

Requirements for Signaling Protocols
<[draft-ietf-nsis-req-06.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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>.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This document defines requirements for signaling across different network environments, where different network environments mean across administrative and technology domains. Signaling is mainly though for QoS such as [\[RSVP\]](#), however in recent year several other applications of signaling have been defined such as signaling for MPLS label distribution [\[RSVP-TE\]](#). 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 memo present the assumptions and the aspects not considered within scope before listing the requirements grouped according to areas such as architecture and design goals, signaling flows, layering, performance, flexibility, security, and mobility.

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 [RFC 2119](#).

Table of Contents

Brunner (Editor)

Informational

[Page 1]

Status of this Memo.....	1
Abstract.....	1
Table of Contents.....	1
1 Introduction.....	2
2 Terminology.....	3
3 Problem Statement and Scope.....	5
4 Assumptions and Exclusions.....	6
4.1 Assumptions and Non-Assumptions.....	6
4.2 Exclusions.....	7
5 Requirements.....	9
5.1 Architecture and Design Goals.....	9
5.2 Signaling Flows.....	11
5.3 Messaging.....	13
5.4 Control Information.....	14
5.5 Performance.....	16
5.6 Flexibility.....	17
5.7 Security.....	18
5.8 Mobility.....	20
5.9 Interworking with other protocols and techniques.....	20
5.10 Operational.....	21
6 Security Considerations.....	22
7 Informative References.....	22
8 Acknowledgments.....	22
9 Author's Addresses.....	22
10 Appendix: Scenarios/Use cases.....	23
10.1 Terminal Mobility.....	23
10.2 Cellular Networks.....	25
10.3 UMTS access.....	26
10.4 Wired part of wireless network.....	27
10.5 Session Mobility.....	29
10.6 QoS reservations/negotiation from access to core network.....	30
10.7 QoS reservation/negotiation over administrative boundaries....	30
10.8 QoS signaling between PSTN gateways and backbone routers.....	31
10.9 PSTN trunking gateway.....	32
10.10 Application request end-to-end QoS path from the network....	34

[1](#) Introduction

This document 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 define a conceptual model of signaling on an abstract level. Additionally, we describe the entities involved in signaling and typical signaling paths. In the appendix are a list of use cases and scenarios where an NSIS protocol could be applied. It is though of helping derive the requirements and to test the requirements

against use cases.

QoS and signaling for QoS is a pretty large field with a lot of interaction with other protocols, mechanisms, applications etc. However, it is not the only field where signaling is used in the

Brunner (Editor)

Informational

[Page 2]

Internet. Even if this requirement documents mainly used QoS as the sample application other application must be possible and are also addressed.

There are several areas related to networking aspects which are incomplete, for example, interaction with host and site multi-homing, use of anycast services, and so on. These issues should be considered in any future analysis work.

2 Terminology

We try to list the most often used terms in the document. However, don't be too religious about it, they are not meant to prescribe any solution in the document. All of them need refined definitions in follow-up documents.

Resource Management Function (RMF): An abstract concept, representing the management of resources in a domain or a node. This includes admission control and resource allocation.

NSIS Domain (ND): Administrative domain where an NSIS protocol signals for a resource or set of resources.

NSIS Entity (NE): The function within a node, which implements an NSIS protocol.

NSIS Forwarder (NF): NSIS Entity on the path between a NI and NR, which may interact with local resource management function (RMF) for this purpose. NSIS Forwarder also propagates NSIS signaling further through the network. It is responsible for interpreting the signaling carrying the user parameters, optionally inserting or modifying the parameters according to domain network management policy.

NSIS Initiator (NI): NSIS Entity that initiates NSIS signaling for a network resource based on user or application requirements. This can be located in the end system, but may reside elsewhere in network.

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).

Higher Layers: The higher layer (transport protocol and application) functions that request QoS or any other service from the network

layer. The trigger for the request might be generated within the end system, or the trigger might be provided by some entity within the network (e.g. application proxy or policy server).

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 Segment: The segment of a path within a single domain/subdomain.

Control Information: the information that governs for instance the QoS treatment to be applied to a flow or aggregate, including the service class, flow administration, and any associated security or accounting information.

