Internet Engineering Task Force
INTERNET-DRAFT
draft-fu-nsis-mobility-00.ps

NSIS
X. Fu, H. Schulzrinne, H. Tschofenig
U. Goettingen/Columbia U./Siemens
June 23, 2003
Expires: December 2003

Mobility Support in NSIS

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

To view the list Internet-Draft Shadow Directories, see http://www.ietf.org/shadow.html.

Copyright Notice

Abstract

Mobility support was one of the shortcommings of RSVP and a number of proposals have been made in the past to improve various features.

This document attempts to determine what the expected functionality is, how various problems can be solved and what implications are caused by certain design decisions.

Contents

1 Introduction

The NSIS working group is currently working on the two-layer architecture with a generic NTLP protocol and various NSLP applications (e.g. QoS and NAT/Firewall traversal). The NSIS Framework [1] describes the functionality of the individual layers in detail.

Mobility support is one of the items on a list of desired features for future QoS signaling protocols. Unfortunately, mobility support for a generic signaling protocol still causes a lot of confusion. To start the discussion in the working group the expected functionality has to be discussed and the implications for the protocol design evaluated. A recently published MIP QoS requirements document [2] lists some of the requirements.

Interactions of RSVP signaling with Mobile IP have previously been investigated, e.g., in [3, 4]. In RSVP, as signaling sessions are identified by IP addresses, it becomes difficult to address host mobility networks [5, 6]. When a mobile node (MN) moves, the state established along the previous route remains until it times out after multiples of the soft-state interval, typically after more than a minute. With even modest mobility, large amounts of "*obsoleted*" state may cause inefficient resource allocation. In addition, IP-in-IP encapsulation in Mobile IP (MIP) causes RSVP messages sent to the MN also be encapsulated as messages with a new protocol number in its outer header, thus concealed to the RSVP nodes along the path and unable to perform signaling tasks correctly.

This document attempts to determine what the expected functionality is, how various problems can be solved and what implications are caused by certain design decisions, to deal with mobility support in NSIS.

2 Terminology

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 [7].

In addition, this document frequently uses the following terms or abbreviations most of which are defined in Mobile IP, SEAMOBY and NSIS related documents:

Home Address (HoA)

Home Agent (HA)

Foreign Agent (FA)

Care-of-Address (CoA)

Correspondent Node (CN)

Mobile Node (MN)

Access Router (AR)

Mobile IP (MIP): Mobile IPv4 (MIPv4) or Mobile IPv6 (MIPv6)

Hierarchical Mobile IPv6 (HMIPv6)

Fast Handovers: proposals for fast handover enhancements to MIP, such as Fast Handovers for Mobile IPv6 (FMIPv6) and Low latency Handoffs in Mobile IPv4.

Localized Mobilitiy Management (LMM): micro-mobility proposals such as HMIPv6 and Mobile IPv4 Regional Registration.

Regional Care-of-Address (RCoA)

On-Link Care-of-Address (LCoA)

Mobility Anchor Point (MAP)

Gateway Foreign Agent (GFA)

Standard routing: standard IP packets without routing header or IP-in-IP encapsulation are routed according to the destination address field in the IP header

Normal routing / direct routing: routing data packets with routing header (kind of ”source routing”) or standard routing, i.e., without using (mobility related) tunnels.

Routing with Mobile IP tunnels: packets with IP-in-IP encapsulation are routed according to the destination address field in the outer header. These packets are encapsulated at the entry of a Mobile IP tunnel and decapsulated at the exit of a Mobile IP tunnel.

Next Steps in Signaling (NSIS)

NSIS Entity (NE)

NSIS Transport-Layer Protocol (NTLP)

NSIS Signaling-Layer Protocol (NSLP)

Peer discovery

Cross-over Router (CR): the NSIS-aware node which is the nearest to the MN but locates in the common path where new and old paths converge after a handoff.

3 Mobility Support in NSIS: Problem Statement

Efforts like [8, 9, 10, 3, 1, 11] have been made to analyse mobility issues in this area and result in some interesting discussions. What has not emerged is a consensus for defining what problems would be expected, what functionalities need to be contained in NSIS to tackle mobility support and how various problems can be solved. This section attempts to make a summary of problems and next section discusses possible solutions.

