NSIS Working Group
Internet-Draft
Expires: December 16, 2003

M. Brunner
M. Stiemerling
M. Martin
NEC
H. Tschofenig
Siemens
H. Schulzrinne
Columbia U.

June 17, 2003

# NSIS NAT/FW NSLP: Problem Statement and Framework
# draft-brunner-nsis-midcom-ps-00

## Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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 December 16, 2003.

## Copyright Notice



## Abstract

This memo presents the problems for using the Next Steps in Signaling base protocol for firewall/NAT traversal commonly refered to middlebox traversal.

Table of Contents

## 1. Introduction

Even though the NSIS WG (Next Steps in Signaling) has as a primary application the signaling for QoS in mind, other types of applications should be possible.

In this draft, we look into the scenario, framework, problems, and issues of using a signaling protocol for middlebox traversal, where a middlebox in most cases is a Network Address Translator (NAT) or firewall.

One of the requirements in NSIS [1] is that the NTLP signaling protocol must be independent of the service requested. The thinking definitely goes into the direction to request end-to-end or edge-to-edge QoS from IP networks. However, the service might be "open me" the data path through all the firewalls through the network to host X". Also this type of service is running end-to-end.

See also [10] and [11] for proposals to use RSVP or CASP for NAT and Firewall traversal.

## 2. Terminology and Abbreviations

Sender-/Receiver Initiated Signaling

Sender-initiated: NAT bindings and firewall rules are created immediately when the "path" message hits the nsis nodes. With "path" message we refer to the signaling message traveling from the data sender towards the data receiver.

Receiver-initiated: NAT bindings and firewall rules are created when the "resv" message returns from the other end. With "resv" message we refer to a signaling message on the reverse path, this means from the receiver to the sender (i.e. backwards routed).

Note that these definitions have nothing to do with number of roundtrips, who performs authorization etc.

Firewalls vs. Security Gateway: As discussed in Section 3 different types of firewalls exist. This document focuses on firewalls, which perform packet filtering, and possibly application level filtering and does not address IPsec based security gateways.

Middlebox:

From [13]: "A middlebox is defined as any intermediate device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and a destination host."

The term middlebox in context of this document and in NSIS refers to firewalls and NATs only. Other types of middlebox are currently outside the scope.

The following abbreviations are used in various figures throughout the document:

- MB - Middlebox
- FW - Firewall
- S - Data Sender
- R - Data Receiver

## 3. What problem should be solved?

The term firewall and middlebox in general raises different expectations about the functionality provided by such a device. Different groups have worked on the problem of securing access to a network which different procedures with the help of different protocols. From an abstract point of view two different mechanisms for restricting access to a network can be differentiated:

- Packet Filters
- Cryptographically protected data traffic

Within this document we assume that packet filters are installed at devices along the path. These packet filters typically consist of a 5 tuple (src/dst ip address, transport protocol, src/dst port). Some devices entitled as firewalls only accept traffic after cryptographic verification (i.e. IPsec protected data traffic). Particularly for network access scenarios either link layer or network layer data protection is common. Hence we do not address these types of devices (referred as security gateways) since per-flow signaling is rather uncommon in this environment. For a discussion of network access authentication and associated scenarios the reader is referred to the PANA working group (see. [9] ).

In mobility scenarios an often experienced problem is the traversal of a security gateway at the edge of the corporate network. Network administrators often rely on the policy that only authenticated data traffic is allowed to enter the network. A problem statement for the traversal of these security gateways in the context of Mobile IP can be found at [8] ).

The goal of NSIS FW/NAT signaling therefore focuses on packet filter installation, due to the nature of the path-coupled discovery procedure and signaling message delivery. Discovering security gateways, which was also mentioned as an application for NSIS signaling, for the purpose of executing an IKE to create an IPsec SA, is already solved without requiring NSIS.

Installing packet filters provides some security but has some weaknesses, which heavily depend on the type of packet filter installed. A packet filter cannot prevent an adversary to inject traffic (due to the IP spoofing capabilities). This type of attack might not be particular helpful if the packet filter is a standard 5 tuple which is very restrictive. If packet filter installation, however, allows specifying a rule, which restricts only the source IP address, then IP spoofing allows transmitting traffic to an arbitrary address. NSIS aims to provide path-coupled signaling and therefore

an adversary is somewhat restricted in the location from which attacks can be performed. Some trust is therefore assumed from nodes and networks along the path.

## 4. Basic NSIS Usage for NAT/FW traversal

The basic high-level picture for using NSIS for NAT and firewall traversals is that at end host there are applications running, which need to communicate. Potentially, there are some application level intermediate servers. For example a SIP proxy would be such an application server. But also Gaming servers etc could be though of. Naturally, none, one, or more of these application server instances are possible.

End hosts use the NSIS NATFW NSLP for opening firewall pinholes and for creating NAT bindings. Therefore it is necessary that each firewall and each NAT involved in the signaling communication needs to run an NSIS daemon. There might be several NATs and FWs in various combinations possible on a path between two hosts. The reader is referred to Section 5 where different scenarios are presented.

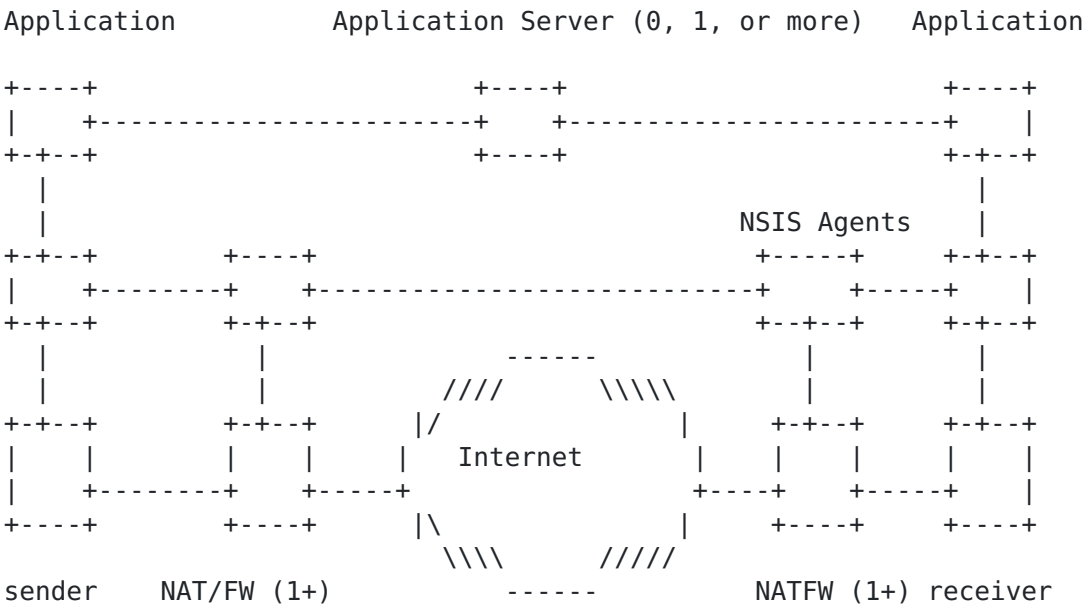


Figure 1: Generic View on NSIS in a NAT / Firewall case

## 5. Scenarios for Protocol Functionality

This section introduces several scenarios for middleboxes in the Internet. These middleboxes are firewalls or different flavours of NATs, like NAPT. Combination of both in the same device are possible as well.

Each section introduces a different scenario for a different set of middleboxes and their ordering within the topology.

### 5.1 Firewall traversal

The following scenario shows two end hosts behind a firewall but connected via the public Internet. The application can somehow trigger firewall traversal (e.g. via an API call) at the NSIS agent at the local host. The NSIS agent then signals this request to the next NSIS aware node and therefore to the receiver. Each firewall in the middle must provide traversal service in order to permit the NSIS message to reach the other end host.

