling style with end-to-end addressed messages. The end-to-end addressed message contains the source and the destination IP addresses of the data flow. The messages are intercepted along the path by NSIS nodes interested in these messages (by using Router Alert Options). The GIMPS specification refers to this as the Datagram mode (D-mode).
- o The second mode (called Connection mode or C-mode) is used when NSIS nodes are directly addressed. This mode assumes that the discovery procedure has already finished (or the address of the receiving node is known via other means) and information about the node is already available.

From the previous description it might be apparent that an important part of the NTLP is its discovery mechanism. Without knowing the next NSIS aware node discovery (either implicit or explicit) is necessary. Providing security for a discovery message is difficult, particularly if standard security protocols should be used. Combining discovery with signaling message delivery is, from a

signaling point of view, faster, but security protection is a lot harder. Currently, the GIMPS specification says that D-Mode does not provide security protection. TLS and IPsec are suggested for C-Mode message protection.

Figure 1 shows the explicit discovery mechanism. Because it is assumed that an NSIS node is unaware of the topology, it is difficult to protect the discovery procedure against all threats. For example, the querying node might not be able to tell whether a responding node is truly the next NSIS node along the path. Furthermore, the querying node might not know the identity of the responding node and hence authentication cannot provide a sufficient guarantee that this node is an authorized NSIS node. Hence, some authorization mechanism has to exist in the routing infrastructure and in the entire system to ensure that nodes along the path act according to their prescribed roles. Such mechanisms might not exist in ad hoc networks. Unauthorized entities located along the path are able to harm NSIS signaling and some NSIS applications, such as NAT/Firewall NSLP and QoS NSLP.

As an example, an adversary along the path not authorized to participate in NSIS signaling observes the NSIS signaling messages and the subsequent data traffic. The adversary is able to learn which IP traffic is allowed to pass the firewall and might learn which QoS treatment a given flow will receive.

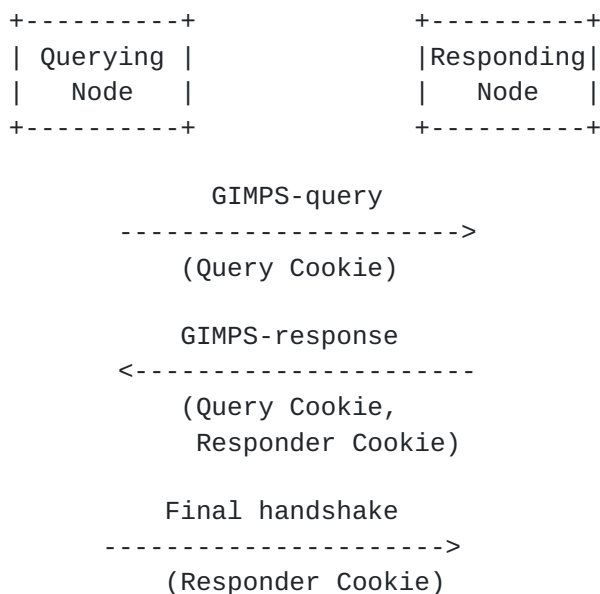


Figure 1: GIMPS discovery mechanism

The discovery mechanism shown in Figure 1 only presents the high-level details. It allows the querying node to learn the IP address of the responding node. Additional functionality, such as discovering NSIS-unaware NATs between these two nodes is under discussion.

The usage of two cookies is somewhat unusual and requires explanation. The responder cookie is used to prevent denial of service attacks in the classical sense as used by other protocols, such as SCTP or IKE. The query cookie has to ensure that an adversary does not redirect the discovery message to another NSIS node. This is guaranteed by providing a cookie by the querying node and by returning the same cookie in the response. This mechanism prevents off-path adversaries from flooding the querying node with GIMPS-responses. The querying node uses this cookie to match a request with a pending response. Furthermore, transmitting the query cookie from the responding node to the querying node after a security association is established between the two ensures that the responder has actually participated in the discovery exchange (i.e., the discovery procedure is bound to the subsequent exchange).

