

Internet Engineering Task Force (IETF)
Request for Comments: 5978
Category: Informational
ISSN: 2070-1721

J. Manner
Aalto University
R. Bless
KIT
J. Loughney
Nokia
E. Davies, Ed.
Folly Consulting
October 2010

Using and Extending the NSIS Protocol Family

Abstract

This document gives an overview of the Next Steps in Signaling (NSIS) framework and protocol suite created by the NSIS Working Group during the period of 2001-2010. It also includes suggestions on how the industry can make use of the new protocols and how the community can exploit the extensibility of both the framework and existing protocols to address future signaling needs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5978>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. The NSIS Architecture	3
3. The General Internet Signaling Transport	6
4. Quality-of-Service NSLP	11
5. NAT/Firewall Traversal NSLP	12
6. Deploying the Protocols	13
6.1. Deployment Issues Due to Use of RAO	14
6.2. Deployment Issues with NATs and Firewalls	15
6.3. Incremental Deployment and Workarounds	15
7. Security Features	16
8. Extending the Protocols	16
8.1. Overview of Administrative Actions Needed When Extending NSIS	17
8.2. GIST	17
8.2.1. Use of Different Message Routing Methods	18
8.2.2. Use of Different Transport Protocols or Security Capabilities	18
8.2.3. Use of Alternative Security Services	19
8.2.4. Query Mode Packet Interception Schemes	19
8.2.5. Use of Alternative NAT Traversal Mechanisms	20
8.2.6. Additional Error Identifiers	20
8.2.7. Defining New Objects To Be Carried in GIST	21
8.2.8. Adding New Message Types	21
8.3. QoS NSLP	21
8.4. QoS Specifications	22
8.5. NAT/Firewall NSLP	23
8.6. New NSLP Protocols	23
9. Security Considerations	26
10. Acknowledgements	27
11. References	27
11.1. Normative References	27
11.2. Informative References	28

1. Introduction

The Next Steps in Signaling (NSIS) Working Group was formed in November 2001 to develop an Internet signaling protocol suite that would attempt to remedy some of the perceived shortcomings of solutions based on the Resource ReSerVation Protocol (RSVP), e.g., with respect to mobility and Quality-of-Service (QoS) interoperability. The initial charter was focused on QoS signaling as the first use case, taking RSVP as the background for the work. In May 2003, middlebox traversal was added as an explicit second use case. The requirements for the new generation of signaling protocols are documented in [RFC3726], and an analysis of existing signaling protocols can be found in [RFC4094].

The design of NSIS is based on a two-layer model, where a general signaling transport layer provides services to an upper signaling application layer. The design was influenced by Bob Braden's document entitled "A Two-Level Architecture for Internet Signaling" [TWO-LEVEL].

This document gives an overview of the NSIS framework and protocol suite at the time of writing (2010), provides an introduction to the use cases for which the current version of NSIS was designed, describes how to deploy NSIS in existing networks, and summarizes how the protocol suite can be enhanced to satisfy new use cases.

2. The NSIS Architecture

The design of the NSIS protocol suite reuses ideas and concepts from RSVP but essentially divides the functionality into two layers. The lower layer, the NSIS Transport Layer Protocol (NTLP), is in charge of transporting the higher-layer protocol messages to the next signaling node on the path. This includes discovery of the next-hop NSIS node, which may not be the next routing hop, and different transport and security services depending on the signaling application requirements. The General Internet Signaling Transport (GIST) [RFC5971] has been developed as the protocol that fulfills the role of the NTLP. The NSIS protocol suite supports both IP protocol versions, IPv4 and IPv6.

The actual signaling application logic is implemented in the higher layer of the NSIS stack, the NSIS Signaling Layer Protocol (NSLP). While GIST is only concerned with transporting NSLP messages hop-by-hop between pairs of signaling nodes, the end-to-end signaling functionality is provided by the NSLP protocols if needed. Not all NSLP protocols need to perform end-to-end signaling. The current protocols have features to confine the signaling to a limited part of the path (such as the interior of a domain). Messages transmitted by

GIST on behalf of an NSLP are identified by a unique NSLP identifier (NSLPID) associated with the NSLP. Two NSLP protocols are currently specified: one concerning Quality-of-Service signaling [RFC5974] and one to enable NAT/firewall traversal [RFC5973].

NSIS is primarily designed to provide the signaling needed to install state on nodes that lie on the path that will be taken by some end-to-end flow of data packets; the state installed should facilitate or enhance some characteristic of the data flow. This is typically achieved by routing signaling messages along the same path (known as "path-coupled signaling") and intercepting the signaling message at NSIS-capable nodes. However, the NSIS architecture is designed to be flexible, and the routing of signaling messages is controlled by the Message Routing Method (MRM) that is applied to the signaling messages. The initial specifications define two MRMs:

- o the basic Path Coupled MRM designed to drive signaling along the path that will be followed by the data flow, and
- o an alternative Loose End MRM, which is applicable for preconditioning the state in firewalls and Network Address Translation (NAT) middleboxes when data flow destinations lie behind this sort of middlebox. Without preconditioning, these middleboxes will generally reject signaling messages originating outside the region 'protected' by the middlebox and where the destination is located.

Parameters carried by each signaling message drive the operation of the relevant transport or signaling application. In particular, the messages will carry Message Routing Information (MRI) that will allow the NSIS nodes to identify the data flow to which the signaling applies. Generally, the intercepted messages will be reinjected into the network after processing by the NSIS entities and will be routed further towards the destination, possibly being intercepted by additional NSIS-capable nodes before arriving at the flow endpoint.

As with RSVP, it is expected that the signaling message will make a complete round trip either along the whole end-to-end path or a part of it if the scope of the signaling is limited. This implements a two-phase strategy in which capabilities are assessed and provisional reservations are made on the outbound leg; these provisional reservations are then confirmed and operational state is installed on the return leg. Unlike RSVP, signaling is normally initiated at the source of the data flow, making it easier to ensure that the signaling follows the expected path of the data flow, but can also be receiver initiated as in RSVP.

A central concept of NSIS is the Session Identifier (SID). Signaling application states are indexed and referred to through the SID in all the NSLPs. This decouples the state information from IP addresses, allowing dynamic IP address changes for signaling flows, e.g., due to mobility: changes in IP addresses do not force complete teardown and re-initiation of a signaling application state; they force merely an update of the state parameters in the NSLP(s), especially the MRI.

At the NTLP (GIST) layer, the SID is not meaningful by itself, but is used together with the NSLP identifier (NSLPID) and the Message Routing Information (MRI). This 3-tuple is used by GIST to index and manage the signaling flows. Changes of routing or dynamic IP address changes, e.g., due to mobility, will require GIST to modify already established Messaging Associations (MAs) that are used to channel NSLP messages between adjacent GIST peers in order to satisfy the NSLP MRI for each SID.