The difference between this scenario and the following is that only firewalls are on the path, but no NATs. This has specific implication concerning the path-coupled signaling.

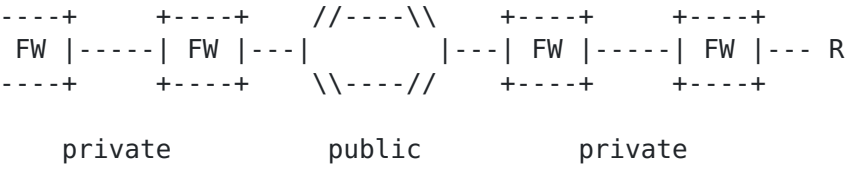


FW: Firewall
S: Data Sender
R: Data Receiver

Figure 2: Firewall Traversal Scenario

### 5.2 NAT with two private networks

This scenario deals with NATs on both ends of the network. Therefore, each application instance is behind a NAT and is connected to the public Internet (see Figure 3 ).

The case where more than one MB on each side ("only" two are shown in the figure) is present must be taken into account. This aspect

introduces more topology problems.


MB: Middlebox
S: Data Sender
R: Data Receiver

Figure 3: NAT with two private networks Scenario

Data traffic from the sender to the receiver has to traverse all four middleboxes on the path and all four middleboxes must be configured properly to allow subsequent data packets to flow. The sender has to know the IP address of the receiver in advance, i.e. before any NSIS message can be sent. Or more general the NSIS Initiator must know the IP addresses of the NSIS Responder, otherwise he cannot send a single NSIS signaling message towards the responder. Note that this IP address is not the private IP address of the responder. Instead a NAT binding (including an public IP address) has to be obtained from the NAT which subsequently allows packets hitting the NAT to be forwarded to the receiver within the private address realm. This in general requires further support from an application layer protocol for the purpose of discovering and exchanging information. The receiver might have a number of ways to learn its public IP address and port number and might need to signal this information to the sender using the application level signaling protocol.

### 5.3 NAT with private network on sender side

This scenario shows an application instance at the sending node which is behind one ore more FW/NATs. The receiver is located in the public internet.

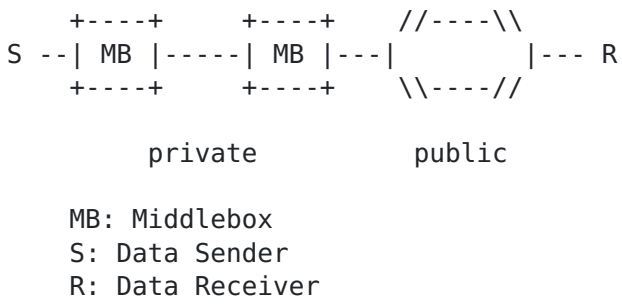


Figure 4: NAT with private network on sender scenario

The traffic from the sender to the receiver has to traverse only middleboxes on the sender's side. The receiver has a public IP address and therefore the procedure is simple. The sender sends its signaling message directly to the receiver whereby it is intercepted by the middleboxes along the path.

Note that the data sender does not necessarily knows whether the receiver is behind a NAT or not, and so, it is the receiving side that has to detect it. As described later NSIS can also provide help for this procedure.

## 5.4 NAT with private network on receiver side

The application instance receiving data is behind one or more NATs.


Figure 5: NAT with private network on receiver Scenario

First, the sender must determine the public IP address of the receiver.

One possibility is that an application level protocol is used. In this case, the receiver must first find out its public IP addresses

at the middlebox on its side. This information about IP address and port numbers could be signalled somehow to the actual sender directly or indirectly via a third party. In the scenario, this means the receiver has to determine its public IP address (NAT binding) and register this address with the third party.

The sender can start NSIS signaling after he has received information about the receiver's address and port number.

Note that it is part of the solution design where to terminate the signaling messages.

### 5.5 Both end hosts are in same private network behind NATs

This is a special case, where the main problem is to detect that both nodes are within the same network behind a NAT. This scenario primarily addresses performance aspects.

Sender and receiver are both within a private address space and even the address space might be the same. Figure 6 shows the ordering of NATs. This is a common configuration in several networks. For instance after two companies merge each network uses the same private IP address space.

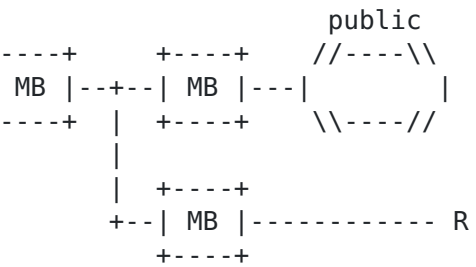


MB: Middlebox
S: Data Sender
R: Data Receiver

Figure 6: NAT to public, receiver in same private network Scenario

The middleboxes are twice-NATs, i.e. they map the IP addresses and port numbers on both sides, private and public interface.

From a protocol point of view, this means that the protocol must be

robust enough to at least not break with this scenario.

In the worst case, both sender and receiver obtain a public IP address at the NAT and the communication path is not optimal anymore.

### 5.6 IPv4/v6 NAT with two private networks

This scenario combines the usage case mentioned in Section 5.2 with the IPv4 to IPv6 transition scenario, i.e. using Network Address and Protocol Translators (NAT-PT).

The difference to the other scenarios lies in the use of IPv6 - IPv4 address translation, which happens in both directions. Additionally, the base NTLP must take care of this case for its own functionality of forwarding messages between NSIS peers.

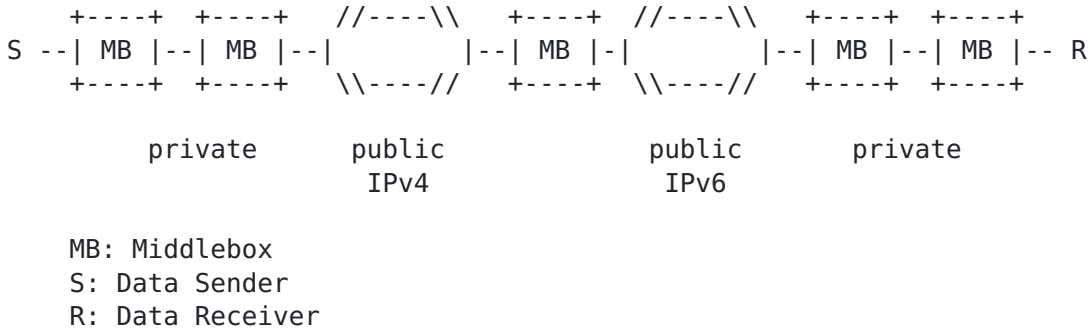


Figure 7: IPv4/v6 NAT with two private networks

## 6. Trust Relationship and Authorization

Trust relationships and authorization are very important for the protocol machinery. Trust and authorization closely related to each other in the sense that a certain degree of trust is required to authorize a particular action. If the action is "create/delete packet filters" then authorization is very important due to the nature of a firewall.

It is not particular surprising that differences exist between authorization in a QoS signaling environment and firewall signaling. As elaborate in [6] the establishment of a financial relationship is very important for QoS signaling whereas for firewall signaling is not directly of interest.

In the subsequent sections different trust relationships will be described which appear in firewall signaling environments. Peer-to-peer trust relationships are those, which are used in QoS signaling today and seem to be the simplest. However, there are reasons to believe that this is not the only type of trust relationship found in today's networks.