Once the next NSIS node is known, a messaging association can be established between these two nodes using C-mode. The same procedure is repeated again and again for the C-Mode until the last GIMPS node is reached. Note that the NSIS signaling does not necessarily need to terminate at the data flow receiver. The data flow receiver might not be NSIS capable, and some other node along the path (e.g., the access router) might act on his behalf.

The GIMPS protocol itself is only executed between NSIS peers, and they also implement the signaling application. There are no GIMPS nodes along the path that do not contain an upper layer signaling application. This is, both an architectural principle and a technical protocol design simplification. As with other protocols, such as Diameter, security mechanisms at the "lower-layer" prevent certain attacks at both layers between two NSIS nodes and allow standard channel security mechanisms to be used.

2.2 NAT/Firewall NSLP

Currently, the NAT/Firewall NSLP description (see [[I-D.ietf-nsis-nslp-natfw](#)]) mostly analyses the different problems and challenges, describes trust relationships and motivates the different scenarios where the protocol is used.

Unlike other protocols, little information is actually carried in the NSLP beyond the information carried at the NTLP: information about a created NAT binding, as well as lifetime and signaling information (such as protocol headers and error messages). Information about the flow identifier and the session identifier is carried in the NTLP. Currently no additional security payloads at the NSLP layer are specified.

The most valuable part of these information elements is the flow identifier (in most cases a 5-tuple but in some cases not completely known to the sender and/or the receiver at the time of transmitting a message). As an example, a data sender might indicate which source port, protocol type and source IP address has to be used, but it cannot know the public IP address, of the NAT binding yet since it is up to the protocol execution to establish and learn this NAT binding.

It is useful to distinguish between two signaling modes:

The first mode (CREATE) is the traditional way of creating a NAT binding by sending a message from the data sender along the path to the data receiver. Figure 2 shows a message exchange for this signaling mode.

The second mode (RESERVE) is used when a data receiver is behind a NAT and wants to establish a NAT binding to allow incoming data traffic. Figure 3 shows this mode. It was necessary to introduce this mode, because path-coupled signaling in the traditional sense is not immediately applicable.

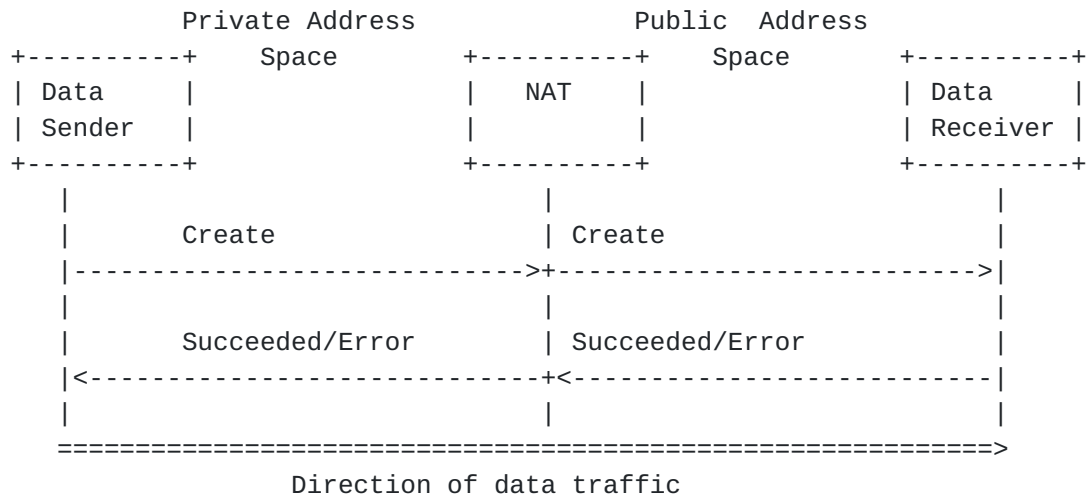


Figure 2: CREATE Mode

With the CREATE mode shown in Figure 2 the data sender (which happens to be the NSIS initiator in this case) sends a message to request a NAT binding to be created. The message is targeted to the data receiver (or even to any node in the Internet), which returns a success or failure message. The data sender learns about the new NAT binding, as a consequence.

