

Requirements for Signaling Protocols

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This document defines requirements for signaling across different network environments, such as across administrative and/or technology domains. Signaling is mainly considered for Quality of Service (QoS) such as the Resource Reservation Protocol (RSVP). However, in recent years, several other applications of signaling have been defined. For example, signaling for label distribution in Multiprotocol Label Switching (MPLS) or signaling to middleboxes. To achieve wide applicability of the requirements, the starting point is a diverse set of scenarios/use cases concerning various types of networks and application interactions. This document presents the assumptions before listing the requirements. The requirements are grouped according to areas such as architecture and design goals, signaling flows, layering, performance, flexibility, security, and mobility.

Table of Contents

1.	Introduction	5
1.1.	Keywords	5
2.	Terminology.	5
3.	Problem Statement and Scope.	6
4.	Assumptions and Exclusions	8
4.1.	Assumptions and Non-Assumptions.	8
4.2.	Exclusions	9
5.	Requirements	10
5.1.	Architecture and Design Goals.	11
5.1.1.	NSIS SHOULD Provide Availability Information on Request	11
5.1.2.	NSIS MUST be Designed Modularly.	11
5.1.3.	NSIS MUST Decouple Protocol and Information. . .	12
5.1.4.	NSIS MUST Support Independence of Signaling and Network Control Paradigm	12
5.1.5.	NSIS SHOULD be Able to Carry Opaque Objects. . .	12
5.2.	Signaling Flows.	12
5.2.1.	The Placement of NSIS Initiator, Forwarder, and Responder Anywhere in the Network MUST be Allowed.	12
5.2.2.	NSIS MUST Support Path-Coupled and MAY Support Path-Decoupled Signaling	13
5.2.3.	Concealment of Topology and Technology Information SHOULD be Possible	13
5.2.4.	Transparent Signaling Through Networks SHOULD be Possible	13
5.3.	Messaging.	13
5.3.1.	Explicit Erasure of State MUST be Possible . . .	13
5.3.2.	Automatic Release of State After Failure MUST be Possible	14
5.3.3.	NSIS SHOULD Allow for Sending Notifications Upstream	14
5.3.4.	Establishment and Refusal to set up State MUST be Notified.	15
5.3.5.	NSIS MUST Allow for Local Information Exchange .	15
5.4.	Control Information.	16
5.4.1.	Mutability Information on Parameters SHOULD be Possible	16
5.4.2.	It SHOULD be Possible to Add and Remove Local Domain Information	16
5.4.3.	State MUST be Addressed Independent of Flow Identification	16
5.4.4.	Modification of Already Established State SHOULD be Seamless.	16
5.4.5.	Grouping of Signaling for Several Micro-Flows MAY be Provided.	17

5.5.	Performance.	17
5.5.1.	Scalability.	17
5.5.2.	NSIS SHOULD Allow for Low Latency in Setup . . .	18
5.5.3.	NSIS MUST Allow for Low Bandwidth Consumption for the Signaling Protocol	18
5.5.4.	NSIS SHOULD Allow to Constrain Load on Devices .	18
5.5.5.	NSIS SHOULD Target the Highest Possible Network Utilization.	18
5.6.	Flexibility.	19
5.6.1.	Flow Aggregation	19
5.6.2.	Flexibility in the Placement of the NSIS Initiator/Responder.	19
5.6.3.	Flexibility in the Initiation of State Change. .	19
5.6.4.	SHOULD Support Network-Initiated State Change. .	19
5.6.5.	Uni / Bi-directional State Setup	20
5.7.	Security	20
5.7.1.	Authentication of Signaling Requests	20
5.7.2.	Request Authorization.	20
5.7.3.	Integrity Protection	20
5.7.4.	Replay Protection.	21
5.7.5.	Hop-by-Hop Security.	21
5.7.6.	Identity Confidentiality and Network Topology Hiding	21
5.7.7.	Denial-of-Service Attacks.	21
5.7.8.	Confidentiality of Signaling Messages.	22
5.7.9.	Ownership of State	22
5.8.	Mobility	22
5.8.1.	Allow Efficient Service Re-Establishment After Handover	22
5.9.	Interworking with Other Protocols and Techniques	22
5.9.1.	MUST Interwork with IP Tunneling	22
5.9.2.	MUST NOT Constrain Either to IPv4 or IPv6.	23
5.9.3.	MUST be Independent from Charging Model.	23
5.9.4.	SHOULD Provide Hooks for AAA Protocols	23
5.9.5.	SHOULD work with Seamless Handoff Protocols. . . .	23
5.9.6.	MUST Work with Traditional Routing	23
5.10.	Operational.	23
5.10.1.	Ability to Assign Transport Quality to Signaling Messages	23
5.10.2.	Graceful Fail Over	24
5.10.3.	Graceful Handling of NSIS Entity Problems. . . .	24
6.	Security Considerations.	24
7.	Acknowledgments.	24
8.	Appendix: Scenarios/Use Cases.	26
8.1.	Terminal Mobility.	26
8.2.	Wireless Networks.	28
8.3.	An Example Scenario for 3G Wireless Networks	29
8.4.	Wired Part of Wireless Network	31

8.5.	Session Mobility	33
8.6.	QoS Reservation/Negotiation from Access to Core Network.	34
8.7.	QoS Reservation/Negotiation Over Administrative Boundaries	34
8.8.	QoS Signaling Between PSTN Gateways and Backbone Routers	35
8.9.	PSTN Trunking Gateway.	36
8.10.	An Application Requests End-to-End QoS Path from the Network.	38
8.11.	QoS for Virtual Private Networks	39
	8.11.1. Tunnel End Points at the Customer Premises . . .	39
	8.11.2. Tunnel End Points at the Provider Premises . . .	39
9.	References	40
	9.1. Normative References	40
	9.2. Informative References	40
10.	Authors' Addresses	41
11.	Full Copyright Statement	42

1. Introduction

This document is the product of the Next Steps in Signaling (NSIS) Working Group. It defines requirements for signaling across different network environments. It does not list any problems of existing signaling protocols such as [RSVP].

In order to derive requirements for signaling it is necessary to first have an idea of the scope within which they are applicable. Therefore, we list use cases and scenarios where an NSIS protocol could be applied. The scenarios are used to help derive requirements and to test the requirements against use cases.

The requirements listed are independent of any application. However, resource reservation and QoS related issues are used as examples within the text. However, QoS is not the only field where signaling is used in the Internet. Signaling might also be used as a communication protocol to setup and maintain the state in middleboxes [RFC3234].

This document does not cover requirements in relation to some networking areas, in particular, interaction with host and site multihoming. We leave these for future analysis.

1.1. Keywords

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 BCP 14, RFC 2119 [KEYWORDS].

2. Terminology

We list the most often used terms in the document. However, they cannot be made precise without a more complete architectural model, and they are not meant to prescribe any solution in the document. Where applicable, they will be defined in protocol documents.

NSIS Entity (NE): The function within a node, which implements an NSIS protocol. In the case of path-coupled signaling, the NE will always be on the data path.

NSIS Forwarder (NF): NSIS Entity between a NI and NR, which may interact with local state management functions in the network. It also propagates NSIS signaling further through the network.

NSIS Initiator (NI): NSIS Entity that starts NSIS signaling to set up or manipulate network state.

NSIS Responder (NR): NSIS Entity that terminates NSIS signaling and can optionally interact with applications as well.

Flow: A traffic stream (sequence of IP packets between two end systems) for which a specific packet level treatment is provided. The flow can be unicast (uni- or bi-directional) or multicast. For multicast, a flow can diverge into multiple flows as it propagates toward the receiver. For multi-sender multicast, a flow can also diverge when viewed in the reverse direction (toward the senders).

Data Path: The route across the networks taken by a flow or aggregate, i.e., which domains/subdomains it passes through and the egress/ingress points for each.

Signaling Path: The route across the networks taken by a signaling flow or aggregate, i.e., which domains/subdomains it passes through and the egress/ingress points for each.

Path-coupled signaling: A mode of signaling where the signaling messages follow a path that is tied to the data packets. Signaling messages are routed only through nodes (NEs) that are in the data path.

Path-decoupled signaling: Signaling with independent data and signaling paths. Signaling messages are routed to nodes (NEs) which are not assumed to be on the data path, but which are (presumably) aware of it. Signaling messages will always be directly addressed to the neighbor NE, and the NI/NR may have no relation at all with the ultimate data sender or receiver.