Provisioning: the act of actually allocating resources to a flow or aggregate of flows, may include mechanisms such as LSP initiation for MPLS, packet scheduler configuration within a router, and so on. The mechanisms depend on the overall technology and application being used within the domain, and can be from static to dynamic.

Subdomain: a network within an administrative domain using a uniform technology, e.g., a single QoS provisioning function to provision resources.

QoS Technology: a generic term for a set of protocols, standards and mechanisms that can be used within a QoS domain/subdomain to manage the QoS provided to flows or aggregates that traverse the domain. Examples might include MPLS, DiffServ, and so on. A QoS technology is associated with certain QoS provisioning techniques.

Resource: something of value in a network infrastructure to which rules or policy criteria are first applied before access is granted. Examples of resources include the buffers in a router and bandwidth on an interface.

Service: something provided by an 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.

Sender-initiated signaling protocol: A sender-initiated signaling protocol is a protocol where the NSIS Initiator initiates the signaling on behalf of the sender of the data. What this means is that resource management functions are processed from the data sender towards the data receiver. However, the triggering instance is not specified.

Receiver-initiated signaling protocol: A receiver-initiated protocol (e.g., [\[RSVP\]](#)) is a protocol, where the NSIS Responder on behalf of

the receiver of the user data initiates the reservation. What this means is that resource management functions are processed from the data receiver back towards the data sender. However, the triggering instance is not specified.

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.

A basic goal should be to re-use these wherever possible, and to focus requirements work at an early stage on those areas where a new solution is needed (e.g. an especially simple one). We also try to avoid defining requirements related to internal implementation aspects.

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 and 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 resources, 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 resources requested by them, and also provides feedback information to the higher layers, which might be used by transport layer rate management or adaptive applications.

2. Something that assists in managing resources 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 Initiator is that it responds to request at the end of a signaling path.

4. The path segment 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 receiving indications about success or failure in response.

Now concentrating more on the overall end to end (multiple domain) aspects, in particular:

1. The NSIS Initiator need not be located at an end system, and the NSIS Forwarders are not assumed to be located on the flow's data path. However, they must be able to identify the ingress and egress points for the flow's data path as it traverses the NSIS domain. Any signaling protocol must be able to find the appropriate NSIS Forwarder.
2. We 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 apply to the signaling between them.
3. Any domain may contain Resource Management Function (e.g. traffic engineering, admission control, policy and so on). These are assumed to interact with the NSIS Initiator, Responder, and Forwarders using standard mechanisms.
4. The placement of the NSIS Initiators and NSIS Forwarders is not fixed.

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 the 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. 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, provisioning can be carried out in a local implementation-dependent way without further signaling, whereas in the case of remote NSIS Forwarders, a provisioning protocol might be needed to control the routers along the path. This provisioning 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 an

application-layer issue and IETF protocols such as SIP etc. can be used.

3. Where the signaling does cover several NSIS domains or subdomains, we do not exclude that different signaling protocols are used in each path segment. 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 protocols or various flavors of the same protocol. This implies the need for the translation of information into domain specific format as well.

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 it. There might be standardized services, the definition can be negotiated together with a contract, the service definition is published at a Web page, etc.

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, etc.) The discovery of the NSIS entities has security implications that need to be addressed properly. These implications largely depend on the chosen protocol. 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 to operate with networks using standard ("normal") L3 routing. Where "normal" is not specified more exactly on purpose.

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 for provisioning within a domain/subdomain

are not considered. However, NSIS can be used for signaling within a domain/subdomain performing provisioning. 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 guaranties. 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 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. Protection of non-signaling messages is outside the scope of the protocol

The protection of non-signaling messages (including data traffic following a reservation) is not directly considered by a signaling protocol. The protection of data messages transmitted along the provisioned path is outside the scope of a signaling protocol. Regarding data traffic there is an interaction with accounting (metering) and edge routers might require packets to be integrity protected to be able to securely assign incoming data traffic to a particular user.

Additionally there might be an interaction with IPSec protected traffic experiencing service-specific treatment and the established state created due to signaling. One example of such an interaction is the different flow identification with and without IPSec protection.

Many security properties are likely to be application specific and may be provided by the corresponding application layer protocol.

8. 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.

