Internet Engineering Task Force NSIS
Internet Draft
H. Tschofenig, D. Kroeselberg
Siemens AG

draft-ietf-nsis-threats-01.txt
23 January 2003
Expires: August 2003

Security Threats for NSIS

STATUS OF THIS MEMO

Abstract

This threats document provides a detailed analysis of the security threats relevant for the NSIS working group. It motivates and helps to understand various security considerations in the NSIS Requirements, Framework and Protocol proposals. This document does not describe vulnerabilities of specific NSIS protocols.

1 Introduction

Section 1.1 introduces the overall process of addressing the security work done in the NSIS working group. Section 1.2 gives a high-level picture of the different network parts, which are traversed by NSIS signaling. Each part is characterized by a different set of requirements and different trust relationships. The threats described in Section 2 can be assigned to these individual parts.

1.1 NSIS Security Process

Whenever a new protocol has to be developed or existing protocols have to be modified their security threats should be evaluated. The process of securing protocols in separated into individual steps. To address security in the NSIS working group a number of documents have been produced:

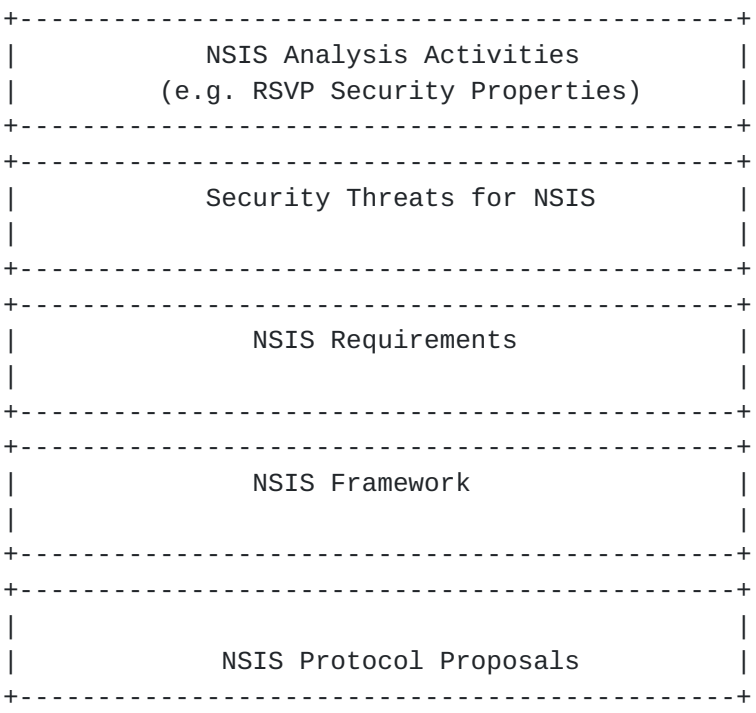

Figure 1: NSIS Security related Documents

All the documents depicted in Figure 1 contribute to the NSIS security approach. The purpose of each of these documents is briefly described below to give the reader insight into the development process.

NSIS Analysis Activities:

The primary goal of the NSIS analysis activity is the investigation of existing approaches in the area of quality of service signaling protocols. Several of the published approaches directly identify security threats and requirements, whereas other threats and requirements can be derived from the different scenarios in which these protocols are used. For instance, [1] points to the reduced complexity if RSVP is used without multicast support. This modification also results in simplified security requirements. In [2], security issues in some example configurations are given. In [3], the security properties of RSVP are described. Furthermore an analysis of the interaction between RSVP and Mobile IP is provided by Michael Thomas in [4] and an analysis of existing QoS protocols is described in [5].

NSIS Requirements:

To address the security threats relevant for NSIS described in this document, security requirements have been specified as part of the NSIS Requirements document [6]. In addition to these requirements [6] describes basic scenarios where the NSIS signaling protocol might be deployed.

NSIS Framework:

Signaling information to a number of devices located in different parts in the network with different trust assumptions and possible interactions with a large number of other protocols require some framework thoughts, which is especially true for security. In [7] a security framework is provided for NSIS.

NSIS Protocol:

Finally there are documents describing concrete protocol proposals. These proposals either rely on existing security mechanisms or develop their own if the existing mechanisms cannot counter all relevant security threats or if they are inappropriate for other reasons. In practice, a protocol proposal might use established security mechanisms or protocols for basic protection, but is likely to require some additional protection mechanisms, or a combination of both for

enhanced security.

Note that the process of developing the above-mentioned documents is not linear. Instead it takes various iterations to reach a satisfactory NSIS security solution.

Security Threats for NSIS:

This document identifies the basic threats that need to be addressed by the NSIS signaling protocol design. In addition, although the base protocol might be secure, some extensions may cause problems when used in a particular environment. Furthermore it is necessary to investigate the context in which a signaling protocol is used and the architecture where it is integrated. As an example of such interaction accounting and charging are taken into account in this document, since without an appropriate integration of the two it is difficult to deploy any NSIS solution. This interaction is also subject of the NSIS framework and some aspects are discussed in [7].

1.2 Relevant communication models

Independent of the threat scenarios described in Section 2 signaling messages traverse different network parts, which demand different security means. The difference in security protection is mainly caused by the fact that the NSIS signaling messages cross trust boundaries where different trust relationships are prevalent. Often a categorization into first-peer/last-peer, intra-domain and inter-domain communication is applicable (see Figure 2). Depending on the concrete security requirements end-to-end security protection across trust boundaries might be required for certain scenarios but is usually not easily addressable by standard means. The main properties of the listed network parts are briefly described in this section and the threat scenarios of Section 2 are classified accordingly. Figure 2 depicts a typical end-to-end communication scenario including an access part between the NSIS end entities and the nearest NSIS hops, respectively. This "first-peer communication" commonly comes with specific security requirements, especially important for properly addressing security in mobile scenarios. Differences in the trust relationship and the required security for first-peer communication, compared to other parts of the signaling path, might exist.