Service: A generic something provided by one entity and consumed by another. It can be constructed by allocating resources. The network can provide it to users or a network node can provide it to packets.

3. Problem Statement and Scope

We provide in the following a preliminary architectural picture as a basis for discussion. We will refer to it in the following requirement sections.

Note that this model is intended not to constrain the technical approach taken subsequently, simply to allow concrete phrasing of requirements (e.g., requirements about placement of the NSIS Initiator.)

Roughly, the scope of NSIS is assumed to be the interaction between the NSIS Initiator, NSIS Forwarder(s), and NSIS Responder including a protocol to carry the information, and the syntax/semantics of the information that is exchanged. Further statements on assumptions/exclusions are given in the next Section.

The main elements are:

1. Something that starts the request for state to be set up in the network, the NSIS Initiator.

This might be in the end system or within some other part of the network. The distinguishing feature of the NSIS Initiator is that it acts on triggers coming (directly or indirectly) from the higher layers in the end systems. It needs to map the services requested by them, and also provides feedback information to the higher layers, which might be used by transport layer algorithms or adaptive applications.

2. Something that assists in managing state further along the signaling path, the NSIS Forwarder.

The NSIS Forwarder does not interact with higher layers, but interacts with the NSIS Initiator, NSIS Responder, and possibly one or more NSIS Forwarders on the signaling path, edge-to-edge or end-to-end.

3. Something that terminates the signaling path, the NSIS Responder.

The NSIS responder might be in an end-system or within other equipment. The distinguishing feature of the NSIS Responder is that it responds to requests at the end of a signaling path.

4. The signaling path traverses an underlying network covering one or more IP hops. The underlying network might use locally different technology. For instance, QoS technology has to be provisioned appropriately for the service requested. In the QoS example, an NSIS Forwarder maps service-specific information to technology-related QoS parameters and receives indications about success or failure in response.

5. We can see the network at the level of domains/subdomains rather than individual routers (except in the special case that the domain contains one link). Domains are assumed to be administrative entities. So security requirements might apply differently for the signaling between the domains and within a domain. Both cases we deal with in this document.

4. Assumptions and Exclusions

4.1. Assumptions and Non-Assumptions

1. The NSIS signaling could run end-to-end, end-to-edge, or edge-to-edge, or network-to-network (between providers), depending on what point in the network acts as NSIS initiator, and how far towards the other end of the network the signaling propagates. In general, we could expect NSIS Forwarders to become more 'dense' towards the edges of the network, but this is not a requirement. For example, in the case of QoS, an over-provisioned domain might contain no NSIS Forwarders at all (and be NSIS transparent); at the other extreme, NSIS Forwarders might be placed at every router. In the latter case, QoS provisioning can be carried out in a local implementation-dependent way without further signaling, whereas in the case of remote NSIS Forwarders, a protocol might be needed to control the routers along the path. This protocol is then independent of the end-to-end NSIS signaling.
2. We do not consider 'pure' end-to-end signaling that is not interpreted anywhere within the network. Such signaling is a higher-layer issue and IETF protocols such as SIP etc. can be used.
3. Where the signaling does cover several domains, we do not exclude that different signaling protocols are used in each domain. We only place requirements on the universality of the control information that is being transported. (The goals here would be to allow the use of signaling protocols, which are matched to the characteristics of the portion of the network being traversed.) Note that the outcome of NSIS work might result in various flavors of the same protocol.
4. We assume that the service definitions a NSIS Initiator can ask for are known in advance of the signaling protocol running. For instance in the QoS example, the service definition includes QoS parameters, lifetime of QoS guarantee etc., or any other service-specific parameters.

There are many ways service requesters get to know about available services. There might be standardized services, the definition can be negotiated together with a contract, the service definition is published in some on-line directory (e.g., at a Web page), and so on.

5. We assume that there are means for the discovery of NSIS entities in order to know the signaling peers (solutions include static configuration, automatically discovered, or implicitly runs over

the right nodes along the data path, etc.). The discovery of the NSIS entities has security implications that need to be addressed properly. For some security mechanisms (i.e., Kerberos, pre-shared secret) it is required to know the identity of the other entity. Hence the discovery mechanism may provide means to learn this identity, which is then later used to retrieve the required keys and parameters.

6. NSIS assumes layer 3 routing and the determination of next data node selection is not done by NSIS.

4.2. Exclusions

1. Development of specific mechanisms and algorithms for application and transport layer adaptation are not considered, nor are the protocols that would support it.
2. Specific mechanisms (APIs and so on) for interaction between transport/applications and the network layer are not considered, except to clarify the requirements on the negotiation capabilities and information semantics that would be needed of the signaling protocol.
3. Specific mechanisms and protocols for provisioning or other network control functions within a domain/subdomain are not considered. The goal is to reuse existing functions and protocols unchanged. However, NSIS itself can be used for signaling within a domain/subdomain.

For instance in the QoS example, it means that the setting of QoS mechanisms in a domain is out of scope, but if we have a tunnel, NSIS could also be used for tunnel setup with QoS guarantees. It should be possible to exploit these mechanisms optimally within the end-to-end context. Consideration of how to do this might generate new requirements for NSIS however. For example, the information needed by a NSIS Forwarder to manage a radio subnetwork needs to be provided by the NSIS solution.

4. Specific mechanisms (APIs and so on) for interaction between the network layer and underlying provisioning mechanisms are not considered.
5. Interaction with resource management or other internal state management capabilities is not considered. Standard protocols might be used for this. This may imply requirements for the sort of information that should be exchanged between the NSIS entities.

6. Security implications related to multicasting are outside the scope of the signaling protocol.
 7. Service definitions and in particular QoS services and classes are out of scope. Together with the service definition any definition of service specific parameters are not considered in this document. Only the base NSIS signaling protocol for transporting the service information are addressed.
 8. Similarly, specific methods, protocols, and ways to express service information in the Application/Session level are not considered (e.g., SDP, SIP, RTSP, etc.).
 9. The specification of any extensions needed to signal information via application level protocols (e.g., SDP), and the mapping to NSIS information are considered outside of the scope of NSIS working group, as this work is in the direct scope of other IETF working groups (e.g., MMUSIC).
 10. Handoff decision and trigger sources: An NSIS protocol is not used to trigger handoffs in mobile IP, nor is it used to decide whether to handoff or not. As soon as or in some situations even before a handoff happened, an NSIS protocol might be used for signaling for the particular service again. The basic underlying assumption is that the route comes first (defining the path) and the signaling comes after it (following the path). This doesn't prevent a signaling application at some node interacting with something that modifies the path, but the requirement is then just for NSIS to live with that possibility. However, NSIS must interwork with several protocols for mobility management.
 11. Service monitoring is out of scope. It is heavily dependent on the type of the application and or transport service, and in what scenario it is used.
5. Requirements

This section defines more detailed requirements for a signaling solution, respecting the problem statement, scoping assumptions, and terminology considered earlier. The requirements are in subsections, grouped roughly according to general technical aspects: architecture and design goals, topology issues, parameters, performance, security, information, and flexibility.

Two general (and potentially contradictory) goals for the solution are that it should be applicable in a very wide range of scenarios, and at the same time be lightweight in implementation complexity and resource consumption requirements in NSIS Entities. We use the terms

'access' and 'core' informally in the discussion of some particular requirements to refer to deployment conditions where particular protocol attributes, especially performance characteristics, have special importance. Specifically, 'access' refers to lower capacity networks with fewer users and sessions. 'Core' refers to high capacity networks with a large number of users and sessions.

One approach to this is that the solution could deal with certain requirements via modular components or capabilities, which are optional to implement or use in individual nodes.

5.1. Architecture and Design Goals

This section contains requirements related to desirable overall characteristics of a solution, e.g., enabling flexibility, or independence of parts of the framework.

5.1.1. NSIS SHOULD Provide Availability Information on Request

NSIS SHOULD provide a mechanism to check whether state to be setup is available without setting it up. For the resource reservation example this translates into checking resource availability without performing resource reservation. In some scenarios, e.g., the mobile terminal scenario, it is required to query, whether resources are available, without performing a reservation on the resource.

5.1.2. NSIS MUST be Designed Modularly

A modular design allows for more lightweight implementations, if fewer features are needed. Mutually exclusive solutions are supported. Examples for modularity:

- Work over any kind of network (narrowband versus broadband, error-prone versus reliable, ...). This implies low bandwidth signaling, and elimination of redundant information MUST be supported if necessary.
- State setup for uni- and bi-directional flows is possible.
- Extensible in the future with different add-ons for certain environments or scenarios.
- Protocol layering, where appropriate. This means NSIS MUST provide a base protocol, which can be adapted to different environments.