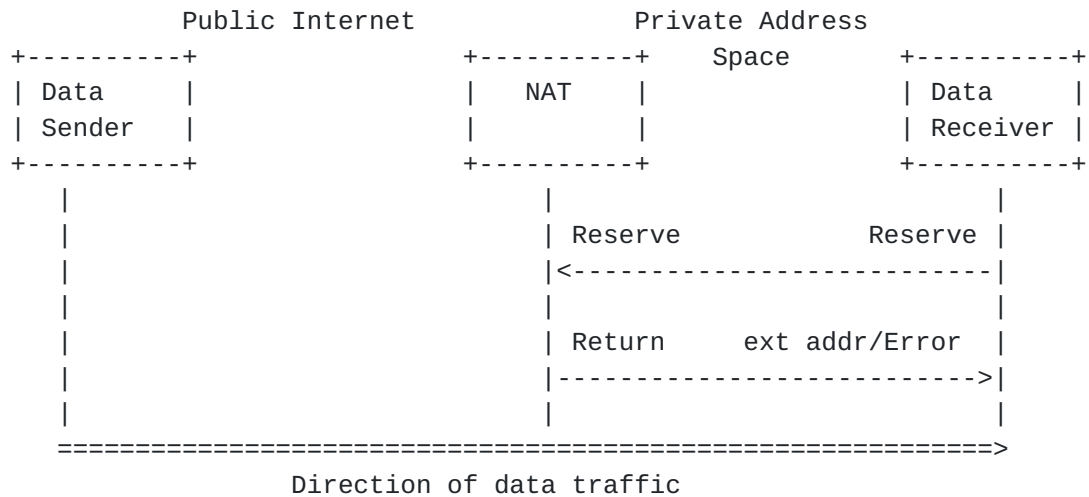


Figure 3: RESERVE Mode

With the RESERVE mode shown in Figure 3 the data receiver behind a NAT creates a NAT binding. This allows data traffic from a node on the Internet to be received. Please note that the RESERVE message is sent in the opposite direction of the data traffic. The RESERVE mode is, in some sense, not path-coupled, since the data receiver starts

signaling but, on the other hand, the data sender will send the data traffic to the IP address (and port) allocated at the NAT.

It should be noted that in [[I-D.ietf-nsis-fw](#)] the RESERVE mode currently requires an additional CREATE message from the data sender to the NAT to activate the binding. This issue is still in discussion.

3. Challenges

This section highlights some of the challenges discovered. Further details can be found in NAT/Firewall NSLP [[I-D.ietf-nsis-nslp-natfw](#)].

3.1 Security for NAT vs. Firewall Traversal

As we tried to motivate in [Section 1](#), the creation of NAT bindings is security sensitive but not to the same degree as firewall traversal. Existing proposals for NAT traversal typically do not use a signaling protocol. Instead regular data traffic from the internal to the external network is used. This is also true for IPv4/IPv6 transition mechanisms in case of automatic tunneling. A typical threat against a NAT device is flooding by an adversary that allocates a large number of NAT bindings. If the dynamically allocated NAT bindings are selected from a limited pool of available bindings (in particular if a NAT instead of a NAPT is used) then this might be a real threat. For a NAPT this threat does not seem to be dangerous enough to require special purpose signaling protocols. As a minor note, STUN [[RFC3489](#)] and TURN [[I-D.rosenberg-midcom-turn](#)] are signaling protocols, but they do not provide additional security for the NAT device when allocating NAT bindings.

If security should be provided for creating NAT bindings, then authentication might be useful in cases of misuse (e.g., allocation of too many NAT bindings). More interesting is, however, authorization. In most networks today every node is automatically authorized to create NAT bindings. To support mobility it is possible either to allocate a new NAT binding (approach used today) or to update the state. Updating a NAT binding is security sensitive, since an adversary can modify an existing NAT binding in order to redirect traffic to a third-party victim, to the adversary itself, or even to a black hole. Flooding a third-party entity might be particularly dangerous if the data sender is streaming a large amount of data (possibly over a wireless interface).