If signaling messages are not exchanged end-to-end and only parts of the signaling path are affected, some threats may not be relevant.

To further refine the above differentiation based on network parts that NSIS signaling may traverse, we consider trust relationships between

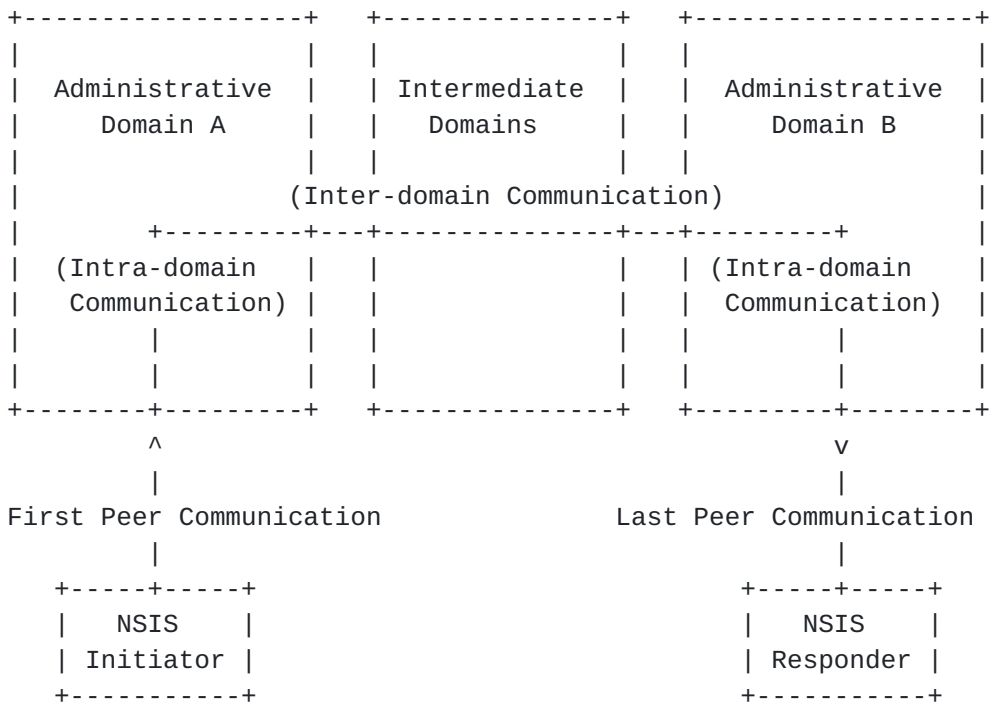

Figure 2: Involved Network Parts

NSIS hops. Additional threats may apply to NSIS communication where one entity involved is an end-entity (initiator or responder) and the other entity is any intermediate hop not being the first peer. This is typically called end-to-middle scenario. The motivation for including this configuration stems for example from the SIP [8] protocol. Any intermediate SIP proxy may request a SIP end entity (UA) to authenticate, countering a number of specific security threats. Such functionality in general seems to be useful for intermediaries at the borders of trust domains that signaling messages need to traverse. Intermediate NSIS hops as well may have to deal with specific security threats that do not (directly) relate to end-entities. Between such intermediate hops, other such NSIS hops will typically be in the signaling path. This scenario is called middle-to-middle. A generic example are two NSIS hops at the border of their respective trust domains with some form of trust relation. NSIS messages between these hops may have to traverse one or more intermediate untrusted hops. Figure 3 illustrates these additional scenarios. The first-peer case discussed further above is covered by the peer-to-peer trust relationships between end entity and closest hop, respectively.

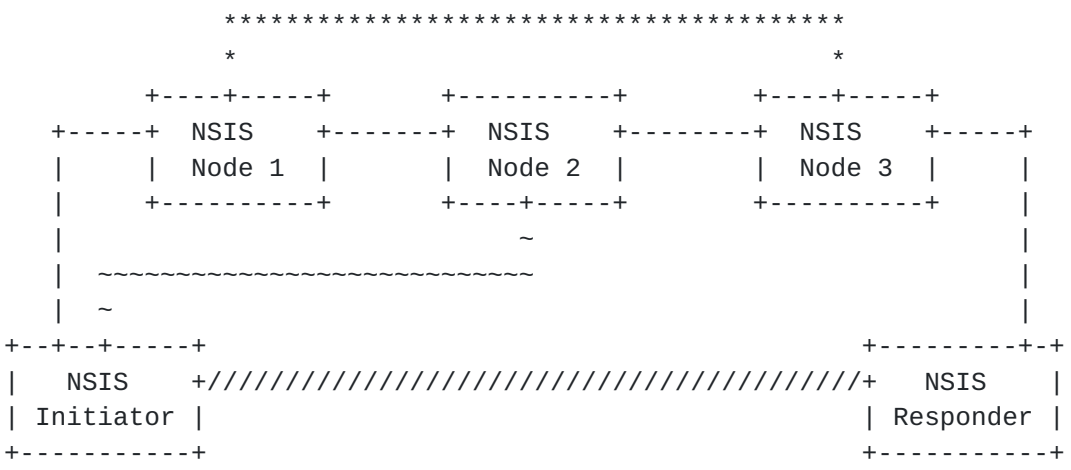

Figure 3: Trust Relationships

First-Peer Communication:

First peer communication refers to the peer-to-peer interaction between a signaling message originator, the NSIS Initiator (NI), and the first NSIS aware entity along the path. Assumptions about the threats, security requirements and the available trust relationships may be difficult here. To illustrate this, in many mobility environments it is difficult to assume the existence of a pre-established security association directly available for NSIS peers involved in first-peer communication, as these peers cannot be assumed to have any relation between each other in advance. For enterprise networks, in contrast, the situation is different. Usually there is a fairly strong (pre-established) trust relationship between the peers. Enterprise network administrators usually have some degree of freedom to select the appropriate security protection and to enforce it. The choice of selecting a security mechanism is therefore often influenced by the already available infrastructure. Per-session negotiation of security mechanisms is therefore often not required (which, in contrast, is required for the mobility

case).

For first-peer communication, especially threats related to initial security association setup, replay attacks, lack of confidentiality, denial of service, integrity violation, identity spoofing and fraud are applicable.

End-to-Middle Communication:

End-to-middle interaction in signaling may be required to e.g. grant end-entities access to, or specific services in trust domains different from the one the first peer belongs to. Threats, in addition to these already discussed for first-hop communication, may be untrusted intermediate NSIS hops that maliciously alter NSIS signaling. These threats are still relevant if security mechanisms are in place between the NSIS hops, but terminate at each hop (e.g. IPsec hop-by-hop protection).

Intra-Domain Communication:

After having been verified at the first peer, an NSIS signaling message traverses the network within the same administrative domain the first peer belongs to. Since the request has already been authenticated and authorized threats are different to those described above in a). To differentiate first-peer communication with the intra-domain communication (i.e. communication internally within one administrative domain) we assume that no end hosts have direct access to the internal network nodes, except the first peer. We furthermore assume that NSIS peers within the same administrative domain have at least some sort of trust relationship.

Inter-Domain Communication:

The threat assumptions between the borders of different administrative domains largely depend on accounting procedures (and therefore business relationships) in case of QoS signaling, which is an important example application of NSIS signaling. If one domain transmits forged QoS reservations (for example stating a higher QoS reservation than a aggregated number of user did) to the next domain then the originating domain may also have to pay for the reservation. Hence in this case, there is no real benefit for the first network domain to forge a QoS reservation. If an end host is directly charged by domains different to the first peer's domain, then such an attack may be quite a reasonable threat. However, security protection of messages transmitted between

different administrative domains is still necessary to tackle attacks like spoofing, integrity violation, or denial of service between these domains, e.g. to allow for proper accounting. In case of securing signaling messages between administrative domains, the number of domains is usually rather limited (compared to first-peer communication) which causes fewer problems for the key management.

Signaling information other than QoS service parameters such as policy rules in case of middlebox communication demands different assumptions for inter-domain communication. Trust assumptions and business relationships are of particular importance for their communication.

If signaling messages are transparent in the core network (i.e. the are not intercepted and processed in the core network) then the signaling message communication effectively takes place between access networks. This might place a burden on the key management infrastructure because of the global PKI requirements. Hence this can be seen as a serious deployment threat since it might be unacceptable for an access network service provider to perform processing (QoS reservations, policy rule installation at firewalls) due to unprotected incoming signaling messages.

End-to-End Communication:

Providing end-to-end signaling message protection for NSIS would cause difficulties for authentication and key establishment procedures. It would furthermore limit the flexibility of a signaling protocol in general. Functionality such as terminating at an arbitrary location along the path, delegating a signaling message exchange to other nodes, etc. would be difficult to achieve in a secure fashion. Protecting signaling messages end-to-end (in addition to peer-to-peer security) is in our opinion rarely required. This is based on the observation that end-to-end issues like charging and payment selection (i.e. which user has to pay for which part of a QoS reservation) are already securely negotiated by preceding upper layer protocols (for example SIP). Information carried within an NSIS signaling protocol for the purpose of charging is therefore assumed opaque to the NSIS protocol itself. Note that this observation makes some assumptions about the charging model and about the existence of a protocol interaction with AAA, QoS and an application layer protocol.

It is however possible to imagine a charging solution that requires end-to-end protection of information delivered within

the NSIS signaling protocol. The following example requires some sort of end-to-end protection: Alice wants Bob to pay for a QoS reservation (reverse charging). Bob wants to be assured that the QoS signaling message he receives was transmitted by Alice because he is only willing to pay for particular users and not for everyone. Hence Bob requires Alice to protect the reservation request.

Regarding end-to-end security one additional issue needs to be addressed - delegation. Whenever a signaling is addressed end-to-end and an arbitrary node along the path acts as a proxy on behalf of the other endpoint a delegation mechanism is required to provide secure interaction. This obviously leads to additional complexity in the area of end-to-end security, as an additional set of threats becomes relevant.

Middle-to-middle:

We do not explicitly consider the middle-to-middle case here, as this is already covered by either intra- or inter-domain communication depending on the location of the involved entities.

2 Threat Scenarios

This section provides threat scenarios that are applicable to signaling protocols. Note that some threat scenarios use the term user instead of NSIS Initiator. This is mainly because security protocols allow a differentiation between entities being hosts and users (based on the identities used).