### 6.1 Peer-to-Peer Trust Relationship

Starting with the simplest scenario it is assumed that neighboring nodes trust each other. They required security association to authenticate a signaling message and to protect it is either available (manual configuration) or dynamically established with the help of an authentication and key exchange protocol. If nodes are located closely together it is assumed that security association establishment is easier than establishing it between far distant node. It is, however, difficult to describe this relationship generally due to the different usage scenarios and environments. Authorization heavily depends on the participating entities but for this scenario it is assumed that neighboring entities are trust each other to an extend that the packet filter creation and deletion is allowed. Note that Figure 8 does not illustrate the trust relationship between the end host and the access network, which is dynamically established as part of the network access authentication procedure as motivated in Section 1 .

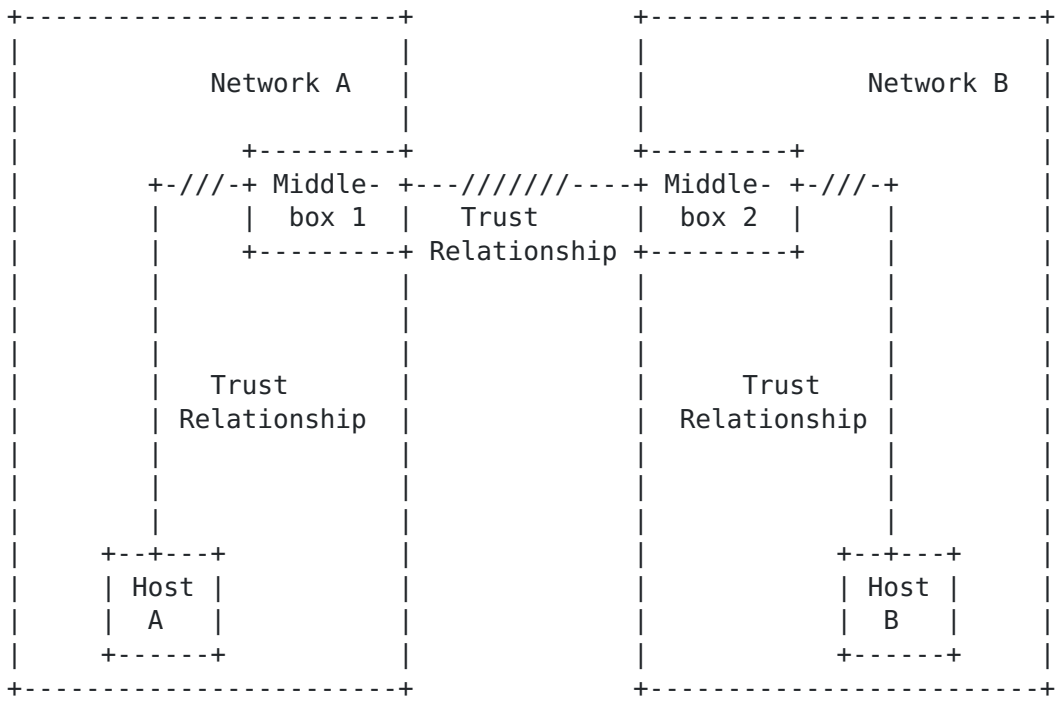


Figure 8: Peer-to-Peer Trust Relationship

## 6.2 Intra-Domain Trust Relationship

In larger corporations often more than one firewall is used to protect different departments. In many cases the entire enterprise is controlled by a security department, which gives instructions to the department administrators. In such a scenario a peer-to-peer trust-relationship might be prevalent. Sometimes however it might be necessary to preserve authentication and authorization information within the network. As a possible solution a centralized approach could be used whereby an interaction between the individual middleboxes and a central entity (for example a policy decision point - PDP) takes place. As an alternative individual firewalls could exchange the authorization decision to another firewalls within the same trust domain. Individual middleboxes within an administrative domain should exploit their trust relationship instead of requesting authentication and authorization of the signaling initiator again and again. Thereby complex protocol interaction is avoided. This provides both a performance improvement without a security disadvantage since a single administrative domain can be seen as a single entity. Figure 9 illustrates a network structure which uses a centralized entity.

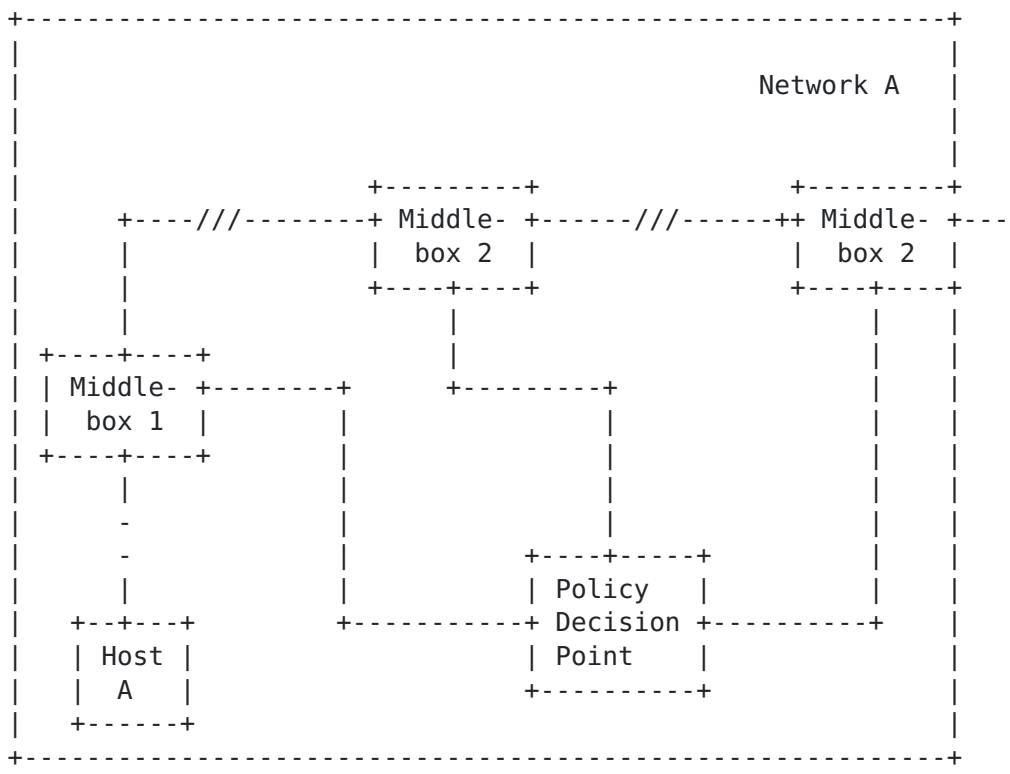


Figure 9: Intra-domain Trust Relationship

## 6.3 End-to-Middle Trust Relationship

In some scenarios a simple peer-to-peer trust relationship between participating nodes is not sufficient. Network B might require additional authorization of the signaling message initiator. If authentication and authorization information is not attached to the initial signaling message then the signaling message arriving at Middlebox 2 would cause an error message to created which indicates the additional authorization requirement. In many cases the signaling message initiator is already aware of the additional required authorization before the signaling message exchange is executed. Replay protection is a requirement for authentication to the non-neighboring firewall which might be difficult to accomplish without adding additional roundtrips to the signaling protocol (e.g. by adding a challenge/response type of message exchange).

Figure 10 shows the slightly more complex trust relationships in this scenario.