9. Similarly, specific methods, protocols, and ways to express service information in the Application/Session level are not considered (e.g., SDP, SIP, RTSP, etc.).

10. The specification of any extensions needed to signal information via application level protocols (e.g. SDP), and the mapping on 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).

11. 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 situation 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.

12. 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 framework, 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 lightweight in implementation complexity and resource requirements in nodes. One approach to this is that the solution could deal with certain requirements via modular components or capabilities, which are optional to implement in individual nodes.

In order to prioritize the various requirements we define different 'parts of the network'. In the different parts of the network a particular requirement might have a different priority.

The parts of the networks we differentiate are the host-to-first router, the access network, and the core network. The host to first router part includes all the layer 2 technologies to access to the Internet. In many cases, there is an application and/or user running on the host initiating signaling. The access network can be characterized by low capacity links, medium speed IP processing capabilities, and it might consist of a complete layer 2 network as well. The core network characteristics include high-speed forwarding capacities and inter-domain issues. These divisions between network types are not strict and do not appear in all networks, but where they do exist they may influence signaling requirements and will be highlighted as necessary.

5.1 Architecture and Design Goals

Brunner (Editor)

Informational

[Page 9]

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 MUST be applicable for different technologies.

The signaling protocol MUST work with various QoS and non-QoS technologies. The basic information exchanged over the signaling protocol MUST be in such detail and quantity that it is useful for various technologies.

5.1.2 SHOULD provide resource availability information on request

NSIS SHOULD provide a mechanism to check whether resources are available without performing a 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.

5.1.3 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 redundant information MUST be supported if necessary.
- Uni- and bi-directional reservations are 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.

5.1.4 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 gained flexibility in the information transported 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.

5.1.5 NSIS MUST reuse existing provisioning

Brunner (Editor)

Informational

[Page 10]

Reuse existing functions and protocols for provisioning within a domain/subdomain unchanged. (Motivation: 'Don't re-invent the wheel'.)

5.1.6 NSIS MUST support independence of signaling and provisioning paradigm

The signaling MUST be independent of the paradigm and mechanism of provisioning. 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.

5.1.7 NSIS MUST be application independent

The signaling protocol MUST be independent of the application. The control information SHOULD be application independent, because we look into network level signaling.

The requirement relates to the way the signaling interacts with upper layer functions (users, applications, and QoS administration), and lower layer technologies.

Opaque application information MAY get transported in the signaling message, without being handled in the network. Development and deployment of new applications SHOULD be possible without impacting the network infrastructure.

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, middle boxes, management systems), in which direction.

5.2.1 The placement of NSIS Initiator, Forwarder, Responder MUST be free

The protocol MUST work in various scenarios such as host-to-network-to-host, edge-to-edge, (e.g., just within one providers 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 No constraint MUST be posed the signaling and NSIS Forwarders to be in the data path.

There is a set of scenarios, where signaling is not on the data path. The NSIS Forwarder being in the data path is one extreme case

and useful in many cases. Therefore the case of having NSIS entities on the data path only **MUST** be supported.

There are going to be cases where a centralized entity will take a decision about service requests. In this case, there is no need to have the data follow the signaling path, or have the signaling follow the data path.

There are going to be cases without a centralized entity managing resources and the signaling will be used as a tool for resource management. For various reasons (such as efficient use of expensive bandwidth), one will want to have fine-grained, fast, and very dynamic control of the resources in the network.

There are going to be cases where there will be neither signaling nor a centralized entity (over-provisioning). Nothing has to be done anyway.

One can capture the requirement with the following, different wording: If one views the domain as a virtual router then NSIS signaling used between those virtual routers **MUST** follow the same path as the data.

Routing the signaling protocol along an independent path is desired by network operators/designers. Ideally, the capability to route the protocol along an independent path would give the network designer/operator the option to manage bandwidth utilization through the topology.

There are other possibilities as well. An NSIS protocol **MUST** accept all of these possibilities with a strong focus on the on-path signaling.

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 traverse 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 flow-related signaling 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 release of resources MUST be possible

When a resource reservation is no longer necessary, e.g. because the application terminates, or because a mobile host experienced a hand-off, it MUST be possible to explicitly release resources. In general explicit release enhances the overall network utilization.