2.1 MITM Attacks

Security protection of protocols is often separated into two steps. The first step provides entity authentication and key establishment whereas the second step provides message protection using the previously established security association. The first step usually tends to be more expensive than the second which is also the main reason for separation. If messages are transmitted very infrequently then these two steps are collapsed into a single and usually rather costly step. One such example is e-mail protection via S/MIME. A good example for an efficient two-step approach is provided by IPsec [9]. We use this separation to cover the different threats in more detail. The first paragraph describes security threats where two peers do not already share a security association, or do not use security mechanisms at all. The next paragraph describes threats which are applicable when a security association is already established. Finally a denial of service attack is described which is applicable to a signaling message when no

separation between SA establishment and signaling protection takes place.

Various security threat are caused by a protocol performing dynamic node discovery. These threats include Denial of Service attacks, which are among other threats described in Section 2.9. Note that the threats are largely independently of the discovery procedure (path discovery, next peer discovery or topology discovery).

1. Attacks during NSIS SA Establishment

   During the process of establishing a security association an adversary fools the signaling message initiator with respect to the entity to which it has to authenticate. The man-in-the-middle adversary is able to modify signaling messages to mount e.g. DoS attacks. In addition, it may be able to terminate NSIS messages of the Initiator and inject messages to a peer itself, therefore acting as the peer to the initiator and as the initiator to the peer. This results in the initiator wrongly believing that it talks to the "real" network whereas it is actually attached to an adversary. For this attack to be successful, pre-conditions have to hold which are described with the following two cases:

   - Missing Authentication

     The first case addresses missing authentication between the neighboring peers: Without authentication a NI, NR or NF is unable to detect an adversary. However in some cases protection available might be difficult to accomplish in a practical environment either because the next peer is unknown, because of misbelieved trust relationships in parts of the network or because of the inability to establish proper security protection (inter-domain signaling messages, dynamic establishment of a security association, etc.). If one of the communication endpoints is unknown then for some security mechanisms it is either not possible or very difficult to apply appropriate security protection. Sometimes network administrators use intra-domain signaling messages without proper security. Such a configuration would then allow an adversary on a compromised non-NSIS aware node to interfere with nodes running an NSIS signaling protocol. Note that this type of threat goes beyond a threat caused by malicious NSIS nodes (described in Section 2.8).

   - Unilateral Authentication

In case of a unilateral authentication the NSIS entity that does not authenticate its peer is unable to discover the man-in-the-middle adversary. Although authentication of signaling messages should take place between each peer participating in the protocol operation special attention is given here to first-peer communication. Unilateral authentication between end hosts and the first peer is still common today, but certainly opens up many possibilities for MITM attackers impersonating either the end host or the (administrative domain represented by the) first peer.

The two threats described above are a general problem of network access without appropriate authentication, not only for an NSIS signaling protocol. Obviously there is a strong need to correctly address them in a future NSIS protocol. The signaling protocols addressed by NSIS are different to other protocols where only two entities are involved. Note, that especially first-peer authentication is important, as the impacts of a security breach likely reach beyond the directly involved entities (or even beyond a local network).

Finally it should be noted that the signaling protocol should be considered as a peer-to-peer protocol where the roles of initiator and responder can be reversed at any time. This leads to the conclusion that unilateral authentication is not very useful for such a protocol. However there might be a need to have some form of asymmetry in the authentication process whereby one entity uses a different authentication mechanism than the other one. As an example the combination of symmetric and asymmetric cryptography should be mentioned.

- Weak Authentication

This threat addresses weak authentication mechanisms whereby information transmitted during the NSIS SA establishment process may leak passwords and/or may allow offline dictionary attacks. This threat is not specific to NSIS signaling protocols but may also be applicable and countermeasures must be taken.

2. Attacks during NSIS SA Usage

Once a security association is establish (and used to protect signaling messages) basic attacks are prevented. However, a malicious NSIS node is still able to perform various attacks as described in Section 2.8. Replay attacks, which can be a problem when a NSIS node crashes, restarts and performs state

re-establishment. Proper re-synchronization capability of the security mechanism must therefore address this problem.

3. Combining Signaling and SA Establishment

This threat covers an attack which allows an adversary to flood an NSIS node with bogus signaling messages to cause a denial of service attack.

When a signaling message arrives at a NSIS aware network element some processing is required. If this message contains security objects such as digital signatures and not security association is already available then some processing is required for the cryptographic verification. Since NSIS signaling should not require several roundtrips between two NSIS peers it is difficult to provide DoS protection mechanisms commonly found in authentication and key agreement protocols. If signaling messages furthermore aim to be idempotent and no security association should be created then some cryptographic mechanisms should be used with precaution (for example public key cryptography).

Additionally to the threat described above an incoming signaling message might require time consuming processing (computations, state maintenance, timer setting, etc) and communication with third-party nodes including policy servers, LDAP servers, etc. If an adversary is able to transmit a large number of signaling message (for example with QoS reservation requests) with invalid credentials then the verifying node may not be able to process further reservation messages by legitimate users.

2.2 Eavesdropping and Traffic Analysis

This threat cases covers adversaries which are able to eavesdrop signaling messages but are unable to actively participate in signaling message exchange (i.e. passive adversary). The collected signaling packets may serve for the purpose of traffic analysis or to later mount replay attacks as described in the Section 2.3. The eavesdropper might learn QoS parameters, communication patterns, policy rules for firewall traversal, policy information, application identifiers, user identities, NAT bindings and more.

2.3 Adversary being able to replay signaling messages

This threat scenario covers the case where an adversary eavesdrops and collects signaling messages and replays them at a latter point in time (or at a different place, or uses parts of them at a different place or

in a different way - e.g. cut and paste attacks). Without proper replay protection an adversary might mount man-in-the-middle, denial of service and theft of service attacks.

A more difficult attack that may cause problems even in case of replay protection requires the adversary to crash an NSIS aware node to loose state information (sequence numbers, security associations, etc.) and to be able to replay old signaling messages. This attack addresses re-synchronization deficiencies.