In the past the following issues with mobility support in NSIS have been discovered:

Issue 1 - Indexing state information: If state information, which is stored at routers along the path, is indexed based on the Care of Address (CoA) then it is unaccessible after most handover procedures. With the handover procedure the mobile node obtains a new Care of Address.

Issue 2 - Double Reservation: State information should not be established twice between the cross-over router (CR) and the CN.

Issue 3 - State Update: Updating state information along the entire path might be necessary, for example, due to the selection of the flow identifier. If an application uses the IPv6 Flow Label (3-tuple) for the flow identifier then end host mobility requires updating the flow identitifer along the entire path.

Issue 4 - Keeping Signaling Local: End-to-end signaling is typically expensive. In some cases it might be desireable to keep signaling messages local. This might require to either add a ”local only” flag (or something like scoping) or to have the cross-over router (or the MAP) to decide by itself whether to stop signaling or to forward the signaling message to the corresponding node.

Issue 5 - Sender-Initiated Reservations: Sender-initiated reservations seem to be more efficient in mobile environments. Until now, the NSIS working group has not yet discussed which approach (sender- or receiver-initiated reservations; or even both) should be supported by a future signaling protocol. In some cases receiver-initiated reservations are better (e.g. Firewall/NAT traversal) whereas in other cases it might even be required to have more than one roundtrip (e.g. when price distribution has to be considered).

Issue 6 - Partial State Release: It might be desireable to delete previously established state on the “*obsoleted*” path (i.e. along the path between the cross-over router and the old access router). In mobility scenarios it is desireable to delete these states as soon as possible, particularly for fast mobility. Soft-state timeouts might be too inefficient.

3.1 Synchronization Issues

Issue 7 - Synchronziation Problems: Partial state release might cause synchronization problems or race conditions in some cases. For example, if an MN moves rapidly from one AR0 to another AR1 afterwards to a third AR2, local repair messages may be initiated along AR1 and AR2 subsequently, but the message along AR1 may arrive later than that along AR2. As illustrated in Fig. 1, this may cause the states be incorrectly updated (step 7 in CR) and released (step 8 in AR2).

Another example of this problem is a ping-pong effect where a mobile node switches between two neighboring ARs, possibly causing states be released or updated incorrectly.

3.2 Tunnel Issues

In addition to direct routing between a mobile node (MN) and the corresponding node (CN), Mobile IP and various local mobility management (LMM) [12, 13, 14] and fast handoff [15, 16] schemes may add one or more IP-in-IP encapsulation tunnel(s) between the home agent (HA) and the corresponding node (CN), as summarized in the appendix. Thus, data packets may take one of several possible routes:

- For data packets from the MN towards the CN:

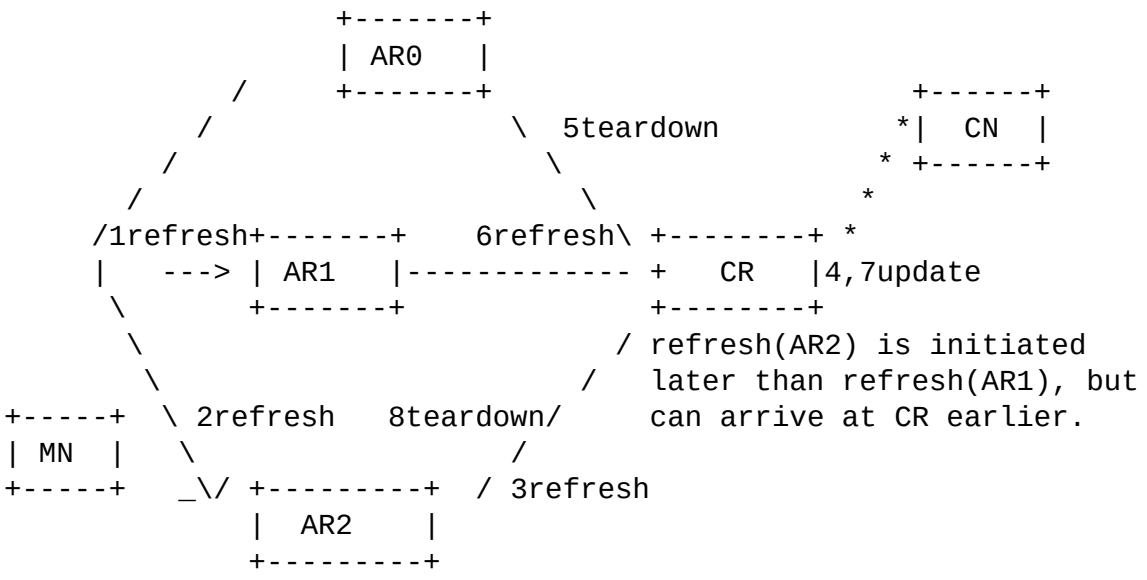

Figure 1: Fast Movement Issue: an Example

1a) direct routing, or

1b) (in case of reverse tunnel [17, 6]) reverse tunnel segment from the MN to the HA plus a direct routing segment from the HA to the CN.

- For data packets from the CN towards the MN:

  2a) optimized path from CN to MN (using routing header through the CoA);

  2b) non-optimized path which consists of a normal routing segment from CN to HA and another tunnel segment from HA to FA/MN;

  2c) for LMM schemes, GFA/MAP divides the tunnel segment from HA to FA/MN in 2b) into two segments, a tunnel from HA to GFA/MAP plus another tunnel from GFA/MAP to FA/MAP;

  2d) for fast handover schemes [15, 16], an additional tunnel is inserted in 2b): the tunnel between old and new access routers.

For (2c) and (2d) the exit of a tunnel is immediately the entry of next tunnel.

IP-in-IP encapsulation in Mobile IP, for example in (1b), (2b), (2c) and (2d), causes each signaling message in end-to-end addressing schemes like RSVP to be encapsulated inside a tunnel message. The tunnel hides the RSVP message and make intermediate nodes between the tunnel entry and the tunnel exit invisible.

The following issues can be identified for the interworking between Mobile IP and NSIS:

Issue 8 - Non-Unaware NSIS Tunnel Endpoints: If the two end points of a tunnel do not support NSIS then it is impossible to provide signaling services for any NSIS node inside the tunnel. In order to (for example) make a QoS reservation for tunnels, it is necessary to have the tunnel end points to support NSIS.

Issue 9 - Selective Tunnel Reservations: When an application triggers a per-flow QoS reservation then in some cases it is desired to trigger a QoS reservation also for the tunnel. In some cases it is, however, not desireable to support a QoS reservation for a tunnel. It must therefore be possible for NSIS to decide whether a reservation for a tunnel is desired. An end-to-end reservation must not automatically be extended to the tunnel since the tunnel might also serve as a means for aggreagation.

Issue 10 - Discovery Procedure: As previously stated NSIS signaling messages are hidden inside a tunnel. It is therefore necessary to trigger a separate discover and signaling message exchange for the tunneled region. The signaling messages must either be prevented to ”enter” the tunnel, need to experience special treatment after encapsulation or require different addressing (i.e. addressing the tunneled reservation towards the tunnel exit).

3.3 NSLP/NTLP Interaction Issue

Issue 11 - NSLP/NTLP Interworking: The NTLP should be able to notify the NSLP to update state (by initializing NSLP refresh/teardown messages appropriately). An open issue is, however, how and what information the NSLP can expect from NTLP, or directly from the routing interface.

3.4 Routing Interface Issue

Issue 12 - NSIS/Routing API: Mobility reflects different route change (due to changes in the binding cache, for example in HA, CN, FA/MN, GFA/MAP) or due to creating, updating or deleting tunnels. An API is required for the communication between the operating system and the NSIS daemon, upon which NSIS is able to trigger the necessary action.

4 Possible Solutions

To deal with the issues summarized in Section 3, we describe some possible approaches to address them.

4.1 NTLP and Discovery

First, issues 1-2 can be resolved by a unique session identifier independent on flow identifier (which reflects the traffic information), to index the state information. When the MN moves, the NTLP only changes the flow identifier of the session, without changing the session identifier. This change takes place when an NTLP node detects the introduction or release of MIP tunnels or simply a route change in the MN or the CN. Then it triggers the next-hop discovery mechanism to determine the new next signaling node and updates its information in the NTLP state according to the unique session identifier.

