Internet Draft

Robert Hancock
Eleanor Hepworth
Andrew McDonald
Siemens/Roke Manor Research

Document: draft-hancock-nsis-sender-receiver-00.txt
Expires: April 2003 October 2002

Sender and Receiver Orientation Issues in NSIS
draft-hancock-nsis-sender-receiver-00.txt

Status of this Memo

Abstract

The NSIS working group is considering protocols for signaling for resources for a traffic flow along its path in the network. The requirements for such signaling are being developed in [2] and a framework in [3].

It is clear from existing work that there are many interrelated issues with NSIS signaling, concerning the respective roles of the two ends of the communication path. These issues include route finding, authorisation, state management requirements, localization of negotiation, and so on. The wide variety of problems involved hinders progress in deciding what approach NSIS should adopt. This Internet Draft attempts to provide a summary of these issues and suggests a way of structuring further analysis. It is not expected that this document should have a long term existence.

Conventions used in this document

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 [4].

Table of Contents

1. Introduction, Terminology, and Scope

Unless otherwise stated, this document follows the terminology given in the current NSIS framework [3].

1.1 Data and Signaling Flows

For the bulk of this document, we are concerned with path-coupled signaling for a single unidirectional flow, as shown in Figure 1 (additional considerations are given in section 4). The node that is sending the user data packets is called the 'sender' and the node

sinking them the 'receiver'; these packets pass through one or more routers.

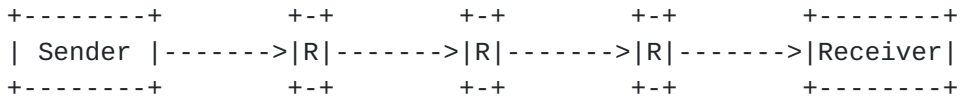

Figure 1: Sender and Receiver

In the case of path-coupled NSIS signaling, there are signaling nodes (NSIS entities) along the data path. The NSIS initiator (NI) notionally controls the signaling (e.g. at application request), whereas the NSIS responder (NR) terminates the signaling at the far end; there may be one of more NSIS forwarders (NF) between the two.

The NI and NR do not have to be colocated with sender and receiver (e.g. they could be at first/last hop access routers); nor do they have to be the same 'way round' as the sender and receiver. This leads to two different cases for analysis. Figure 2 shows the 'sender initiated' case, and Figure 3 shows the 'receiver initiated' case.

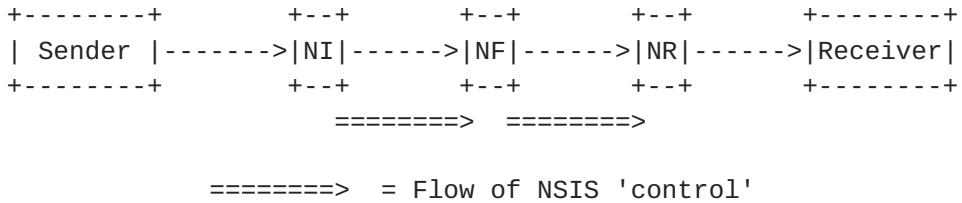

Figure 2: Sender Initiation

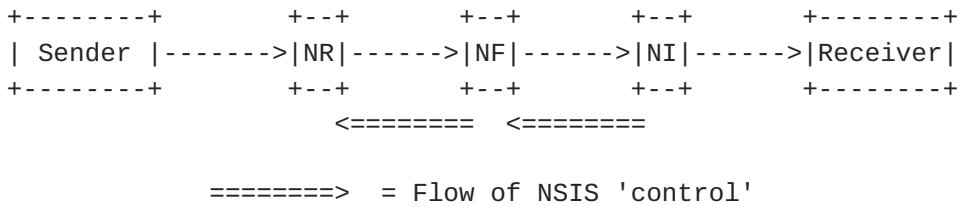

Figure 3: Receiver Initiation

One of the basic open issues in NSIS is whether one or both of these models should be supported, and in either case, what is the real difference in functionality between the NI and NR; to put it another

way, how 'directional' is the relationship between NSIS entities (which up to now has not really been defined).

