

Securing the Next Steps in Signaling (NSIS) Protocol Suite

Hannes Tschofenig

Siemens AG, Corporate Technology
Otto-Hahn-Ring 6, Munich 81739, Germany
Fax: +49 89 636 48000, E-mail: hannes.tschofenig@siemens.com

Xiaoming Fu

University of Göttingen, Institute for Informatics
Lotzestr. 16-18, Göttingen 37083, Germany
Fax: +49 551 39 14403, E-mail: fu@cs.uni-goettingen.de

Abstract: The Next Steps in Signaling (NSIS) protocol suite represents an extensible framework for enabling various signaling applications over IP-based networks. The framework consists of two layers that need different types of security protection: the lower layer mainly deals with the discovery of adjacent peers, establishment of channel security to protect the delivery of signaling messages between two peers, while the upper layer provides the signaling application specific functionalities. Different security properties are desired at the two layers with stronger authorization functionality at the signaling application layer. In this paper we examine how various security vulnerabilities can be utilized by an adversary, including eavesdropping, man-in-the-middle attacks, fraud and denial of service attacks. Moreover, we describe how to protect against a number of selected security threats and highlight some security challenges that require further research.

Keywords: QoS Signaling; RSVP; Next Steps in Signaling (NSIS); General Internet Signaling Transport (GIST); Security; AAA.

Reference to this paper should be made as follows: Tschofenig, H. and Fu, X. (2005) 'Securing the Next Steps in Signaling (NSIS) Protocol Suite', *International Journal of Internet Protocol Technology*, Vol. 1, No. 4, pp. xxx-xxx.

Biographical notes: Hannes Tschofenig received his Diplom degree in Computer Science from University of Klagenfurt, Austria in 2001. Since then he is a Research Scientist at Corporate Technology, Siemens AG, Germany. His primary research interests lie in network security, with a focus on mobile communications and QoS contexts. He is Secretary of the NSIS working group and Chair of ECRIT working group of the IETF. He is a co-author of RFC 3726 'Requirements for Signaling Protocols', RFC 4081 'Security Threats for Next Steps in Signaling (NSIS)', RFC 4230 'RSVP Security Properties', and RFC 4279 'Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)', as well as an editor of 'NAT/Firewall NSIS Signaling Layer Protocol (NSLP)', 'Design of the MOBIKE Protocol', 'Protocol for Carrying Authentication for Network Access (PANA)', 'Carrying Location Objects in RADIUS', among others.

Xiaoming Fu received his Ph.D. in Computer Science from Tsinghua University, China in 2000. He joined the Institute for Informatics, University of Göttingen, Germany as an Assistant Professor (C1) in 2002. Prior to that he was member of research staff at Technical University Berlin, Germany. He is a co-author of RFC 4094 and about 40 refereed papers. His research interests encompass network architectures, protocols and distributed systems design, implementation, verification and performance analysis, particularly on signaling and QoS, security, overlay networks, wireless and mobile networking. In these areas he currently leads several EU- and industry-sponsored research projects and actively contributes to the Internet community.

The Internet is continuously evolving from a traditional packet-switching network to a network with a number of architectural changes, for example, due to the rise of QoS services (1; 2) and stateful middleboxes (3). Subsequently, there has been a general need to employ signaling protocols to install, maintain and remove control states in network nodes which are related to end-to-end communications. Several proposals have been submitted to the IETF to provide such functionality, one of the pioneering work of which is the Resource Reservation Protocol (RSVP) (4).

The strengths of RSVP are that it provides a soft-state management mechanism, a modular design to accommodate QoS reservation state setup among multicast and unicast networks, and a separation between signaling and routing (5). However, RSVP has several key weaknesses, for example, in its message delivery mechanism, protocol complexity, inadequate support in extensibility (to support other signaling applications than just QoS signaling), and inability to support mobility (6). Furthermore, as pointed out in (7) and in Section 3 of this paper, when QoS (or any other stateful middlebox) service provisioning is concerned, the control plane (signaling) and data plane functions can suffer from various attacks, e.g., denial of service attacks on reservation setup or consuming legitimately reserved resources by injecting false packets. Here, security of control plane state installation, maintenance and removal, namely the operations of a signaling protocol, is critical for the success deployment of such new services and applications. However, security was only patched to RSVP later (8; 9; 10) and RSVP lacks a solid security framework especially for end-to-end addressed signaling messages (e.g., PATH) (7; 11; 12). In particular, discovery and signaling message delivery are combined into a single protocol step in RSVP. This design decision turns out to make it difficult to provide proper security protection using existing security protocols as well as to extend the protocol into other deployment environments. In addition, authentication and key management is not adequately addressed. For example, only manual configuration is supported for the Integrity object (8), which protects the entire RSVP signaling message. Authorization aspects are provided to some degree (see discussions in Section 3.6) but do not interwork with today's AAA infrastructure (such as Diameter (13) or RADIUS (14)) and roaming environments.

The Next Steps in Signaling (NSIS) working group of the IETF is working on a signaling protocol suite aiming to overcome the shortcomings of RSVP while learning as much as possible from its experiences. Realizing today's Internet has become a hostile environment, and security functions in general contribute a considerable amount of complexity and performance impacts to the design of a protocol, security aspects have been analyzed in the NSIS design from early stages of the protocol work (see, for example, (15)). We believe addressing security will exhibit

as an important prerequisite for NSIS's success. However, to the best of our knowledge, there is no single paper systematically summarizing these aspects and investigating the possible overall security design space for the NSIS protocol suite.

In this paper we describe NSIS communication models and analyze various threats for the NSIS protocol suite. In addition, we describe how to deal with a selected set of security threats for NSIS protocols, especially in some specific application domains.

2 The NSIS Protocol Suite: Preliminaries

2.1 NSIS Overview

The NSIS protocol suite is designed to comprise two functional layers (16). The lower layer (the NSIS Signaling Transport Layer Protocol, or NTLP) provides a generic delivery service for different signaling applications, based on the General Internet Signaling Transport (GIST) (17) protocol. Signaling applications reside in the upper layer (the NSIS Signaling Layer Protocols, or NSLPs). Current NSLPs include signaling applications for QoS resource reservation (QoS-NSLP) (18) and NAT/Firewall traversal (NAT/FW-NSLP) (19). An introduction of the NSIS protocol suite is given in (20). In addition to a separation between signaling applications and generic delivery services, NSIS differs from RSVP in several other key design choices:

- Reuse of existing transport and security protocols, instead of designing a new transport protocol and later adding protocol properties such as reliability and security,
- (Whenever possible) decoupling discovery of the next signaling node from delivery of signaling messages between neighboring signaling nodes, and
- Introduction of a session identifier, which is carried in signaling messages to uniquely identify installed state, instead of reusing the flow identifier as done in RSVP.

Through GIST, NTLP functions as a universal “messenger” in routing and delivering signaling messages for all signaling applications. First, GIST installs and maintains a “routing” state used to direct signaling messages forwards or backwards along the flow path. In addition, a “messaging” state is maintained at the NTLP level for multiple signaling sessions to reuse existing transport associations between neighboring NSIS peers. These two types of state are different from signaling application (NSLP)-specific states, such as QoS reservation states or NAT bindings.