The following design restrictions were imposed for the first phase of the protocol suite. They may be lifted in the future, and new functionality may be added into the protocols at some later stage.

- o Initial focus on MRMs for path-coupled signaling: GIST transports messages towards an identified unicast data flow destination based on the signaling application request, and does not directly support path-decoupled signaling, e.g., QoS signaling to a bandwidth broker or other off-path resource manager. The framework also supports a Loose End MRM used to discover GIST nodes with particular properties in the direction of a given address; for example, the NAT/firewall NSLP uses this method to discover a NAT along the upstream data path.
- o No multicast support: Introducing support for multicast was deemed too much overhead, considering the currently limited support for global IP multicast. Thus, the current GIST and the NSLP specifications consider unicast flows only.

The key documents specifying the NSIS framework are:

- o Requirements for Signaling Protocols [RFC3726]
- o Next Steps in Signaling: Framework [RFC4080]
- o Security Threats for NSIS [RFC4081]

The protocols making up the suite specified by the NSIS Working Group are documented in:

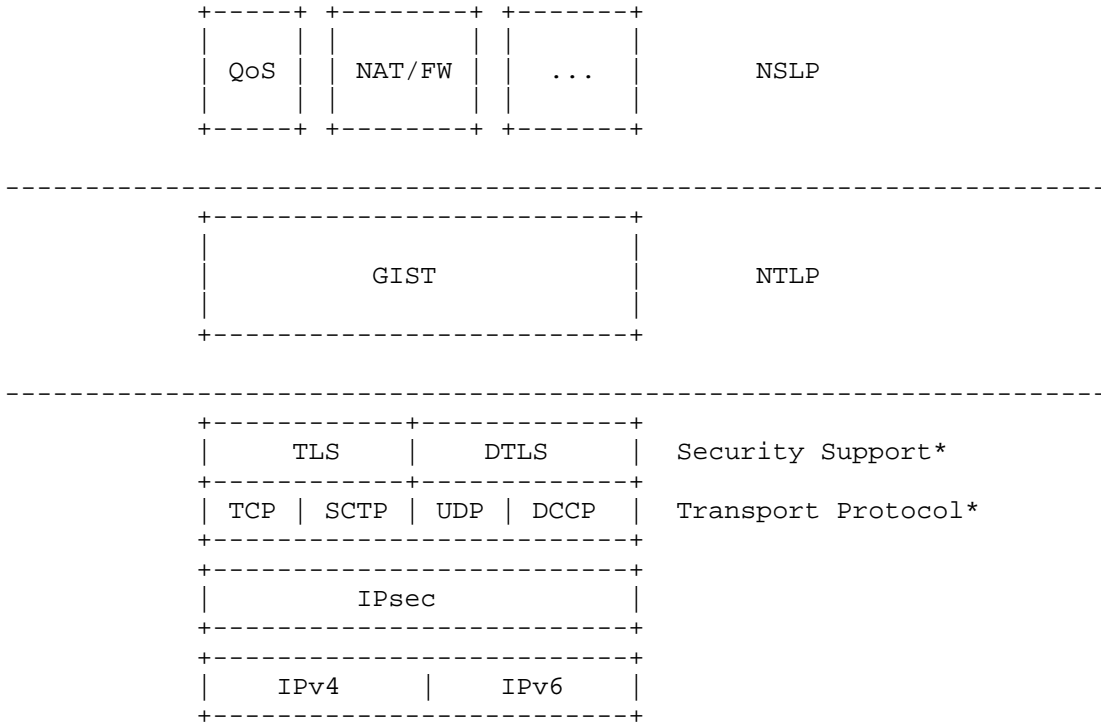
- o The General Internet Signaling Transport protocol [RFC5971]

- o Quality-of-Service NSLP (QoS NSLP) [RFC5974]
- o The QoS specification template [RFC5975]
- o NAT/firewall traversal NSLP [RFC5973]

The next three sections provide a brief survey of GIST, the QoS NSLP, and the NAT/firewall NSLP.

3. The General Internet Signaling Transport

The General Internet Signaling Transport (GIST) [RFC5971] provides signaling transport and security services to NSIS Signaling Layer Protocols (NSLP) and the associated signaling applications. GIST does not define new IP transport protocols or security mechanisms but rather makes use of existing protocols, such as TCP, UDP, TLS, and IPsec. Applications can indicate the desired transport attributes for the signaling flow, e.g., unreliable or reliable, and GIST then chooses the most appropriate transport protocol(s) to satisfy the requirements of the flow. GIST will normally use UDP if unreliable signaling is adequate, TCP if reliability is required, and TLS over TCP for secure (and reliable) signaling flows, but there exist extensibility provisions within GIST that will allow alternatives to be specified in the future. The NSIS layered protocol stack is shown in Figure 1.



* The Security Support and Transport Protocol layers show some possible protocols that could be used to transport NSIS messages. To provide authentication and/or integrity protection support, the transport protocol has to be paired with a suitable security mechanism, e.g., TCP with TLS, or Datagram Congestion Control Protocol (DCCP) with DTLS.

Figure 1: The NSIS protocol stack

GIST divides up the data flow's end-to-end path into a number of segments between pairs of NSIS-aware peer nodes located along the path. Not every router or other middlebox on the path needs to be NSIS aware: each segment of the signaling path may incorporate several routing hops. Also not every NSIS-aware node necessarily implements every possible signaling application. If the signaling for a flow requests services from a subset of the applications, then only nodes that implement those services are expected to participate as peers, and even some of these nodes can decline to operate on a particular flow if, for example, the additional load might overload the processing capability of the node. These characteristics mean that incremental deployment of NSIS capabilities is possible both with the initial protocol suite, and for any future NSLP applications

that might be developed. The following paragraphs describe how a signaling segment is set up to offer the transport and security characteristics needed by a single NSLP.

When an NSLP application wants to send a message towards a flow endpoint, GIST starts the process of discovering the next signaling node by sending a Query message towards the destination of the related data flow. This Query carries the NSLP identifier (NSLPID) and Message Routing Information (MRI), among others. The MRI contains enough information to control the routing of the signaling message and to identify the associated data flow. The next GIST node on the path receives the message, and if it is running the same NSLP, it provides the MRI to the NSLP application and requests it to make a decision on whether to peer with the querying node. If the NSLP application chooses to peer, GIST sets up a Message Routing State (MRS) between the two nodes for the future exchange of NSLP data. State setup is performed by a three-way handshake that allows for negotiation of signaling flow parameters and provides counter-measures against several attacks (like denial-of-service) by using cookie mechanisms and a late state installation option.