It is the purpose of this document to gather together some of the information about this subject and propose a way forward.

1.2 Status of Existing Protocols

The principle existing path-coupled signaling protocol is RSVP [5]. RSVP is commonly described as 'receiver initiated', although there are some subtleties in this categorization.

From the point of view of the act of resource reservation, RSVP is clearly receiver initiated, in that the receiver is responsible for generating the RESV message which actually defines the QoS that the receiver wants for incoming traffic. This RESV message can also be accompanied with security-related policy information to support the request (see [6]). The primary motivation behind adopting receiver initiation for resource reservation appears to have been multicast support, as described in [7].

On the other hand, key elements of RSVP operation (RESV routing and route change detection) depend on the PATH message which is generated by the sender, and can be seen as triggering the RESV message (at least the first one). This sender-generated message also contains QoS-related information (and can even contain policy elements).

We can therefore see RSVP as containing several functions, some sender oriented and some receiver oriented. It might be that this distinction should be carried over into a successor protocol.

1.3 Protocol Layering Assumptions

The working assumption in the NSIS group is the signaling protocol should be 'layered' in two parts (see section 4 of [3] for more details), and this is consistent with several protocol proposals, such as CSTP/ALSP [8] and others too provocative to mention here.

In this document, we refer to these layers as follows:
*) The 'NSIS Base Protocol' (NBP), handling message routing aspects specific to path-coupled signaling; it may include transport-layer-like functionality (reliability, congestion control and so on) or be layered on an existing transport protocol.
*) 'A Signaling Application Protocol' (ASAP), a 'placeholder' for one of many possible protocols which handle particular signaling applications (QoS, middlebox control, and so on).

2. Constraints on Sender/Receiver Orientation

Depending on the particular NSIS function (or specific signaling application function) under consideration, it may be much easier to implement it in a sender or receiver 'oriented' way. This section summarizes these various constraints or influences.

2.1 Signaling Message Routing

Regardless of the particular signaling application in question, path-coupled signaling requires the capability of message routing along the path from sender to receiver. It appears that there are only two methods for the signaling protocol to acquire awareness of the route:
*) Using a PATH mechanism similar to RSVP.
*) Using local topology information (e.g. from a routing protocol, or local configuration).

Signaling message routing, which is a function of the NBP layer, should therefore be sender oriented, possibly with the ability to use additional information sources if available.

A related question is whether signaling messages need to be routed with or against the data flow (or both). (So far as we can tell, the NBP layer only sends and receives messages over a single NSIS hop, so the question only applies to the ASAP layer. It applies both to 'real' signaling application messages and probably also to application-specific error notifications.) If messages need to be routed against the data flow, this has implications for the need to store reverse-path message routing state at intermediate nodes.

The conclusion therefore seems to be that the NBP layer should be able to operate in a sender-oriented mode, but what state it needs to store depends on ASAP layer requirements.

2.2 User Application Triggering

Ultimately, the NSIS signaling is supporting the requirements of some user application (e.g. a VoIP or other media capability). It is likely that sometimes, only one 'party' will have a clear view on what to request, e.g. what is the appropriate QoS, or even what are the flow identification characteristics (port numbers or flow labels may be allocated only at the sender).

Even if both ends know, still one end probably knows first and communicates the information via upper layer exchanges; therefore, fixing sender or receiver orientation for NSIS signaling may impose additional roundtrip delays compared to an 'optimised' solution.

The constraints here are probably both
*) Signaling application specific, and
*) User application specific.

2.3 Renegotiation

There has been some discussion (requirement 5.6.3 of [2] and section 3.3.2 of [3]) of the need for flexibility in which entities can renegotiate aspects of a reservation - for example, whether the sender or receiver should be able to do this, or the initiator or responder, or whether it should be possible from within the network.