Since a NAT binding has a life-time, it is necessary to refresh it continually. This mechanism provides a self-healing property, since a new data packet (or a new signaling message) causes either the creation of a new binding or the refresh of the old one.

In contrast, firewall pinholing is more security sensitive. Creating or deleting packet filters might easily violate the security policy of a network and might allow an adversary to mount a number of attacks. Only authorized entities are typically allowed to modify packet filters. This requires proper authorization. Authentication will also most likely be required, since typically the authenticated identity is used for computing the authorization decision. As

(communication within an administrative domain), inter-domain communication (communication between administrative domains) and finally the communication with the end hosts and the attached network. NSIS, unlike RSVP, does not necessarily need to be executed between the true data sender and data receiver. It is possible to use NSIS within a single administrative domain only. The impact on security of using NSIS in such diverse environments is that different security protocols, not just one, need to be supported. Some architectures use Kerberos, others rely on special authentication and key exchange protocols, and again others rely on public-key-based mechanisms.

It is highly desirable to provide some flexibility for the authentication and key exchange protocol.

For some NSLPs, such as a quality of service signaling protocol, it is desirable to execute the authorization procedure at an entity where the user is known (typically its home network). This typically implies that the authentication and key exchange protocol is also terminated at the same entity.

3.4 Mobility, Sender Invariance, and Authorization Problems

The NAT/Firewall NSLP establishes state at possibly several entities between the NSIS initiator and the NSIS responder. Providing authentication of the signaling initiator to each individual node along the path might be possible but not particularly useful, since the initiator is most likely unknown to some middlebox along the path. Hence, authentication per se does not solve the security problem.

If authentication is only provided to some entities along the path (most likely to the neighboring NSIS nodes), then information about the initiator of the session is known to some NSIS nodes (except for the session identifier, which does not change along the path and over the lifetime of a session). Now, with the introduction of mobility it might be possible that intermediate NSIS nodes need some assurance that a particular sender is the owner of a session. No other entity should be allowed to modify state since this would allow certain attacks. In some respect this issue is similar to the authorization property in Mobile IP where the mobility binding needs to be protected against unauthorized modifications.

It seems that the property of "sender Invariance" is required in this case: "A party is assured that the source of the communication has remained the same as the one that started the communication, although the actual identity of the source is not important to the recipient."

The interactions among security, mobility, session ownership, and authorization are subject to ongoing discussions in [[I-D.manyfolks-signaling-protocol-mobility](#)].

3.5 Dependencies among QoS, NAT, and Firewall Signaling

Routing asymmetry has to be considered for firewalls but is not applicable to NAT-only signaling. In the presence of NATs, we are always sure that the forward path and the backward path are same with regard to the NAT boxes, since the NAT forces the IP packets to flow through these devices. But, in the presence of firewalls, the forward and the backward routes may be different. A solution needs to focus on the more difficult case where the routes are different. In the forward direction some rules are established in the traversed firewalls. In the reverse direction, if a different route is taken, the packets might be blocked by some other firewall.

It is important to study the relationship between NSIS signaling and other application protocols (such as SIP) and also between different NSIS signaling applications themselves. Different NATs and firewalls can be found along the path, and the worst case needs to be assumed. As we argue in NAT-FW [[I-D.ietf-nsis-nslp-natfw](#)], it is always possible with mobility that an end host finds itself located behind a NAT. Before NSIS can start, the NSIS initiator needs to know the destination IP address, since this is an integral part of path-coupled signaling. This might, however, already assume some application layer signaling exchange. The IP address information exchanged during this exchange might, however, be wrong due to the presence of a NAT. In some scenarios (e.g., receiver behind a NAT) NSIS signaling might need to start beforehand. With NSIS QoS signaling it is also necessary to avoid breaking this type of signaling application. The NSIS NAT/Firewall NSLP does not aim to learn topology information but rather to create NAT bindings and firewall pinholes and to make information about the NAT binding available to the end host. In a short memo on NAT handling in NSIS (see [[nat-memo](#)]) we argue that it is necessary to incorporate a mechanisms for learning the presence of NSIS unaware nodes into the GIMPS discovery procedure. Additionally, it is necessary to modify the flow identifier of the QoS signaling message at the corresponding NAT device along the path to reflect the address change. An NSIS initiator should not need to know where this address translation takes place; this would require topology information. Providing the necessary flow identifier modification in addition to the installation of a QoS reservation would be useful.