If a transport connection is required and needs to provide for reliable or secure signaling, like TCP or TLS/TCP, a Messaging Association (MA) is established between the two peers. An MA can be reused for signaling messages concerning several different data flows, i.e., signaling messages between two nodes are multiplexed over the same transport connection. This can be done when the transport requirements (reliability, security) of a new flow can be met with an existing MA, i.e., the security and transport properties of an existing MA are equivalent or better than what is requested for a potential new MA.

For path-coupled signaling, we need to find the nodes on the data path that should take part in the signaling of an NSLP and invoke them to act on the arrival of such NSLP signaling messages. The basic concept is that such nodes along a flow's data path intercept the corresponding signaling packets and are thus discovered automatically. GIST places a Router Alert Option (RAO) in Query message packets to ensure that they are intercepted by relevant NSIS-aware nodes, as in RSVP.

Late in the development of GIST, serious concerns were raised in the IETF about the security risks and performance implications of extensive usage of the RAO [RAO-BAD]. Additionally, evidence was discovered indicating that several existing implementations of RAO were inconsistent with the (intention of the) standards and would not support the NSIS usage. There were also concerns that extending the need for RAO recognition in the fast path of routers that are

frequently implemented in hardware would delay or deter implementation and deployment of NSIS. Eventually, it was decided that NSIS would continue to specify RAO as its primary means for triggering interception of signaling messages in intermediate nodes on the data path, but the protocol suite would be published with Experimental status rather than on the Standards Track while deployment experience was gathered. More information about the use of RAO in GIST can be found in [GIST-RAO]. Also, the deployment issues that arise from the use of RAO are discussed in Section 6.1.

Alternative mechanisms have been considered to allow nodes to recognize NSIS signaling packets that should be intercepted. For example, NSIS nodes could recognize UDP packets directed to a specific destination port as Query messages that need to be intercepted even though they are not addressed to the intercepting node. GIST provides for the use of such alternatives as a part of its extensibility design. NSIS recognizes that the workload imposed by intercepting signaling packets could be considerable relative to the work needed just to forward such packets. To keep the necessary load to a minimum, NSIS provides mechanisms to limit the number of interceptions needed by constraining the rate of generation and allowing for intentional bypassing of signaling nodes that are not affected by particular signaling requests. This can be accomplished either in GIST or in the NSLP.

Since GIST carries information about the data flow inside its messages (in the MRI), NAT gateways must be aware of GIST in order to let it work correctly. GIST provides a special object for NAT traversal so that the actual translation is disclosed if a GIST-aware NAT gateway provides this object.

As with RSVP, all the state installed by NSIS protocols is "soft-state" that will expire and be automatically removed unless it is periodically refreshed. NSIS state is held both at the signaling application layer and in the signaling transport layer, and is maintained separately. NSLPs control the lifetime of the state in the signaling application layer by setting a timeout and sending periodic "keep alive" messages along the signaling path if no other messages are required. The MAs and the routing state are maintained semi-independently by the transport layer, because MAs may be used by multiple NSLP sessions, and can also be recreated "on demand" if the node needs to reclaim resources. The transport layer can send its own "keep alive" messages across a MA if no NSLP messages have been sent, for example, if the transport layer decides to maintain a heavily used MA even though there is no current NSLP session using it. Local state can also be deleted explicitly when no longer needed.

If there is a change in the route used by a flow for which NSIS has created state, NSIS needs to detect the change in order to determine if the new path contains additional NSIS nodes that should have state installed. GIST may use a range of triggers in order to detect a route change. It probes periodically for the next peer by sending a GIST Query, thereby detecting a changed route and GIST peer. GIST monitors routing tables and the GIST peer states, and it notifies NSLPs of any routing changes. It is then up to the NSLPs to act appropriately, if needed, e.g., by issuing a refresh message. The periodic queries also serve to maintain the soft-state in nodes as long as the route is unchanged.

In summary, GIST provides several services in one package to the upper-layer signaling protocols:

- o Signaling peer discovery: GIST is able to find the next-hop node that runs the NSLP being signaled for.
- o Multiplexing: GIST reuses already established signaling relationships and messaging associations to next-hop peers if the signaling flows require the same transport attributes.
- o Transport: GIST provides transport with different attributes -- namely, reliable/unreliable and secure/unsecure.
- o Security: If security is requested, GIST uses TLS to provide an encrypted and integrity-protected message transport to the next signaling peer.
- o Routing changes: GIST detects routing changes, but instead of acting on its own, it merely sends a notification to the local NSLP. It is then up to the NSLP to act.
- o Fragmentation: GIST uses either a known Path MTU for the next hop or limits its message size to 576 bytes when using UDP or Query mode. If fragmentation is required, it automatically establishes an MA and sends the signaling traffic over a reliable protocol, e.g., TCP.
- o State maintenance: GIST establishes and then maintains the soft-state that controls communications through MAs between GIST peers along the signaling path, according to usage parameters supplied by NSLPs and local policies.

4. Quality-of-Service NSLP

The Quality-of-Service (QoS) NSIS Signaling Layer Protocol (NSLP) establishes and maintains state at nodes along the path of a data flow for the purpose of providing some forwarding resources for that flow. It is intended to satisfy the QoS-related requirements of RFC 3726 [RFC3726]. No support for QoS architectures based on bandwidth brokers or other off-path resource managers is currently included.

The design of the QoS NSLP is conceptually similar to RSVP, RFC 2205 [RFC2205], and uses soft-state peer-to-peer refresh messages as the primary state management mechanism (i.e., state installation/refresh is performed between pairs of adjacent NSLP nodes, rather than in an end-to-end fashion along the complete signaling path). The QoS NSLP extends the set of reservation mechanisms to meet the requirements of RFC 3726 [RFC3726], in particular, support of sender- or receiver-initiated reservations, as well as, a type of bidirectional reservation and support of reservations between arbitrary nodes, e.g., edge-to-edge, end-to-access, etc. On the other hand, there is currently no support for IP multicast.

A distinction is made between the operation of the signaling protocol and the information required for the operation of the Resource Management Function (RMF). RMF-related information is carried in the QSPEC (QoS Specification) object in QoS NSLP messages. This is similar to the decoupling between RSVP and the IntServ architecture, RFC 1633 [RFC1633]. The QSPEC carries information on resources available, resources required, traffic descriptions, and other information required by the RMF. A template for QSPEC objects is defined in [RFC5975]. This provides a number of basic parameter objects that can be used as a common language to specify components of concrete QoS models. The objects defined in [RFC5975] provide the building blocks for many existing QoS models such as those associated with RSVP and Differentiated Services. The extensibility of the template allows new QoS model specifications to extend the template language as necessary to support these specifications.