5.1.3. NSIS MUST Decouple Protocol and Information

The signaling protocol MUST be clearly separated from the control information being transported. This provides for the independent development of these two aspects of the solution, and allows for this control information to be carried within other protocols, including application layer ones, existing ones or those being developed in the future. The flexibility gained in the transport of information allows for the applicability of the same protocol in various scenarios.

However, note that the information carried needs to be standardized; otherwise interoperability is difficult to achieve.

5.1.4. NSIS MUST Support Independence of Signaling and Network Control Paradigm

The signaling MUST be independent of the paradigm and mechanism of network control. E.g., in the case of signaling for QoS, the independence of the signaling protocol from the QoS provisioning allows for using the NSIS protocol together with various QoS technologies in various scenarios.

5.1.5. NSIS SHOULD be Able to Carry Opaque Objects

NSIS SHOULD be able to pass around opaque objects, which are interpreted only by some NSIS-capable nodes.

5.2. Signaling Flows

This section contains requirements related to the possible signaling flows that should be supported, e.g., over what parts of the flow path, between what entities (end-systems, routers, middleboxes, management systems), in which direction.

5.2.1. The placement of NSIS Initiator, Forwarder, and Responder Anywhere in the Network MUST be Allowed

The protocol MUST work in various scenarios such as host-to-network-to-host, edge-to-edge, (e.g., just within one provider's domain), user-to-network (from end system into the network, ending, e.g., at the entry to the network and vice versa), and network-to-network (e.g., between providers).

Placing the NSIS Forwarder and NSIS Initiator functions at different locations allows for various scenarios to work with the same protocol.

5.2.2. NSIS MUST Support Path-Coupled and MAY Support Path-Decoupled Signaling.

The path-coupled signaling mode MUST be supported. NSIS signaling messages are routed only through nodes (NEs) that are in the data path.

However, there is a set of scenarios, where signaling is not on the data path. Therefore, NSIS MAY support the path-decoupled signaling mode, where signaling messages are routed to nodes (NEs), which are not assumed to be on the data path, but which are aware of it.

5.2.3. Concealment of Topology and Technology Information SHOULD be Possible

The NSIS protocol SHOULD allow for hiding the internal structure of a NSIS domain from end-nodes and from other networks. Hence an adversary should not be able to learn the internal structure of a network with the help of the signaling protocol.

In various scenarios, topology information should be hidden for various reasons. From a business point of view, some administrations don't want to reveal the topology and technology used.

5.2.4. Transparent Signaling Through Networks SHOULD be Possible

It SHOULD be possible that the signaling for some flows traverses path segments transparently, i.e., without interpretation at NSIS Forwarders within the network. An example would be a subdomain within a core network, which only interpreted signaling for aggregates established at the domain edge, with the signaling for individual flows passing transparently through it.

In other words, NSIS SHOULD work in hierarchical scenarios, where big pipes/trunks are setup using NSIS signaling, but also flows which run within that big pipe/trunk are setup using NSIS.

5.3. Messaging

5.3.1. Explicit Erasure of State MUST be Possible

When state along a path is no longer necessary, e.g., because the application terminates, or because a mobile host experienced a hand-off, it MUST be possible to erase the state explicitly.

5.3.2. Automatic Release of State After Failure MUST be Possible

When the NSIS Initiator goes down, the state it requested in the network SHOULD be released, since it will most likely no longer be necessary.

After detection of a failure in the network, any NSIS Forwarder/Initiator MUST be able to release state it is involved in. For example, this may require signaling of the "Release after Failure" message upstream as well as downstream, or soft state timing out.

The goal is to prevent stale state within the network and add robustness to the operation of NSIS. So in other words, an NSIS signaling protocol or mechanisms MUST provide means for an NSIS entity to discover and remove local stale state.

Note that this might need to work together with a notification mechanism. Note as well, that transient failures in NSIS processing shouldn't necessarily have to cause all state to be released immediately.

5.3.3. NSIS SHOULD Allow for Sending Notifications Upstream

NSIS Forwarders SHOULD notify the NSIS Initiator or any other NSIS Forwarder upstream, if there is a state change inside the network. There are various types of network changes for instance among them:

Recoverable errors: the network nodes can locally repair this type error. The network nodes do not have to notify the users of the error immediately. This is a condition when the danger of degradation (or actual short term degradation) of the provided service was overcome by the network (NSIS Forwarder) itself.

Unrecoverable errors: the network nodes cannot handle this type of error, and have to notify the users as soon as possible.

Service degradation: In case the service cannot be provided completely but only partially.

Repair indication: If an error occurred and it has been fixed, this triggers the sending of a notification.

Service upgrade available: If a previously requested better service becomes available.

The content of the notification is very service specific, but it is must at least carry type information. Additionally, it may carry the location of the state change.

The notifications may or may not be in response to a NSIS message. This means an NSIS entity has to be able to handle notifications at any time.

Note however, that there are a number of security consideration needs to be solved with notification, even more important if the notification is sent without prior request (asynchronously). The problem basically is, that everybody could send notifications to any NSIS entity and the NSIS entity most likely reacts on the notification. For example, if it gets an error notification it might erase state, even if everything is ok. So the notification might depend on security associations between the sender of the notification and its receiver. If a hop-by-hop security mechanism is chosen, this implies also that notifications need to be sent on the reverse path.

5.3.4. Establishment and Refusal to Set Up State MUST be Notified

A NR MUST acknowledge establishment of state on behalf of the NI requesting establishment of that state. A refusal to set up state MUST be replied with a negative acknowledgement by the NE refusing to set up state. It MUST be sent to the NI. Depending on the signaling application the (positive or negative) notifications may have to pass through further NEs upstream. Information on the reason of the refusal to set up state MAY be made available. For example, in the resource reservation example, together with a negative answer, the amount of resources available might also be returned.

5.3.5. NSIS MUST Allow for Local Information Exchange

The signaling protocol MUST be able to exchange local information between NSIS Forwarders located within one single administrative domain. The local information exchange is performed by a number of separate messages not belonging to an end-to-end signaling process. Local information might, for example, be IP addresses, notification of successful or erroneous processing of signaling messages, or other conditions.

In some cases, the NSIS signaling protocol MAY carry identification of the NSIS Forwarders located at the boundaries of a domain. However, the identification of edge should not be visible to the end host (NSIS Initiator) and only applies within one administrative domain.

5.4. Control Information

This section contains requirements related to the control information that needs to be exchanged.

5.4.1. Mutability Information on Parameters SHOULD be Possible

It is possible that nodes modify parameters of a signaling message. However, it SHOULD be possible for the NSIS Initiator to control the mutability of the signaled information. For example, the NSIS Initiator should be able to control what is requested end-to-end, without the request being gradually mutated as it passes through a sequence of nodes.

5.4.2. It SHOULD be Possible to Add and Remove Local Domain Information

It SHOULD be possible to add and remove local scope elements. Compared to Requirement 5.3.5 this requirement does use the normal signaling process and message exchange for transporting local information. For example, at the entrance to a domain, domain-specific information is added information is added, which is used in this domain only, and the information is removed again when a signaling message leaves the domain. The motivation is in the economy of re-using the protocol for domain internal signaling of various information pieces. Where additional information is needed within a particular domain, it should be possible to carry this at the same time as the end-to-end information.

5.4.3. State MUST be Addressed Independent of Flow Identification

Addressing or identifying state MUST be independent of the flow identifier (flow end-points, topological addresses). Various scenarios in the mobility area require this independence because flows resulting from handoff might have changed end-points etc. but still have the same service requirement. Also several proxy-based signaling methods profit from such independence, though these are not chartered work items for NSIS.

5.4.4. Modification of Already Established State SHOULD be Seamless

In many case, the established state needs to be updated (in QoS example upgrade or downgrade of resource usage). This SHOULD happen seamlessly without service interruption. At least the signaling protocol should allow for it, even if some data path elements might not be capable of doing so.

5.4.5. Grouping of Signaling for Several Micro-Flows MAY be Provided

NSIS MAY group signaling information for several micro-flows into one signaling message. The goal of this is the optimization in terms of setup delay, which can happen in parallel. This helps applications requesting several flows at once. Also potential refreshes (in case of a soft state solution) might profit from grouping.