This is probably a question which depends on the ASAP layer. If additional flexibility has to be supported for renegotiation compared to initial reservation setup, then this will be an additional source of complexity. Note that some of the motivation for this flexibility is (presumably) to allow localized renegotiation, which is also discussed in section 2.5.

2.4 'Service' Authorization

When any 'resource' is being requested from the network, in some cases the use of this resource must be authorised (or somehow verified to be compatible with a network's internal policy requirements).

It is a hard question to work out how authorisation approaches might impact on the sender/receiver orientation aspects. For example, it is possible that current inter-provider peering agreements would favour a 'sender-initiated' authorisation approach, since typically the traffic originator 'pays' for traffic. On the other hand, in mobile environments, the mobile user may be prepared to authorise a resource request for both directions; a firewall application may only accept resource requests from one side.

Therefore, the service authorisation constraints on sender/receiver orientation are both
*) Signaling application dependent, and
*) Network policy dependent (although it may be the case that for any given signaling application, there is a single 'natural' authorisation direction). Indeed, even for a single path, the network policy may change at provider boundaries.

One reason why sender/receiver authorisation has an impact on signaling flows is the state management aspects while a request is being authorised end to end. For example, Figure 4 shows a 'initiator authorised' signaling flow: messages flowing in the direction NI-->NF-->NR can carry their own authorisation data (they could even

carry it idempotently/statelessly), which could allow very simple authorisation processing at intermediate nodes.

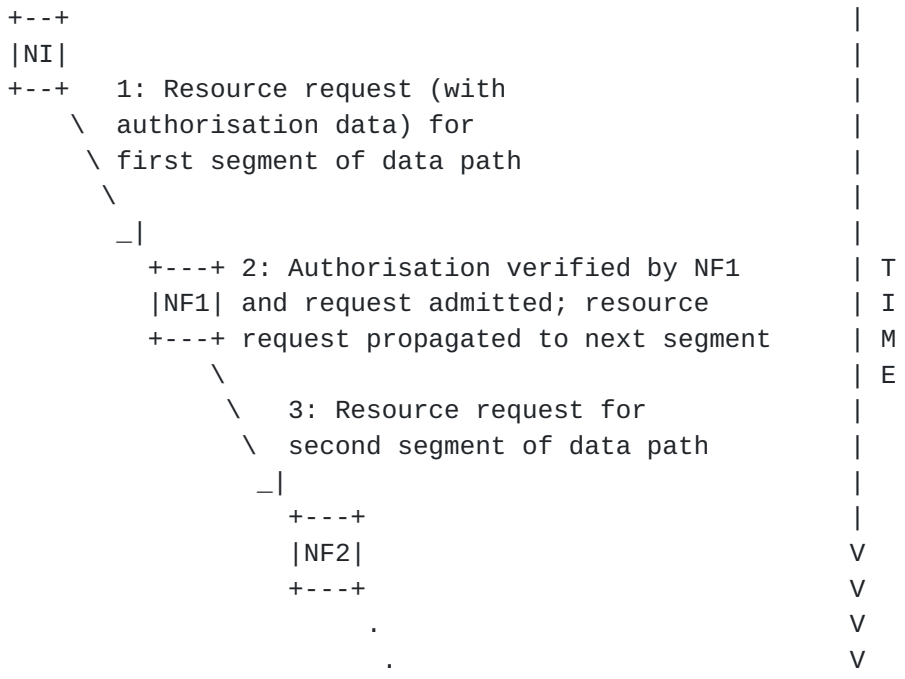

Figure 4: Message Flow for Initiator Authorisation

However, the 'responder authorised' situation is more complex, since the actual authorisation data has to come from the remote end of the signaling exchange, and intermediate nodes may have to retain state waiting for this to arrive, as shown in Figure 5.