One question is therefore how much NAT handling needs to be incorporated into the NTLP and into every NSLP to support proper signaling behavior. If NAT handling is added to the QoS signaling,

then it automatically inherits the authorization applied to QoS signaling. Should this procedure be extended to Firewall signaling? Does the right to make a QoS reservation imply the right to traverse the firewall?

3.6 End-to-end security

Securing the communication between neighboring NAT/FW NSLPs with a chain of trust is a convenient assumption that allows simplified signaling message processing. However, it might not always be applicable, especially between two arbitrary networks. We assume that NATs and firewalls are typically located in the access networks and are typically not found in the core network. Hence, two observations follow:

- o The two access networks (and the firewalls/NATs in these networks) do not trust each other or they might not even know of each other.
- o The end host might have a trust relationship with the local access network that allows it to create firewall pinholes. However, it cannot be assumed that the end host of one network is able to create packet filters at another network. In the example of Figure 5 Host A is not authorized to create pinholes at Middlebox 2. A trust relationship exists only between Host B and Middlebox 2. This scenario represents a scenario in which two employees of two companies want to communicate through their corporate network firewalls. Company B trusts its own employees but not employees of company A.

In [[I-D.ietf-nsis-nslp-natfw](#)] we describe three possible approaches to tackle this problem. None of these three approaches is without drawbacks. We have chosen the one approach that assumes the signaling message is sent end-to-end and each end host contributes its part to the authorization decision. Furthermore, we have to assume that the NSIS signaling message is allowed to bypass the firewall (without installing a packet filter at this stage of the protocol) to reach the other end host.

Based on Figure 5 the authorization steps can be described as follows: Host A starts with the NSIS signaling message exchange and has to authenticate itself to Middlebox 1. Middlebox 1 authorizes Host A to create a pinhole based on the existing trust relationship. Then the signaling message is forwarded along the path and intercepted by Middlebox 2. No trust relationship between Middleboxes 1 and 2 or between Host A and Middlebox 2 exists. Hence, Middlebox 2 does not authorize the pinhole creation. For more restrictive firewalls an error message is returned to Host A, but in our scenario the NSIS signaling message is forwarded to Host B (the final destination of the signaling message exchange). Host B

verifies that the signaling message is provided from a trusted device, might already expect an incoming message based on some application layer signaling exchange with Host A, and returns a response message. Middlebox 2 authorizes Host B's request for pinhole creation due to the existing trust relationship. The message travels back to Host A, which receives a positive confirmation that the signaling message exchange is successful. Host A can start transmitting data packets.

Please note that if Middlebox 2 is actually a NAT (instead of a firewall) then the scenario for a receiver behind a NAT is applicable, which allows Host B to perform signaling locally without the above-described complications. This is one of the other solutions described in [[I-D.ietf-nsis-nslp-natfw](#)].