GIST defines two modes of delivering signaling messages. The first mode, known as Datagram mode (D-mode), follows an RSVP signaling style by using end-to-end addressed messages. The end-to-end addressed message contains the source and the destination IP addresses

of the data flow. The messages are intercepted along the path by NSIS nodes interested in these messages (by using the Router Alert Option (24)). The second mode, so-called Connection mode (C-mode), is used when NSIS nodes are directly addressed. This mode assumes that a discovery procedure should have been performed (or the address of the receiving node is known via other means, e.g., by manual configuration or a new discovery mechanism). The default discovery mechanism is based on a Query-Response message exchange using D-mode encapsulation, which afterwards allows establishing a C-mode messaging association (upon which NSIS signaling messages can be delivered). Any NSIS node that implements the desired NSLP functionality, upon receipt of a Query D-mode message, will respond with a Response message to the D-mode message sender. Therefore, GIST in D-mode for NSIS message delivery essentially makes no much difference from RSVP: both provide unreliable transport for delivering signaling messages, maintaining soft state for routing of the signaling messages. Therefore, the focus here is on securing GIST C-mode message delivery. However, unlike in RSVP, GIST does not employ two-way signaling message exchange, nor does it install signaling application specific state (QoS reservation in particular), since they are related to specific signaling applications, such as QoS reservation, firewall pinhole or NAT binding configuration. These function are now part of the NSLP layer functionality, e.g., a Reserve message possibly followed by a Response. Fig. 1 compares these two signaling approaches (RSVP and NSIS) according to the necessary function components. More discussions about the difference are given in (20).

We also discuss to some extent how to secure D-mode discovery process as it forms a basis of initializing the C-mode operations.

2.2 NSIS Signaling Scenarios

An example NSIS signaling scenario is shown in Fig. 2. Each NSIS node may store messaging association state information about its peers. A node, the NSIS Initiator (NI), initiates the signaling exchange, while some nodes along the signaling path, called NSIS Forwarders (NFs), intercept and then forward signaling messages, and the NSIS Responder (NR) terminates the signaling.

Fig. 2 also shows that not all routers along the data path need to be NSIS aware nor do all NSIS nodes necessarily support all signaling applications. For a particular NSIS session, nodes not supporting the desired signaling application are bypassed. In this example, messages of signaling application type A will be delivered between Router 2 and the edge node, without being processed in Router 3.

The NSIS protocol suite is envisioned to support various signaling applications that need to install and manipulate these application-specific states as well as message routing and transport states at signaling nodes along the data flow path through the network. Unlike many other protocols, which work in an end-to-end fashion and involve no in-

termediate nodes, NSIS is a protocol suite for distributed state establishment with the ability to interact with intermediaries. The set of nodes participating in the chain in the signaling communication can change over time, e.g., due to re-routing events and possibly node mobility. As such, the communication between an NSIS Initiator and an NSIS Responder (i.e., the two signaling ends) is more complex. Furthermore, for deployment reasons it is likely that the proxy signaling functionality (as developed with RSVP proxy (21)) will be desirable, whereby the data receiver and or the data sender are not NSIS aware.

As illustrated in Fig. 3, the entire NSIS end-to-end communication path can be split into different parts, namely first-peer, last-peer, intra-domain, or inter-domain.

For example, in some cases it may be necessary to allow an NSIS signaling node to explicitly authorize a non-adjacent NSIS node. Each of these models can represent some different impacts on security requirements.

First-Peer and Last-Peer Communication: First-peer (and last-peer as a variant) communication can vary from one access scenario to another. For example, in an enterprise network scenario, there can be a pre-established security association between the NSIS entities for a first- or last-peer communication. In a roaming scenario, it is difficult to assume a pre-established key between a mobile node and the attached network (referred as visited network) since the mobile node will most likely be unknown to the visited network. Generally, in a hostile environment, it would be desirable to perform a proper authentication and key exchange to establish the necessary pre-condition for providing channel security.

Intra-Domain Communication: When an NSIS signaling message travels between NSIS nodes belonging to the same administrative domain, authorization and key management can be simpler. However, security protection is still desired to prevent non-NSIS nodes from interfering with the NSIS-aware signaling nodes.

Inter-Domain Communication: Inter-domain communication will be necessary, e.g., when a QoS signaling request needs to traverse multiple domains, or when one network wants to signal a QoS reservation towards a neighboring domain. Note that the difference between these two cases is mainly related to the granularity of the QoS reservation request. In the former case a per-flow reservation needs to be made, whereas in the latter case most likely a reservation for an aggregated flow is desired. Neighboring domains can establish the necessary infrastructure, such as key distribution and contractual aspects, to subsequently secure signaling messages.

In short, the NSIS protocol suite needs to be able to protect signaling messages and payloads between adjacent NSIS nodes, between non-adjacent NSIS nodes over multiple NSIS hops (such as required by middle-to-middle

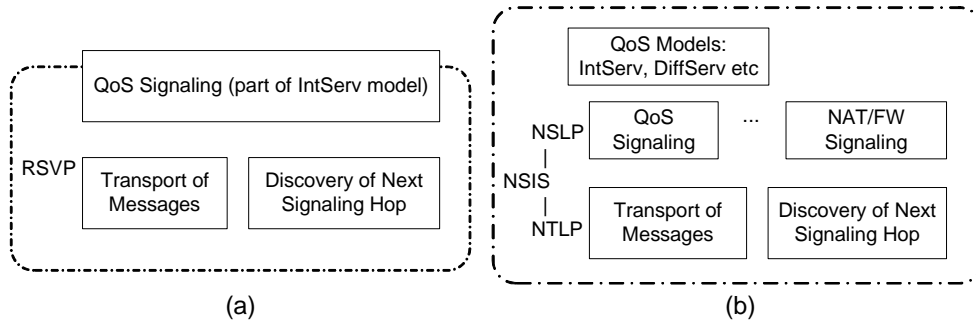


Figure 1: RSVP vs. NSIS: different approaches of signaling functionality composition




Figure 2: An NSIS signaling scenario between a host and an edge node

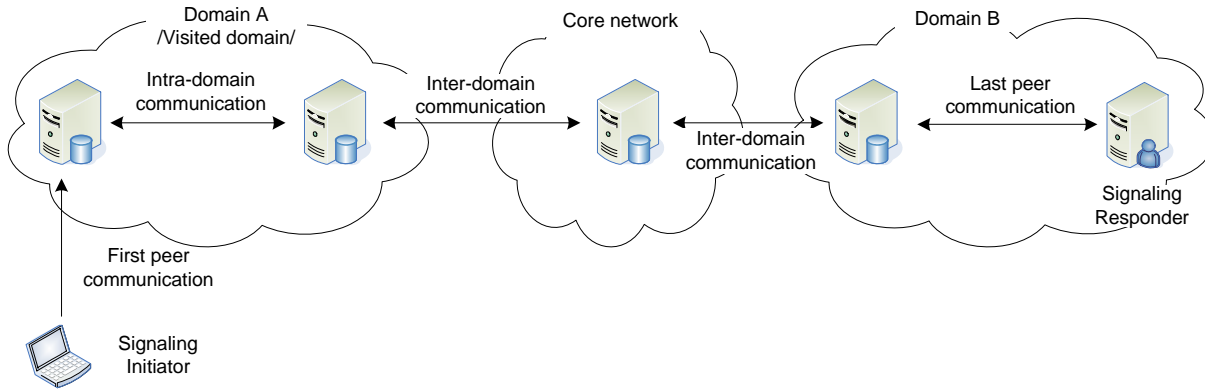


Figure 3: An end-to-end signaling environment

protection of payloads), end-to-middle or middle-to-end protection and end-to-end protection. End-to-middle (or middle-to-end) protection is necessary when intermediate NSIS nodes are not allowed to eavesdrop or to modify certain objects, or are required to cryptographically verify a certain payload. As an example, the protected exchange of signaling messages (or selected payloads) from an ingress towards an egress node of an administrative domain might be desired as described in RMD (22). End-to-end protection of the entire signaling message might not be possible (or useful) since intermediate NSIS nodes need to add, inspect, modify or delete objects in this message. However, some objects may need to be protected between the signal-

ing initiator and the signaling responder (i.e., end-to-end) or between non-neighboring NSIS nodes.