5.3.2 Automatic release of resources after failure SHOULD be possible

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

After detection of a failure in the network, any NSIS Forwarder/Initiator MUST be able to release a reservation 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 of reservations.

The goal is to prevent stale state within the network and adds 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.

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/severe congestion: In case the service cannot be provided completely but partially only.

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. E.g., if it gets an error notification it might teardown the reservation, 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 Feedback about success of service request MUST be provided

A request for service MUST be answered at least with yes or no. However, it may be useful in case of a negative answer to also get a description of what amount of resources a request is possible. So an opaque element MAY be included into the answer. The element heavily depends on the service requested.

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, severe congestion notification, notification of successful or erroneous processing of signaling messages.

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

Brunner (Editor)

Informational

[Page 14]

It SHOULD be possible for the NSIS initiator to control the mutability of the signaled information. This prevents from being changed in a non-recoverable way. 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 domains. This implies that in case of changes made on the parameters, the original requested ones must still be available.

Note that we do not require anything about particular parameters being changed.

Additionally, note that a provider or that particular services requested, can still influence the provisioning but in the signaling message the request should stay the same.

5.4.2 SHOULD possible to add and remove local domain information

It SHOULD be possible for the Resource Management Function 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. E.g., at the entrance to a domain domain-specific 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-use 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.

5.4.4 Modification of already reserved resources SHOULD be seamless

In many case, the reservation needs to be updated (up or downgrade). 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-flow 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 **MUST 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 and each reservation **MUST** be handled separately from each other.

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 a must anyway. However, depending on the scenario the question to which extends the protocol must be scalable.

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 overweight 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 a reservation. This applies for end-systems setting up several reservations. Some servers might be expected to setup a large number of reservations.

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

And Scalability in CPU use **MUST** be achieved on end terminals in case of many reservations at the same time and intermediate nodes mainly in the core.

5.5.2 NSIS SHOULD allow for low latency in setup

NSIS **SHOULD** allow for low latency setup of reservations. This is only needed in scenarios, where reservations are in 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 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, and this are only examples, include message aggregation, by ignoring signaling message, header compression, or minimizing functionality. The framework may choose any method as long as the requirement is met.

5.5.5 NSIS SHOULD target highest possible network utilization

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 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

NSIS MUST be flexible in placing an NSIS Initiator. The NSIS Initiator might be the sender or the receiver of content. But also network-initiated reservations MUST be allowed in various scenarios such as PSTN gateways, some VPNs, and mobility.

5.6.3 Flexibility in the initiation of re-negotiation

Brunner (Editor)

Informational

[Page 17]

The NSIS Initiator or the NSIS Responder SHOULD be able to initiate a re-negotiation or change the reservation due to various reasons, such as local resource shortage (CPU, memory on end-system) or a user changed application preference/profiles.

5.6.4 SOULD support network-initiated re-negotiation

NSIS SHOULD support network-initiated re-negotiation. This is used in cases, where the network is not able to further guarantee resources and want to e.g. downgrade a reservation.

5.6.5 Uni / bi-directional reservation

Both unidirectional as well as bi-direction reservations SHOULD be possible. With bi-directional reservations we mean here reservations having the same end-points. But the path in the two directions does not need to be the same.

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

5.7 Security

This section discusses security-related requirements. For a discussion of security threats see [SEC-THR]. The NSIS protocol MUST provide means for security, but it MUST be allowed that nodes implementing NSIS signaling do not need 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 Resource Authorization

The signaling protocol MUST provide means to authorize resource 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 resource request. Authorization prevents reservations by unauthorized entities, reservations 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 reservations. 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

Brunner (Editor)

Informational

[Page 18]

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

Hop-by-Hop security **SHOULD** be supported. It is a well known and proven concept in Quality-of-Service and other signaling protocols that allows intermediate nodes that actively participate in the protocol to modify the messages as it is required by processing rule. 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 initiator and the 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 location privacy

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.

Location privacy **MAY** be supported. It is an issue for the initiator who triggers the signaling protocol. In some scenarios the initiator may not be willing to reveal location information to the responder as part of the signaling procedure.