The QoS NSLP supports different QoS models because it does not define the QoS mechanisms and RMF that have to be used in a domain. As long as a domain knows how to perform admission control for a given QSPEC, QoS NSLP actually does not care how the specified constraints are enforced and met, e.g., by putting the related data flow in the topmost of four Diffserv classes or by putting it into the third highest of twelve Diffserv classes. The particular QoS configuration used is up to the network provider of the domain. The QSPEC can be seen as a common language to express QoS requirements between different domains and QoS models.

In short, the functionality of the QoS NSLP includes:

- o Conveying resource requests for unicast flows
- o Resource requests (QSPEC) that are decoupled from the signaling protocol (QoS NSLP)
- o Sender- and receiver-initiated reservations, as well as bidirectional
- o Soft-state and reduced refresh (keep-alive) signaling
- o Session binding, i.e., session X can be valid only if session Y is also valid
- o Message scoping, end-to-end, edge-to-edge, or end-to-edge (proxy mode)
- o Protection against message re-ordering and duplication
- o Group tear, tearing down several sessions with a single message
- o Support for rerouting, e.g., due to mobility
- o Support for request priorities and preemption
- o Stateful and stateless nodes: stateless operation is particularly relevant in core networks where large amounts of QoS state could easily overwhelm a node
- o Reservation aggregation

The protocol also provides for a proxy mode to allow the QoS signaling to be implemented without needing all end-hosts to be capable of handling NSIS signaling.

The QSPEC template supports situations where the QoS parameters need to be fine-grained, specifically targeted to an individual flow in one part of the network (typically the edge or access part) but might need to be more coarse-grained, where the flow is part of an aggregate (typically in the core of the network).

5. NAT/Firewall Traversal NSLP

The NAT/firewall traversal NSLP [RFC5973] lets end-hosts interact with NAT and firewall devices in the data path. Basically, it allows for a dynamic configuration of NATs and/or firewalls along the data path in order to enable data flows to traverse these devices without

being obstructed. For instance, firewall pinholes could be opened on demand by authorized hosts. Furthermore, it is possible to block unwanted incoming traffic on demand, e.g., if an end-host is under attack.

Configurations to be implemented in NAT and firewall devices signaled by the NAT/firewall NSLP take the form of a (pattern, action) pair, where the pattern specifies a template for packet header fields to be matched. The device is then expected to apply the specified action to any passing packet that matches the template. Actions are currently limited to ALLOW (forward the packet) and DENY (drop the packet). The template specification allows for a greater range of packet fields to be matched than those allowed for in the GIST MRI.

Basically, NAT/firewall signaling starts at the data sender (NSIS Initiator) before any actual application data packets are sent. Signaling messages may pass several middleboxes that are NAT/firewall NSLP aware (NSIS Forwarder) on their way downstream and usually hit the receiver (being the NSIS Responder). A proxy mode is also available for cases where the NAT/firewall NSLP is not fully supported along the complete data path. NAT/firewall NSLP is based on a soft-state concept, i.e., the sender must periodically repeat its request in order to keep it active.

Additionally, the protocol also provides functions for receivers behind NATs. The receiver may request an external address that is reachable from outside. The reserved external address must, however, be communicated to the sender out-of-band by other means, e.g., by application level signaling. After this step the data sender may initiate a normal NAT/firewall signaling in order to create firewall pinholes.

The protocol also provides for a proxy mode to allow the NAT/firewall signaling to be implemented without needing all end-hosts to be capable of handling NSIS signaling.

6. Deploying the Protocols

The initial version of the NSIS protocol suite is being published with the status of Experimental in order to gain deployment experience. Concerns over the security, implementation, and administrative issues surrounding the use of RAO are likely to mean that initial deployments occur in "walled gardens" where the characteristics of hardware in use are well-known, and there is a high level of trust and control over the end nodes that use the protocols. This section addresses issues that need to be considered in a deployment of the NSIS protocol suite.

First of all, NSIS implementations must be available in at least some of the corresponding network nodes (i.e., routers, firewalls, or NAT gateways) and end-hosts. That means not only GIST support, but also the NSLPs and their respective control functions (such as a resource management function for QoS admission control, etc.) must be implemented. NSIS is capable of incremental deployment and an initial deployment does not need to involve every node in a network domain. This is discussed further in Section 6.3. There are a number of obstacles that may be encountered due to broken implementations of RAO (see Section 6.1) and due to firewalls or NATs that drop NSIS signaling packets (see Section 6.2).

Another important issue is that applications may need to be made NSIS-aware, thereby requiring some effort from the applications' programmers. Alternatively, it may be possible to implement separate applications to control, e.g., the network QoS requests or firewall pinholes, without needing to update the actual applications that will take advantage of NSIS capabilities.

6.1. Deployment Issues Due to Use of RAO

The standardized version of GIST depends on routers and other middleboxes correctly recognizing and acting on packets containing RAO. There are a number of problems related to RAO that can obstruct a deployment of NSIS:

- o Some implementations do not respond to RAO at all.
- o Some implementations respond but do not distinguish between the RAO parameter values in IPv4 packets or reject anything except 0 (in which case, only the value 0 can be used).
- o The response to RAO in a GIST Query mode packet, which is sent using the UDP transport, is to dispatch the packet to the UDP stack in the intercepting node rather than to a function associated with the RAO parameter. Since the node will not normally have a regular UDP receiver for these packets they are dropped.
- o The major security concern with RAO in NSIS is that it provides a new vector for hosts to mount a (distributed) denial-of-service (DDoS) attack on the control plane of routers on the data path. Such attacks have occurred, and it is therefore normal for service providers to prohibit "host-to-router" signaling packets such as RSVP or NSIS from entering their networks from customer networks. This will tend to limit the deployment of NSIS to "walled gardens" unless a suitable mitigation of the DDoS threat can be found and deployed.

In order to deploy NSIS effectively, routers and other hardware need to be selected and correctly configured to respond to RAO and dispatch intercepted packets to the NSIS function.

A further obstacle results from the likelihood that IPv4 packets with IP options of any kind will be filtered and dropped by firewalls and NATs. In many cases, this is the default behavior so that explicit configuration is needed to allow packets carrying the RAO to pass through. The general inclination of domain administrators is to deny access to packets carrying IP options because of the security risks and the additional load on the routers in the domain. The situation with IPv6 may be easier, as the RAO option in IPv6 is better defined, but the security concerns remain.

Deployment issues are discussed at more length in Appendix C of the GIST specification [RFC5971].

