

NSIS Extensibility Model
draft-loughney-nsis-ext-02.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

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

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

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

This Internet-Draft will expire on September 7, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document discusses the Next Steps in Signaling extensibility model. This model is based upon a two-layer model, where there is a transport layer and a signaling application model. This two-layer provides the ability to develop new signaling applications, while retaining the use of a common transport layer. This document will serve as guidance on how the NSIS architecture can be extended.

Table of Contents

1.	Requirements notation	3
2.	Introduction	3
2.1.	NSIS Extensibility Types	3
3.	GIST Extensibility	4
3.1.	NSLP Identifiers	4
3.2.	Message Routing Methods	4
3.3.	Protocol Indicators	5
3.4.	Router Alert Values	5
4.	NSLP Extensibility	6
4.1.	Common Functionality Among Signaling Applications	7
5.	QoS Model Extensibility	7
6.	IANA Considerations	7
7.	Security Considerations	7
8.	Acknowledgements	7
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	8
	Author's Address	9
	Intellectual Property and Copyright Statements	10

1. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[1\]](#).

2. Introduction

The Next Steps in Signaling Framework NSIS Framework [\[2\]](#) details a basic two-layer framework for signaling on the Internet. The document decomposes signaling into a two-layer model, into a generic transport layer and specific signaling layers.

This model allows for an extensible model for different signaling needs on the the Internet. Currently, the NSIS working group is working on two main signaling applications - QoS signaling [\[3\]](#) and NAT/Firewall signaling [\[4\]](#). It is expected that there will be more signaling applications.

The NSIS Transport Layer Protocol, GIST (General Internet Signaling Transport) GIST [\[5\]](#) defines a basic protocol for routing and transport of per-flow signaling along the path taken by that flow through the network; managing the underlying transport and security protocols.

Above GIST can one or more NSIS Signaling Layer protocols, which can signal for things such as QoS, firewall control and NAT signaling. QoS NSLP [\[3\]](#), NAT/FW NSLP [\[4\]](#). These signaling applications manage their state by using the services that the GIST provides them for signaling.

This two layer approach allows for signaling applications to be developed indepently of the transport. As it is likely that the functionality entities for different signaling applications will be distinct, not all path elements support NSIS signaling will require that a specific signaling application is present - only those nodes that will be maintaining some signaling application state need to support the signaling application.

2.1. NSIS Extensibility Types

Generally, NSIS protocols can be extended in multiple ways. This secition is and overview of the mechansims used. The extensibility rules are based upon the procedures by which IANA assigns values: "Standards Action" (as defined in [\[IANA\]](#)), "IETF Action", "Expert Review", and "Organization/Vendor Private", defined below.

Extensions subject to "IETF Action" require either a Standards Track RFC, Experimental RFC or an Information RFC.

Extensions subject to "Expert Review" refer to values that are to be reviewed by an Expert designated by the IESG. The code points from these ranges are typically used for experimental extensions; such assignments MUST be requested by either Experimental or Information RFCs that document their use and processing, and the actual assignments made during the IANA actions for the document. Values from "Expert Review" ranges MUST be registered with IANA.

"Organization/Vendor Private" ranges refer to values that are enterprise-specific. In this way, different enterprises, vendors, or Standards Development Organizations (SDOs) can use the same code point without fear of collision.

In the NSIS protocols, experimental code points are allocated for experimentation, usually within closed networks, as explained in [RFC 3692](#).^[7] If these experiments yield useful results, it is assumed that they will be formally allocated by one of the above mechanisms.

[3. GIST Extensibility](#)

GIST is Major extensibility options for GIST are:

[3.1. NSLP Identifiers](#)

Each signaling application requires one or more NSLPIDs (different NSLPIDs may be used to distinguish different classes of signaling node, for example to handle different aggregation levels or different processing subsets). An NSLPID must be associated with a unique RAO value. IETF Action is required to allocate a new NSLP Identifier.

[3.2. Message Routing Methods](#)

GIST allows the idea of multiple Message Routing Methods (MRM). The message routing method is indicated in the leading 2 bytes of the MRI object. GIST allocates 2 bits for experimental Routing Methods, for use in closed networks for experimentation purposes. Standards Action is required to allocate new Routing Methods.

Experimental NSLPs are required to be able to use experimental MRMs, and the experimental NSLP is not supported, then the MRM is ignored. The expectation is that the experimental MRM is used within a closed network, for experimental purposes, as explained in [RFC 3692](#).^[7]

3.3. Protocol Indicators

The GIMPS design allows the set of possible protocols to be used in a messaging association to be extended. Every new mode of using a protocol is given by a Protocol Indicator, which is used as a tag in the Node Addressing and Stack Proposal objects. New protocol indicators require IETF Action. Allocating a new protocol indicator requires defining the higher layer addressing information in the Node Addressing Object that is needed to define its configuration.

3.4. Router Alert Values

Router Alert Option (RAO) values are allocated on the basis of IETF consensus. However, new RAO values SHOULD NOT be allocated for each new NSLP. Careful consideration needs to be exercised when choosing to allocate a new RAO value. This section discusses some considerations on how to choose if an existing RAO option should be chosen or a new RAO should be allocated for an NSLP

The use of the RAO is the primary mechanism to indicate that an GIST message should be intercepted by a particular node. There are two basic reasons why a GIST node might wish not to intercept a particular message. The first reason would be because the message is for a signaling application that the node does not process. The second reason would be because the node is processes signaling messages at the aggregate level, not for individual flow, even though the signaling application is present on the node. However, these reasons do not preclude a node processing several RAO values, implying it supports several different signaling applications.