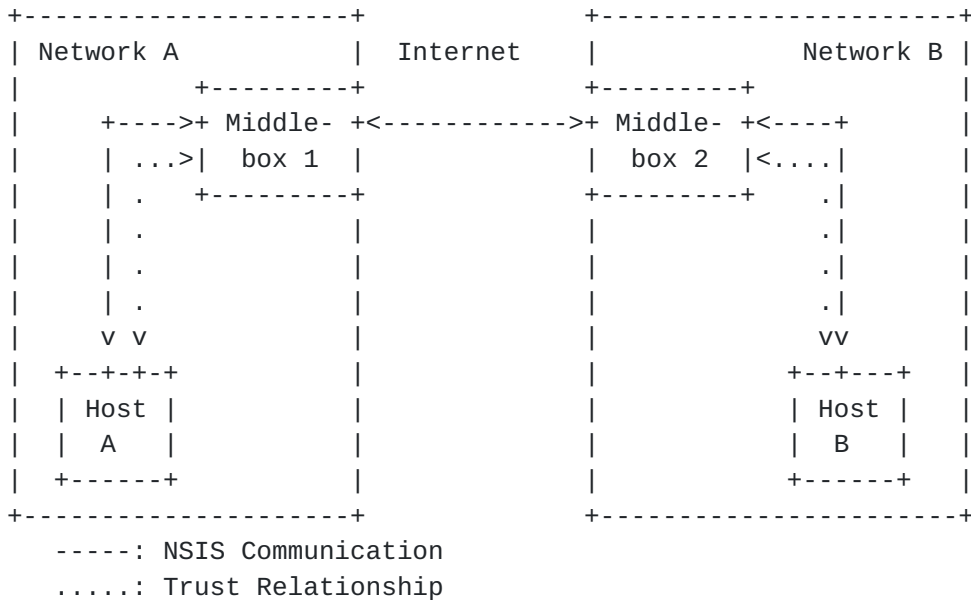


Figure 5: Authorization Problems

Without a proper binding of the NSIS to application signaling, Host B might suddenly receive an NSIS signaling message that indicates a firewall pinhole has to be created. Host B does not know which end host requested this NAT binding nor for which reason. Hence it might be reasonable to think about providing end-to-end security (via a binding between NSIS signaling and the application signaling) as an option to provide the receiving node stronger guarantees about the entity requesting certain actions. It has to be noted that other NSIS signaling scenarios in which intermediate nodes start (or terminate) NSIS signaling on behalf of the end hosts might be much more difficult to deploy along with end-to-end security.

The main questions raised by this section are whether the described observations are correct and whether it seems possible to make the assumption that an NSIS signaling message be allowed to traverse the packet filter firewall. Furthermore, it needs to be studied whether end-to-end security provides better properties.

It is worth noting that the observation such a need of this application layer signaling to NSIS signaling binding is raised in [[I-D.aoun-nsis-nslp-natfw-migration](#)].

3.7 Asymmetry of Security Protocols

Some security protocols operate asymmetrically, which leads to unpleasant consequences for the NSIS protocol suite. The Transport Layer Security protocol (TLS) [[RFC2246](#)], IKEv2 [[I-D.ietf-ipsec-ikev2](#)], and also custom security protocols (such as those provided with RSVP Identity Representation and, for example, Kerberos [[RFC3182](#)]). NAT/Firewall NSLP signaling messages travel along a path between the NSIS initiator and the NSIS responder containing a number of entities that act in different roles. Due to the routing asymmetry it is necessary to start signaling from both end hosts (one signaling exchange for each data flow direction). In IKEv2 the security properties of the initiator and the responder are different with respect to denial of service protection and support for the Extensible Authentication Protocol (EAP) [[I-D.ietf-eap-rfc2284bis](#)]. This is also the case if an end host wants run TLS with unilateral authentication (NSIS entity in the network to the end host) with upper layer client-side authentication. This type of exchange might be typical for QoS signaling, since authorization has to be executed at entities other than those executing the security protocol. From a deployment point of view it is simpler to have public key based authentication of the network to the user than to support a client-side PKI. Such a client-side PKI is, however, necessary when the roles are reversed. Figure 6 shows this problem graphically. Unfortunately, TLS cannot reverse its roles and cannot reuse the session cache for the reverse direction. This problem was also observed in the context of SIP, where [[I-D.ietf-sip-connect-reuse](#)] provides a solution to reuse an established TCP or TLS connection that was established based on a SIP REGISTER before a SIP INVITE (or similar message) is used. NSIS, however, has no provision to support separate 'registration' and a 'end-to-end' signaling message exchanges due to the path-coupled property.