6.2. Deployment Issues with NATs and Firewalls

NAT gateways and firewalls may also hinder initial deployment of NSIS protocols for several reasons:

- o They may filter and drop signaling traffic (as described in Section 6.1) to deny access to packets containing IP options.
- o They may not permit "unsolicited" incoming GIST Query mode packets. This behavior has been anticipated in the design of the protocols but requires additional support to ensure that the middleboxes are primed to accept the incoming queries (see [RFC5974] and [RFC5973]).
- o NATs that are not aware of the NSIS protocols will generally perform address translations that are not coordinated with the NSIS protocols. Since NSIS signaling messages may be carrying embedded IP addresses affected by these translations, it may not be possible to operate NSIS through such legacy NATs. The situation and workarounds are discussed in Section 7.2.1 of [RFC5971].

6.3. Incremental Deployment and Workarounds

NSIS is specifically designed to be incrementally deployable. It is not required that all nodes on the signaling and data path are NSIS aware. To make any use of NSIS, at least two nodes on the path need to be NSIS aware. However, it is not essential that the initiator and receiver of the data flow are NSIS aware. Both the QoS and NAT/firewall NSLPs provide "proxy modes" in which nodes adjacent to the initiator and/or receiver can act as proxy signaling initiator or

receiver. An initiator proxy can monitor traffic and, hopefully, detect when a data flow of a type needing NSIS support is being initiated. The proxies can act more or less transparently on behalf of the data flow initiator and/or receiver to set up the required NSIS state and maintain it while the data flow continues. This capability reduces the immediate need to modify all the data flow endpoints before NSIS is viable.

7. Security Features

Basic security functions are provided at the GIST layer, e.g., protection against some blind or denial-of-service attacks, but note that introduction of alternative MRMs may provide attack avenues that are not present with the current emphasis on the path-coupled MRM. Conceptually, it is difficult to protect against an on-path attacker and man-in-the-middle attacks when using path-coupled MRMs, because a basic functionality of GIST is to discover as yet unknown signaling peers. Transport security can be requested by signaling applications and is realized by using TLS between signaling peers, i.e., authenticity and confidentiality of signaling messages can be assured between peers. GIST allows for mutual authentication of the signaling peers (using TLS means such as certificates) and can verify the authenticated identity against a database of nodes authorized to take part in GIST signaling. It is, however, a matter of policy that the identity of peers is verified and accepted upon establishment of the secure TLS connection.

While GIST is handling authentication of peer nodes, more fine-grained authorization may be required in the NSLP protocols. There is currently an ongoing work to specify common authorization mechanisms to be used in NSLP protocols [NSIS-AUTH], thus allowing, e.g., per-user and per-service authorization.

8. Extending the Protocols

This section discusses the ways that are available to extend the NSIS protocol suite. The Next Steps in Signaling (NSIS) Framework [RFC4080] describes a two-layer framework for signaling on the Internet, comprising a generic transport layer with specific signaling-layer protocols to address particular applications running over this transport layer. The model is designed to be highly extensible so that it can be adapted for different signaling needs.

It is expected that additional signaling requirements will be identified in the future. The two-layer approach allows for NSLP signaling applications to be developed independently of the transport protocol. Further NSLPs can therefore be developed and deployed to meet these new needs using the same GIST infrastructure, thereby

providing a level of macro-extensibility. However, the GIST protocol and the two signaling applications have been designed so that additional capabilities can be incorporated into the design should additional requirements within the general scope of these protocols need to be accommodated.

The NSIS framework is also highly supportive of incremental deployment. A new NSLP need not be available on every NSIS-aware node in a network or along a signaling path in order to start using it. Nodes that do not (yet) support the application will forward its signaling messages without complaint until it reaches a node where the new NSLP application is deployed.

One key functionality of parameter objects carried in NSIS protocols is the so-called "Extensibility flags (A/B)". All the existing protocols (and any future ones conforming to the standards) can carry new experimental objects, where the A/B flags can indicate whether a receiving node must interpret the object, or whether it can just drop them or pass them along in subsequent messages sent out further on the path. This functionality allows defining new objects without forcing all network entities to understand them.

8.1. Overview of Administrative Actions Needed When Extending NSIS

Generally, NSIS protocols can be extended in multiple ways, many of which require the allocation of unique code point values in registries maintained by IANA on behalf of the IETF. This and the following sections provide an overview of the administrative mechanisms that might apply. The extensibility rules defined below are based upon the procedures by which IANA assigns values: "IESG Approval", "IETF Review", "Expert Review", and "Private Use" (as specified in [RFC5226]). The appropriate procedure for a particular type of code point is defined in one or other of the NSIS protocol documents, mostly [RFC5971].

In addition to registered code points, all NSIS protocols provide code points that can be used for experimentation, usually within closed networks, as explained in [RFC3692]. There is no guarantee that independent experiments will not be using the same code point!

8.2. GIST

GIST is extensible in several aspects covered in the subsections below. In these subsections, there are references to document sections in the GIST specification [RFC5971] where more information can be found. The bullet points at the end of each subsection specify the formal administrative actions that would need to be carried out when a new extension is standardized.

More generally, as asserted in Section 1 of the GIST specification, the GIST design could be extended to cater for multicast flows and for situations where the signaling is not tied to an end-to-end data flow. However, it is not clear whether this could be done in a totally backwards-compatible way, and this is not considered within the extensibility model of NSIS.

8.2.1. Use of Different Message Routing Methods

Currently, only two message routing methods are supported (Path Coupled MRM and Loose End MRM), but further MRMs may be defined in the future. See Sections 3.3 and 5.8 of the GIST specification [RFC5971]. One possible additional MRM under development is documented in [EST-MRM]. This MRM would direct signaling towards an explicit target address other than the (current) data flow destination and is intended to assist setting up of state on a new path during "make-before-break" handover sequences in mobile operations. Note that alternative routing methods may require modifications to the firewall traversal techniques used by GIST and NSLPs.

- o New MRMs require allocation of a new MRM-ID either by IETF review of a specification or expert review [RFC5971].

8.2.2. Use of Different Transport Protocols or Security Capabilities

The initial handshake between GIST peers allows a negotiation of the transport protocols to be used. Currently, proposals exist to add DCCP [GIST-DCCP] and the Stream Control Transmission Protocol (SCTP) [GIST-SCTP] transports to GIST; in each case, using Datagram TLS (DTLS) to provide security. See Sections 3.2 and 5.7 of the GIST specification [RFC5971]. GIST expects alternative capabilities to be treated as selection of an alternative protocol stack. Within the protocol stack, the individual protocols used are specified by MA Protocol IDs that are allocated from an IANA registry if new protocols are to be used. See Sections 5.7 and 9 of the GIST specification [RFC5971].