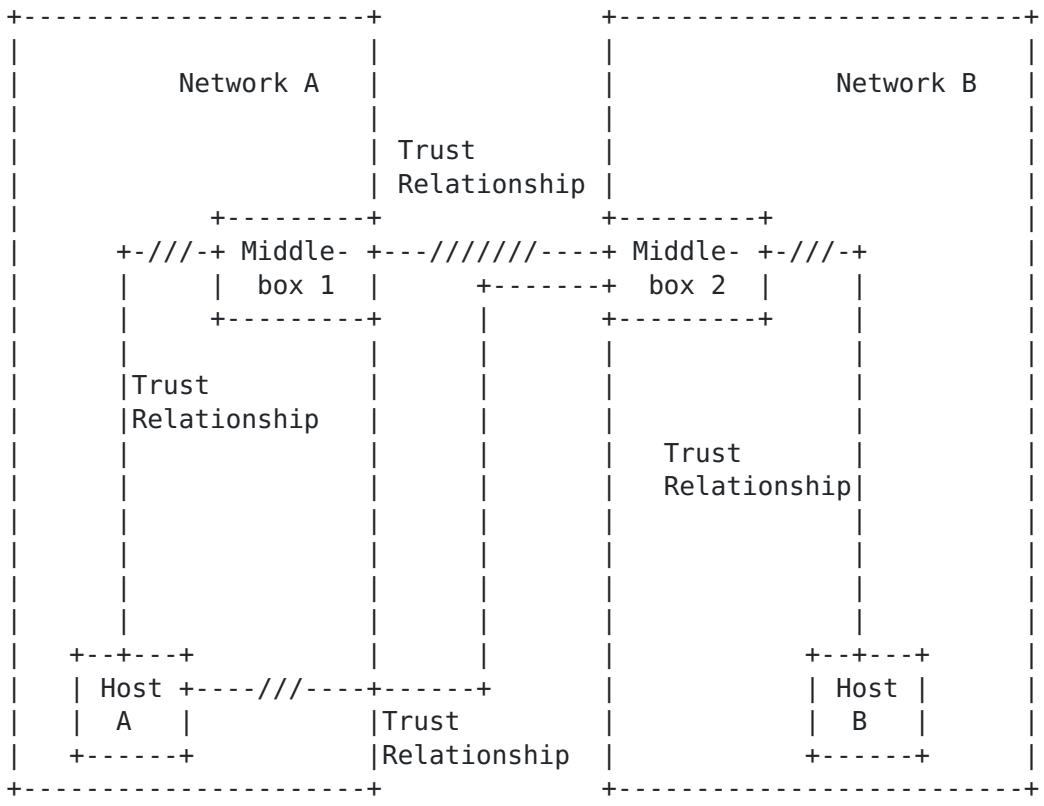


Figure 10: End-to-Middle Trust Relationship

## 7. Problems and Challenges

This section describes a number of problems which have to be addressed for NSIS NAT/Firewall. There are also of relevance for other NSLP protocols.

### 7.1 Missing Network-to-Network Trust Relationship

Peer-to-peer trust relationship, as shown in Figure 8 , is a very convenient assumption that allows simplified signaling message processing. However, it might not always be applicable. Especially the trust relationship between two arbitrary access networks (over a core network where signaling messages are not interpreted) does possibly not exist because of the large number of networks and the unwillingness of administrators to have other network operators to create holes in their firewalls without proper authorization. Hence in the following scenario we assume a somewhat different message processing and show three possible approaches to tackle the problem. None of these three approaches is without drawbacks or constraining assumptions.

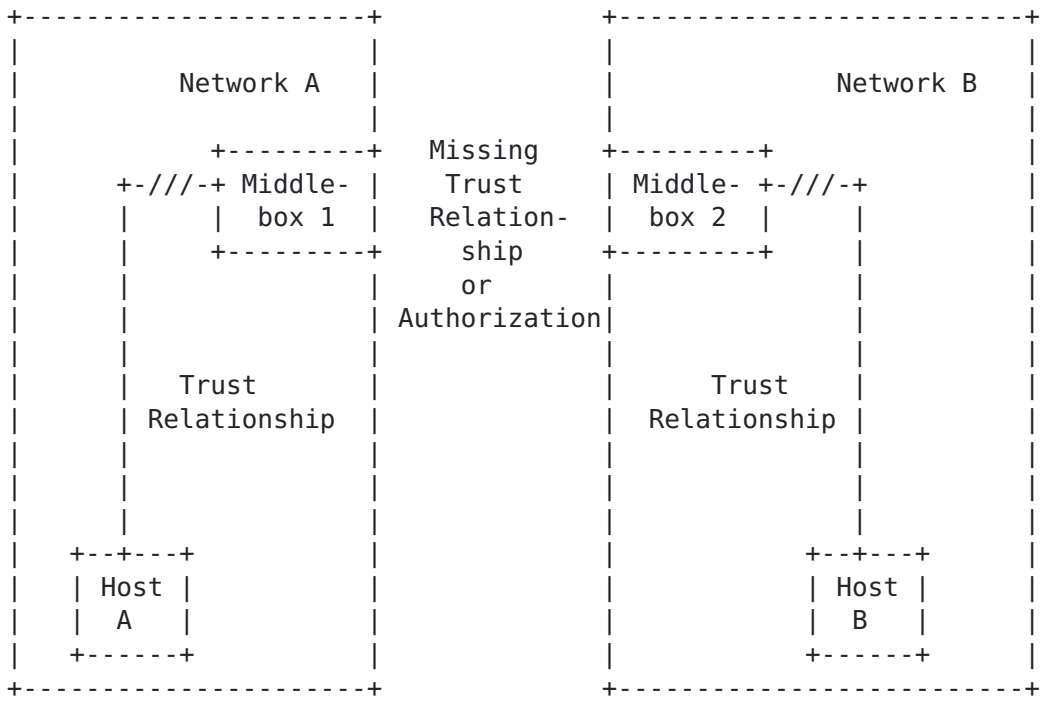


Figure 11: Missing Network-to-Network Trust Relationship

Figure 11 illustrates a problem whereby an external node is not allowed to manipulate (create, delete, query, etc.) packet filters at a firewall. Opening pinholes is only allowed for internal nodes or with a certain authorization permission. Hence the solution alternatives focus on establishing the necessary trust with cooperation of internal nodes. We have identified three possible approaches of tackling the problem which are described in Appendix D.

## 7.2 End-to-end significance

Also in the case of NAT/firewall traversals, we need to have the end-to-end significance since more than one NAT/Firewall might be in the path between a data sender and a data receiver.

## 7.3 Relationship with routing

The data path is following the "normal" routes. The NAT/FW devices along the data path are those providing the service. In this case the service is something like "open a pinhole" or even more general "allow for connectivity between two communication partners". The benefit of using path-coupled signaling is that the NSIS NATFW NSLP does not need to take care where middleboxes can be found and in

which order they appear.

Creating NAT bindings modifies routing of data packets between end points. This is unlike other NSIS NSLPs, which do not interfere with routing - instead they only follow the path of the data packets.

### 7.4 Dynamic state installation and maintenance

For NAT/Firewall traversal, the lifetime of a NAT binding or a packet filter must be provided and needs to be continuously refreshed. So specifically for short-lived flows signaling for pinholes and NAT bindings is preferable. The capability to specify a lifetime for a NAT binding provides some advantages to what exists today where unknown NAT binding lifetimes can lead to unexpected protocol actions.

For more static behavior both NAT bindings and pinholes can be provisioned statically and no signaling is used. For static state other mechanisms than an NSIS signaling protocol might be preferable. Most time this is a matter of configuration of a middlebox using a management protocol such as SNMP or CLI.

### 7.5 Affected Parts of the Network

NATs and Firewalls tend to be located at the edge of the network, whereby other signaling applications effect all nodes along the path. One typical example is QoS signaling where all networks along the path must provide QoS in order to achieve true end-to-end QoS. In the NAT/Firewall case, only some of the domains/nodes are affected (typically access networks), whereas most parts of the networks and nodes are unaffected (e.g. the core network).

This fact raises some questions. Should an NSIS NTLP node intercept every signaling message independently of the upper layer signaling application or should it be possible to make the discovery procedure more intelligent to skip nodes. These questions are also related to the question whether NSIS NAT/FW should be combined with other NSIS signaling applications.