3 Security Threats

Based on the discussions in Section 2, we analyze the various threats to the NSIS protocol suite in more detail. For the convenience of discussion we follow the definition in (7) and differentiate adversaries into three types: *InsiderNSIS*, *OutsiderOnPath*, and *OutsiderOffPath*. These terms refer to participating NSIS nodes, traversed non-NSIS nodes along the signaling path, and nodes not along the signaling

path, respectively. Note that this section does not assume any particular security solutions for NSIS.

3.1 Eavesdropping and Traffic Analysis

An *OutsiderOnPath* and an *InsiderNSIS* adversary can eavesdrop NSIS signaling messages, and use the collected signaling packets to perform traffic analysis, for example, learning QoS parameters, communication patterns, policy rules for firewall traversal, policy information, application identifiers, user identities, NAT bindings, authorization objects, network configuration and performance information, etc. Based on collected information, an eavesdropper can also perform replay attacks (see Section 3.2).

An adversary's capability to eavesdrop on signaling messages may violate a user's preference for privacy, particularly if unprotected authentication or authorization information (including policies and profile information) is exchanged. Note, for certain objects or messages it is highly desirable to permit actively participating intermediate NSIS nodes to inspect either all or most of the signaling message payloads, in order to allow these nodes to perform the desired protocol processing. Hence, *InsiderNSIS* nodes are also considered as a potential threat source for the NSIS protocol suite.

3.2 Replay Attacks

An *OutsiderOnPath* and an *InsiderNSIS* adversary can replay (or reorder) eavesdropped signaling messages. Making the assumption of *InsiderNSIS* adversaries makes the protocol interaction very difficult and is likely not to be useful for many signaling applications due to the nature of the signaling interaction. With a QoS signaling, for example, a mobile host might request a certain QoS treatment from a QoS routers in the access network. The access network might need to forward the QoS signaling message towards the destination address and will therefore be charged by its neighboring network accordingly, too. The existence of such an authorization relationship between neighboring entities and the utilized charging model does not make it beneficial for the access network node to replay any QoS reservation requests for end hosts.

Without data origin authentication, integrity, replay and optional confidentiality protection of signaling messages an adversary can impact the functionality of a network considerably. For example, a denial of service attack can be launched by replaying past QoS signaling requests to install faked QoS reservations. Replaying refresh messages can also be exploited to prevent soft state timeouts. This threat may be particularly severe in mobile environments.

3.3 Man-in-the-Middle Attacks

The discovery phase and the establishment of a messaging association is particularly vulnerable to man-in-the-middle (MITM) attacks. In the discovery phase, an *OutsiderOnPath* adversary can inject a bogus response message, forcing the querying node to start a messaging asso-

ciation establishment procedure with either an adversary (or with another NSIS node which is not on the signaling path). Note that an adversary located in a broadcast medium, such as Ethernet or Wireless LAN, where where ARP spoofing is possible is referred as an *OutsiderOnPath* entity.

Clearly, for end-to-end addressed messages such attacks are possible, particularly if the adversary is located along the path and able to intercept the discovery message which traverses the adversary. The MITM adversary can redirect signaling messages to another legitimate NSIS node that might not even be located at the signaling path.

Besides, without a proper authorization procedure, an NSIS node implementing the functionality of NSLP *X* can pretend to implement the functionality of NSLP *Y*. Therefore, it is necessary to have a trusted third party to warrant for the capabilities of individual routers. Furthermore, a malicious non-NSIS node can be detected with the corresponding security mechanisms. A legitimate NSIS node, which is not the next NSIS node along the path, cannot be detected without using topology knowledge as part of the authorization decision. A more detailed discussions of the problem associated with this type of authorization can be found in (23).

In the example shown in Fig. 4 we investigate the NSLP peer discovery exchange used in context of C-Mode messaging. The subsequent paragraph provides a high-level and abstract description of the exchange. The details of the discussed proposal can be found in (17). Hence, this message is addressed towards the data receiver's IP address will be intercepted by an intermediate NSIS aware node.

An *OutsiderOnPath*, *OutsiderOffPath* or a malicious *InsiderNSIS* may respond to the discovery with a Query-Response with its own IP address as the address of the next NSIS aware node along the path. When the adversary is located in a same physical network as the legitimate NSIS node, the attacker may succeed if its "response" reaches the querying node faster than the legitimate response. Without any additional information the querying node has to rely on the data returned by the Query-Response message. Then a messaging association is established with an entity at a given IP address (IP_x) in step (3). The adversary may then again establish a messaging association with the next NSIS node to forward the signaling message. Note that the adversary can just modify the Query-Response message, forcing the querying node to establish a messaging association with any other NSIS node that is not even along the data path.

As a variant of this attack, an *OutsiderOffPath* adversary can also flood a node with bogus discovery reply messages, even if it is not able to eavesdrop transmitted discovery requests. If the discovery message sender accidentally accepts one of those bogus messages then a MITM-attack as described in Fig. 4 is possible. It should be noted that the process is self-healing since the discovery process is periodically performed. If an adversary is unable to mount this attack with every discovery message, the correct next

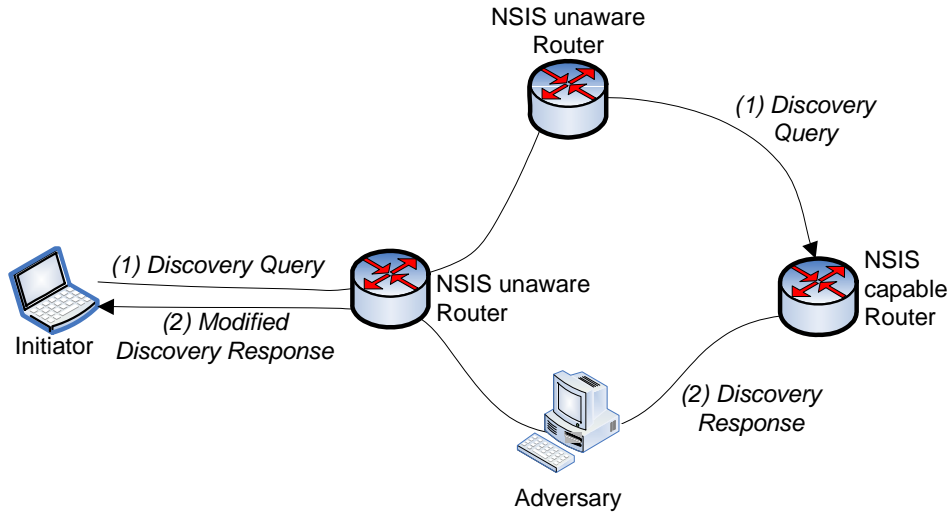


Figure 4: MITM Attack during the Discovery Exchange

NSIS node along the path will be discovered again. However, this could result in a ping-pong effect on the underlying NSIS state (and also the message routing and delivery behaviors).