- o Use of an alternative transport protocol or security capability requires allocation of a new MA-Protocol-ID either by IETF review of a specification or expert review [RFC5971].

8.2.3. Use of Alternative Security Services

Currently, only TLS is specified for providing secure channels with MAs. Section 3.9 of the GIST specification [RFC5971] suggests that alternative protocols could be used, but the interactions with GIST functions would need to be carefully specified. See also Section 4.4.2 of the GIST specification [RFC5971].

- o Use of an alternative security service requires allocation of a new MA-Protocol-ID either by IETF review of a specification or expert review [RFC5971].

8.2.4. Query Mode Packet Interception Schemes

GIST has standardized a scheme using RAO mechanisms [GIST-RAO] with UDP packets. If the difficulties of deploying the RAO scheme prove insuperable in particular circumstances, alternative interception schemes can be specified. One proposal that was explored for GIST used UDP port recognition in routers (rather than RAO mechanisms) to drive the interception of packets. See Section 5.3.2 of the GIST specification [RFC5971]. Each NSLP needs to specify membership of an "interception class" whenever it sends a message through GIST. A packet interception scheme can support one or more interception classes. In principle, a GIST instance can support multiple packet interception schemes, but each interception class needs to be associated with exactly one interception scheme in a GIST instance, and GIST instances that use different packet interception schemes for the same interception class will not be interoperable.

Defining an alternative interception class mechanism for incorporation into GIST should be considered as a very radical step, and all alternatives should be considered before taking this path. The main reason for this is that the mechanism will necessarily require additional operations on every packet passing through the affected router interfaces. A number of considerations should be taken into account:

- o Although the interception mechanism need only be deployed on routers that actually need it (probably for a new NSLP), deployment may be constrained if the mechanism requires modification to the hardware of relevant routers and/or needs to await modification of the software by the router vendor.
- o Typically, any packet fields to be examined should be near the header of the packet so that additional memory accesses are not needed to retrieve the values needed for examination.

- o The logic required to determine if a packet should be intercepted needs to be kept simple to minimize the extra per-packet processing.
- o The mechanism should be applicable to both IPv4 and IPv6 packets.
- o Packet interception mechanisms potentially provide an attack path for denial-of-service attacks on routers, in that packets are diverted into the "slow path" and hence can significantly increase the load on the general processing capability of the router. Any new interception mechanism needs to be carefully designed to minimize the attack surface.

Packet interception mechanisms are identified by an "interception class" which is supplied to GIST through the Application Programming Interface for each message sent.

- o New packet interception mechanisms will generally require allocation of one or more new Interception-class-IDs. This does not necessarily need to be placed in an IANA registry as it is primarily used as a parameter in the API between the NSLPs and GIST and may never appear on the wire, depending on the mechanism employed; all that is required is consistent interpretation between the NSLPs and GIST in each applicable node. However, if, as is the case with the current RAO mechanism [GIST-RAO], the scheme distinguishes between multiple packet interception classes by a value carried on the wire (different values of RAO parameter for the RAO mechanism in GIST), an IANA registry may be required to provide a mapping between interception classes and on-the-wire values as discussed in Section 6 of [GIST-RAO].

8.2.5. Use of Alternative NAT Traversal Mechanisms

The mechanisms proposed for both legacy NAT traversal (Section 7.2.1 of the GIST specification [RFC5971]) and GIST-aware NAT traversal (Section 7.2.2 of the GIST specification [RFC5971]) can be extended or replaced. As discussed above, extension of NAT traversal may be needed if a new MRM is deployed. Note that there is extensive discussion of NAT traversal in the NAT/firewall NSLP specification [RFC5973].

8.2.6. Additional Error Identifiers

Making extensions to any of the above items may result in having to create new error modes. See Section 9 and Appendix A.4.1 - A.4.3 of the GIST specification [RFC5971].

- o Additional error identifiers require allocation of new error code(s) and/or subcode(s) and may also require allocation of Additional Information types. These are all allocated on a first-come, first-served basis by IANA [RFC5971].

8.2.7. Defining New Objects To Be Carried in GIST

The A/B (extensibility) flags in each signaling object carried in NSIS protocols enable the community to specify new objects applicable to GIST that can be carried inside a signaling session without breaking existing implementations. See Appendix A.2 of the GIST specification [RFC5971]. The A/B flags can also be used to indicate in a controlled fashion that a certain object must be understood by all GIST nodes, which makes it possible to probe for the support of an extension. One such object already designed is the "Peering Information Object (PIO)" [PEERING-DATA] that allows a Query message to carry additional peering data to be used by the recipient in making the peering decision.

- o New objects require allocation of a new Object Type ID either by IETF review of a specification or through another acceptable published specification [RFC5971].

8.2.8. Adding New Message Types

Major modifications could be made by adding additional GIST message types and defining appropriate processing. It might be necessary to define this as a new version of the protocol. A field is provided in the GIST Common Header containing the version number. GIST currently has no provision for version or capability negotiation that might be needed if a new version was defined.

- o New GIST Message Types require allocation of a new GIST Message Type ID either by IETF review of a specification or expert review [RFC5971].

8.3. QoS NSLP

The QoS NSLP provides signaling for QoS reservations on the Internet. The QoS NSLP decouples the resource reservation model or architecture (QoS model) from the signaling. The signaling protocol is defined in Quality-of-Service NSLP (QoS NSLP) [RFC5974]. The QoS models are defined in separate specifications, and the QoS NSLP can operate with one or more of these models as required by the environment where it is used. It is anticipated that additional QoS models will be developed to address various Internet scenarios in the future. Extensibility of QoS models is considered in Section 8.4.

The QoS NSLP specifically mentions the possibility of using alternative Message Routing Methods (MRMs), apart from the general ability to extend NSLPs using new objects with the standard A/B extensibility flags to allow them to be used in new and old implementations.

There is already work to extend the base QoS NSLP and GIST to enable new QoS signaling scenarios. One such proposal is the Inter-Domain Reservation Aggregation aiming to support large-scale deployment of the QoS NSLP [RESV-AGGR]. Another current proposal seeks to extend the whole NSIS framework towards path-decoupled signaling and QoS reservations [HYPATH].

8.4. QoS Specifications

The QoS Specification template (QSPEC) is defined in [RFC5975]. This provides the language in which the requirements of specific QoS models are described. Introduction of a new QoS model involves defining a new QSPEC. In order to have a new QSPEC allocated by IANA, there must be an acceptable published specification that defines the specific elements within the QSPEC used in the new model. See [RFC5975] for details.

The introduction of new QoS models is designed to enable deployment of NSIS-based QoS control in specific scenarios. One such example is the Integrated Services Controlled Load Service for NSIS [CL].