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 computation complexity to verify a resource request prior to authenticating the requesting entity. Additionally the signaling protocol and the used security mechanisms **SHOULD NOT** require large

resource consumption (for example main memory or other additional message exchanges) before a successful authentication was 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. To prevent this from happening, confidentiality of the signaling message in a hop-by-hop manner MAY be provided. Note that the protection can be provided on a hop-by-hop basis for most message payloads since it is required that entities which actively participating in the signaling protocol must be able to read and eventually modify the content of the signaling messages.

5.7.9 Ownership of a reservation

When existing reservations have to be modified then there is a need to use a reservation identifier to uniquely identify the established state. A signaling protocol MUST provide the appropriate security protection to prevent other entities to modify state without having the proper ownership.

5.8 Mobility

5.8.1 Allow efficient service re-establishment after handover

Handover is an essential function in wireless networks. After handover, the reservation 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:

5.9.1 MUST interwork with IP tunneling

IP tunneling for various applications MUST be supported. More specifically tunneling for IPsec tunnels are of importance as discussed in [Section 4.2](#). This mainly impacts the identification of flows. Using IPsec parts of information used for flow identification (e.g. transport protocol information and ports) may not be accessible due to encryption.

5.9.2 The solution MUST NOT constrain either to IPv4 or IPv6

5.9.3 MUST be independent from charging model

Brunner (Editor)

Informational

[Page 20]

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

5.9.4 SHOULD provide hooks for AAA protocols

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

5.9.5 SHOULD interwork with seamless handoff protocols

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

5.9.6 MAY interwork with non-traditional routing

NSIS assumes traditional routing, but networks, which do non-traditional L3 routing, should not break it.

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 for possible denial-of-service attacks, this requirement gives the network operator a mean, but also the obligation, to trade-off between signaling latency and the impact (from the signaling messages) on devices within his/her 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 how an NSIS signaling system behaves if some kind of request it sent stays without answer. The basic transport protocol to be used between adjacent NSIS units **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 peers **SHOULD** be able to detect the 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 Informative References

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

[RSVP-TE] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.

8 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), Vlori 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.

9 Author's Addresses

Marcus Brunner (Editor)
NEC Europe Ltd.
Network Laboratories

Kurfuersten-Anlage 34
D-69115 Heidelberg
Germany
E-Mail: brunner@ccrle.nec.de

Brunner (Editor)

Informational

[Page 22]

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@icn.siemens.de

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munchen
Germany
Email: Hannes.Tschofenig@mchp.siemens.de

10 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.

10.1 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 change 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 reservation 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

Brunner (Editor)

Informational

[Page 24]

- 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 reservations

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 reservations 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.

10.2 Cellular Networks

In this scenario, the user is using the packet service of a 3rd generation cellular system, e.g. UMTS. The region between the End Host and the edge node connecting the cellular network to another QoS domain (e.g. the GGSN in UMTS or the PDSN in 3GPP2) is considered to be a single QoS domain.

The issues in such an environment regarding QoS include:

1) Cellular systems provide their own QoS technology with specialized parameters to co-ordinate the QoS provided by both the radio access and wired access network. For example, in a UMTS network, one aspect of GPRS is that it can be considered as a QoS technology; provisioning of QoS within GPRS is described mainly in terms of calling UMTS bearer classes. This QoS technology needs to

be invoked with suitable parameters when higher layers trigger a request for QoS, and this therefore involves mapping the requested IP QoS onto these UMTS bearer classes. This request for resources might be triggered by IP signaling messages that pass across the

cellular system, and possibly other QoS domains, to negotiate for network resources. Typically, cellular system specific messages invoke the underlying cellular system QoS technology in parallel with the IP QoS negotiation, to allocate the resources within the cellular system.

2) The placement of NSIS Initiators and NSIS Forwarders (terminology in the framework given here). The NSIS Initiator could be located at the End Host (triggered by applications), the GGSN/PDSN, or at a node not directly on the data path, such as a bandwidth broker. In the second case, the GGSN/PDSN could either be acting as a proxy on behalf of an End Host with little capabilities, and/or managing aggregate resources within its QoS domain (the UMTS core network). The IP signaling messages are interpreted by the NSIS Forwarders, which may be located at the GGSN/PDSN, and in any QoS sub-domains within the cellular system.