Some of this information can be encoded in the RAO value field, which then allows messages to be filtered on the fast path. There is a tradeoff between two approaches here, whose evaluation depends on whether the processing node is specialised or general purpose:

Fine-Grained: The signaling application (including specific version) and aggregation level are directly identified in the RAO value. A specialised node which handles only a single NSLP can efficiently ignore all other messages; a general purpose node may have to match the RAO value in a message against a long list of possible values.

Coarse-Grained> RAO values are allocated are used on common applications or sets of applications (such as 'All QoS Signaling Applications'). This speeds up the processing in a general purpose node, but a specialised node may have to carry out further processing on the GIST common header to identify the precise messages it needs to consider.

These considerations imply that the RAO value should not be tied directly to the NSLPID, but should be selected for the application on broader considerations of likely deployment scenarios. Note that the exact NSLP is given in the GIMPS common header, and some implementations may still be able to process it on the fast path. The semantics of the node dropping out of the signaling path are the same however the filtering is done. Which is to say that the RAO does not have to have a one-to-one relation to a specific NSLPID, the RAO must be uniquely derivable from the NSLPID.

There is a special consideration in the case of the aggregation level. In this case, whether a message should be processed depends on the network region it is in (specifically, the link it is on). There are then two basic possibilities:

All routers have essentially the same algorithm for which messages they process, i.e. all messages at aggregation level 0. However, messages have their aggregation level incremented on entry to an aggregation region and decremented on exit.

Router interfaces are configured to process messages only above a certain aggregation level and ignore all others. The aggregation level of a message is never changed; signaling messages for end to end flows have level 0, but signaling messages for aggregates are generated with a higher level.

The first technique requires aggregating/deaggregating routers to be configured with which of their interfaces lie at which aggregation level, and also requires consistent message rewriting at these boundaries. The second technique eliminates the rewriting, but requires interior routers to be configured also. It is not clear what the right trade-off between these options is.

4. NSLP Extensibility

An NSLP roughly should correspond to a class of signaling application, which requires some state maintenance along a network path. Signaling applications should be generic enough to allow for state manipulation for a common set of functions. This allows for an architecture which allows for flexible network deployment, without over-burdening nodes with extra signaling applications.

New NSLPs should be created when there is a new signaling application. Creating new NSLPs which only slightly modify existing NSLPs is not recommended as it will increase deployment complexity (common nodes would need to support multiple NSLPs for similar functionality).

4.1. Common Functionality Among Signaling Applications

While NSIS has adopted a two-layer signaling approach, in practice, there is much in common between different NSLPs. Efforts should be made to ensure some high-level of compatibility among signaling applications, which could be reused by some implementations in order to combine multiple signaling applications into one implementation, but that is an implementation decision.

5. QoS Model Extensibility

The QoS NSLP provides signaling for QoS reservations on the Internet. The QoS NSLP decouples the resource reservation model or architecture from the signaling. The QoS specification is defined in QSpec [6]. New QoS Models require IETF action, which defines the elements within the QSpec. See QSpec [6] for details.

A key part of the QoS model is support a common language, which can be shared among several QOSMs. These QSPEC parameters ensure a certain level of interoperability of QOSMs. Optional QSPEC parameters support the extensibility of the QoS NSLP to other QOSMs in the future. The node initiating the NSIS signaling adds an Initiator QSPEC that must not be removed, thereby ensuring the intention of the NSIS initiator is preserved along the signaling path.

6. IANA Considerations

This document outlines the basic rules for extending NSIS protocols. This instructions IANA on allocation policies for NSIS protocols.

7. Security Considerations

This document is an informational document, outlining the extensibility model of the NSIS protocol suite. As such, this document does not impact the security of the Internet directly.

8. Acknowledgements

This document borrows some ideas and some text from [RFC3936](#) [8], Procedures for Modifying the Resource reSerVation Protocol (RSVP).

Robert Hancock provided text for much of the GIST section. Claudia Keppler have provided feedback on this draft.

Allison Mankin and Bob Braden suggest that this draft be worked on.

9. References

9.1. Normative References

- [2] Hancock, R., "Next Steps in Signaling: Framework", [draft-ietf-nsis-fw-07](#) (work in progress), December 2004.
- [4] Stiemerling, M., "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", [draft-ietf-nsis-nslp-natfw-09](#) (work in progress), February 2006.
- [5] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signaling Transport", [draft-ietf-nsis-ntlp-09](#) (work in progress), February 2006.
- [3] Manner, J., "NSLP for Quality-of-Service signalling", [draft-ietf-nsis-qos-nslp-09](#) (work in progress), February 2006.
- [6] Ash, J., "QoS-NSLP QSPEC Template", [draft-ietf-nsis-qspec-08](#) (work in progress), December 2005.
- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

9.2. Informative References

- [7] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.
- [8] Kompella, K. and J. Lang, "Procedures for Modifying the Resource reSerVation Protocol (RSVP)", [BCP 96](#), [RFC 3936](#), October 2004.

Author's Address

John Loughney
Nokia
Itamerenkatu 11-13
Helsinki 00180
Finland

Phone: +358504836242

Email: john.loughney@nokia.com

Intellectual Property Statement

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

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

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

Disclaimer of Validity

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

Copyright Statement

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

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