The conclusion from this part of the discussion is that:
*) Either the initiator or the responder might be responsible for authorisation aspects (depending on the discussion above), but
*) If the responder is responsible, the NBP will have to handle messages in both directions, and intermediate nodes will have to handle more local state storage.

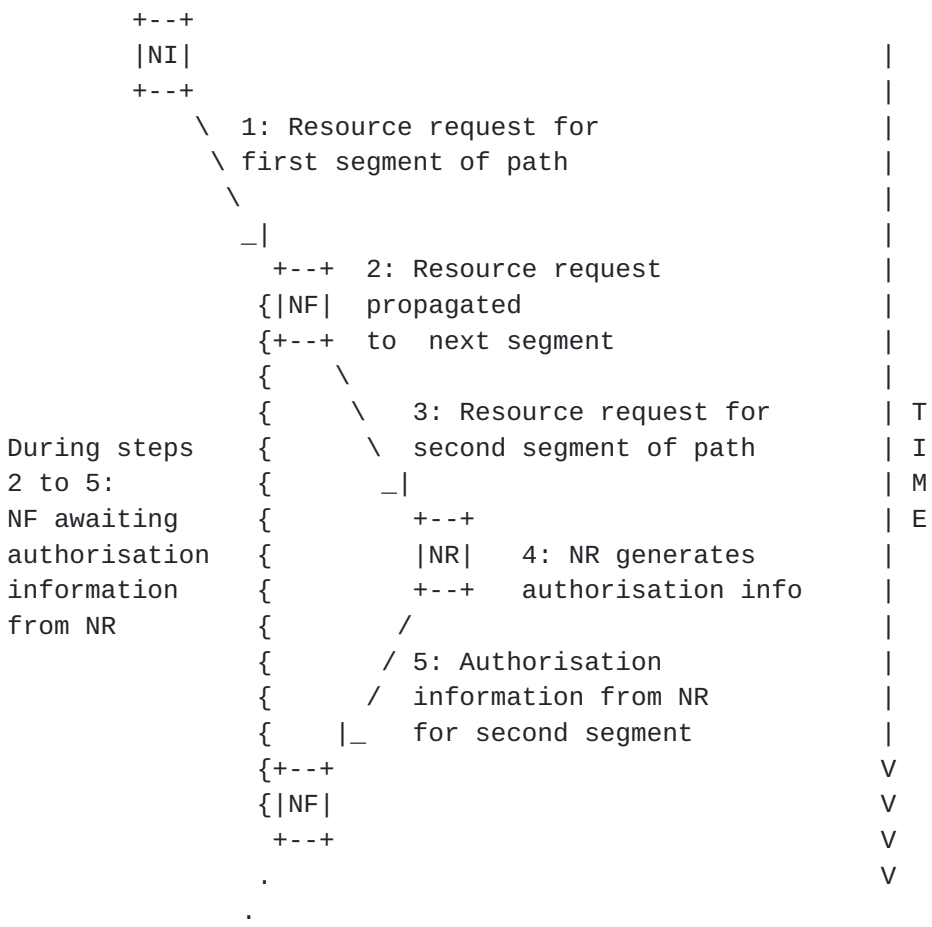

Figure 5: Message Flows for Responder Authorisation

2.5 Localized Signaling Support

Technical approaches for localization of signaling have already been discussed in the context of RSVP, for example in [9] and [10]. There are several reasons why it may be desirable to localize the scope of some aspect of the signaling, such as:

*) Only one endpoint may be generally NSIS aware (e.g. because the other endpoint has no motivation to implement it, or because it is a legacy device).

*) Only one endpoint may be aware of the specific ASAP which is relevant.

*) One endpoint may be mobile and wish to manage aspects of its reservations locally to improve handover performance.

Regardless of the motivation, the end result is that in some scenarios, an endpoint will probably wish to carry out both sender and receiver oriented signaling over some local region of the network, i.e. for incoming and outgoing packets for a bi-directional session. Ideally this would be done both:

*) for the NBP layer (although we have said in 2.1 that this is hard), and
*) for the ASAP layer.