### 7.6 Traversing NSIS unaware domains

Signaling of QoS information even works if NSIS (or QoS NSLP) unaware domains are traversed. The thinking behind this is that we hope to get the best even if traversing unaware domains. Although it might not produce the desired effect from a Quality of Service point of view it is still possible for NSIS messages to reach the intended end host.

With firewalls the situation is somewhat different. An NSIS unaware firewall should actually reject such a request. Since firewalls typically implement the policy "default = deny" the traversal of NSIS signaling messages must be We believe this is easily possible by normal firewall functionality. So this does not seam to be a real problem in most cases. But this is a deployment problem, since all firewalls along the path must be NSIS aware in order to get an open path. Which packet filters are required to allow NSIS signaling messages itself to pass the firewall depends on the NSIS signaling message. Since RSVP signaling messages are addressed end-to-end (in case of the path message) it is necessary to create a packet filter, which allows IP datagrams using protocol 46 to pass. For signaling protocols, which perform peer-to-peer addressing, addressing of a specific port needs to be allowed (assuming that a transport protocol is used and that the firewall or NAT is NSIS aware).

For NATs this is more problematic since signaling messages are forwarded (at least in one direction), but with a changed IP address and changed port numbers. The content of the NSIS signaling message is, however, unchanged. This can lead to unexpected results. NSIS unaware NATs must be detected in order to let all entities involved take care of that situation and in order to work correctly. Such a "legacy" NAT detection procedure can be done during the NSIS discover procedure itself.

Based on experience it was discovered that routers unaware of the Router Alert IP option [RFC 2113] discarded packets. This is certainly a problem for NSIS signaling.

### 7.7 Authentication and Authorization

Since a firewall has security functionality, strong authentication and authorization means MUST be provided.

For NATs security is not a major concern, but might play a role in the perceived security measure of some administrators. For NAT sometimes address depletion is mentioned as a threat.

### 7.8 Directional Property

A firewall has a directional property. Hosts are sitting behind a firewall, or hosts are in the intra-net. Others are outside the firewall. So from a security point of view, the way NSIS signaling messages enters the NSIS agent of a firewall (see Figure 2) might be important, because different policies might apply for authentication and admission control.

Also for NATs there is a natural direction from the private to the

public address space. Only after establishing the NAT binding packets can flow in both directions. NAT bindings are therefore typically created by data traffic originating from the internal network.

Most of the time hosts inside the firewall-protected domain are more trusted than external hosts. However, based on changes in the network architecture and the corporate policy not even this might be true anymore. Nevertheless it would imply that the data sender and the data receiver might need to tell their respective firewalls that it should open a pinhole. In general it is inappropriate to perform operations on the firewall from the outside; particularly without sufficient authorization privileges.

### 7.9 Routing Asymmetry

Routing asymmetry [7] is a general problem for path-coupled signaling. Similarly to path-coupled QoS signaling firewall signaling also has to be aware of the routing asymmetry although the routing asymmetry might not be large within the local networks where firewalls are typically located. For signaling NAT bindings this issue comes with a different flavor since an established NAT binding changes the path of the data packets. Hence a data receiver might still be able to send NSIS signaling messages to create a NAT binding, although they travel the previously "wrong" path.

### 7.10 Addressing

Also a more general problem of NATs is the addressing of the end-point. NSIS signaling message have to be addressed to the other end host to follow data packets subsequently sent. Therefore a public IP address of the receiver has to be known. When NSIS signaling messages contain IP addresses of the sender and the receiver in the signaling message payloads, then an NSIS agent must modify them. This is one of the cases, where a NSIS aware NATs is also helpful for other types of signaling applications e.g. QoS signaling.

### 7.11 NTLP/NSLP NAT Support

It must be possible for NSIS NATs along the path to change NTLP and/or NSLP message payloads , which carry IP address and port information. This functionality includes the support of providing mid-session and mid-path modification of these payloads. As a consequence these payloads must not be reordered, integrity protected and/or encrypted in a non peer-to-peer fashion (e.g. end-to-middle, end-to-end protection). Ideally these mutable payloads must be marked (e.g. a protected flag) to assist NATs in their effort of

adjusting these payloads.

### 7.12 Route changes

The effect of route changes are more severe than in other signaling applications since a firewall pinhole and NAT binding is needed before further communication can takes place. This is true for both NSIS signaling and for subsequent data traffic. If a route changes and NSIS signaling messages do not configure NSIS NATs and firewalls along the new path then the communication is temporarily interrupted. This is naturally a big problem for networks where routes frequently change e.g. ad-hoc networks or in case of fast mobility. In these cases either refresh messages and/or triggers have to provide a mechanism for fast reaction.

### 7.13 Combining Middlebox and QoS signaling

In many cases, middlebox and QoS signaling has to be combined at least logically. Hence it was suggested to combine them into a single signaling message or to tie them together with the help of the same session identifier. This, however, has some disadvantages such as: - NAT/FW NSLP signaling affects a much small number of NSIS nodes along the path (for example compared to the QoS signaling). - NAT/FW signaling might show different signaling patters (e.g. required end-to-middle communication). - The refresh intervals are likely to be different. - The number of error cases increase as different signaling applications are combined into a single message. The combination of error cases has to be considered.

### 7.14 Difference between sender- and receiver-initiated signaling

For NAT/FW signaling there seems to be little difference between sender- and receiver- initiated signaling messages. Some other characteristics of QoS signaling protocols are not applicable (e.g. the adspec object) to the NAT/FW context. It seems that a full roundtrip is always required if the protocol aims to be generic enough.

### 7.15 Inability to know the scenario

In Section 5 a number of different scenarios are presented. In some scenario NSIS signaling is fairly easy whereas in others it is quite complex. Additionally different trust relationships exist between networks along the path, which might require interaction with the end host or a different signaling behavior. However, the user (or the NSIS agent initially) typically does not know which scenario is currently applicable. To make things worse, the scenario might actually change with moving networks, adhoc networks or with mobility

in general. Hence NSIS signaling must assume the worst case and cannot put responsibility to the user to know which scenario is currently applicable. As a result, it might be necessary to perform a "discovery" periodically such that the NSIS agent at the end host has enough information to decide which scenario is currently applicable. This additional messaging, which might not be necessary in all cases, eats performance, bandwidth and adds complexity. Additional information by the user can provide information to assist this "discovery" process but cannot replace it.

Some protocols already aim to provide a solution for an end host to learn something about the topology such as STUN [3]. To some extend these protocols can help NSIS NAT/FW signaling.

## 8. Security Considerations

Security is of major concern specifically if the middlebox is a firewall. General threats to signaling have been discussed in [2]. These apply here as well. Additionally, the draft discusses some problems concerning security for that specific purpose.

Normative References

[1] Brunner et al., M., "Requirements for Signaling Protocols", DRAFT draft-ietf-nsis-req-07.txt, March 2003.

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

[3] 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.

[4] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A. and A. Rayhan, "Middlebox communication architecture and framework", RFC 3303, August 2002.

Informative References

[5] Manner, J., Suihko, T., Kojo, M., Liljeberg, M. and K. Raatikainen, "Localized RSVP", DRAFT draft-manner-lrsvp-00.txt, November 2002.

[6] Tschofenig, H., Buechli, M., Van den Bosch, S. and H. Schulzrinne, "NSIS Authentication, Authorization and Accounting Issues", draft-tschofenig-nsis-aaa-issues-01 (work in progress), March 2003.

[7] Amini, L. and H. Schulzrinne, "Observations from router-level internet traces", DIMACS Workshop on Internet and WWW Measurement, Mapping and Modelin Jersey) , Februar 2002.