A key feature provided by defining the QSPEC template is support of a common language for describing QoS requirements and capabilities, which can be reused by any QoS models intending to use the QoS NSLP to signal their requirements for traffic flows. The commonality of the QSPEC parameters ensures a certain level of interoperability of QoS models and reduces the demands on hardware that has to implement the QoS control. Optional QSPEC parameters support the extensibility of the QoS NSLP to other QoS models in the future; new QSPEC parameters can be defined in the document that specifies a new QoS model. See Sections 4.4 and 7 of [RFC5975].

The QSPEC consists of a QSPEC version number, QSPEC objects, plus specification of processing and procedures that can be used to build many QoS models. The definition of a QSPEC can be revised without necessarily changing the version if the changes are functionally backwards compatible. If changes are made that are not backwards compatible, then a new QSPEC version number has to be assigned. Note that a new QSPEC version number is not needed just because additional QSPEC parameters are specified; new versions will be needed only if the existing functionality is modified. The template includes version negotiation procedures that allow the originator of an NSLP

message to retry with a lower QSPEC version if the receiver rejects a message because it does not support the QSPEC version signaled in the message. See Section 3.2 of [RFC5975].

- o Creation of a new, incompatible version of an existing QSPEC requires allocation of a new QSPEC version number that is documented in a permanent and readily available public specification. See [RFC5975].
- o Completely new QSPECs can also be created. Such new QSPECs require allocation of a QSPEC type that is documented in a permanent and readily available public specification. Values are also available for local or experimental use during development. See [RFC5975].
- o Additional QSPEC procedures can be defined requiring allocation of a new QSPEC procedure number that is documented in a permanent and readily available public specification. Values are also available for local or experimental use during development. See [RFC5975].
- o Additional QSPEC parameters and associated error codes can be defined requiring a permanent and readily available public specification document. Values are also available for local or experimental use during development. See [RFC5975].

8.5. NAT/Firewall NSLP

The NAT/firewall signaling can be extended broadly in the same way as the QoS NSLP by defining new parameters to be carried in NAT/firewall NSLP messages. See Section 7 of [RFC5973]. No proposals currently exist to fulfill new use cases for the protocol.

8.6. New NSLP Protocols

Designing a new NSLP is both challenging and easy.

New signaling applications with associated NSLPs can be defined to work in parallel or replace the applications already defined by the NSIS working group. Applications that fit into the NSIS framework will be expected to use GIST to provide transport of signaling messages and appropriate security facilities that relieve the application designer of many "lower-level" problems. GIST provides many important functions through the API that it exposes to the code of the signaling application layer, and allows the signaling application programmer to offload various tasks to GIST, e.g., the channel security, transport characteristics, and signaling node discovery.

Yet, on the other hand, the signaling application designer must take into account that the network environment can be dynamic, both in terms of routing and node availability. The new NSLP designer must take into account at least the following issues:

- o Routing changes, e.g., due to mobility: GIST sends network notifications when something happens in the network, e.g., peers or routing paths change. All signaling applications must be able to handle these notifications and act appropriately. GIST does not include logic to figure out what the NSLP would want to do due to a certain network event. Therefore, GIST gives the notification to the application, and lets it make the right decision.
- o GIST indications: GIST will also send other notifications, e.g., if a signaling peer does not reply to refresh messages, or a certain NSLP message was not successfully delivered to the recipient. NSLP applications must also be able to handle these events. Appendix B in the GIST specification discusses the GIST-NSLP API and the various functionality required, but implementing this interface can be quite challenging; the multitude of asynchronous notifications that can arrive from GIST increases the implementation complexity of the NSLP.
- o Lifetime of the signaling flow: NSLPs should inform GIST when a flow is no longer needed using the SetStateLifetime primitive. This reduces bandwidth demands in the network.
- o NSLP IDs: NSLP messages may be multiplexed over GIST MAs. The new NSLP needs to use a unique NSLPID to ensure that its messages are delivered to the correct application by GIST. A single NSLP could use multiple NSLPIDs, for example, to distinguish different classes of signaling nodes that might handle different levels of aggregation of requests or alternative processing paths. Note that unlike GIST, the NSLPs do not provide a protocol versioning mechanism. If the new NSLP is an upgraded version of an existing NSLP, then it should be distinguished by a different NSLPID.
 - * A new generally available NSLP requires IESG approval for the allocation of a new NSLP ID [RFC5971]
- o Incremental deployment: It would generally be unrealistic to expect every node on the signaling path to have a new NSLP implemented immediately. New NSLPs need to allow for this. The QoS and NAT/firewall NSLPs provide examples of techniques such as proxy modes that cater for cases where the data flow originator and/or receiver does not implement the NSLP.

- o Signaling Message Source IP Address: It is sometimes challenging for an NSLP originating a signaling message to determine the source IP address that should be used in the signaling messages, which may be different from the data flow source address used in the MRI. This challenge occurs either when a node has multiple interfaces or is acting as a proxy for the data flow originator (typically expected to occur during the introduction of NSIS when not all nodes are NSIS enabled). A proxy signaling flow originator generally needs to know and use the correct data flow source IP address, at least initially. As discussed in Section 5.8.1.2 of [RFC5971], the signaling flow originator may choose to alter the source IP address after the initial Query message has established the flow path in order that ICMP messages are directed to the most appropriate node. In the proxy case, the data flow originator would be unaware of the signaling flow, and ICMP messages relating to the signaling would be meaningless if passed on to the data flow originator. Hence, it is essential that an NSLP is aware of the position and role of the node on which it is instantiated and has means of determining the appropriate source address to be used and ensuring that it is used on signaling packets.
- o New MRMs: GIST currently defines two Message Routing Methods, and leaves the door open for new ideas. Thus, it is possible that a new NSLP also requires a new MRM; path-decoupled routing being one example.
- o Cooperation with other NSLPs: Some applications might need resources from two or more different classes in order to operate successfully. The NSLPs managing these resources could operate cooperatively to ensure that such requests were coordinated to avoid wasting signaling bandwidth and prevent race conditions.

It is essential that the security considerations of a new NSLP are carefully analyzed. NSIS NSLPs are deployed in routers as well as host systems; a poorly designed NSLP could therefore provide an attack vector for network resources as well as end systems. The NSLP must also support authorization of users and must allow the use of the GIST authentication and integrity protection mechanisms where users deem them to be necessary.

The API between GIST and NSLPs (see Appendix B in [RFC5971]) is very important to understand. The abstract design in the GIST specification does not specify the exact messaging between GIST and the NSLPs but gives an understanding of the interactions, especially what kinds of asynchronous notifications from GIST the NSLP must be prepared to handle: the actual interface will be dependent on each implementation of GIST.