In practice, the mechanism for localizing signaling will be some kind of proxy, and the difficulty in the NBP layer is precisely the difficulty in locating the proxy using purely local signaling. Given the proxy location, however, the ASAP layer signaling between it and the end point then suffers from all the same constraints related to sender/receiver orientation as in the end to end case.

2.6 Protocol - Protocol Interactions

As well as operating locally (in isolation), NSIS signaling will have to interact with other protocols, such as RSVP in other parts of the network. Also, several NSIS deployment scenarios consider NSIS interacting with itself in a 'layered' style, or end-to-end NSIS using edge-to-edge signaling for intradomain provisioning (see for examples sections 3.2 and 7 of [3]).

In these circumstances, NSIS is at least partly at the mercy of these other protocols or other instances of itself, to be initiated and to respond in a compatible way at the protocol interworking boundary. In particular, to interwork with RSVP, NSIS signaling may have to be able to operate in compatible way (e.g. receiver oriented for reservation).

2.7 Multicast Support

Multicast support is the primary justification for the receiver orientation of the reservation signaling in RSVP. The reason is that this naturally allows for progressive state merging from large numbers of receivers back towards the senders, thereby allowing better scalability. For the most general multicast case, this conclusion seems unchallenged (although restricted multicast scenarios, such SSM [11] or multicast with homogeneous receivers, other options may be possible).

Multicast support is not an initial requirement for NSIS protocol work. However, in the future, it might be desirable to extend parts of NSIS to support multicast signaling applications, in which case particular sorts of receiver orientation should not be permanently excluded.

2.8 Something Unpleasant about NAT

The existence of NATs poses some special problems for signaling protocols, since they change the header information in packets

downstream from the sender in a way which may not be predictable before the data flow along the path is actually active (e.g. if dynamic address sharing is taking place).

The consequence of this is that, even if we would naturally imagine a certain signaling operation being controlled from the receiver, this may not be possible because the receiver does not know how to refer to the flow in the first place. Therefore, the signaling has to at least involve the sender as well, probably in cooperation with the receiver (and NAT) as well.

2.9 Summary

The overall conclusion of this section is that there are all sorts of reasons why:
*) Sender orientiation may be required for some functions or in some scenarios;
*) Receiver orientation may be required for other functions or other scenarios;
*) Sender and receiver orientiation have different costs and complexities (e.g. in state management or latency) associated with them.

The choice between sender and receiver orientation therefore appears as a classic rock and hard place dilemma, especially given the natural desire to build a solution that is not overwhelmed by complexity or option negotiation.

3. Possible Approaches

This section presents three possible approaches to resolving this conundrum.

3.1 Fix on One Paradigm

Initially, the most attractive possibility would be to fix on a single paradigm and impose it throughout the NSIS work.

However, it seems impossible to imagine that a single paradigm will support all the requirements and scenarios under discussion; even the baseline RSVP approach, summarized in 1.2, covers only some of the possibilities, and in some scenarios simpler sender-only solutions are possible. A wider set of options might also make incremental deployment (which could be a critical issue) more achievable.

3.2 Allow Both Paradigms

The opposite approach is to allow everything - all aspects of NSIS - to be both sender and receiver oriented. The basic danger here is of overwhelming the NSIS protocols with excessive complexity, since they may well have to operate differently depending on which direction they are working in. It would also make it more difficult to implement a minimal subset of NSIS for particularly constrained environments.

Even if the NSIS protocols could be specified and implemented, the variety of options would pose some operational problems. It might be that both sender and receiver would attempt to initiate the signaling protocol and cause a protocol collision (or indeed that neither of them would). The necessary remedy for this would be to introduce yet another component of the NSIS protocol, to negotiate which end should take the initiative.