2.4 Missing Protection of Authorization Information

Authorization is an important step for providing resources such as QoS reservations, NAT bindings and pin-holed firewalls. Authorization information might be delivered to the NSIS participating entities in a number of ways.

One such approach is to use a successful authorization step done by a different protocol in a later NSIS signaling message by providing some sort of token. The functionality and structure of such an authorization token for RSVP is described in [10] and in [11].

The interaction between different protocols based on authorization tokens, however, requires some care. Using such an authorization token it is possible to link state information between different protocols. Returning an unprotected authorization token to the end host might allow an adversary (for example an eavesdropper) to steal resources. An adversary might also use the token to learn communication patters. An untrustworthy end host might also modify the token content.

Other authorization mechanisms might depend on availability of sufficient funds and therefore real-time information. Deployment threats of this kind are described in Section 2.14. The Session/Reservation Ownership problem can also be considered as an authorization problem. Details are described in Section 2.11. In enterprise networks authorization is often coupled with membership to a particular class user of users/groups. This type of information can either be delivered as part of the authentication and key agreement procedure or has to be retrieved via separate protocols from other entities. If an adversary manages to modify information relevant for determining authorization or the outcome of the authorization process itself then theft of service might be the consequence.

2.5 Identity Spoofing

Identity spoofing relevant for NSIS appears in two flavors: First, identity spoofing can appear during the establishment of a security association if based on a weak authentication mechanism.

Eve, acting as an adversary, claims to be the registered user Alice by spoofing the identity of Alice. Thereby Eve causes the network to charge Alice for the consumed network resources. This type of attack is possible if authentication is done based on a simple username identifier (i.e. in absence of cryptographic authentication) or if authentication is provided for hosts and multiple users have access to a single host. This attack could also be classified as theft of service.

Second, an adversary is able to perform identity spoofing on transmitted data packets. This type of attack is often labeled as IP spoofing. Since most NSIS signaling messages contain some sort of flow identifier for which a certain behavior is performed (e.g. particular flow experiences QoS treatment or is allowed to bypass a firewall, etc.) an adversary could mount an attack by modifying the flow identifier of a signaling message. The following example tries to show an adversary using identity spoofing of the first category:

An adversary is able to exploit the established flow identifiers (required for QoS and Midcom specific signaling protocols). Some identifiers such as IP addresses, transport protocol identifiers, port numbers, flow labels (see [12] and [13]) and others are communicated in these protocols. Modification of these flow identifiers cause quality of service reservations or policy rules at middleboxes to be either ineffective or beneficial for adversaries.

The following paragraph describes a possible threat caused by identity spoofing of transmitted data traffic. The spoofed identity is thereby the source IP addresses. For this attack to be successful accounting records are collected based on the source IP address and not on a SPI due to IPSec protection. After the network receives a properly protected reservation request, transmitted by the legitimate user Alice, Traffic Selectors are installed at the corresponding devices (for example edge router). These Traffic Selectors are used for flow identification and allow to match data traffic originated from a given source address to be assigned to a particular QoS reservation. The adversary Eve now spoofs the IP address of the Alice. Additionally Alice's host may be crashed by the adversary as a result of a denial of service attack or lost connectivity for example because of mobility reasons. If both nodes are located at the same link and use the same IP address then obviously a duplicate IP address will be detected. Assuming that only Eve is present at the link then she is able to receive and transmit data (for example RTP data traffic), which receives preferential QoS treatment based on the previous reservation. Depending on the installed Traffic Selector granularity Eve might have more possibilities to exploit the QoS reservation or a pin-holed firewall. Assuming the soft state paradigm, where periodical refresh messages are required, the absence of Alice will not be detected until the next signaling message appears and forces Eve to respond with a protected signaling message. Again this issue is

not only applicable to QoS traffic but the existence of QoS reservation causes more difficulties since this type of traffic is more expensive. The same procedure is also applicable to a Middlebox communication protocol.

The ability for an adversary to inject data traffic which matches a certain Traffic Selector established by a legitimate user often requires the ability to also receive the data traffic. This is, however, only true if the Traffic Selector consists of values which contain addresses used for routing. If we imagine to use attributes for a Traffic Selector where such a property is not required then identity spoofing and injecting traffic is much easier. An adversary can use a nearly arbitrary endpoint identifier to experience the desired result. Obviously the endpoint identifiers are still not irrelevant since the messages have to travel the same path through the network. DiffServ marking of IP packets is such an example and others can be constructed very easily.

Data traffic marking based on DiffServ is such an example. Whenever an ingress router uses only marked incoming data traffic for admission control procedures then various attacks are possible. These problems are known in the DiffServ community for a long time and documented in various DiffServ related documents. The IPSec protection of DiffServ Code Points is described in Section 6.2 of [14]. Related security issues (for example denial of service attacks) are described in Section 6.1 of the same document.

2.6 Adversary being able to inject/modify messages

This type of threat addresses integrity violations whereby an adversary modifies signaling messages (e.g. by acting as a man-in-the-middle attacker) to cause an unexpected network behavior. Possible actions an adversary might consider for its attack are reordering, delaying, dropping, injecting and modifying.

An adversary may inject a signaling message requesting a large amount of resources (possibly using a different user identity). Other resource requests could then be rejected. In combination with identity spoofing it is also possible accomplish fraud. This attack is only successful in absence of signaling message protection.

Some directly related threats are described in Section 2.8, 2.5, 2.8 and 2.9.