Messages transmitted by GIST on behalf of an NSLP are identified by a unique NSLP identifier (NSLPID). NSLPIDs are 16-bit unsigned numbers taken from a registry managed by IANA and defined in Section 9 of the GIST specification [RFC5971].

A range of values (32704-32767) is available for Private and Experimental use during development. Any new signaling application that expects to be deployed generally on the Internet needs to use the registration procedure "IESG Approval" in order to request allocation of unique NSLPID value(s) from the IANA registry. There is additional discussion of NSLPIDs in Section 3.8 of the GIST specification.

9. Security Considerations

This document provides information to the community. It does not itself raise new security concerns.

However, any extensions that are made to the NSIS protocol suite will need to be carefully assessed for any security implications. This is particularly important because NSIS messages are intended to be actively processed by NSIS-capable routers that they pass through, rather than simply forwarded as is the case with most IP packets. It is essential that extensions provide means to authorize usage of capabilities that might allocate resources and recommend the use of appropriate authentication and integrity protection measures in order to exclude or adequately mitigate any security issues that are identified.

Authors of new extensions for NSIS should review the analysis of security threats to NSIS documented in [RFC4081] as well as considering whether the new extension opens any new attack paths that need to be mitigated.

GIST offers facilities to authenticate NSIS messages and to ensure that they are delivered reliably. Extensions must allow these capabilities to be used in an appropriate manner to minimize the risks of NSIS messages being misused and must recommend their appropriate usage.

If additional transport protocols are proposed for use in association with GIST, an appropriate set of compatible security functions must be made available in conjunction with the transport protocol to support the authentication and integrity functions expected to be available through GIST.

10. Acknowledgements

This document combines work previously published as two separate documents: "What is Next Steps in Signaling anyway - A User's Guide to the NSIS Protocol Family" written by Roland Bless and Jukka Manner, and "NSIS Extensibility Model" written by John Loughney.

Max Laier, Nuutti Varis and Lauri Liuhto have provided reviews of the "User's Guide" and valuable input. Teemu Huovila also provided valuable input on the later versions.

The "Extensibility Model" borrowed some ideas and some text from RFC 3936 [RFC3936], "Procedures for Modifying the Resource ReSerVation Protocol (RSVP)". Robert Hancock provided text for the original GIST section, since much modified, and Claudia Keppler has provided feedback on this document, while Allison Mankin and Bob Braden suggested that this document be worked on.

11. References

11.1. Normative References

- [RFC3726] Brunner, M., "Requirements for Signaling Protocols", RFC 3726, April 2004.
- [RFC4080] Hancock, R., Karagiannis, G., Loughney, J., and S. Van den Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.
- [RFC4081] Tschofenig, H. and D. Kroeselberg, "Security Threats for Next Steps in Signaling (NSIS)", RFC 4081, June 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5971] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.
- [RFC5973] Stiemerling, M., Tschofenig, H., Aoun, C., and E. Davies, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", RFC 5973, October 2010.
- [RFC5974] Manner, J., Karagiannis, G., and A. McDonald, "NSIS Signaling Layer Protocol (NSLP) for Quality-of-Service Signaling", RFC 5974, October 2010.

- [RFC5975] Ash, G., Bader, A., Kappler, C., and D. Oran, "QSPEC Template for the Quality-of-Service NSIS Signaling Layer Protocol (NSLP)", RFC 5975, October 2010.

11.2. Informative References

- [CL] Kappler, C., Fu, X., and B. Schloer, "A QoS Model for Signaling IntServ Controlled-Load Service with NSIS", Work in Progress, April 2010.
- [EST-MRM] Bless, R., "An Explicit Signaling Target Message Routing Method (EST-MRM) for the General Internet Signaling Transport (GIST) Protocol", Work in Progress, June 2010.
- [GIST-DCCP] Manner, J., "Generic Internet Signaling Transport over DCCP and DTLS", Work in Progress, June 2007.
- [GIST-RAO] Hancock, R., "Using the Router Alert Option for Packet Interception in GIST", Work in Progress, November 2008.
- [GIST-SCTP] Fu, X., Dickmann, C., and J. Crowcroft, "General Internet Signaling Transport (GIST) over Stream Control Transmission Protocol (SCTP) and Datagram Transport Layer Security (DTLS)", Work in Progress, May 2010.
- [HYPATH] Cordeiro, L., Curado, M., Monteiro, E., Bernardo, V., Palma, D., Racaru, F., Diaz, M., and C. Chassot, "GIST Extension for Hybrid On-path Off-path Signaling (HyPath)", Work in Progress, February 2008.
- [NSIS-AUTH] Manner, J., Stiemerling, M., Tschofenig, H., and R. Bless, "Authorization for NSIS Signaling Layer Protocols", Work in Progress, July 2008.
- [PEERING-DATA] Manner, J., Liuhto, L., Varis, N., and T. Huovila, "Peering Data for NSIS Signaling Layer Protocols", Work in Progress, February 2008.
- [RAO-BAD] Rahman, R. and D. Ward, "Use of IP Router Alert Considered Dangerous", Work in Progress, October 2008.
- [RESV-AGGR] Doll, M. and R. Bless, "Inter-Domain Reservation Aggregation for QoS NSLP", Work in Progress, July 2007.

- [RFC1633] Braden, B., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, January 2004.
- [RFC3936] Kompella, K. and J. Lang, "Procedures for Modifying the Resource reSerVation Protocol (RSVP)", BCP 96, RFC 3936, October 2004.
- [RFC4094] Manner, J. and X. Fu, "Analysis of Existing Quality-of-Service Signaling Protocols", RFC 4094, May 2005.
- [TWO-LEVEL] Braden, R. and B. Lindell, "A Two-Level Architecture for Internet Signaling", Work in Progress, November 2002.

Authors' Addresses

Jukka Manner
Aalto University
Department of Communications and Networking (Comnet)
P.O. Box 13000
FIN-00076 Aalto
Finland

Phone: +358 9 470 22481
EMail: jukka.manner@tkk.fi
URI: <http://www.netlab.tkk.fi/~jmanner/>

Roland Bless
Institute of Telematics, Karlsruhe Institute of Technology (KIT)
Zirkel 2, Building 20.20
P.O. Box 6980
Karlsruhe 76049
Germany

Phone: +49 721 608 6413
EMail: bleess@kit.edu
URI: <http://tm.kit.edu/~bleess>

John Loughney
Nokia
955 Page Mill Road
Palo Alto, CA 94303
USA

Phone: +1 650 283 8068
EMail: john.loughney@nokia.com

Elwyn Davies (editor)
Folly Consulting
Soham
UK

EMail: elwynd@folly.org.uk
URI: <http://www.folly.org.uk>