[8] Adrangi, F. and H. Levkowetz, "Problem Statement: Mobile IPv4 Traversal of VPN Gateways", draft-ietf-mobileip-vpn-problem-statement-req-02.txt (work in progress), April 2003.

[9] Ohba, Y., Das, S., Patil, P., Soliman, H. and A. Yegin, "Problem Space and Usage Scenarios for PANA", draft-ietf-pana-usage-scenarios-06 (work in progress), April 2003.

[10] Shore, M., "The TIST (Topology-Insensitive Service Traversal) Protocol", DRAFT draft-shore-tist-prot-00.txt, May 2002.

[11] Tschofenig, H., Schulzrinne, H. and C. Aoun, "A Firewall/NAT Traversal Client for CASP", DRAFT draft-tschofenig-nsis-casp-midcom-01.txt, March 2003.

[12] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M. and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[13] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.

Authors' Addresses

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 0
EMail: lopez@ccrle.nec.de
URI:

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

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

Henning Schulzrinne
Columbia University, Dept. of Computer Science
1214 Amsterdam Avenue
New York NY 10027
USA

Phone:
EMail: schulzrinne@cs.columbia.edu
URI: http://www.cs.columbia.edu/~hgs/

## Appendix A. Interworking of SIP with NSIS NATFW NSLP

This document aims at pinpointing the problems of using SIP in nowadays networks, focusing on the problems derived of NAT's, Firewalls and multi-path communications. It is intended to fit in a scenario description that shows the necessity of NSIS, as well as depicting it's requirements. However, note that there are a number of other solutions available. For example the IETF Midcom working group is working on [4].

### A.1 The Session Initiation Protocol

[12] describes the Session Initiation Protocol, an application-layer control protocol that can establish, modify, and terminate multimedia sessions. This often involves several flows for video and voice, which are transported over new connections. These use of dynamically allocated ports which results in protocol complexity which can not be handled by nowadays NAT's and Firewalls.

Session initiation when one or both of the users is behind a NAT is also not possible, given the impossibility to address a private IP over the internet. Moreover, network deployments often allow for different paths per connection and direction, making the setup of the middle boxes even more complicated.

The following figure depicts a typical SIP connection:

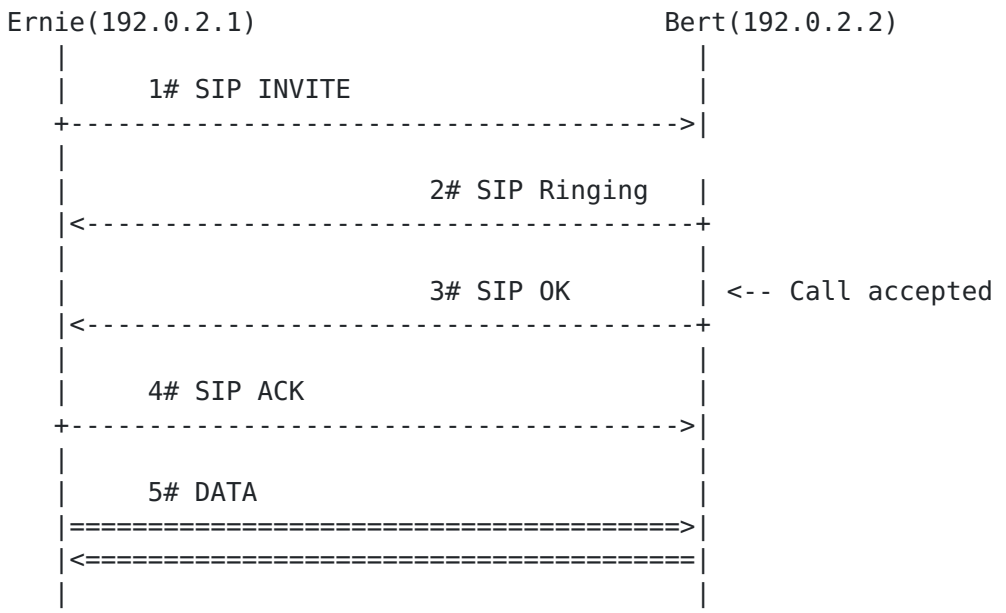


1# SIP Invite (192.0.2.1:? -> 192.0.2.2:SIP): I Listen on 192.0.2.1:1000 Ernie invites Bert to the conference, and informs it's awaiting media data on port 1000.

2# SIP Ringing (192.0.2.2:SIP -> 192.0.2.1:?): Ringing Bert's phone The ringing simply inplies that there's something sip aware on Berts side, and that it's ringing his phone

3# SIP OK (192.0.2.2:SIP -> 192.0.2.1:?): Call accepted, I listen on 192.0.2.2:2000 This OK means that the Bert took the phone off hook, and thus accepted the call. It also informs Ernie that Bert is awaiting his media data at port 2000

4# SIP ACK (192.0.2.1:? -> 192.0.2.2:SIP): All is fine, start transmitting. ACK means the ports are accepted and the call can start in the slected data ports on both sides.

5# DATA (192.0.2.1:? -> 192.0.2.2:2000 and 192.0.2.2:? -> 192.0.2.1:1000): Voice,image, video.. This is the actual data being transmited.

In the above example, SIP is used successfully to establish a communication, which includes negotiating the data ports for the actual transmission. Unfortunatelly, this scheme will not work for more complex setups.

Let's now consider one firewall in the data path, be it on Ernie's or Bert's network, or elsewhere in the middle. We assume that the

firewall is allowing traffic directed to the SIP port. As to the rest of the ports, a typical setup involves outgoing connections being allowed, and incoming connections being dropped, except for those already established. That is, we allow packets to go out and their replies to come in, but disable all other traffic.

In this case, the connection is as follows, for the case of a firewall on Ernie's network:

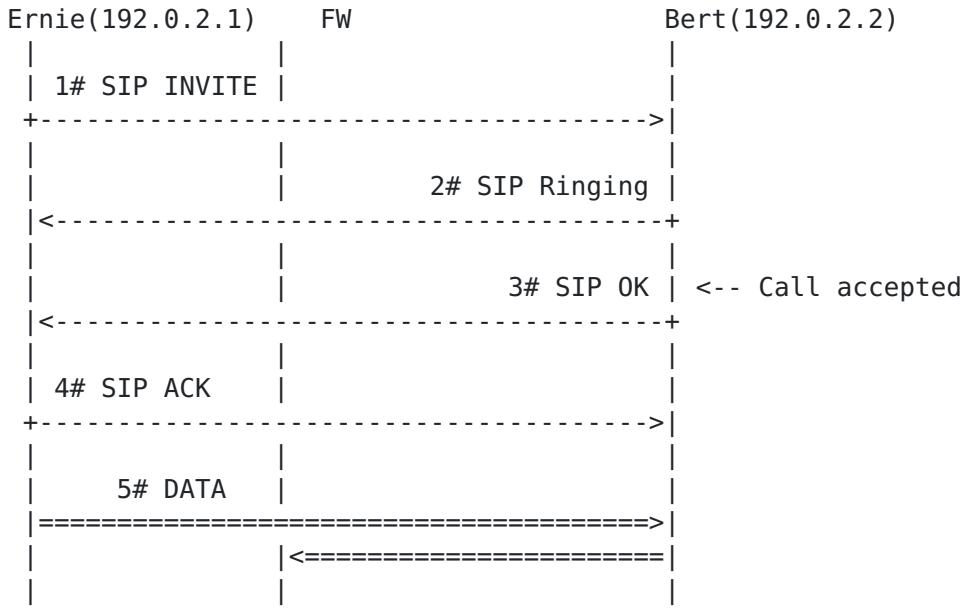


Notice how the SIP messages #1 and #4 traverse the firewall, because they are outbound, and how 2# and 3# traverse it too, because they are replies to the connection established at 1#.