3) Initiation of IP-level QoS negotiation. IP-level QoS re-negotiation may be initiated by either the End Host, or by the network, based on current network loads, which might change depending on the location of the end host.

4) The networks are designed and mainly used for speech communication (at least so far).

Note that in comparison to the previous scenario emphasis is much less on the mobility aspects, because mobility is mainly handled on the lower layer.

10.3 UMTS access

The UMTS access scenario is shown in Figure 1. The Proxy-Call State Control Function/Policy Control Function (P-CSCF/PCF) is the outbound SIP proxy of the visited domain, i.e. the domain where the mobile user wants to set-up a call. The Gateway GPRS Support Node (GGSN) is the egress router of the UMTS domain and connects the UMTS access network to the Edge Router (ER) of the core IP network. The P-CSCF/PCF communicates with the GGSN via the COPS protocol. The User Equipment (UE) consists of a Mobile Terminal (MT) and Terminal Equipment (TE), e.g. a laptop.



[Page 26]

+----+ +-----+ +----+

Figure 1: UMTS access scenario

In this scenario the GGSN has the role of Access Gate. According to 3GPP standardization, the PCF is responsible for the policy-based control of the end-user service in the UMTS access network (i.e. from UE to GGSN). In the current UMTS release R.5, the PCF is part of the P-CSCF, but in UMTS R.6 the interface between P-CSCF and PCF may evolve to an open standardized interface. In any case the PCF has all required QoS information for per-flow admission control in the UMTS access network (which it gets from the P-CSCF and/or GGSN). Thus the PCF would be the appropriate entity to host the functionality of NSIS Initiator, initiating the "NSIS" QoS signaling towards the core IP network. The PCF/P-CSCF has to do the mapping from codec type (derived from SIP/SDP signaling) to IP traffic descriptor. SDP extensions to explicitly signal QoS information are useful to avoid the need to store codec information in the PCF and to allow for more flexibility and accurate description of the QoS traffic parameters. The PCF also controls the GGSN to open and close the gates and to configure per-flow policers, i.e. to authorize or forbid user traffic.

The NSIS Forwarder is (of course) not part of the standard UMTS architecture. However, to achieve end-to-end QoS a NSIS Forwarder is needed such that the PCF 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 performs per-flow admission control into these pipes. In this way, a connection can be made between two UMTS access networks, and hence, end-to-end QoS can be achieved. In this case the NSIS Initiator and NSIS Forwarder are clearly two separate entities.

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 inter-connection of UMTS networks over an IP backbone.

10.4 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).