3.3 Choose Separately for Each Protocol Component

A third way is to select between sender and receiver orientation independently for each component; provided the inter-component interactions can be controlled, this should then allow better fitting of protocol behavior to the constraints identified above.

Specifically, we could imagine the following:

The NBP layer would be (universally) sender oriented, the same way as the RSVP PATH message (possibly also allowing for other peer discovery mechanisms and proxy usage).

The ASAP layer would be either sender or receiver oriented, depending on the signaling application in question. There might even be different variants for different deployment scenarios (e.g. a sender-oriented intra-domain QoS signaling application, which worked with a receiver-oriented inter-domain counterpart at domain boundaries).

The operation of the NBP and ASAP layers would be interdependent to some extent. The dependencies would include:

*) A receiver-oriented ASAP would suffer from (at least) a single end-to-end delay, waiting for the NBP layer to complete establishing the signaling path. However, this delay is probably an unavoidable consequence of whatever constraints meant the ASAP was receiver-oriented in the first place.

*) The NBP might unnecessarily store reverse-path state for a purely sender-oriented ASAP (in other words, one which required no receiver-to-sender messages). This could be fine tuned by allowing the ASAP to invoke the NBP in a mode which didn't store such state.

3.4 Implications of a Layered Choice

Splitting the responsibility in this way and leaving the selection to the ASAP layer represents quite a significant shift in thinking compared to current protocols. There are therefore some dangers.

The first danger is of excessive flexibility. On the other hand, the flexibility is a consequence of the NSIS requirements and constraints. This approach does allow simpler solutions in particular environments (e.g. for specific ASAP layers).

The split decision probably has implications for the way state is managed between the layers, especially where different layers are in different protocol states in the interior of the network. This clearly needs further analysis.

If the choice between sender and receiver initiation is really a matter for the ASAP layer, the implication is that the messages visible in the NBP should be somewhat neutral in content. The existing NSIS framework (section 4.3.2 of [3]) may be too specific in this regard. Also, the basic NI/NF/NR concepts may have to be split depending on the NBP/ASAP layer.

4. Additional Considerations

The adoption of a split approach for sender/receiver orientation could have some implications for other aspects of NSIS-related work beyond the basic unicast path-coupled case. These are summarized here.

4.1 Bidirectional Reservations

NSIS work (especially requirements work) has discussed the case of 'bidirectional' reservations, in other words, signaling for both directions of a point-to-point data flow. The baseline approach for this feature (see section 3.2.7 of [3]) is to simply combine a pair of unidirectional reservations, which is then covered by the previous discussion.

However, a 'true' bi-directional reservation (integrating the signaling for each direction) would also be interesting in some applications. Topologically, this would only be possible over a path segment that was symmetrically routed.

Following the split layer approach of section 3.3, it seems that asking for bi-directional protocol within the NBP layer is not meaningful, since in general, even if the route is symmetric, NBP

layer procedures have to operate asymmetrically while finding this out. However, it could be possible for the NBP layer to detect this symmetry (i.e. correlate the routes for incoming and outgoing flows) and provide this as an enhanced service interface to the ASAP layer.

Whether the ASAP layer can or must use this capability to set up a bi-directional reservation using that interface is probably very much dependent on the signaling application and possibly scenario in question. It seems likely that the logical behavior (to do with state management, message sequences and so on) is the same as just a sender and receiver initiated reservation; however, the sending and reception of the messages in pairs might enable more efficient local processing.

4.2 Path-Decoupled Signaling

Although NSIS does not currently have path-decoupled signaling in its scope, it is worth pointing out here some issues that may be special related to sender/receiver aspects in the path-decoupled case.

The main issue with path-decoupled signaling is that once the signaling endpoints are not on the data path, it is no longer an unambiguous topological decision to categorize one of them as being related to the sender and the other to the receiver (see Figure 6).