However, the network need not know that a relationship between the grouped flows exists. There MUST NOT be any transactional semantic associated with the grouping. It is only meant for optimization purposes.

5.5. Performance

This section discusses performance requirements and evaluation criteria and the way in which these could and should be traded off against each other in various parts of the solution.

Scalability is always an important requirement for signaling protocols. However, the type of scalability and its importance varies from one scenario to another.

Note that many of the performance issues are heavily dependent on the scenario assumed and are normally a trade-off between speed, reliability, complexity, and scalability. The trade-off varies in different parts of the network. For example, in radio access networks low bandwidth consumption will outweigh the low latency requirement, while in core networks it may be reverse.

5.5.1. Scalability

NSIS MUST be scalable in the number of messages received by a signaling communication partner (NSIS Initiator, NSIS Forwarder, and NSIS Responder). The major concern lies in the core of the network, where large numbers of messages arrive.

It MUST be scalable in number of hand-offs in mobile environments. This mainly applies in access networks, because the core is transparent to mobility in most cases.

It MUST be scalable in the number of interactions for setting up state. This applies for end-systems setting up several states. Some servers might be expected to setup a large number of states.

Scalability in the amount of state per entity MUST be achieved for NSIS Forwarders in the core of the network.

Scalability in CPU usage MUST be achieved on end terminals and intermediate nodes in case of many state setup processes at the same time.

Specifically, NSIS MUST work in Internet scale deployments, where the use of signaling by hosts becomes universal. Note that requirement 5.2.4 requires the functionality of transparently signaling through networks without interpretation. Additionally, requirement 5.6.1 lists the capability to aggregate. Furthermore, requirement 5.5.4 states that NSIS should be able to constrain the load on devices. Basically, the performance of the signaling MUST degrade gracefully rather than catastrophically under overload conditions.

5.5.2. NSIS SHOULD Allow for Low Latency in Setup

NSIS SHOULD allow for low latency setup of states. This is only needed in scenarios where state setups are required on a short time scale (e.g., handover in mobile environments), or where human interaction is immediately concerned (e.g., voice communication setup delay).

5.5.3. NSIS MUST Allow for Low Bandwidth Consumption for the Signaling Protocol

NSIS MUST allow for low bandwidth consumption in certain access networks. Again only small sets of scenarios call for low bandwidth, mainly those where wireless links are involved.

5.5.4. NSIS SHOULD Allow to Constrain Load on Devices

The NSIS architecture SHOULD give the ability to constrain the load (CPU load, memory space, signaling bandwidth consumption and signaling intensity) on devices where it is needed. One of the reasons is that the protocol handling should have a minimal impact on interior (core) nodes.

This can be achieved by many different methods. Examples include message aggregation, header compression, minimizing functionality, or ignoring signaling in core nodes. NSIS may choose any method as long as the requirement is met.

5.5.5. NSIS SHOULD Target the Highest Possible Network Utilization

This requirement applies specifically to QoS signaling.

There are networking environments that require high network utilization for various reasons, and the signaling protocol SHOULD to its best ability support high resource utilization while maintaining appropriate service quality.

In networks where resources are very expensive (as is the case for many wireless networks), efficient network utilization for signaling traffic is of critical financial importance. On the other hand there are other parts of the network where high utilization is not required.

5.6. Flexibility

This section lists the various ways the protocol can flexibly be employed.

5.6.1. Flow Aggregation

NSIS MUST allow for flow aggregation, including the capability to select and change the level of aggregation.

5.6.2. Flexibility in the Placement of the NSIS Initiator/Responder

NSIS MUST be flexible in placing an NSIS Initiator and NSIS Responder. The NSIS Initiator might be located at the sending or the receiving side of a data stream, and the NSIS Responder naturally on the other side.

Also network-initiated signaling and termination MUST be allowed in various scenarios such as PSTN gateways, some VPNs, and mobility. This means the NSIS Initiator and NSIS Responder might not be at the end points of the data stream.

5.6.3. Flexibility in the Initiation of State Change

The NSIS Initiator or the NSIS Responder SHOULD be able to initiate a change of state. In the example of resource reservation this is often referred to as resource re-negotiation. It can happen due to various reasons, such as local resource shortage (CPU, memory on end-system) or a user changed application preference/profiles.

5.6.4. SHOULD Support Network-Initiated State Change

NSIS SHOULD support network-initiated state change. In the QoS example, this is used in cases, where the network is not able to further guarantee resources and wants to e.g., downgrade a resource reservation.

5.6.5. Uni / Bi-Directional State Setup

Both unidirectional as well as bi-direction state setup SHOULD be possible. With bi-directional state setup we mean that the state for bi-directional data flows is setup. The bi-directional data flows have the same end-points, but the path in the two directions does not need to be the same.

The goal of a bi-directional state setup is mainly an optimization in terms of setup delay. There is no requirements on constrains such as use of the same data path etc.

5.7. Security

This section discusses security-related requirements. The NSIS protocol MUST provide means for security, but it MUST be allowed that nodes implementing NSIS signaling do not have to use the security means.

5.7.1. Authentication of Signaling Requests

A signaling protocol MUST make provision for enabling various entities to be authenticated against each other using strong authentication mechanisms. The term strong authentication points to the fact that weak plain-text password mechanisms must not be used for authentication.

5.7.2. Request Authorization

The signaling protocol MUST provide means to authorize state setup requests. This requirement demands a hook to interact with a policy entity to request authorization data. This allows an authenticated entity to be associated with authorization data and to verify the request. Authorization prevents state setup by unauthorized entities, setups violating policies, and theft of service. Additionally it limits denial of service attacks against parts of the network or the entire network caused by unrestricted state setups. Additionally it might be helpful to provide some means to inform other protocols of participating nodes within the same administrative domain about a previous successful authorization event.

5.7.3. Integrity Protection

The signaling protocol MUST provide means to protect the message payloads against modifications. Integrity protection prevents an adversary from modifying parts of the signaling message and from mounting denial of service or theft of service type of attacks against network elements participating in the protocol execution.

5.7.4. Replay Protection

To prevent replay of previous signaling messages the signaling protocol MUST provide means to detect old i.e., already transmitted signaling messages. A solution must cover issues of synchronization problems in the case of a restart or a crash of a participating network element.

5.7.5. Hop-by-Hop Security

Channel security between signaling entities MUST be implemented. It is a well known and proven concept in Quality of Service and other signaling protocols to have intermediate nodes that actively participate in the protocol to modify the messages as it is required by processing rules. Note that this requirement does not exclude end-to-end or network-to-network security of a signaling message. End-to-end security between the NSIS Initiator and the NSIS Responder may be used to provide protection of non-mutable data fields. Network-to-network security refers to the protection of messages over various hops but not in an end-to-end manner i.e., protected over a particular network.

5.7.6. Identity Confidentiality and Network Topology Hiding

Identity confidentiality SHOULD be supported. It enables privacy and avoids profiling of entities by adversary eavesdropping the signaling traffic along the path. The identity used in the process of authentication may also be hidden to a limited extent from a network to which the initiator is attached. However the identity MUST provide enough information for the nodes in the access network to collect accounting data.

Network topology hiding MAY be supported to prevent entities along the path to learn the topology of a network. Supporting this property might conflict with a diagnostic capability.

5.7.7. Denial-of-Service Attacks

A signaling protocol SHOULD provide prevention of Denial-of-service attacks. To effectively prevent denial-of-service attacks it is necessary that the used security and protocol mechanisms MUST have low computational complexity to verify a state setup request prior to authenticating the requesting entity. Additionally the signaling protocol and the used security mechanisms SHOULD NOT require large resource consumption on NSIS Entities (for example main memory or other additional message exchanges) before a successful authentication is done.

5.7.8. Confidentiality of Signaling Messages

Based on the signaling information exchanged between nodes participating in the signaling protocol an adversary may learn both the identities and the content of the signaling messages. Since the ability to listen to signaling channels is a major guide to what data channels are interesting ones.

To prevent this from happening, confidentiality of the signaling message in a hop-by-hop manner SHOULD be provided. Note that most messages must be protected on a hop-by-hop basis, since entities, which actively participate in the signaling protocol, must be able to read and eventually modify the signaling messages.

5.7.9. Ownership of State

When existing states have to be modified then there is a need to use a session identifier to uniquely identify the established state. A signaling protocol MUST provide means of security protection to prevent adversaries from modifying state.