Notice now how 5# can go outwards, but Bert can not go through the firewall to reach Ernie's port 1000. The reason is the connection is a new one, and the firewall won't allow it through.

Bert will now get media from Ernie, but Ernie is never going to get anything from Bert. The call is thus considered unsuccessful. The reason is that the application level port negotiation is never acknowledge by the network-transport layer firewall, which doesn't know what to expect. We would still face the same problem if the connection used a SIP Proxy, for it would only translate names into IP addresses.

Let us now assume that we indeed have an application layer firewall,

be it by design, or because we load some sort of SIP module to it. The previous case would now work, since the firewall can now understand the packets going through it and open the necessary ports. Still, we cannot assume that SIP signalization packets and the actual data follow the same path. The following figure shows a likely setup. FW+ stands for one or more firewalls:

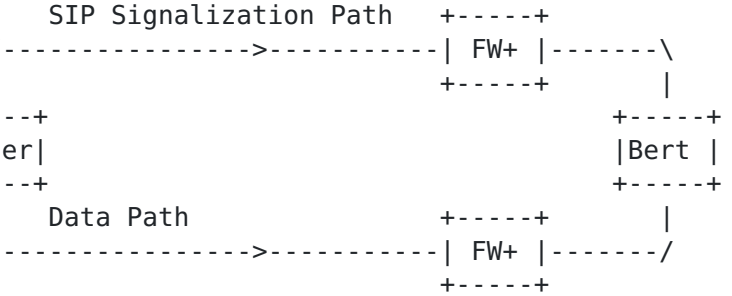


The SIP packets with the information about the listening ports now travels on the SIP Signalization path, and so the firewalls on that path can read them. The Data, though, is traveling through the Data path, and the firewalls in that path never get to see the Invite and Ok packets. They are thus unable to open the ports.

Two issues are arisen here: first, we need on-path signalization unless we already know the path our packets will take; a highly unlikely situation in today's internet. Second, if we patch the firewalls to understand SIP, we will provide any caller with a hole-puncher for the firewall, since SIP is not provisioned with proper authentication mechanism.

It is now clear that tight firewalls prevent SIP from successfully working. There is still another obstacle: NATs.

NATs provide for a link between two different address spaces, typically connecting a private range network to a public range one. As a consequence, connections going from the inside (usually the private range) are translated using the NAT's public interface address, and the replies are routed back. The public side of the network can only see the NATs public interface, and know nothing of the private network inside. This means computers outside the NAT won't be able to address computers inside the NAT.

Let us analyze the SIP example when Ernie is behind a NAT. The following figure depicts a typical session:

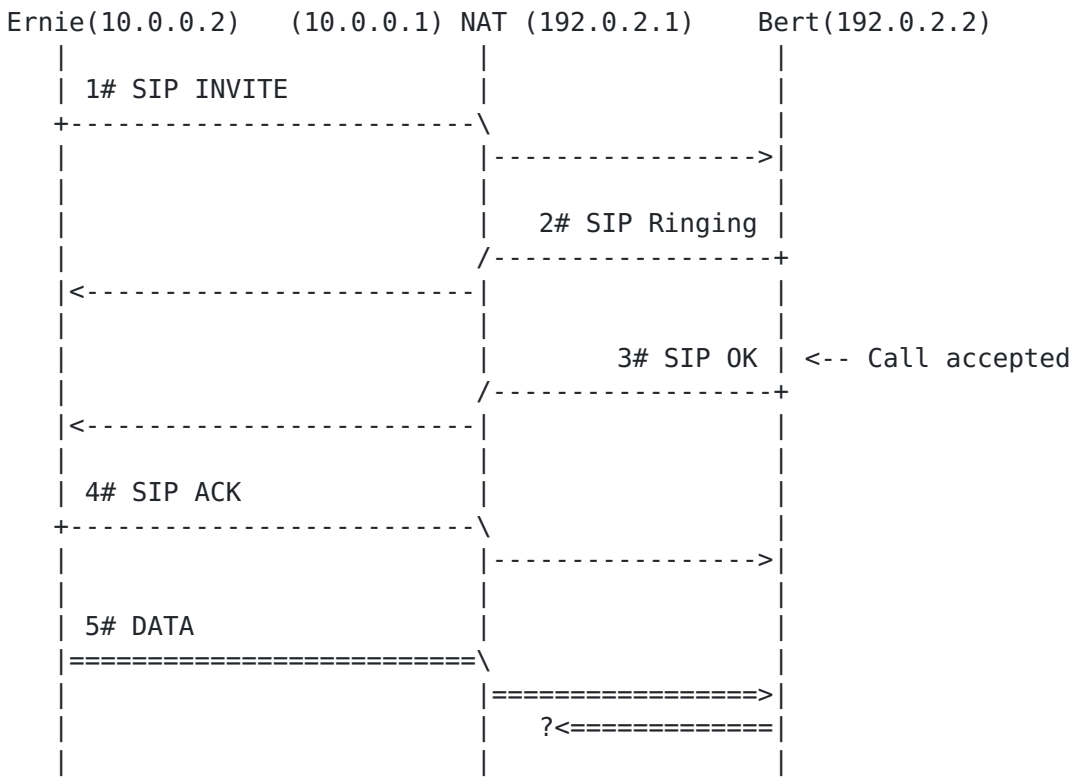


The communication is analogous to the one in the previous examples, except for the fact the NAT is rewriting the source address of the packets as they traverse it.

For instance, packet 1# is going from 10.0.0.2:? towards 192.0.2.2:SIP. The NAT box intercepts the message and puts 192.0.2.1:? as the source address and port, with ? being a dynamically picked port, which might be different from the original one 1# used.

On the way back, Bert is replyinc to the source of the IP packet, that is, 192.0.2.1, and so, when 2# reaches 192.0.2.1, the NAT know it is a reply from 1#, because it established a NAT binding, and this replaces the destination address, 192.0.2.1:? with 10.0.0.2:? and forwards the packet inside the NAT.

As a result, Ernie never knows there is a NAT in his communication path, since he sends and receives packets from 192.0.2.2 normally. This means that the INVITE packet will tell Bert to send data back to 10.0.0.2, a private IP. Once the signalization is finished, and the actual DATA transmission starts, Bert tries to connect to 10.0.0.2, a private IP address, from the internet; The routers don't know how to route this, and the packet is eventually dropped.

One possible solution would be for Ernie to know the NAT exists, and already indicate that it listens on 192.0.2.1, and not 10.0.0.2. That, still would not work, since the NAT binding is not performed at the NAT box.

## A.2 Conclusions

The above examples display the inability to use standard SIP through tight firewalls or NATs, and points at the necessity of a secure on-path protocol to negotiate firewall pinholes and NAT bindings.