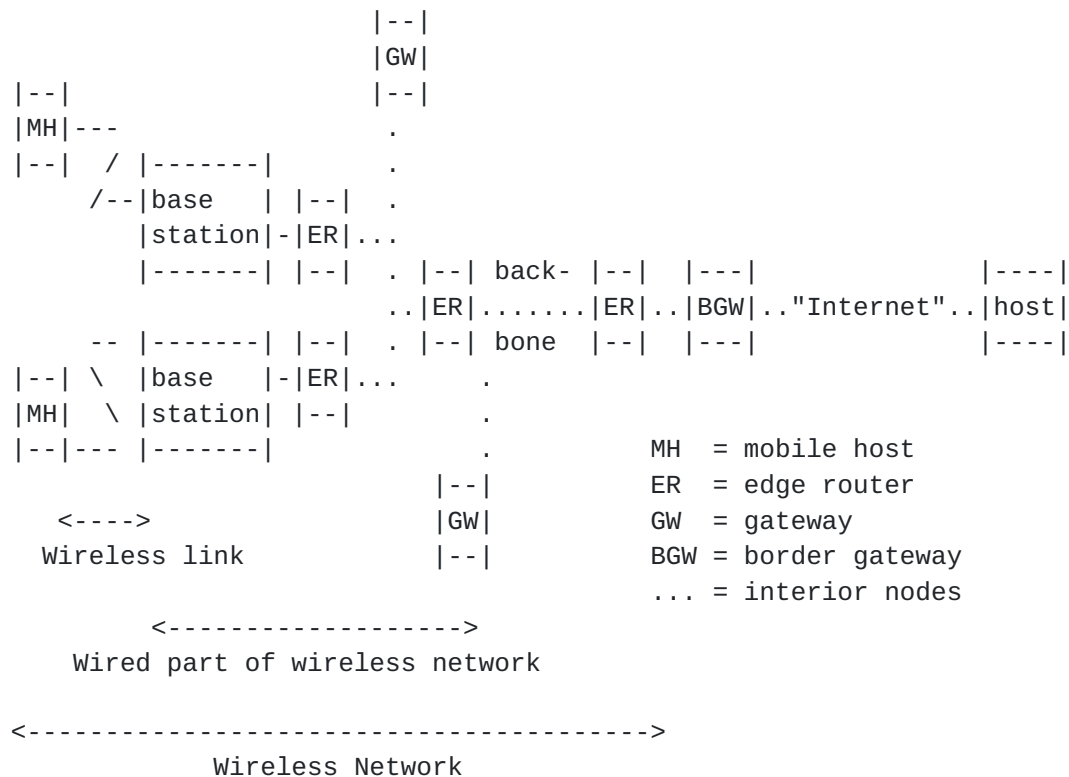


Figure 2. QoS architecture of wired part of wireless network

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

- 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.
- 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 characteristics when compared to traditional IP networks:

1. The network supports a high proportion of real-time traffic. The majority of the traffic transported in the

wired part of the wireless network is speech, which is very sensitive to delays and delay variation (jitter).

2. 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 reservations degrade the performance of the wireless network. Moreover, for maximal utilization of the radio spectrum, frequent handover operations are required.

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

10.5 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. Not the terminal (end-system) has 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

reservations.

2) The trigger of the session move might be the user or any other party involved in the session.

10.6 QoS reservations/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 reservations might be different (long living reservations 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.

10.7 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 as 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 need 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.

10.8 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 endpoints 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.

10.9 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 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.


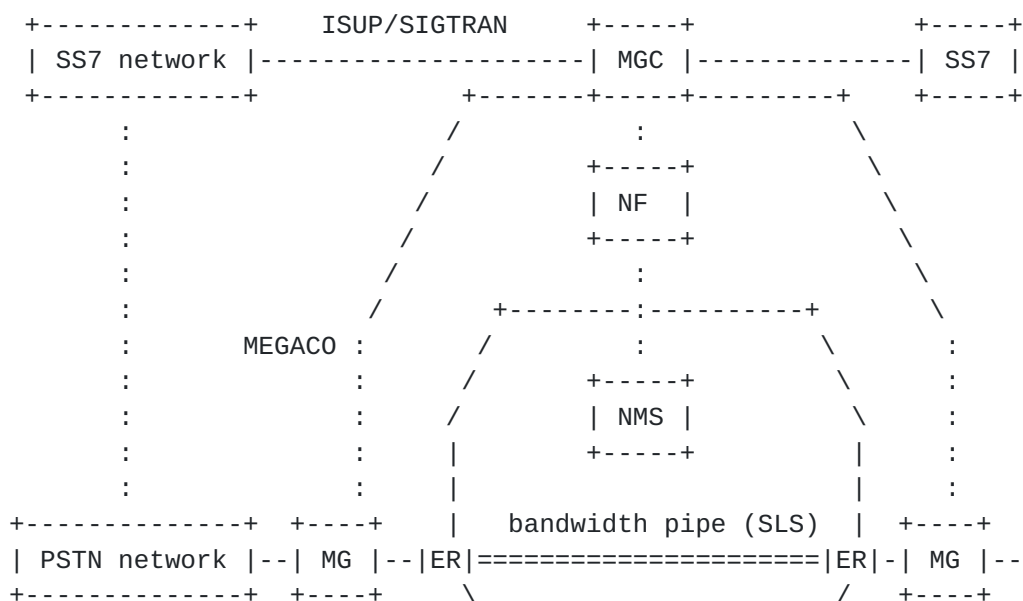


Figure 3: PSTN trunking gateway scenario

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.



\ QoS network /
+-----+

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.

10.10 Application request end-to-end QoS path from the network

This is actually the easiest case, nevertheless might be most often used in terms of number of users. So 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). Additionally, we assume no mobility and standard devices.

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, IP-Sec, 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.

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.

Tunnel end points at the provider premises

In the case where 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 customer's 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.

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.