5.8. Mobility

5.8.1. Allow Efficient Service Re-Establishment After Handover

Handover is an essential function in wireless networks. After handover, the states may need to be completely or partially re-established due to route changes. The re-establishment may be requested by the mobile node itself or triggered by the access point that the mobile node is attached to. In the first case, the signaling MUST allow efficient re-establishment after handover. Re-establishment after handover MUST be as quick as possible so that the mobile node does not experience service interruption or service degradation. The re-establishment SHOULD be localized, and not require end-to-end signaling.

5.9. Interworking with Other Protocols and Techniques

Hooks SHOULD be provided to enable efficient interworking between various protocols and techniques including the following listed.

5.9.1. MUST Interwork with IP Tunneling

IP tunneling for various applications MUST be supported. More specifically IPSec tunnels are of importance. This mainly impacts the identification of flows. When using IPSec, parts of information commonly used for flow identification (e.g., transport protocol information and ports) may not be accessible due to encryption.

5.9.2. MUST NOT Constrain Either to IPv4 or IPv6

5.9.3. MUST be Independent from Charging Model

Signaling MUST NOT be constrained by charging models or the charging infrastructure used.

5.9.4. SHOULD Provide Hooks for AAA Protocols

The NSIS protocol SHOULD be developed with respect to be able to collect usage records from one or more network elements.

5.9.5. SHOULD Work with Seamless Handoff Protocols

An NSIS protocol SHOULD work with seamless handoff protocols such as context transfer and candidate access router (CAR) discovery.

5.9.6. MUST Work with Traditional Routing

NSIS assumes traditional L3 routing, which is purely based on L3 destination addresses. NSIS MUST work with L3 routing, in particular it MUST work in case of route changes. This means state on the old route MUST be released and state on the new route MUST be established by an NSIS protocol.

Networks, which do non-traditional routing, should not break NSIS signaling. NSIS MAY work for some of these situations. Particularly, combinations of NSIS unaware nodes and routing other than traditional one causes some problems. Non-traditional routing includes, for example, routing decisions based on port numbers, other IP header fields than the destination address, or splitting traffic based on header hash values. These routing environments result in the signaling path being potentially different than the data path.

5.10. Operational

5.10.1. Ability to Assign Transport Quality to Signaling Messages

The NSIS architecture SHOULD allow the network operator to assign the NSIS protocol messages a certain transport quality. As signaling opens up the possibility of denial-of-service attacks, this requirement gives the network operator a means, but also the obligation, to trade-off between signaling latency and the impact (from the signaling messages) on devices within the network. From protocol design this requirement states that the protocol messages SHOULD be detectable, at least where the control and assignment of the messages priority is done.

Furthermore, the protocol design must take into account reliability concerns. Communication reliability is seen as part of the quality assigned to signaling messages. So procedures MUST be defined for how an NSIS signaling system behaves if some kind of request it sent stays unanswered. The basic transport protocol to be used between adjacent NSIS Entities MAY ensure message integrity and reliable transport.

5.10.2. Graceful Fail Over

Any unit participating in NSIS signaling MUST NOT cause further damage to other systems involved in NSIS signaling when it has to go out of service.

5.10.3. Graceful Handling of NSIS Entity Problems

NSIS entities SHOULD be able to detect a malfunctioning peer. It may notify the NSIS Initiator or another NSIS entity involved in the signaling process. The NSIS peer may handle the problem itself e.g., switching to a backup NSIS entity. In the latter case note that synchronization of state between the primary and the backup entity is needed.

6. Security Considerations

Section 5.7 of this document provides security related requirements of a signaling protocol.

7. Acknowledgments

Quite a number of people have been involved in the discussion of the document, adding some ideas, requirements, etc. We list them without a guarantee on completeness: Changpeng Fan (Siemens), Krishna Paul (NEC), Maurizio Molina (NEC), Mirko Schramm (Siemens), Andreas Schrader (NEC), Hannes Hartenstein (NEC), Ralf Schmitz (NEC), Juergen Quittek (NEC), Morihisa Momona (NEC), Holger Karl (Technical University Berlin), Xiaoming Fu (Technical University Berlin), Hans-Peter Schwefel (Siemens), Mathias Rautenberg (Siemens), Christoph Niedermeier (Siemens), Andreas Kassler (University of Ulm), Ilya Freytsis.

Some text and/or ideas for text, requirements, scenarios have been taken from an Internet Draft written by the following authors: David Partain (Ericsson), Anders Bergsten (Telia Research), Marc Greis (Nokia), Georgios Karagiannis (Ericsson), Jukka Manner (University of Helsinki), Ping Pan (Juniper), Vloria Rexhepi (Ericsson), Lars Westberg (Ericsson), Haihong Zheng (Nokia). Some of those have actively contributed new text to this document as well.

Another Internet Draft impacting this document has been written by Sven Van den Bosch, Maarten Buchli, and Danny Goderis (all Alcatel). These people contributed also new text.

Thanks also to Kwok Ho Chan (Nortel) for text changes. And finally thanks Alison Mankin for the thorough AD review and thanks to Harald Tveit Alvestrand and Steve Bellovin for the IESG review comments.

8. Appendix: Scenarios/Use Cases

In the following we describe scenarios, which are important to cover, and which allow us to discuss various requirements. Some regard this as use cases to be covered defining the use of a signaling protocol. In general, these scenarios consider the specific case of signaling for QoS (resource reservation), although many of the issues carry over directly to other signaling types.

8.1. Terminal Mobility

The scenario we are looking at is the case where a mobile terminal (MT) changes from one access point to another access point. The access points are located in separate QoS domains. We assume Mobile IP to handle mobility on the network layer in this scenario and consider the various extensions (i.e., IETF proposals) to Mobile IP, in order to provide 'fast handover' for roaming Mobile Terminals. The goal to be achieved lies in providing, keeping, and adapting the requested QoS for the ongoing IP sessions in case of handover. Furthermore, the negotiation of QoS parameters with the new domain via the old connection might be needed, in order to support the different 'fast handover' proposals within the IETF.

The entities involved in this scenario include a mobile terminal, access points, an access network manager, and communication partners of the MT (the other end(s) of the communication association). From a technical point of view, terminal mobility means changing the access point of a mobile terminal (MT). However, technologies might change in various directions (access technology, QoS technology, administrative domain). If the access points are within one specific QoS technology (independent of access technology) we call this intra-QoS technology handoff. In the case of an inter-QoS technology handoff, one changes from e.g., a DiffServ to an IntServ domain, however still using the same access technology. Finally, if the access points are using different access technologies we call it inter-technology hand-off.

The following issues are of special importance in this scenario:

1) Handoff decision

- The QoS management requests handoff. The QoS management can decide to change the access point, since the traffic conditions of the new access point are better supporting the QoS requirements. The metric may be different (optimized towards a single or a group/class of users). Note that the MT or the network (see below) might trigger the handoff.

- The mobility management forces handoff. This can have several reasons. The operator optimizes his network, admission is no longer granted (e.g., emptied prepaid condition). Or another example is when the MT is reaching the focus of another base station. However, this might be detected via measurements of QoS on the physical layer and is therefore out of scope of QoS signaling in IP. Note again that the MT or the network (see below) might trigger the handoff.
- This scenario shows that local decisions might not be enough. The rest of the path to the other end of the communication needs to be considered as well. Hand-off decisions in a QoS domain do not only depend on the local resource availability, e.g., the wireless part, but involve the rest of the path as well. Additionally, decomposition of an end-to-end signaling might be needed, in order to change only parts of it.

2) Trigger sources

- Mobile terminal: If the end-system QoS management identifies another (better-suited) access point, it will request the handoff from the terminal itself. This will be especially likely in the case that two different provider networks are involved. Another important example is when the current access point bearer disappears (e.g., removing the Ethernet cable). In this case, the NSIS Initiator is basically located on the mobile terminal.
- Network (access network manager): Sometimes, the handoff trigger will be issued from the network management to optimize the overall load situation. Most likely this will result in changing the base-station of a single providers network. Most likely the NSIS Initiator is located on a system within the network.

3) Integration with other protocols

- Interworking with other protocol must be considered in one or the other form. E.g., it might be worth combining QoS signaling between different QoS domains with mobility signaling at hand-over.