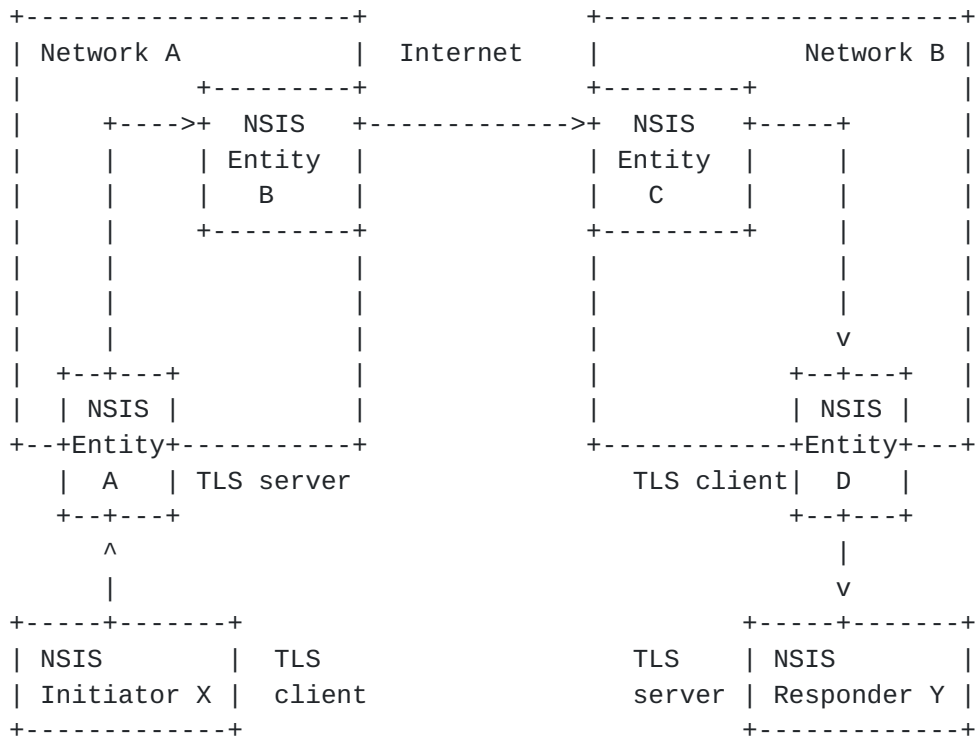


Figure 6: Problems with Asymmetric Protocols

4. Conclusion

This section summarizes and reiterates a few questions addressed in this document:

- o Is it useful to separate the security aspects for NAT and firewall signaling?
- o To what extent is end-to-end security important?
- o How can specific authorization problems be addressed?
- o What can be done with regard to the properties of asymmetric security protocols?
- o Are the proposals in [[I-D.tschofenig-nsis-sid](#)] adequate to address the sender-invariance property for mobility scenarios? This document, for example describes how to reuse concepts like hash-chains and the Purpose-Built Key mechanism [[I-D.bradner-pbk-frame](#)] to provide a mobility solution without a global PKI.

5. Security Considerations

This entire document addresses security issues of path-coupled NAT/Firewall signaling. The main intention is to solicit feedback and comments from the community at an early stage of the protocol development.

6. Contributors

The author would like to thank Richard Graveman for his detailed review.

7. Acknowledgements

The author would like to thank Cedric Aoun, Marcus Brunner, Srinath Thiruvengadam, Martin Stiernerling and Miquel Martin for their time to discuss many NAT/Firewall related issues.

8. References

8.1 Normative References

[I-D.[draft-ietf-nsis-ntlp](#)]

Schulzrinne, H. and R. Hancock, "GIMPS: General Internet Messaging Protocol for Signaling", [draft-draft-ietf-nsis-ntlp-00](#) (work in progress), October 2003, <reference.I-D.[draft-ietf-nsis-ntlp.xml](#)>.

[I-D.[ietf-nsis-nslp-natfw](#)]

Stiemerling, M., Tschofenig, H., Martin, M. and C. Aoun, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", [draft-ietf-nsis-nslp-natfw-02](#) (work in progress), May 2004, <reference.I-D.[ietf-nsis-nslp-natfw.xml](#)>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

[[draft-tschofenig-nsis-sid](#)]

Tschofenig, H., Schulzrinne, H., Hancock, R., McDonald, A. and X. Fu, "Security Implications of the Session Identifier", June 2003.