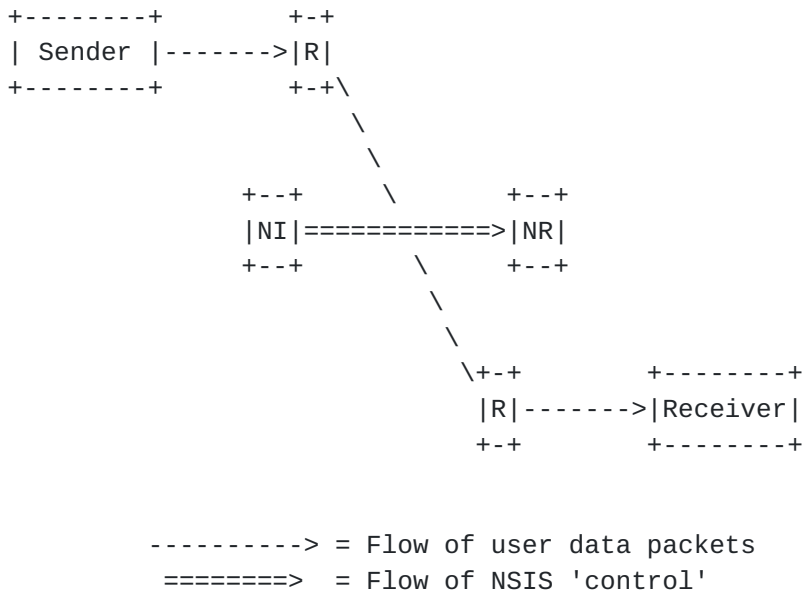

Figure 6: Path-Decoupled Signaling Topology

If there is to be an attempt to re-use path-coupled NSIS signaling in this type of environment, and the signaling depends significantly on

sender and receiver orientation, it will be necessary to work out how to match these concepts in the path-decoupled case.

One approach to this would be to place the responsibility for 'path-orientation' in the NBP layer or its equivalent (which has to be modified anyway for the path-decoupled case to support off-path nodes). This layer will also have to have some more explicit (application layer?) interaction with the data sender and receiver, just to trigger the signaling process in the first place. However, once this is done, the ASAP layer (at least in terms of message exchanges) might operate in almost exactly the same way as in the path-coupled case.

5. Conclusions

This document has no conclusions. However, it proposes a method for reasoning (possibly constructively) about the sender/receiver orientation possibilities. Implications for the requirements and framework, and consequences for path-decoupled signaling, have also been identified.

References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

2 Brunner, M., "Requirements for QoS Signaling Protocols", draft-ietf-nsis-req-04.txt (work in progress), August 2002

3 Freytsis, I., R. E. Hancock, G. Karagiannis, J. Loughney, S. van den Bosch, "Next Steps in Signaling: Framework", draft-ietf-nsis-fw-00.txt (work in progress), October 2002

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

5 Braden, R. et al., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997

6 Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000

7 Braden, R. et al., "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994

8 Braden, R., "A Two-Level Architecture for Internet Signaling", draft-braden-2level-signal-arch-00.txt (work in progress), November 2001 (expired)

9 Gai, S. et al., "RSVP Proxy", draft-ietf-rsvp-proxy-03.txt (work in progress), March 2002

10 Manner, J., et al., "Localized RSVP", draft-manner-lrsvp-00.txt (work in progress), May 2002

11 Bhattacharyya, S. et al., "An Overview of Source-Specific Multicast (SSM)", draft-ietf-ssm-overview-03.txt (work in progress), March 2002

Acknowledgments

The authors would like to thank all their colleagues and fellow participants in the NSIS working group for exposing the complexities and subtleties in this subject area.

Author's Addresses

{Robert Hancock, Eleanor Hepworth, Andrew McDonald}
Roke Manor Research
Old Salisbury Lane
Romsey
Hampshire
SO51 0ZN
United Kingdom
email: {robert.hancock|eleanor.hepworth|andrew.mcdonald}@roke.co.uk

Full Copyright Statement