4) Handover rates

In mobile networks, the admission control process has to cope with far more admission requests than call setups alone would generate. For example, in the GSM (Global System for Mobile communications) case, mobility usually generates an average of one to two handovers

per call. For third generation networks (such as UMTS), where it is necessary to keep radio links to several cells simultaneously (macro-diversity), the handover rate is significantly higher.

5) Fast state installation

Handover can also cause packet losses. This happens when the processing of an admission request causes a delayed handover to the new base station. In this situation, some packets might be discarded, and the overall speech quality might be degraded significantly. Moreover, a delay in handover may cause degradation for other users. In the worst-case scenario, a delay in handover may cause the connection to be dropped if the handover occurred due to bad air link quality. Therefore, it is critical that QoS signaling in connection with handover be carried out very quickly.

6) Call blocking in case of overload

Furthermore, when the network is overloaded, it is preferable to keep states for previously established flows while blocking new requests. Therefore, the resource reservation requests in connection with handover should be given higher priority than new requests for resource reservation.

8.2. Wireless Networks

In this scenario, the user is using the packet services of a wireless system (such as the 3rd generation wireless system 3GPP/UMTS, 3GPP2/cdma2000). The region between the End Host and the Edge Node (Edge Router) connecting the wireless network to another QoS domain is considered to be a single QoS domain.

The issues in such an environment regarding QoS include:

- 1) The wireless networks provide their own QoS technology with specialized parameters to coordinate the QoS provided by both the radio access and wired access networks. Provisioning of QoS technologies within a wireless network can be described mainly in terms of calling bearer classes, service options, and service instances. These QoS technologies need to be invoked with suitable parameters when higher layers trigger a request for QoS. Therefore these involve mapping of the requested higher layer QoS parameters onto specific bearer classes or service instances. The request for allocation of resources might be triggered by signaling at the IP level that passes across the wireless system, and possibly other QoS domains. Typically, wireless network specific messages are invoked to setup the underlying bearer

classes or service instances in parallel with the IP layer QoS negotiation, to allocate resources within the radio access network.

- 2) The IP signaling messages are initiated by the NSIS initiator and interpreted by the NSIS Forwarder. The most efficient placement of the NSIS Initiator and NSIS Forwarder has not been determined in wireless networks, but a few potential scenarios can be envisioned. The NSIS Initiator could be located at the End Host (e.g., 3G User equipment (UE)), the Access Gateway or at a node that is not directly on the data path, such as a Policy Decision Function. The Access Gateway could act as a proxy NSIS Initiator on behalf of the End Host. The Policy Decision Function that controls per-flow/aggregate resources with respect to the session within its QoS domain (e.g., the 3G wireless network) may act as a proxy NSIS Initiator for the end host or the Access Gateway. Depending on the placement of the NSIS Initiator, the NSIS Forwarder may be located at an appropriate point in the wireless network.
- 3) The need for re-negotiation of resources in a new wireless domain due to host mobility. In this case the NSIS Initiator and the NSIS Forwarder should detect mobility events and autonomously trigger re-negotiation of resources.

8.3. An Example Scenario for 3G Wireless Networks

The following example is a pure hypothetical scenario, where an NSIS signaling protocol might be used in a 3G environment. We do not impose in any way, how a potential integration might be done. Terms from the 3GPP architecture are used (P-CSCF, IMS, expanded below) in order to give specificity, but in a hypothetical design, one that reflects neither development nor review by 3GPP. The example should help in the design of a NSIS signaling protocol such that it could be used in various environments.

The 3G wireless access scenario is shown in Figure 1. The Proxy-Call State Control Function (P-CSCF) is the outbound SIP proxy (only used in IP Multimedia Subsystems (IMS)). The Access Gateway is the egress router of the 3G wireless domain and it connects the radio access network to the Edge Router (ER) of the backbone IP network. The Policy Decision Function (PDF) is an entity responsible for controlling bearer level resource allocations/de-allocations in relation to session level services e.g., SIP. The Policy Decision Function may also control the Access Gateway to open and close the gates and to configure per-flow policies, i.e., to authorize or forbid user traffic. The P-CSCF (only used in IMS) and the Access Gateway communicate with the Policy Decision Function, for network

resource allocation/de-allocation decisions. The User Equipment (UE) or the Mobile Station (MS) consists of a Mobile Terminal (MT) and Terminal Equipment (TE), e.g., a laptop.

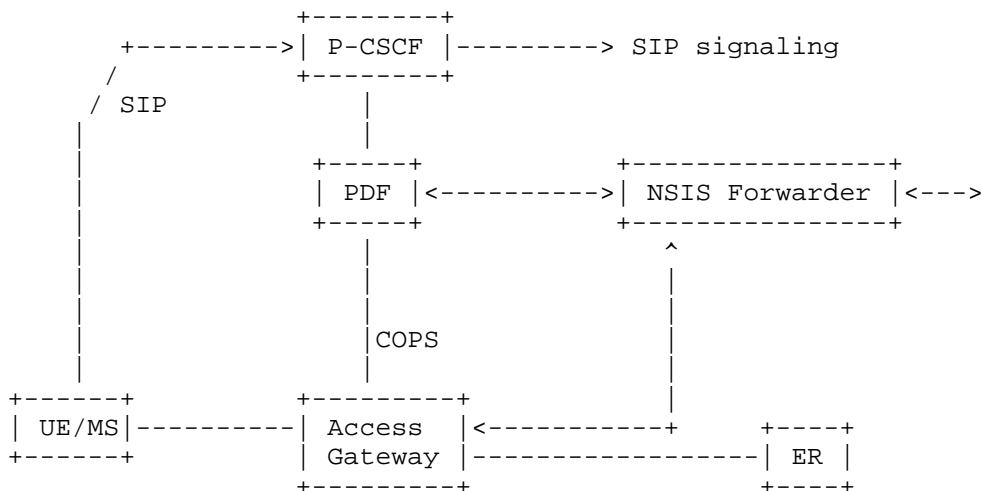


Figure 1: 3G wireless access scenario

The PDF has all the required QoS information for per-flow or aggregate admission control in 3G wireless networks. It receives resource allocation/de-allocation requests from the P-CSCF and/or Access Gateway etc. and responds with policy decisions. Hence the PDF may be a candidate entity to host the functionality of the NSIS Initiator, initiating the NSIS QoS signaling towards the backbone IP network. On the other hand, the UE/MS may act as the NSIS Initiator or the Access Gateway may act as a Proxy NSIS Initiator on behalf of the UE/MS. In the former case, the P-CSCF/PDF has to do the mapping from codec types and media descriptors (derived from SIP/SDP signaling) to IP traffic descriptor. In the latter case, the UE/MS may use any appropriate QoS signaling mechanism as the NSIS Initiator. If the Access Gateway is acting as the Proxy NSIS initiator on behalf of the UE/MS, then it may have to do the mapping of parameters from radio access specific QoS to IP QoS traffic parameters before forwarding the request to the NSIS Forwarder.

The NSIS Forwarder is currently not part of the standard 3G wireless architecture. However, to achieve end-to-end QoS a NSIS Forwarder is needed such that the NSIS Initiators can request a QoS connection to the IP network. As in the previous example, the NSIS Forwarder could manage a set of pre-provisioned resources in the IP network, i.e., bandwidth pipes, and the NSIS Forwarder perform per-flow admission control into these pipes. In this way, a connection can be made

between two 3G wireless access networks, and hence, end-to-end QoS can be achieved. In this case the NSIS Initiator and NSIS Forwarder are clearly two separate logical entities. The Access Gateway or/and the Edge Router in Fig.1 may contain the NSIS Forwarder functionality, depending upon the placement of the NSIS Initiator as discussed in scenario 2 in section 8.2. This use case clearly illustrates the need for an NSIS QoS signaling protocol between NSIS Initiator and NSIS Forwarder. An important application of such a protocol may be its use in the end-to-end establishment of a connection with specific QoS characteristics between a mobile host and another party (e.g., end host or content server).

8.4. Wired Part of Wireless Network

A wireless network, seen from a QoS domain perspective, usually consists of three parts: a wireless interface part (the "radio interface"), a wired part of the wireless network (i.e., Radio Access Network) and the backbone of the wireless network, as shown in Figure 2. Note that this figure should not be seen as an architectural overview of wireless networks but rather as showing the conceptual QoS domains in a wireless network.

In this scenario, a mobile host can roam and perform a handover procedure between base stations/access routers. In this scenario the NSIS QoS protocol can be applied between a base station and the gateway (GW). In this case a GW can also be considered as a local handover anchor point. Furthermore, in this scenario the NSIS QoS protocol can also be applied either between two GWs, or between two edge routers (ER).