## Appendix B. Ad-Hoc networks

Some forms of ad-hoc networks exist where trust in the network is not justified. Figure Figure 16 mainly illustrates the problems of malicious NSIS entities graphically:

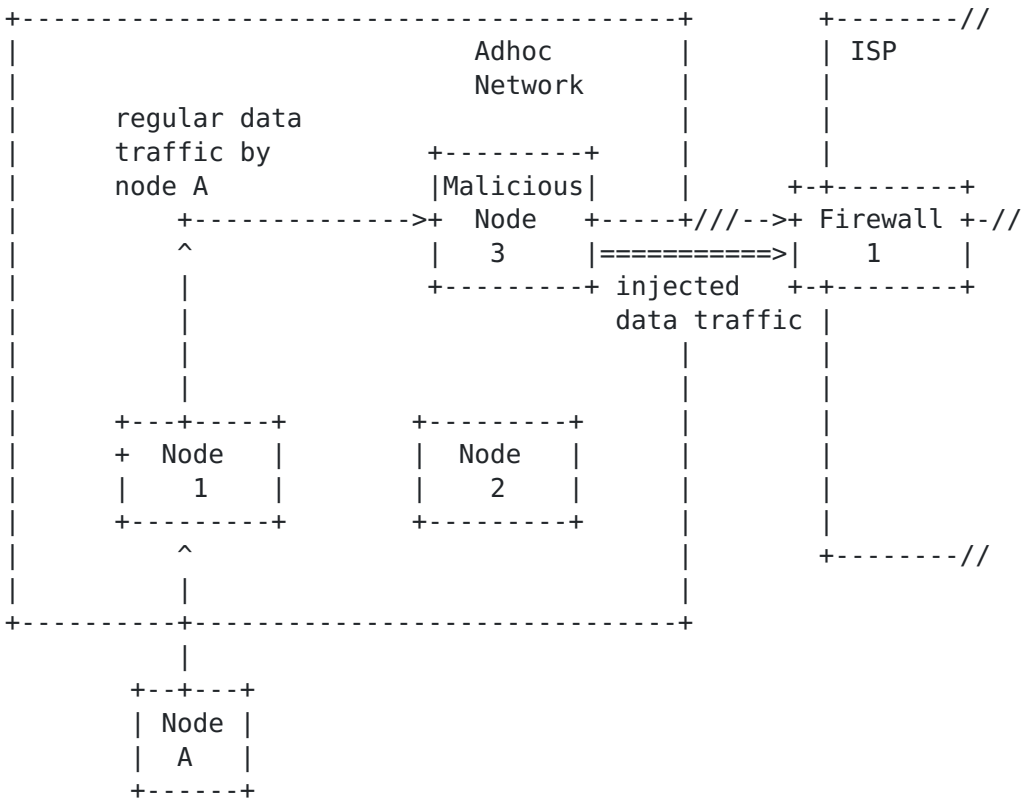


Figure 16: Limits of packet filter security

An ad-hoc networks consists of a number of nodes between the end host (Node A) and the ISP to which Node A wants to get access. Although Node A uses an authentication and key exchange protocol to create a policy rule at the firewall 1 it is still possible for an untrusted node (in this case Node 3) to inject data traffic which will pass Firewall 1 since the data traffic is not authenticated. To prevent this type of threat two approaches are possible. First, a restrictive packet filter limits the capabilities of an adversary. Finally, there is always the option of using data traffic protection.

## Appendix C. Interworking of Security Mechanisms and NSIS NATFW NSLP

TBD

## Appendix D. Solution approaches in case of missing authorization

### D.1 Solution Approach: Local authorization from both end points

The first approach makes use of local authorization from both end points. If Host A sends a signaling message toward the destination to Middlebox 1 the message will perform the desired action in Network A. Middlebox 1 establishes some state information and forwards the signaling message towards Host B. Signaling message protection between the two access networks might be difficult. A missing trust relationship does not necessarily mean that no security association establishment is possible. The lacking trust disallows Middlebox 1 (or indirectly Host A where the signaling message was initiated) to create packet filters at Middlebox 2. We assume that the NSIS signaling message is allowed to pass the firewall then it finally reaches Host B. Due to the missing authorization no packet filter specific state is created. The filters will be installed later after receiving an authorization from Host B. When Host B returns a confirmation or acknowledgement then Middlebox 2 treats it as an authorization and finally triggers filter creation. The message is then forwarded to Middlebox 1, where filters are either already installed or require an additional confirmation. Finally the signaling message is forwarded to Host A, which can be assured that subsequent data traffic can be transmitted end-to-end from Host A to Host B. The same procedure has to be applied again to signal information for the other direction (Host B to Host A).

The following behavior has to be assumed in order for this approach to be applicable:

1. Signaling messages must be allowed to pass firewalls along the path.

2. NSIS signaling must operate in the described manner which could be described as: Install where you have authorization - delay and forward where you have no authorization.

This approach suffers from the following drawbacks:

1. Firewalls which block NSIS signaling from external networks or nodes prevent a successful operation.

2. A full roundtrip is required to signal packet filter information. The NSIS signaling message must therefore provide the capability to route signaling message in both direction which might either require state installation at nodes along the path (route pinning) or a stateless version via record-route. Some risk of DoS protection might exist.

### D.2 Solution Approach: Access Network-Only Signaling

The next approach is based on signaling packet filter information by both hosts into the local access network only. An NSIS allows specifying such a behavior by indicating the signaling endpoint with the help of scoping (for example with domain name or a "local network only" flag). Scoping means that the signaling message although addressed to a particular destination IP address terminates somewhere along the path. If packet filters for both directions have to be installed then the signaling messages have to make packet filter installations up- and downstream along the data path. Similar to proposals in the area of QoS signaling some problems are likely to occur. One such problem is that downstream signaling in general causes problems because of asymmetric routes. In particular it is difficult to determine the firewall where the downstream data traffic will enter a network. The problem of triggering downstream reservations is for example described in [5] . Another problem for example is the placement of a firewall or NAT along the path other than in the access network. This would prevent a successful data exchange.

The following behavior has to be assumed in order for this approach to be applicable:

1. It must be possible to trigger a signaling message exchange for a downstream signaling message exchange at the firewall where the data traffic enters the network.

2. No other firewalls or NATs are present along the path other than in the access network.

This approach suffers from the following drawbacks:

1. To signal policy rules only within the access network (by both end-points) has a number of disadvantage and challenges (see for example [5] ). The complex message processing caused by this approach strongly argues against it although it might sound simple (and even might be simple in restricted environments).

2. Complex topologies might lead to ineffective policy rules (i.e. data traffic hits firewalls hits wrong firewalls).

### D.3 Solution Approach: Authorization Tokens

The last approach is based on some exchanged authorization tokens which are created by an authorized entity (such as the PDP) or by a trusted third party. Both end hosts need to exchange these tokens

with protocols such as SIP or HTTP since these protocols are likely to be allowed to bypass the firewall. The basic idea of this approach is to provide an end host, which requests access to the network, with credentials (referred as authorization tokens). These tokens have to possess some properties, namely:

1. They have to be restrictive by including lifetimes, source and destination identifiers, usage indication and more.

2. They have to provide basic replay protection to prevent unauthorized reuse.

3. The have be cryptographically protected to prevent manipulations.

4. There has to be a mechanism to dynamically create them for a specific reason and to distribute them to the end points.

5. It has to be possible to exchange tokens via a trusted third part in cases where no direct communication between the end hosts is possible (due to NAT).

6. The token can be created locally at the network or by a trusted third party.

An example of a possible signaling communication could have the following structure: After exchanging the tokens between the two end hosts. Host A would include the received authorization token to the signaling message for Network B. When the signaling message arrives at Middlebox 2 then the token is verified by the token-creating entity. In order to prevent parties from reusing the token timestamps (e.g. token creation, token lifetime, etc.) have to be included. Adding IP address information about Host A would create difficulties in relationship with NATs. Information about Host B might be possible to include in order to limit attacks where a token is lost and reused by a different host for a different purpose. The goal is to restrict the usage of the token for a specific session. The content of the token only needs to be verified by the originator of the token since it only has to be verified locally. Since authorization needs to be linked to the authorized actions, which have to be performed on the packets matching the packet filter, the token may include the associated action or a reference to it. The following behavior has to be assumed in order for this approach to be applicable:

1. The exchange of authorization tokens between end-systems must be possible. These protocols must be allowed to pass the firewalls.

2. An end-system must be able to request such an authorization token

at some entity in the local network or at a trusted third party.

This approach suffers from the following drawback:

1. Possibly an additional protocol is required for an end host to request an authorization token from an entity in the local network.

## Intellectual Property Statement



## Full Copyright Statement

## Acknowledgement

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