2.7 Missing Non-Repudiation

Repudiation in this context refers to a problem where one party later denies to have requested a certain action (such as a QoS reservation).

The problem of a missing non-repudiation property appears in two flavors:

>From a service provider point-of-view the following threat may be worth an investigation. A user may deny to have issued reservation request for which it was charged. A service provider may then like to prove that a particular user issued reservation requests.

The same threat can be interpreted from the users point-of-view. A service provider claims to have received a number of reservation requests. The user in question thinks that he never issued those requests and wants to have a proof for correct service usage for a given set of QoS parameters.

In today's telecommunication networks non-repudiation is not provided. The user has to trust the network operator to correctly meter the traffic, collect and merge accounting data and that no unforeseen problems occur. If a signaling protocol is used to establish QoS reservations with a higher volume (for example service level agreements) then it might impact protocol design.

Looking at threats based on missing non-repudiation it must be carefully considered whether non-repudiation is needed. Non-repudiation poses additional requirements on the security mechanisms as it can only be provided through public-key cryptography. As this would often increase the overall cost for security, threats related to missing non-repudiation are only considered relevant for certain specific scenarios but not for the general NSIS scenario.

2.8 Malicious NSIS Entity

Network elements within a domain (intra-domain) experience a different trust relationship with regard to the security protection of signaling messages compared to edge routers. We assume that edge routers have the responsibility to perform cryptographic processing (authentication, integrity and replay protection, authorization and accounting) for signaling message arriving from outside. This prevents signaling messages to appear unprotected within the internal network. If however an adversary manages to take over an edge router then the security of the entire network is affected. An adversary is then able to launch a number of attacks including denial of service, integrity violation, replay attacks etc. In case of policy rule installation a rogue firewall can cause harm to other firewalls by modifying the policy rules accordingly. The chain-of-trust principle applied in the peer-to-peer security protection cannot provide protection against a malicious NSIS node. An adversary with access an NSIS router is then also able to get access to security associations to transmit secured signaling messages. Note that even non peer-to-peer security protection might not be able to

fully prevent this problem. Since an NSIS node might issue signaling message on behalf of someone else (by acting as a proxy) additional problems are the consequence.

An NSIS aware edge router is a critical component that requires strong security protection. A strong security policy applied at edge does not imply that all routers within an intra-domain network do not need to cryptographically verify signaling messages. If the chain-of-trust principle is deployed then the security protection of the entire path (in this case within the network of a single administrative domain) is as strong as the weakest link. In our case the edge router is the most critical component of this network that may also act as a security gateway/firewall for incoming/outgoing traffic. For outgoing traffic this device has to act according to the security policy of the local domain to apply the appropriate security protection.

For an adversary to mount this attack either an existing NSIS aware node along the path has to be successfully attacked or an adversary succeeds to convince another NSIS node to be the next NSIS peer (man-in-the-middle attack).

2.9 Denial of Service Attacks

A number of denial of service attacks can cause NSIS nodes to malfunction. Other attacks that could lead to DoS, such as man-in-the-middle attacks, replay attacks, injection or modification of signaling messages etc., are mentioned throughout this document.

1. Path Finding

   This threat tries to address potential denial of service attacks when the reservation setup is split into two phases i.e. path and reservation (as for example used in receiver based reservation setup). For this example we assume that the node transmitting the path message is not charged for the path message itself and is able to issue a high number of reservation request (possibly in a distributed fashion). Charging is activated only after successful verification of the reservation request. The reservations are however never intended to be successful because of various reasons: the destination node cannot be reached; it is not responding or simply rejects the reservation. An adversary can benefit from the fact that resources are already consumed along the path for various processing tasks including path pinning.

2. Discovery Phase

Signaling information to a large number of entities along a data path requires some sort of discovery. This discovery process is vulnerable to a number of attacks since it is difficult to secure. An adversary can use the discovery mechanisms to convince an entity to signal information to another entity which is not along the data path or to cause the discovery process to fail. In the first case the signaling protocol could be correctly continued with the problem that policy rules are installed at incorrect firewalls or QoS resource reservations take place at the wrong entities. For an end host this means that the protocol failed for unknown reasons.

3. Faked Error/Response messages

An adversary may be able to use false error/response messages as part of a denial of service attack. This could be either at the message signaling protocol level, at the level of each client layer protocol (QoS, Midcom, etc.) or at the transport level protocol. An adversary might cause unexpected protocol behavior or produce denial of service attacks. Especially the discovery protocol shows vulnerabilities with regard to this threat. In case that no separate discovery protocol is used by addressing signaling messages to end hosts only (with a Router Alert Option to intercept message as NSIS aware nodes) then an error message might be used to indicate a path change. Such a design is a combination of a discovery protocol together with a signaling message exchange protocol.

2.10 Disclosing the network topology

In some architectures there is a desire not to reveal the internal network structure (or other related information) to the outside world. An adversary might be able to use NSIS messages for network mapping (e.g. discovering which nodes exist, which use NSIS, what version, what resources are allocated, capabilities of nodes along a paths etc.). Discovery messages, traceroute, diagnostic messages (see [14] for a description of diagnostic message functionality for RSVP), query messages in addition to record route and route objects provide the potential to assist an adversary. Hence the requirement of not disclosing a network topology might conflict with another requirement to provide means for automatically discovering NSIS aware nodes or to provide diagnostic facilities (used for network monitoring and administration).

2.11 Session/Reservation Ownership

Figure 4 shows an NSIS Initiator which established state information at NSIS nodes along the path as part of the signaling procedure. As a result the Access Router1 Router 3 and Router 4 (and other nodes) store session state information including the Session Identifier SID-x.