During the process of establish a security association based on the previous discovery, an adversary can fool the signaling message sender with respect to the entity with whom it has to establish the security association. Without proper authorization the sender may successfully finalize the security protocol run with an adversary. If successful it is able to modify signaling messages to mount DoS attacks or steal services (depending on the used signaling application). In addition, an adversary can terminate the sender’s messages and inject messages to another NSIS peer.

3.4 Denial of Service Attacks

An *OutsiderOnPath*, *OutsiderOffPath* or an *InsiderNSIS* adversary can flood an NSIS node with bogus messages to cause a denial of service (DoS) attack. These bogus messages can be, for example, discovery query requests, idempotent refresh messages, or NSLP-level trigger messages requiring the involvement of a third-party node (such as policy or AAA servers), faked GIST or NSLP error or response messages, or vulnerabilities of the used transport layer protocols (see for example (25)) or even causing a node to perform heavy cryptographic operations. Sometimes even the usage of a Router Alert Option is seen as an opportunity for a DoS attack. If the adversary transmits a large number of such messages, the responding node(s) may be overloaded with high amount of computation and transmission requirements, thus unable to process other messages originated from legitimate entities/users. This can be even more challenging if an NSIS node needs to issue signaling messages on behalf of someone else (by acting as a proxy).

As an example of DoS attacks, an adversary can flood a QoS aware network with receiver-initiated reservation requests, such as RSVP PATH-like messages. If these messages are unauthenticated and contain no authorization information, then state at QoS aware routers along the path will be created until the data responder authorizes the reservation in the reverse direction. In such a scenario an adversary can force the QoS aware routers to store state information. Furthermore, since the state expiration might take some time, the routers might reject legitimate QoS reservation requests due to lack of resources.

Among all participating NSIS nodes, NSIS nodes located at the edges of administrative domains are the most critical entities of NSIS signaling, and they may also act as a security gateway or firewall for incoming and outgoing signaling messages (as well as data traffic). For outgoing traffic these devices have to implement the security policy of the local domain and apply the appropriate security protection.

3.5 Identity Spoofing

Identity spoofing in NSIS operation can occur in three forms: first, spoofing the identifier that is used for computing the authorization decision, for example caused by the usage of weak authentication mechanisms; second, an adversary can modify the flow identifier carried within a signaling message; third, data traffic can be spoofed. The latter two issues refer to aspect that the flow identifier can be seen as an identifier of identifying the QoS reservation on the data plane.

The first type of attack can be mounted by an *OutsiderOnPath*, *InsiderNSIS* and an *OutsiderOffPath* adversary. The latter two attacks require the capability of the *OutsiderOnPath* and the *InsiderNSIS* adversary.

In the first case, Eve, acting as an adversary, may claim to be the registered user Alice by spoofing Alice’s identity.

Eve thereby causes the network to charge Alice for the network resources which are consumed by Eve. This attack could also be classified as theft of service. The choice of selecting an identifier that can be associated to a chargeable entity depends on the deployment environment. Making inappropriate assumptions can easily lead to vulnerabilities, such as using a plaintext identifier in a QoS signaling protocol as the basis for authorization as, for example, provided by the RSVP Policy Object (10) allows a malicious end host or a man-in-the-middle to inject a modified identifier.

In the second case, an adversary may be able to exploit the established flow identifiers (required for QoS and NAT/FW NSLP). These identifiers are, among others, IP addresses, transport protocol type (e.g., UDP, TCP), port numbers, IPv6 flow labels, etc. Modification of these flow identifiers allows adversaries to make a QoS reservation ineffective for the QoS resource requesting entity, to reuse it for its own traffic, or to install arbitrary policy rules at middleboxes for further attacks against end hosts (or the network infrastructure).

In the third case, an adversary may spoof data traffic, i.e., inject data traffic with a flow identifier matching the installed identifier. As described in the previous section, NSIS signaling messages carry a flow identifier, which is associated with a specified behavior (e.g., allow a particular flow to receive certain QoS treatment or for packets to traverse a firewall) based on the type of signaling application. Thus, an adversary may use IP spoofing in order to inject data packets to benefit from previously installed flow identifiers (which is also a theft of service). For example, an adversary might observe a NAT/Firewall NSLP message towards a corporate network firewall. After the signaling message exchange was successful, user Alice is allowed to traverse the company firewall based on the established packet filter to contact her internal mail server. Now, adversary Eve, who monitored the signaling exchange, is able to craft data packets by spoofing some header fields that will allow the packets to pass the corporate firewall. Depending on the exact location of the adversary and the degree of routing asymmetry the adversary might even see the response messages. Here, to actualize this attack, Alice does not need to participate in the exchange of signaling messages. To address this attack one can rely on data origin authentication (e.g., by using IPsec (26)).

3.6 Stealing Authorization Information

Different information is used for computing an authorization decisions, such as the authenticated identity, availability of sufficient funds, roles and traits, number of concurrent sessions, amount of requested bandwidth or other QoS parameters, etc. The specific authorization policies might heavily depend on the usage of the specific signaling application and the specific request, such as a request for making a QoS reservation, the request to change a NAT binding or to allocate a firewall pinhole.

One approach to verify an entity's rights to access these

resources is using the authenticated identity. Another approach is to use an authorization token, e.g., as described in RFC 3182 (10). The functionality and the structure of such an authorization token for RSVP is described in (27; 28). By using such an authorization token, it is possible to bind the authorization decision provided by one protocol (e.g., SIP (31)) to the QoS signaling protocol.

However, if an authorization token is returned to the end host without confidentiality protection, then it might allow an eavesdropper to reuse it (depending on the constraints carried inside the token) to gain the same authorization rights as the legitimate owner of the token. An adversary might, for example, use such a token to make an expensive QoS reservation. Without appropriate protection of the token an end host may want to modify the token content and thereby grant itself more rights. A more detailed discussions of authorization tokens used in the context of RSVP and SIP can be found in (27; 28). A more recent example of authorization tokens is available with the Security Assertion Markup Language (SAML) (29; 30).

3.7 Fraud

Signaling applications (such as QoS NSLP or NAT/FW NSLP) often involve three parties: the user, a network that offers NSLP services such as QoS service support and a third party which authenticates and authorizes the user (AAA server). For a QoS reservation the trusted third party in most cases needs to ensure that the network providing the QoS service actually receives a financial compensation.

The QoS service providing network might intentionally deliver incorrect resource reports about the resource consumption to the user's home network. In a typical AAA setting it would be quite difficult to detect this fact since the user is not involved in the exchange between the AAA client and the AAA server. Furthermore, the price of the QoS reservation might not even be known to the user or might change due to mobility events or the end host.

Another problem may appear when the service provider (or the user) later denies about the existence or some parameters (e.g., volume or price) of a QoS reservation (or other NSLP services) towards the third party. This is often regarded as a problem if non-repudiation is not provided, which may appear in two forms:

Service provider's point-of-view: A user may deny having issued a reservation request, which was actually charged for. The service provider may then want to be able to show that a particular user issued the reservation request in question.

User's point-of-view: A service provider may claim to have received a number of reservation requests from a particular user. The user in question may want to show that such a reservation requests have never been issued and may want to see correct service usage records for a given set of QoS parameters.