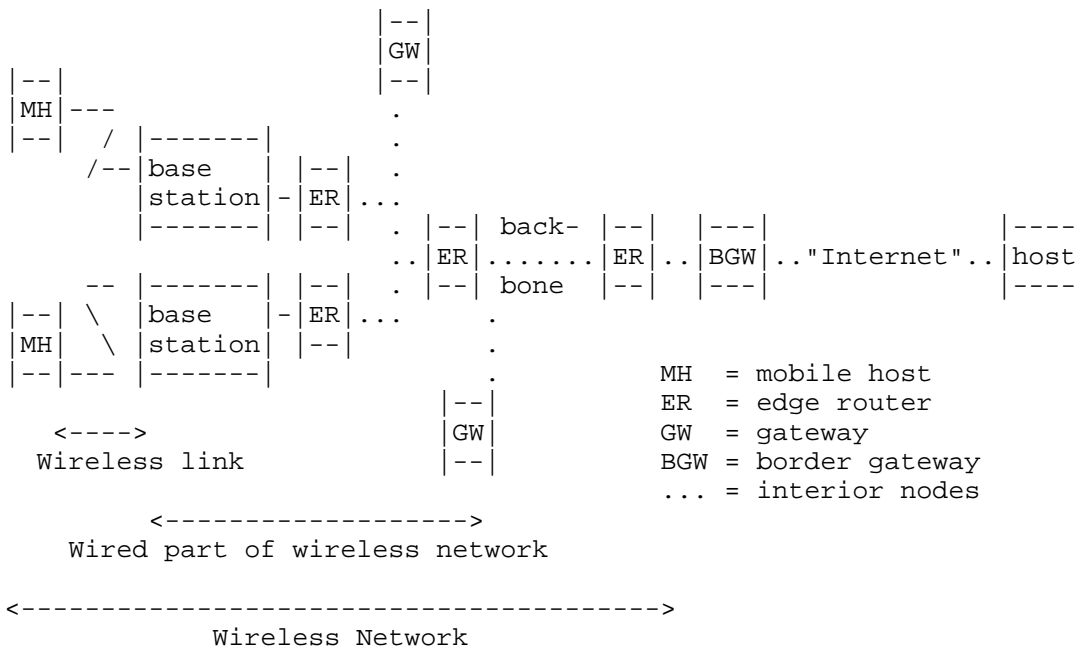


Figure 2. QoS architecture of wired part of wireless network

Each of these parts of the wireless network impose different issues to be solved on the QoS signaling solution being used:

- 1) Wireless interface: The solution for the air interface link has to ensure flexibility and spectrum efficient transmission of IP packets. However, this link layer QoS can be solved in the same way as any other last hop problem by allowing a host to request the proper QoS profile.
- 2) Wired part of the wireless network: This is the part of the network that is closest to the base stations/access routers. It is an IP network although some parts logically perform tunneling of the end user data. In cellular networks, the wired part of the wireless network is denoted as a radio access network.

This part of the wireless network has different requirements for signaling protocol characteristics when compared to traditional IP networks:

- The network must support mobility. Many wireless networks are able to provide a combination of soft and hard handover procedures. When handover occurs, reservations need to be established on new paths. The establishment time has to be as

short as possible since long establishment times for s degrade the performance of the wireless network. Moreover, for maximal utilization of the radio spectrum, frequent handover operations are required.

- These links are typically rather bandwidth-limited.
- The wired transmission in such a network contains a relatively high volume of expensive leased lines. Overprovisioning might therefore be prohibitively expensive.
- The radio base stations are spread over a wide geographical area and are in general situated a large distance from the backbone.

- 3) Backbone of the wireless network: the requirements imposed by this network are similar to the requirements imposed by other types of backbone networks.

Due to these very different characteristics and requirements, often contradictory, different QoS signaling solutions might be needed in each of the three network parts.

8.5. Session Mobility

In this scenario, a session is moved from one end-system to another. Ongoing sessions are kept and QoS parameters need to be adapted, since it is very likely that the new device provides different capabilities. Note that it is open which entity initiates the move, which implies that the NSIS Initiator might be triggered by different entities.

User mobility (i.e., a user changing the device and therefore moving the sessions to the new device) is considered to be a special case within the session mobility scenario.

Note that this scenario is different from terminal mobility. The terminal (end-system) has not moved to a different access point. Both terminals are still connected to an IP network at their original points.

The issues include:

- 1) Keeping the QoS guarantees negotiated implies that the end-point(s) of communication are changed without changing the s.
- 2) The trigger of the session move might be the user or any other party involved in the session.

8.6. QoS Reservation/Negotiation from Access to Core Network

The scenario includes the signaling between access networks and core networks in order to setup and change reservations together with potential negotiation.

The issues to be solved in this scenario are different from previous ones.

- 1) The entity of reservation is most likely an aggregate.
- 2) The time scales of states might be different (long living states of aggregates, less often re-negotiation).
- 3) The specification of the traffic (amount of traffic), a particular QoS is guaranteed for, needs to be changed. E.g., in case additional flows are added to the aggregate, the traffic specification of the flow needs to be added if it is not already included in the aggregates specification.
- 4) The flow specification is more complex including network addresses and sets of different address for the source as well as for the destination of the flow.

8.7. QoS Reservation/Negotiation Over Administrative Boundaries

Signaling between two or more core networks to provide QoS is handled in this scenario. This might also include access to core signaling over administrative boundaries. Compared to the previous one it adds the case, where the two networks are not in the same administrative domain. Basically, it is the inter-domain/inter-provider signaling which is handled in here.

The domain boundary is the critical issue to be resolved. Which of various flavors of issues a QoS signaling protocol has to be concerned with.

- 1) Competing administrations: Normally, only basic information should be exchanged, if the signaling is between competing administrations. Specifically information about core network internals (e.g., topology, technology, etc.) should not be exchanged. Some information exchange about the "access points" of the core networks (which is topology information as well) may be required, to be exchanged, because it is needed for proper signaling.
- 2) Additionally, as in scenario 4, signaling most likely is based on aggregates, with all the issues raise there.

- 3) Authorization: It is critical that the NSIS Initiator is authorized to perform a QoS path setup.
- 4) Accountability: It is important to notice that signaling might be used as an entity to charge money for, therefore the interoperation with accounting needs to be available.

8.8. QoS Signaling Between PSTN Gateways and Backbone Routers

A PSTN gateway (i.e., host) requires information from the network regarding its ability to transport voice traffic across the network. The voice quality will suffer due to packet loss, latency and jitter. Signaling is used to identify and admit a flow for which these impairments are minimized. In addition, the disposition of the signaling request is used to allow the PSTN GW to make a call routing decision before the call is actually accepted and delivered to the final destination.

PSTN gateways may handle thousands of calls simultaneously and there may be hundreds of PSTN gateways in a single provider network. These numbers are likely to increase as the size of the network increases. The point being that scalability is a major issue.

There are several ways that a PSTN gateway can acquire assurances that a network can carry its traffic across the network. These include:

1. Over-provisioning a high availability network.
2. Handling admission control through some policy server that has a global view of the network and its resources.
3. Per PSTN GW pair admission control.
4. Per call admission control (where a call is defined as the 5-tuple used to carry a single RTP flow).

Item 1 requires no signaling at all and is therefore outside the scope of this working group.

Item 2 is really a better informed version of 1, but it is also outside the scope of this working group as it relies on a particular telephony signaling protocol rather than a packet admission control protocol.

Item 3 is initially attractive, as it appears to have reasonable scaling properties, however, its scaling properties only are effective in cases where there are relatively few PSTN GWs. In the

more general case where a PSTN GW reduces to a single IP phone sitting behind some access network, the opportunities for aggregation are reduced and the problem reduces to item 4.

Item 4 is the most general case. However, it has the most difficult scaling problems. The objective here is to place the requirements on Item 4 such that a scalable per-flow admission control protocol or protocol suite may be developed.

The case where per-flow signaling extends to individual IP end-points allows the inclusion of IP phones on cable, DSL, wireless or other access networks in this scenario.

Call Scenario

A PSTN GW signals end-to-end for some 5-tuple defined flow a bandwidth and QoS requirement. Note that the 5-tuple might include masking/wildcarding. The access network admits this flow according to its local policy and the specific details of the access technology.