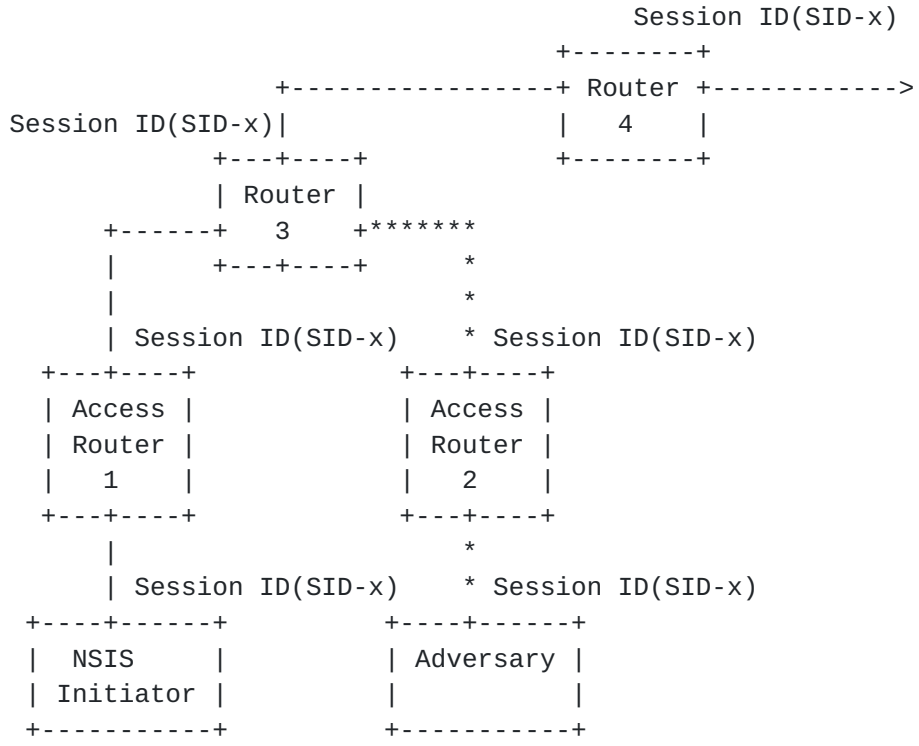

Figure 4: Session/Reservation Ownership

The Session Identifier is included in signaling messages to reference to the established state.

If an adversary was able to obtain the Session Identifier for example by eavesdropping signaling messages it is able to add the same Session Identifier SID-x to a new a signaling message. When the signaling message hits Router3 (as shown in Figure 3) then existing state information can be modified. The adversary can then modify or delete the established reservation causing unexpected behavior for the legitimate user.

The source of the problem is that Router3 (cross-over router) is unable to decide whether the new signaling message was initiated from the owner

of the session/reservation.

To make processing even more difficult it must be mentioned that not only the initial signaling message originator is allowed to signal information during the lifetime of an established session. As part of the protocol any NSIS aware node along the path (and the path might change over time) could be involved in the signaling message exchange and it might be necessary to provide mobility support or to trigger a local repair procedure. Hence if only the initial signaling message originator is allowed to trigger signaling message exchange some protocol behavior will not be possible.

In case that this threat is not addressed an adversary can launch denial of service, theft of service, and various other attacks.

2.12 Security Parameter Exchange/Negotiation

Protocols, which should be useful for a variety of scenarios, tend to have different security requirements. It is often difficult to meet these (sometimes conflicting requirements) with a single security mechanism or a fixed security parameter. Hence often a few selected mechanisms/parameters are supported. Therefore some protocol exchange is required to agree on some security mechanisms/parameters. This protocol exchanged can be the misused by an adversary to mount a downgrading attack by selecting weaker mechanisms than desired. Hence without protecting the negotiation process the security of an NSIS protocol might be as secure as the weakest mechanism if no configuration parameters (for example a security policy disallowing the weakest mechanism, etc.) are used otherwise.

2.13 Attacks against the signaling message transport mechanism

In [15] a two-level architecture is proposed which suggests to split an NSIS protocol into layers: a signaling message transport specific layer and an application specific layer. This architectural assumptions is also considered within the NSIS framework [7]. Most of the threats described in this document are applicable to the application specific part for signaling QoS or middlebox specific information. There are, however, some threats which are applicable to the transport of signaling messages.

Network or transport layer protocols which experience no protected are vulnerable to certain attacks such as header manipulation, DoS, spoofing of identities, session hijacking, unexpected aborts etc.

In case that an existing protocol is used for exchanging NSIS signaling messages then threats known from these protocols are relevant.

2.14 Deployment Threats

This section addresses problems which could appear during the deployment of an NSIS protocol in a real-world environment. Although these problems are theoretically not an obstacle for practical reasons they can represent threats worth a consideration.

Missing Authorization:

Authentication is a very important step for providing the foundation of authorization and accounting. Unlike some other protocols (for example HTTPS) where an authorization verification step is fairly easy (and efficient) QoS and middlebox communication requires more care. First, there is the question what authorization means in the context of NSIS signaling. For quality of service signaling the possible range is broad and could range from pure monetary policies to traditional role-based access control policies. Second, there is a question where this authorization data can be retrieved. Especially in a mobile environment this might be more complicated to securely exchange this information between different network domains. Finally there is an issue of authorization representation (i.e. a language describing authorization policies). If authorization information is exchanged between a large number of networks then this issue deserves further consideration.