The ability to provide non-repudiation places additional requirements on security mechanisms and are often associated with a more complex protocol interaction and additional security mechanisms. Therefore, non-repudiation is not provided with network access authentication in today's networks; the user has to trust the participating network operators to correctly meter the traffic, to collect accounting data, and to ensure that no unforeseen problems occur.

Please note that this issue is also relevant for accessing services in general. The large deployment of wireless LAN hotspots, as one example of a service, also faces the same problems. It is anticipated that a solution applicable for such an environment may be reused for providing non-repudiation protection for NSLPs.

3.8 Disclosing Information about a Network

Some organizations or enterprises may desire not to reveal their internal network structure (or other related information) outside of a closed community. However, an adversary may be able to use NSIS messages for disclosing network topology (e.g., discovering which nodes exist, which use NSIS, what version, what resources are allocated, what capabilities nodes along a path have, etc.). Discovery messages, traceroute, diagnostic messages (see (32) for a description of diagnostic message functionality for RSVP and a Ping tool for NSIS (33)), and query messages, in addition to record route and route objects, provide potential assistance to an adversary. Hence, the requirement of not disclosing a network topology may conflict with other requirements (e.g., to provide diagnostic facilities for network monitoring and administration).

3.9 Session Manipulation

A Session Identifier is included in NSIS signaling messages as a reference to the established state. This helps to simplify mobility and general state manipulation. However, this has a few security implications. The session manipulation issue is a key issue among them (see (34) for a detailed discussion).

Fig. 5 shows an NSIS Initiator that has established state information at NSIS nodes along a path as part of the signaling procedure. As a result, *AccessRouter1*, *Router3*, and *Router4* (and other nodes) have stored session state information including the Session Identifier *SIDx*.

If an adversary were able to obtain the Session Identifier (e.g., by eavesdropping on signaling messages), it would be able to add the same Session Identifier *SIDx* to a new signaling message. When the new signaling message hits *Router3*, existing state information can be modified. The adversary can then modify or delete the established reservation and cause unexpected behavior for the legitimate user. The problem appears at *Router3* (i.e., the cross-over router) that is unable to decide whether the newly received signaling message was initiated from the owner of the session or not.

In addition, nodes other than the initial signaling message originator may be allowed to signal state information during the lifetime of an established session. The flexible design of NSIS with regard to the supported deployment scenarios allows any NSIS-aware node along the path to trigger or terminate a signaling message exchange (e.g., a local repair procedure, or an asynchronous notification). If only the initial signaling message originator were allowed to trigger signaling message exchanges, some enhanced protocol functionality might not be possible. As an example, idempotent refresh messages need to be sent in an end-to-end fashion rather than allowing intermediate nodes to inject them. This tradeoff between signaling protocol flexibility and security become more outstanding for NSIS mobility support (35). Furthermore, the number of NSIS nodes in the flow path may change over the lifetime of a signaling session, e.g., due to node mobility or route changes. Without properly addressing the session ownership problem, an adversary can launch DoS, theft of service and other attacks.

This particular threat reflects an important issue that needs to be considered in many aspects of the protocol design, namely who is authorized to modify established state at various nodes in the network. Even though authorization for the establishment of resources, such as a QoS resource, might need to be provided only between adjacent nodes the aspect of session ownership appears in various places of the NSIS protocol suite. It is also one of the critical issues that need to be addressed in mobile environments (35), in the context of QoS signaling and reservation collisions (36) and a concern needs to be resolved for the data sender behind the NAT signaling scenario in the NAT/Firewall NSLP (19).

A tradeoff analysis is required for each NSLP to evaluate whether the complexity of a solution for the session ownership problem is justified compared to the mitigated threats.

4 Towards a Secure NSIS Protocol Suite

The design of the NSIS protocol suite has not yet completed and a number of open issues especially with security aspects still exist. In this section we present some considerations towards a secure NSIS protocol suite.

Since the NSIS protocol suite is split into two layers, the proposed security solution needs to offer security protection for both NTLP and existing NSLPs. With extensibility in mind, security building blocks within NTLP and some general aspects of NSLPs should be designed to serve as a useful facility for designing future NSLPs. Based on the analysis presented in previous sections, we believe the following issues are vital for securing the NSIS protocol suite:

1. Ability to run an authentication and key exchange protocol between neighboring NSIS peers (supporting either symmetric or asymmetric cryptography or even

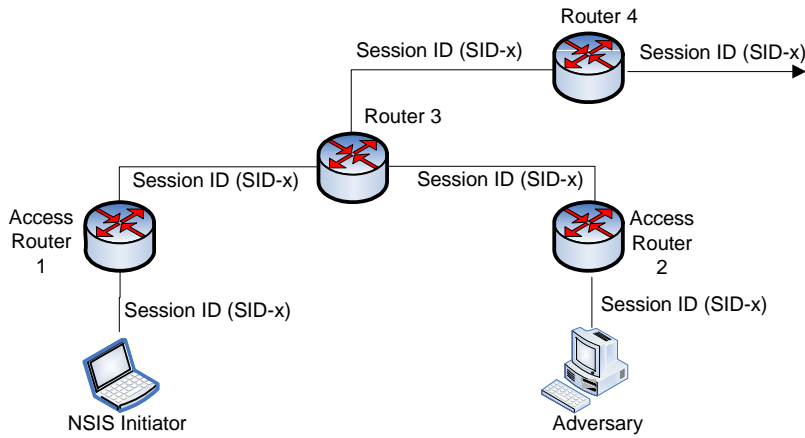


Figure 5: Session/Reservation Ownership

a hybrid version).

2. Security association establishment to provide integrity, confidentiality and replay protection for signaling messages exchanged between neighboring peers.
3. Denial of service protection.
4. Lightweight protection for the discovery mechanism.
5. Authorization of the NTLF signaling peers.
6. Flexible authorization at the NSLP layer including the ability to interwork with the existing AAA infrastructure.

Notably, designing new, custom security protocols is a highly complex and time consuming task. Luckily, many of today's security protocols offer several desired features already: flexible authentication methods, formally verified correctness, offering of DoS protection, ciphersuite negotiation, low number of roundtrips, etc.. Based on this fact, the use of standard security protocols is highly encouraged in securing the NSIS protocol suite. The following subsections discuss some usage scenarios of existing security protocols and some extensions in order to address the above seven issues.

4.1 Effective Authentication and Key Exchange at NTLF Layer

A considerable amount of eavesdropping, replay and man-in-the-middle attacks can be resolved by a same mechanism, namely an effective authentication and key exchange protocol. Usually this is a high cost operation. To avoid expensive cryptographic computations (such as those required by a digital signature) for each individual object, message or signaling session the reuse of an established security association between two NTLF peers will be desirable. As a consequence, the high computational effort

can be amortized with the usage of faster symmetric cryptographic protection for multiple signaling sessions. Fortunately, this is possible to be supported at the NTLF layer, and we discuss several approaches to address it.

One approach is based on the Transport Layer Security (TLS) (37; 38), which provides both a flexible authentication and key exchange protocol framework (i.e., TLS Handshake Layer) and provides protection of the subsequently exchanged application data via the TLS Record Layer. TLS provides both session key establishment based on anonymous, unilateral or mutual authentication. The mutual authentication function provided by TLS can be used for the establishment of a bidirectional NTLF secure messaging association between any two neighboring NSIS peers of the same NSLP type. Additionally, TLS supports additional ciphersuites and enhancements to the protocols, such as Kerberos-based authentication (39), strong password-based authentication (40), pre-shared secret authentication and a mixture of shared secret and public key based authentication (41) and Extensible Authentication Protocol (EAP) (42) support within the TLS Handshake Protocol (43). Note that some enhancements are work in progress and are therefore subject to change.