At the edge router (i.e., border node), the flow is admitted, again with an optional authentication process, possibly involving an external policy server. Note that the relationship between the PSTN GW and the policy server and the routers and the policy server is outside the scope of NSIS. The edge router then admits the flow into the core of the network, possibly using some aggregation technique.

At the interior nodes, the NSIS host-to-host signaling should either be ignored or invisible as the Edge router performed the admission control decision to some aggregate.

At the inter-provider router (i.e., border node), again the NSIS host-to-host signaling should either be ignored or invisible, as the Edge router has performed an admission control decision about an aggregate across a carrier network.

8.9. PSTN Trunking Gateway

One of the use cases for the NSIS signaling protocol is the scenario of interconnecting PSTN gateways with an IP network that supports QoS.

Four different scenarios are considered here.

1. In-band QoS signaling is used. In this case the Media Gateway (MG) will be acting as the NSIS Initiator and the Edge Router (ER) will be the NSIS Forwarder. Hence, the ER should do admission control (into pre-provisioned traffic trunks) for the individual traffic flows. This scenario is not further considered here.
2. Out-of-band signaling in a single domain, the NSIS forwarder is integrated in the Media Gateway Controller (MGC). In this case no NSIS protocol is required.
3. Out-of-band signaling in a single domain, the NSIS forwarder is a separate box. In this case NSIS signaling is used between the MGC and the NSIS Forwarder.
4. Out-of-band signaling between multiple domains, the NSIS Forwarder (which may be integrated in the MGC) triggers the NSIS Forwarder of the next domain.

When the out-of-band QoS signaling is used the Media Gateway Controller (MGC) will be acting as the NSIS Initiator.

In the second scenario the voice provider manages a set of traffic trunks that are leased from a network provider. The MGC does the admission control in this case. Since the NSIS Forwarder acts both as a NSIS Initiator and a NSIS Forwarder, no NSIS signaling is required. This scenario is shown in Figure 3.




Figure 3: PSTN trunking gateway scenario

In the third scenario, the voice provider does not lease traffic trunks in the network. Another entity may lease traffic trunks and may use a NSIS Forwarder to do per-flow admission control. In this case the NSIS signaling is used between the MGC and the NSIS Forwarder, which is a separate box here. Hence, the MGC acts only as a NSIS Initiator. This scenario is depicted in Figure 4.

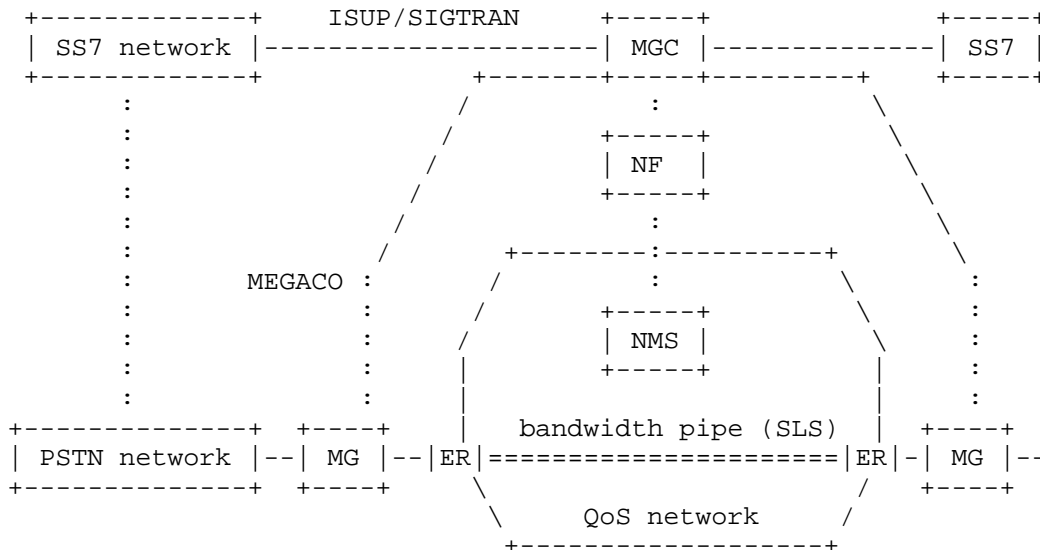


Figure 4: PSTN trunking gateway scenario

In the fourth scenario multiple transport domains are involved. In the originating network either the MGC may have an overview on the resources of the overlay network or a separate NSIS Forwarder will have the overview. Hence, depending on this either the MGC or the NSIS Forwarder of the originating domain will contact the NSIS Forwarder of the next domain. The MGC always acts as a NSIS Initiator and may also be acting as a NSIS Forwarder in the first domain.

8.10. An Application Requests End-to-End QoS Path from the Network

This is actually the conceptually simplest case. A multimedia application requests a guaranteed service from an IP network. We assume here that the application is somehow able to specify the network service. The characteristics here are that many hosts might do it, but that the requested service is low capacity (bounded by the access line). Note that there is an issue of scaling in the number of applications requesting this service in the core of the network.

8.11. QoS for Virtual Private Networks

In a Virtual Private Network (VPN), a variety of tunnels might be used between its edges. These tunnels could be for example, IPSec, GRE, and IP-IP. One of the most significant issues in VPNs is related to how a flow is identified and what quality a flow gets. A flow identification might consist among others of the transport protocol port numbers. In an IP-Sec tunnel this will be problematic since the transport protocol information is encrypted.

There are two types of L3 VPNs, distinguished by where the endpoints of the tunnels exist. The endpoints of the tunnels may either be on the customer (CPE) or the provider equipment or provider edge (PE).

Virtual Private networks are also likely to request bandwidth or other type of service in addition to the premium services the PSTN GW are likely to use.

8.11.1. Tunnel End Points at the Customer Premises

When the endpoints are the CPE, the CPE may want to signal across the public IP network for a particular amount of bandwidth and QoS for the tunnel aggregate. Such signaling may be useful when a customer wants to vary their network cost with demand, rather than paying a flat rate. Such signaling exists between the two CPE routers. Intermediate access and edge routers perform the same exact call admission control, authentication and aggregation functions performed by the corresponding routers in the PSTN GW scenario with the exception that the endpoints are the CPE tunnel endpoints rather than PSTN GWs and the 5-tuple used to describe the RTP flow is replaced with the corresponding flow spec to uniquely identify the tunnels. Tunnels may be of any variety (e.g., IP-Sec, GRE, IP-IP).

In such a scenario, NSIS would actually allow partly for customer managed VPNs, which means a customer can setup VPNs by subsequent NSIS signaling to various end-point. Plus the tunnel end-points are not necessarily bound to an application. The customer administrator might be the one triggering NSIS signaling.

8.11.2. Tunnel End Points at the Provider Premises

In the case were the tunnel end-points exist on the provider edge, requests for bandwidth may be signaled either per flow, where a flow is defined from a customers address space, or between customer sites.

In the case of per flow signaling, the PE router must map the bandwidth request to the tunnel carrying traffic to the destination specified in the flow spec. Such a tunnel is a member of an

aggregate to which the flow must be admitted. In this case, the operation of admission control is very similar to the case of the PSTN GW with the additional level of indirection imposed by the VPN tunnel. Therefore, authentication, accounting and policing may be required on the PE router.

In the case of per site signaling, a site would need to be identified. This may be accomplished by specifying the network serviced at that site through an IP prefix. In this case, the admission control function is performed on the aggregate to the PE router connected to the site in question.

9. References

9.1. Normative References

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[RSVP] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.

10. Authors' Addresses

Marcus Brunner (Editor)
NEC Europe Ltd.
Network Laboratories
Kurfuersten-Anlage 36
D-69115 Heidelberg
Germany

E-Mail: brunner@netlab.nec.de

Robert Hancock
Roke Manor Research Ltd
Romsey, Hants, SO51 0ZN
United Kingdom

E-Mail: robert.hancock@roke.co.uk

Eleanor Hepworth
Roke Manor Research Ltd
Romsey, Hants, SO51 0ZN
United Kingdom

E-Mail: eleanor.hepworth@roke.co.uk

Cornelia Kappler
Siemens AG
Berlin 13623
Germany

E-Mail: cornelia.kappler@siemens.com

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

E-Mail: Hannes.Tschofenig@mchp.siemens.de

11. Full Copyright Statement

Copyright (C) The Internet Society (2004). 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.

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.

Intellectual Property

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.

Acknowledgement

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