In the discovery phase an additional issue of authorization was raised. Whenever a node wants to discover the next NSIS aware node then authentication might not be sufficient. In many cases the IP address or FQDN of a particular router in an unknown network does not add too much trust. An end host for example might want some assurance that this node belongs to a network which has some sort of business relationship which is known and acceptable (from an accounting, charging, security and privacy point of view).

Missing Cost Control:

There is a risk that a large number of service providers with complex roaming agreements create a non-transparent cost-structure. In a traditional subscription-based scenario users are registered with their home networks and use this trust relationship to dynamically establishment other security associations. This is the typical AAA deployment scenario. In these scenarios users do not learn the identity of the access network as part of a regular authentication and key exchange protocol message exchange. The identity of the access network

is possibly never revealed (in a secure fashion). The user is therefore only authenticated to the home network (and hopefully vice versa). When issuing a QoS reservation request to the next NSIS peer (for example in the access network) the end host is typically unaware of the cost of such a reservation. Due to mobility and route changes along the path the cost for an end-to-end QoS reservation might not be transparent for the end host or even become unacceptable.

Today there is no standarized protocol available which allows users to communicate cost limits, to request cost information for network resources or to learn already accumulated costs for a particular reservation.

Especially in mobility environments where an end host is likely to have access to a large number of networks within a short time period cost control is even more complicated.

Some mobility/QoS protocol proposals try to merge existing mobility protocols with QoS signaling (i.e. to apply in-band signaling). Thereby the access network is queried (towards the cross-over router or the MAP) for the possibility making a QoS reservation (without actually making the reservation itself). Without a query mechanism a user cannot take reservation costs into account when choosing between different access networks (or different access routers). Hence the user might not be unable to refuse a more expensive service provider. To allow a user to choose between different providers might be required not only because of the availability of different access technologies (e.g. IEEE 802.1x, Bluetooth, UTRAN) and the different service quality offered but also for cost reasons.

Although real-time notifications of quality of service reservation costs (cost control) to the user are outside the scope of NSIS some interaction might be required.

3 Security Considerations

This entire memo discusses security issues relevant for NSIS. To counter these threats security requirements have been defined and the framework relevant topics have been described. Some additional threats applicable for first peer communication in mobile environments are described in [16].

4 Acknowledgements

We would like to thank (in alphabetical order) Marcus Brunner, Jorge Cuellar, Mehmet Ersue, Xiaoming Fu and Robert Hancock for their comments

to this draft. Jorge and Robert gave us an extensive list of comments and provided information on additional threats.

5 Authors' Addresses

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Hannes.Tschofenig@siemens.com

Dirk Kroeselberg
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Dirk.Kroeselberg@siemens.com

6 Bibliography

[1] X. Fu, C. Kappler, and H. Tschofenig, "Analysis on RSVP regarding multicast," Internet Draft, Internet Engineering Task Force, June 2002. Work in progress.

[2] H. D. Meer et al. , "Analysis of existing qos solutions," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[3] H. Tschofenig, "Rsvp security properties," Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[4] M. Thomas, "Analysis of mobile ip and rsvp interactions," Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[5] J. Manner and X. Fu, "Analysis of existing quality of service signaling protocols," Internet Draft, Internet Engineering Task Force, 2002. **Work in progress.**

[6] M. Brunner, "Requirements for QoS signaling protocols," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[7] R. Hancock, I. Freytsis, G. Karagiannis, J. Loughney, and S. V. den Bosch, "Next steps in signaling: Framework," Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[8] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: session

initiation protocol," RFC 3261, Internet Engineering Task Force, June 2002.

[9] S. Kent and R. Atkinson, "Security architecture for the internet protocol," RFC 2401, Internet Engineering Task Force, Nov. 1998.

[10] L. Hamer, B. Gage, M. Broda, B. Kosinski, and H. Shieh, "Session authorization for RSVP," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[11] L. Hamer, B. Gage, and H. Shieh, "Framework for session set-up with media authorization," Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[12] C. Partridge, "Using the flow label field in IPv6," RFC 1809, Internet Engineering Task Force, June 1995.

[13] J. Rajahalme, A. Conta, B. Carpenter, and S. Deering, "IPv6 flow label specification," Internet Draft, Internet Engineering Task Force, June 2002. Work in progress.

[14] A. Terzis, B. Braden, S. Vincent, and L. Zhang, "RSVP diagnostic messages," RFC 2745, Internet Engineering Task Force, Jan. 2000.

[15] B. Braden and B. Lindell, "A two-level architecture for internet signaling," Internet Draft, Internet Engineering Task Force, Nov. 2001. Work in progress.

[16] J. Kempf and E. Nordmark, "Threat analysis for IPv6 public multi-access links," Internet Draft, Internet Engineering Task Force, June **2002**. **Work in progress.**

Table of Contents

1 Introduction 2
1.1 NSIS Security Process 2
1.2 Relevant communication models 4
2 Threat Scenarios 9
2.1 MITM Attacks 9
2.2 Eavesdropping and Traffic Analysis 12
2.3 Adversary being able to replay signaling messages 12
2.4 Missing Protection of Authorization Information 13
2.5 Identity Spoofing 13
2.6 Adversary being able to inject/modify messages 15
2.7 Missing Non-Repudiation 15
2.8 Malicious NSIS Entity 16
2.9 Denial of Service Attacks 17
2.10 Disclosing the network topology 18
2.11 Session/Reservation Ownership 18
2.12 Security Parameter Exchange/Negotiation 20
2.13 Attacks against the signaling message transport mechanism 20
2.14 Deployment Threats 21
3 Security Considerations 22
4 Acknowledgements 22
5 Authors' Addresses 23
6 Bibliography 23