Another approach is to apply the Internet Key Exchange protocol (IKE) (44), IKEv2 (45) or KINK (46) to support establishment of IPsec security associations, when IPsec is available and chosen for securing signaling messages between two neighboring peers. In addition, the usage of the Host Identity Protocol (HIP) (47; 48; 49) can be used to secure NSIS signaling messages. The rich availability of the mechanisms for IPsec SA establishment allows a flexibility in the use of SAs in different deployment scenarios.

Both approaches require minimal changes to existing security protocols. They can be integrated into an NTLF engine without the need to change the protocol itself, and do not necessarily require changes to the key exchange protocol itself.

For example, after a QoS NSLP node discovers the next QoS NSLP peer along the path, GIST establishes a messag-

ing association with the discovered QoS NLSP node when the C-Mode should be used. As a result the QoS NSLP engine receives the payload provided by the NTLP and additionally security related information via the GIST-NSLP API (such as the authenticated identity and authorization information). This information can then be used by the NSLP as input to the policy engine for the authorization decision (which is detailed in Section 4.5).

4.2 Lightweight Security in Discovery Mechanism

The discovery message exchange is a security sensitive process and additionally very difficult to secure. To prevent adversaries from redirecting messages, a cookie-based mechanism can be used in the discovery procedure. This countermeasures refers to the man-in-the-middle attacks described in Section 3.3 and in Fig. 4.

The cookie-based mechanism (see Fig. 6) can be illustrated as follows.

A Cookie(I) is included in the Discovery-Response message to prevent *OutsiderOffPath* adversaries from flooding the querying node with bogus responses since the initiator can use Cookie(I) to match the response with the request. Since Cookie(I) is a randomly chosen value an *OutsiderOffPath* adversary cannot compute a valid response.

The cookie provided by the responding node (Cookie(R)) is used to prevent DoS attacks in the classical sense as used by other protocols, such as SCTP (54) or IKEv2 (45). Note that the responder must not create per-session state with responding to the Discovery-Query otherwise denial of service vulnerabilities will be introduced. The Responder returns the received Cookie(I) and its own Cookie(R).

Finally, when the initiator receives the Discovery-Response it compares the Cookie(I) value and runs an authentication and key exchange protocol (such as TLS) with the discovered node before establishing a messaging association. To prevent an *OutsiderOnPath* adversary from modifying the Discovery-Response message and adding wrong information about the next NSIS node along the path, Cookie(R) is repeated once channel security is in place. This allows the responder to verify that it has actually participated in the discovery exchange. Thereby the discovery procedure is bound to the subsequent exchange.

A more detailed treatment of the path-coupled discovery procedure security aspects are provided in (17).

4.3 Basic Authorization at NTLP Layer

The goal of computing an authorization decision at the NTLP layer is to ensure that only legitimate NSIS nodes initiate a signaling communication. In a generic signaling environment it is quite difficult for the GIST engine to make a meaningful decisions without consulting the NSLP with respect to signaling application specific functionality. In some deployment environments it is, however, possible for the NTLP layer to perform basic access control operations and to allow only certain nodes from a particular

domain to establish a messaging association. The authenticated identity might be used for computing this authorization decision but it is feasible to utilize authorization certificates (if available). Such authorization decisions at the NTLP layer are particularly useful in intra-domain scenarios as well as in environments, such as enterprise networks, where the communicating peers are known in advance based on pre-configuration.

In order to address one of the man-in-the-middle attack variant, which is described in Section 3.3, where an NSIS node claims to support a certain signaling application, it is necessary to provide information along with the credentials, such as authorization certificates. Alternatively, it may be possible to tie the protocol operation within a previous protocol exchange, e.g., the network access authentication procedure. A more detailed treatment of these issues is provide in (23).

4.4 Flexible Authorization at NSLP Layer

As described in Section 3.6, authorization aspects deserve a special attention. The work on the NAT/Firewall and the QoS NSLP showed the difficulty in properly tackling authorization issues in a generic way for all NSLPs (see (19) for a discussions about NAT/Firewall specific authorization issues and (50; 51; 52) for a discussion of the QoS specific authorization aspects. Individual NSLPs might be able to reuse some building blocks but the authorization handling will be at least in some aspects be different for each NSLP.

The interworking of NSIS protocols with today's AAA protocols such as Diameter (13) and RADIUS (14) is being regarded as an important building block of the network infrastructure.

Although the NTLP layer needs to provide certain authorization functionality, most authorization decisions will be made at the NSLP layer. The decision to successfully authorize QoS reservation requests might be related to the ability of the user (or other another entity requesting a resource reservation) to pay for the preferential treatment. Making an authorization decision to create a NAT binding might likely depend on the traffic direction (network internal traffic towards the Internet vs. traffic from the Internet towards a private network). To create packet filters at a firewall the security policy of the administrative domain will certainly play an important role. This policy will be different in a corporate network, home network and in a 3G network. Modifying established state should only be possible for the entity that created the state. This authorization decision is also known as sender invariance.

An individual NSIS router will, in many cases, be unable to make an authorization decision by itself without consulting third parties. This is particularly the case for an environment where hosts roam from one network to another. Hence, a QoS aware router that receives a QoS reservation request might want to contact the AAA infrastructure to off-load the authorization decision.

In a recent work (53) we describe the integration of the

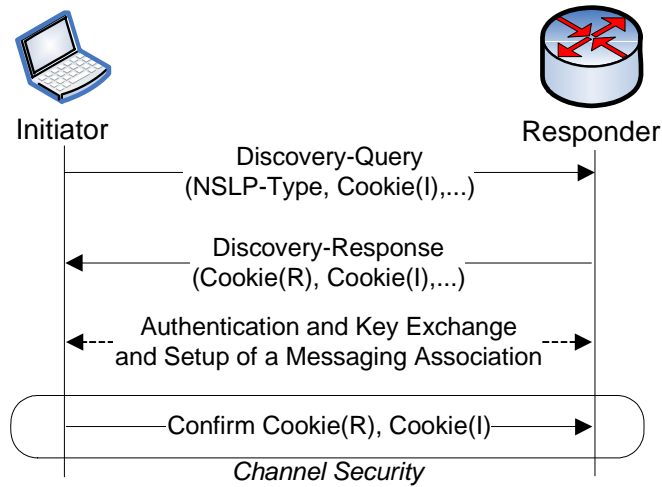


Figure 6: Protection of the discovery procedure in GIST

EAP (42) into NSIS, which attempts to provide flexible advanced authorization mechanism.

4.5 An Example Usage of Securing NSIS

Based on the above considerations with some individual examples, Fig. 7 shows an example of possible overall protocol exchange for the interactions between NSIS, SIP and Diameter. Note that a user will be authorized by its neighboring NSIS node rather than some arbitrary networks along the path towards the end host.