8.2 Informative References

[I-D.[aoun-nsis-nslp-natfw-migration](#)]

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](#)>.

[I-D.[bradner-pbk-frame](#)]

Bradner, S., Mankin, A. and J. Schiller, "A Framework for Purpose-Built Keys (PBK)", [draft-bradner-pbk-frame-06](#) (work in progress), June 2003, <reference.I-D.[bradner-pbk-frame.xml](#)>.

[I-D.[iab-model](#)]

Rescorla, E., "Writing Protocol Models", [draft-iab-model-01](#) (work in progress), May 2004, <reference.I-D.[iab-model.xml](#)>.

[I-D.[ietf-aaa-diameter-cms-sec](#)]

Calhoun, P., Farrell, S. and W. Bulley, "Diameter CMS Security Application", [draft-ietf-aaa-diameter-cms-sec-04](#) (work in progress), March 2002,

<reference.I-D.ietf-aaa-diameter-cms-sec.xml>.

[I-D.ietf-eap-rfc2284bis]

Blunk, L., Vollbrecht, J., Aboba, B., Carlson, J. and H. Levkowitz, "Extensible Authentication Protocol (EAP)", [draft-ietf-eap-rfc2284bis-07](#) (work in progress), December 2003.

[I-D.ietf-ipsec-ikev2]

Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", [draft-ietf-ipsec-ikev2-12](#) (work in progress), January 2004, <reference.I-D.ietf-ipsec-ikev2.xml>.

[I-D.ietf-kink-kink]

Thomas, M. and J. Vilhuber, "Kerberosized Internet Negotiation of Keys (KINK)", [draft-ietf-kink-kink-05](#) (work in progress), January 2003, <reference.I-D.ietf-kink-kink.xml>.

[I-D.ietf-nsis-fw]

Hancock, R., "Next Steps in Signaling: Framework", [draft-ietf-nsis-fw-05](#) (work in progress), October 2003, <reference.I-D.ietf-nsis-fw.xml>.

[I-D.ietf-nsis-qos-nslp]

Bosch, S., "NSLP for Quality-of-Service signaling", [draft-ietf-nsis-qos-nslp-01](#) (work in progress), October 2003, <reference.I-D.ietf-nsis-qos-nslp.xml>.

[I-D.ietf-sip-connect-reuse]

Mahy, R., "Connection Reuse in the Session Initiation Protocol (SIP)", [draft-ietf-sip-connect-reuse-01](#) (work in progress), February 2004, <reference.I-D.ietf-sip-connect-reuse.xml>.

[I-D.manyfolks-signaling-protocol-mobility]

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>.

[I-D.rosenberg-midcom-turn]

Rosenberg, J., "Traversal Using Relay NAT (TURN)", [draft-rosenberg-midcom-turn-04](#) (work in progress), February 2004, <reference.I-D.rosenberg-midcom-turn.xml>.

[I-D.tschofenig-nsis-sid]

Tschofenig, H., "Security Implications of the Session

Identifier", [draft-tschofenig-nsis-sid-00](#) (work in progress), June 2003, <reference.I-D.tschofenig-nsis-sid.xml>.

- [RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", [RFC 2246](#), January 1999, <reference.RFC.2246.xml>.
- [RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998, <reference.RFC.2409.xml>.
- [RFC3182] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T., Herzog, S. and R. Hess, "Identity Representation for RSVP", [RFC 3182](#), October 2001, <reference.RFC.3182.xml>.
- [RFC3489] 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, <reference.RFC.3489.xml>.
- [nat-memo] Tschofenig, H., "Memo about NSIS NAT Handling, available at: <http://www.tschofenig.com/drafts/NSIS-NAT-Handling.txt> (Feb. 2004)", February 2004, <reference.nat-memo>.
- [refs.tist] Shore, M., "The TIST (Topology-Insensitive Service Traversal) Protocol", DRAFT [draft-shore-tist-prot-00.txt](#), May 2002.

Author's Address

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

E-Mail: Hannes.Tschofenig@siemens.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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