First, a service request is sent from the end host Alice to the SIP proxy. In this example, the SIP proxy creates and returns an authorization token to bind the application layer signaling exchange to the subsequent NSIS signaling session (e.g., a sender-initiated reservation). The authorization token is attached to the NSIS signaling message and the message itself is intercepted by the first-hop NSIS NSLP node, “Bob”. Upon the receipt of this token, Bob needs to authorize the QoS request and delegates this responsibility to the Diameter QoS application (50). A Diameter QoS authorization request, which includes authorization information and QoS information, is forwarded to the SIP proxy that created the authorization token for verification. As a response, the authorization decision is returned with a corresponding Diameter message. If the Diameter response is positive and the admission control process allows then the Traffic Control engine is contacted to install the QoS reservation and Bob forwards the NSIS QoS message further along the path. Finally, accounting and possibly credit control procedures are initiated to keep track of the consumed resources, to control the re-authorization policy and to ensure that the user’s credit limit is not exhausted.

5 Summary and Outlook

This paper describes security threats and lists the most important aspects of security protection for the NSIS protocol suite. Through identifying open issues in previous QoS signaling protocols, some critical design decisions have been made for NSIS especially with respect to security.

We believe that the integration of security in the design of NSIS has helped to develop a sound and reasonable framework. Many protocols developed so far have experienced critical shortcomings due to a lack of security considerations (see an analysis of QoS signaling protocols in (6)).

The main specifications of the NSIS protocol suite are getting closer to their final completion and implementation work is ongoing (20). We expect that these activities will reveal further details about performance and the practical usage of security mechanisms. Security will therefore continue to play an important role in further NSIS protocol development since a large fraction of the total performance will go into security processing. Particularly the usage of NSIS in mobile environments demands further investigations and will remain a hot research topic due to performance constraints and the need for optimizations. (35) provides an overview of ongoing work in this field.

We also plan to assess the deployment aspects of the NSIS protocol suite and their impacts on existing and future network architectures (e.g., 3GPP, 3GPP2, WiMax, ETSI-TISPAN, ITU-T NGN).

Acknowledgements

We would like to thank members of the IETF NSIS working group for the fruitful discussions and anonymous reviewers for their helpful comments. In particular we would like to thank Richard Graveman, Robert Hancock, Dirk Kroesberg, John Loughney, Allison Mankin, Andreas Pashalidis, Henning Peters, Henning Schulzrinne and Tseno Tsenov.

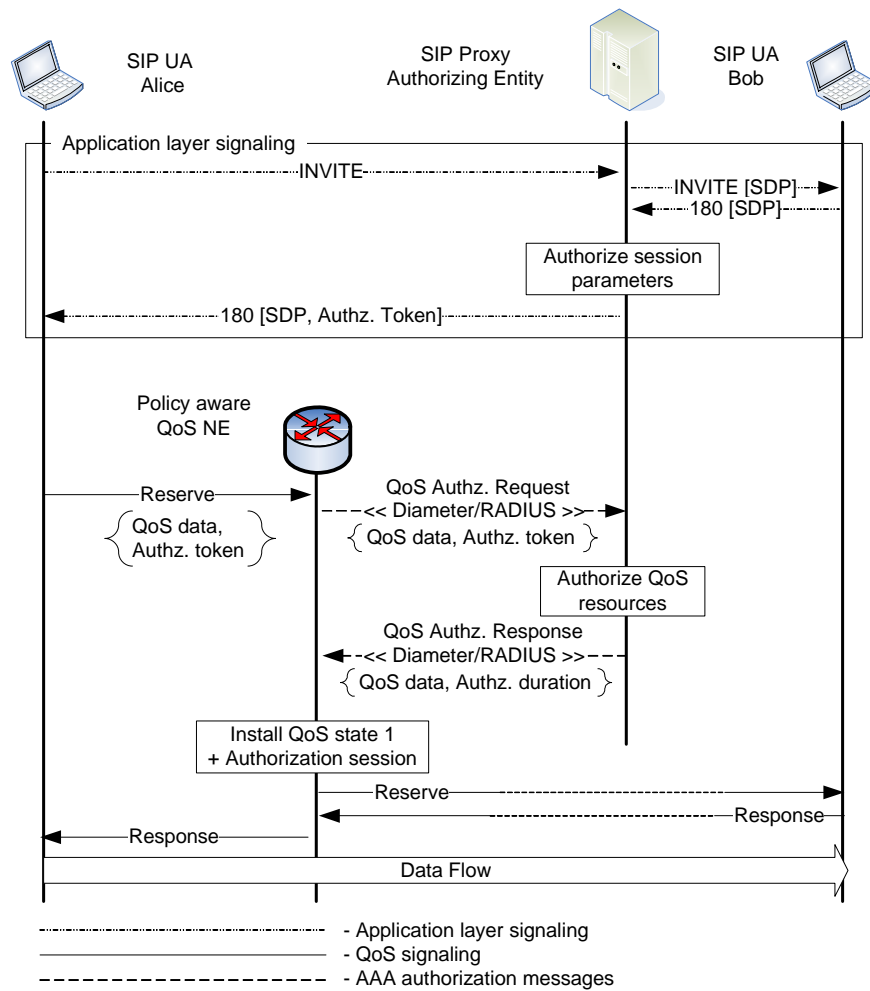


Figure 7: QoS Authorization with SIP and Diameter Interaction

REFERENCES

- [1] Braden, R., Clark, D., Shenker, S.: Integrated Services in the Internet Architecture: an Overview. RFC 1633 (Informational) (1994)
- [2] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W.: An Architecture for Differentiated Service. RFC 2475 (Informational) (1998) Updated by RFC 3260.
- [3] Kempf, J., Austein, R., IAB: The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture. RFC 3724 (Informational) (2004)
- [4] Braden, R., Zhang, L., Berson, S., Herzog, S., Jamin, S.: Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification. RFC 2205 (Proposed Standard) (1997) Updated by RFCs 2750, 3936.
- [5] Zhang, L., Deering, S., Estrin, D., Shenker, S., Zappala, D.: RSVP: A New Resource ReSerVation Protocol. IEEE Network **7** (1993) 8–18
- [6] Manner, J., Fu, X.: Analysis of Existing Quality-of-Service Signaling Protocols. RFC 4094 (Informational) (2005)
- [7] Wu, T.L., Wu, S., Gong, F.: Securing QoS: Threats to RSVP Messages and Their Countermeasures. In: Proc of IWQoS (1999)
- [8] Baker, F., Lindell, B., Talwar, M.: RSVP Cryptographic Authentication. RFC 2747 (Proposed Standard) (2000) Updated by RFC 3097.
- [9] Herzog, S.: RSVP Extensions for Policy Control. RFC 2750 (Proposed Standard) (2000)
- [10] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T., Herzog, S., Hess, R.: Identity Representation for RSVP. RFC 3182 (Proposed Standard) (2001)
- [11] Talwar, V., Nahrstedt, K.: Securing RSVP for Multimedia Applications. In: Proc of ACM Multimedia Workshop, Los Angeles, CA, USA (2000)
- [12] Tschofenig, H., Graveman, R.: RSVP Security Properties. RFC 4230 (Informational) (2005)

- [13] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., Arkko, J.: Diameter Base Protocol. RFC 3588 (Proposed Standard) (2003)
- [14] Rigney, C., Willens, S., Rubens, A., Simpson, W.: Remote Authentication Dial In User Service (RADIUS). RFC 2865 (Proposed Standard) (2000)
- [15] Tschofenig, H., Kroesenberg, D.: Security Threats for Next Steps in Signaling (NSIS). RFC 4081 (Informational) (2005)
- [16] Hancock, R., Karagiannis, G., Loughney, J., Van den Bosch, S.: Next Steps in Signaling (NSIS): Framework. RFC 4080 (Informational) (2005)
- [17] Schulzrinne, H., Hancock, R.: GIST: General Internet Signaling Transport. Internet draft (draft-ietf-nsis-ntlp-08), work in progress (2005)
- [18] Manner, J., Karagiannis, G., McDonald, A., Van den Bosch, S.: NSLP for Quality-of-Service signaling. Internet draft (draft-ietf-nsis-qos-nslp-08), work in progress (2005)
- [19] Stiemerling, M., Tschofenig, H., Martin, M., Aoun, C.: NAT/Firewall NSIS Signaling Layer Protocol (NSLP). Internet draft (draft-ietf-nsis-nslp-natfw-08), work in progress (2005)
- [20] Fu, X., Schulzrinne, H., Bader, A., Hogrefe, D., Kappler, C., Karagiannis, G., Tschofenig, H., Van den Bosch, S.: NSIS: A New Extensible IP Signaling Protocol Suite. IEEE Communications Magazine **43** (2005) 134–141
- [21] Gai, S., Dutt, D., Elfassy, N., Bernet, Y.: RSVP Proxy. Internet draft (draft-ietf-rsvp-proxy-03), work in progress (2002)
- [22] Bader, A., Westberg, L., Karagiannis, G., Kappler, C., Phelan, T.: RMD-QOSM – The Resource Management in DiffServ QOS Model. Internet draft (draft-ietf-nsis-rmd-03), work in progress (2005)
- [23] Aoun, C., Davies, E., Tschofenig, H.: Securing Middlebox Discovery for Path-Directed Signaling in the Internet. In: Proc. of the 5th Workshop on Applications and Services in Wireless Networks (ASWN) (2005)
- [24] Katz, D.: IP Router Alert Option. RFC 2113 (Proposed Standard) (1997)
- [25] Gont, F.: ICMP attacks against TCP. Internet draft (draft-gont-tpm-icmp-attacks-05), work in progress (2005)
- [26] Kent, S., Atkinson, R.: Security Architecture for the Internet Protocol. RFC 2401 (Proposed Standard) (1998) Updated by RFC 3168.
- [27] Hamer, L.N., Gage, B., Kosinski, B., Shieh, H.: Session Authorization Policy Element. RFC 3520 (Proposed Standard) (2003)
- [28] Hamer, L.N., Gage, B., Shieh, H.: Framework for Session Set-up with Media Authorization. RFC 3521 (Informational) (2003)
- [29] Maler, E., Philpott, R., Mishra, P.: Bindings and Profiles for the OASIS Security Assertion Markup Language (SAML) V1.1. September 2003.
- [30] Maler, E., Philpott, R., Mishra, P.: Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V1.1. September 2003.
- [31] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., Schooler, E.: SIP: Session Initiation Protocol. RFC 3261 (Proposed Standard) (2002) Updated by RFCs 3265, 3853.
- [32] Terzis, A., Braden, B., Vincent, S., Zhang, L.: RSVP Diagnostic Messages. RFC 2745 (Proposed Standard) (2000)
- [33] Dickmann C., Juchem, I., Willert, S., Fu, X.: A stateless Ping tool for simple tests of GIMPS implementations. Internet draft (draft-juchem-nsis-ping-tool-02), work in progress (2005)
- [34] Tschofenig, H., Schulzrinne, H., Hancock, R., McDonald, A., Fu, X.: Security Implications of the Session Identifier. Internet draft (draft-tschofenig-nsis-id-00), work in progress (2003)
- [35] Lee, S., Jeong, S., Tschofenig, H., Fu, X., Manner, J.: Applicability Statement of NSIS Protocols in Mobile Environments. Internet draft (draft-ietf-nsis-applicability-mobility-signaling-03), work in progress (2005)
- [36] McDonald, A., Hancock, R., Tschofenig, H., Kappler, C.: A Quality of Service NSLP for NSIS. Internet draft (draft-mcdonald-nsis-qos-nslp-00), work in progress (2003)
- [37] Dierks, T., Allen, C.: The TLS Protocol Version 1.0. RFC 2246 (Proposed Standard) (1999) Updated by RFC 3546.
- [38] Dierks, T., Rescorla, E.: The TLS Protocol Version 1.1. Internet draft (draft-ietf-tls-rfc2246-bis-13), work in progress (2005)
- [39] Medvinsky, A., Hur, M.: Addition of Kerberos Cipher Suites to Transport Layer Security (TLS). RFC 2712 (Proposed Standard) (1999)
- [40] Taylor, D., Wu, T., Mavrogianopoulos, N.: Using SRP for TLS Authentication. Internet draft (draft-ietf-tls-srp-10), work in progress (2005)

- [41] Eronen, P., Tschofenig, H.: Pre-Shared Key Ciphersuites for Transport Layer Security (TLS). RFC 4279 (Proposed Standard) (2005)
- [42] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., Levkowetz, H.: Extensible Authentication Protocol (EAP). RFC 3521 (Proposed Standard) (2004)
- [43] Funk, P., Blake-Wilson, S., Smith, N., Tschofenig, H., Hardjono, T.: TLS Inner Application Extension (TLS/IA). Internet draft (draft-funk-tls-inner-application-extension-01), work in progress (2005)
- [44] Harkins, D., Carrel, D.: The Internet Key Exchange (IKE). RFC 2409 (Proposed Standard) (1998)
- [45] Kaufman, C.: Internet Key Exchange (IKEv2) Protocol. Internet draft (draft-ietf-ipsec-ikev2-17), work in progress (2004)
- [46] Thomas, M.R., Vilhuber, J.: Kerberized Internet negotiation of keys (KINK). Internet draft (draft-ietf-kink-kink-11), work in progress (2005)
- [47] Moskowitz, R., Nikander, P.: Host Identity Protocol Architecture. Internet draft (draft-ietf-hip-arch-03), work in progress (2005)
- [48] Moskowitz, R., Nikander, P., Jokela, P., Henderson, T.: Host Identity Protocol. Internet draft (draft-ietf-hip-base-04), work in progress (2005)
- [49] Jokela, P., Moskowitz, R., Nikander, P.: Using ESP transport format with HIP. Internet draft (draft-ietf-hip-esp-01), work in progress (2005)
- [50] Alfano, F., McCann, P., Tschofenig, H., Tsenov, T.: Diameter Quality of Service Application. Internet draft (draft-alfano-aaa-qosprot-05), work in progress (2005)
- [51] Tschofenig, H., Buechli, M., Van den Bosch, S., Schulzrinne, H., Chen, T.: QoS NSLP Authorization Issues. Internet draft (draft-tschofenig-nsis-qos-authz-issues-00), work in progress (2003)
- [52] Tschofenig, H., Buechli, M., Van den Bosch, S., Schulzrinne, H.: NSIS Authentication, Authorization and Accounting Issues. Internet draft (draft-tschofenig-nsis-aaa-issues-01), work in progress (2003)
- [53] Tsenov, T., Tschofenig, H., Fu, X., Körner, E.: Advanced Authentication and Authorization for Quality of Service Signaling. In: Proc. of 1st IEEE/CREATENET SECURECOM Workshop on Security and QoS in Communication Networks (Sec-QoS) (2005)
- [54] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L., Paxson, V.: Stream Control Transmission Protocol. RFC 2960 (Proposed Standard) (2000)