Issue 3 — releasing of existing state in the old path, by default, can be achieved by the state timer expiration. However, as the state timer is relatively long, keeping state in these nodes may be inefficient. Alternatively, we propose to use local explicit teardown messages. A reasonable place for initiating such a teardown message is the cross-over router (CR), i.e., the node where the old and new paths merge after a route has changed. To help the release of the obsoleted states and determine the behavior of the cross-over router, we further introduce a next-hop branch identifier. Once a new next NSIS node is determined, the NTLP state in the current associates it with a new next-hop branch identifier (to represent the new branch). Then a refresh message is sent through the new branch to establish the necessary states, until the cross-over router with a same NTLP session state is reached. After that a teardown message assigned with the

old branch identifier can be sent reversely towards another end point and terminated by finding a different branch identifier for the same session state.

To reflect the change of flow identifier for NTLP sessions, it is also necessary to refresh the common path (e.g., the path between the CR and the HA, or between the CN and the HA). This needs to differentiate from the CN as data source to the MN as data source.

- For the MN as data source: the route change in the MN triggers a refresh towards the CN. A refresh is then sent along the newly-discovered branch towards the CN; after the CR is determined, the CR sends a teardown along the reverse path of last-recently-used branch, and the refresh is forwarded on toward the CN.

- For the CN as data source: the route change in the CN (in route optimization case), the HA or the GFA/MAP/... (other cases — typically a tunnel segment is created or modified and thus causes a route change in such nodes for the path from the CN to the MN) triggers a refresh towards the MN. The state in first segment (between the route change point to the CR) is refreshed with the updated flow identifier; then a local repair takes place from CR: a refresh is sent along the newly-discovered branch and a teardown is sent along the last-recently-used branch.

  There are additionally two issues with this local repair process. First, a teardown message arriving at MN (along the obsoleted branch), if arrives earlier than the refresh message (along the new branch), can incorrectly make a decision to release the state in the MN. There are two possible solutions: 1) only allow the MN to initiate the release of states in the obsoleted branch, but this solution does not work if the MN looses its connection to the old AR at all; or 2) still let CR to initiate release, but mandate the MN not to remove state upon an earlier-arrived teardown message; or 3) let CR to initiate release after a period of time (e.g., an RTT) after initiating the refresh in the new branch. Our suggestion is 2), as it neither potentially remains "orphan" states nor introduces additional states in the routers.

  Second issue concerns with delivery of teardown messages. In case of peer-to-peer addressing of NSIS signaling messages, the above description works. However, in case of end-to-end addressing of NSIS signaling messages, default routing in the CR may cause the teardown messages to follow the same path as the refresh messages traverse (i.e., in the new path), however these teardown messages need to be sent through the obsoleted branch. One solution is to specify the logical outgoing interface — logical interface handle (LIH), same as in RSVP [18] — for each branch and use it for message delivery.

To deal with issue 7 requires effective sequencing of branches. Our proposal is to let the CR always assign a next branch id different from the current branch id. A refresh with an existing session identifier will cause a state update (and accordingly cause a release) only when the branch identifier in the refresh is larger than that in the current state.

Issue 8 can be resolved by mandatory NTLP support in the tunnel end points, as well as the CN and the MN. When an NSIS node detects the introduction of an MIP tunnel, if signaling into the tunnel is desired, it can issue a discovery towards the tunnel exit, not to the destination address of the end-to-end reservation. For the NSIS node inside the tunnel, it can do that in the same way.

Issue 11 and issue 12 require API functionality within NSIS and the operating system (notification of mobility route change, etc.).

Issues 4-6, 8-9 and 10 are covered/resolved while presenting themselves or discussing other issues. Additionally, issue 5 requires to perform reverse signaling during the local repair for receiver-initiated signaling

services.

In the example in Fig. 2, an MN acting as the data sender communicates with a CN. When it changes its AR to that of a new network, it is associated with a new CoA (nCoA) and the flow identifier in the signaling message is changed. The session identifier, however, remains intact so that only a single state is held in the NSIS nodes along the path from the CR to the CN. After the CR is determined, it can issue a teardown message to release states towards the MN along the reverse path that former signaling messages traverse. Note the refresh messages arriving at the CR will be forwarded on towards the signaling target, to refresh existing states to reflect the change in flow identifier.

Fig. 3 illustrates an example for CN as a data source.

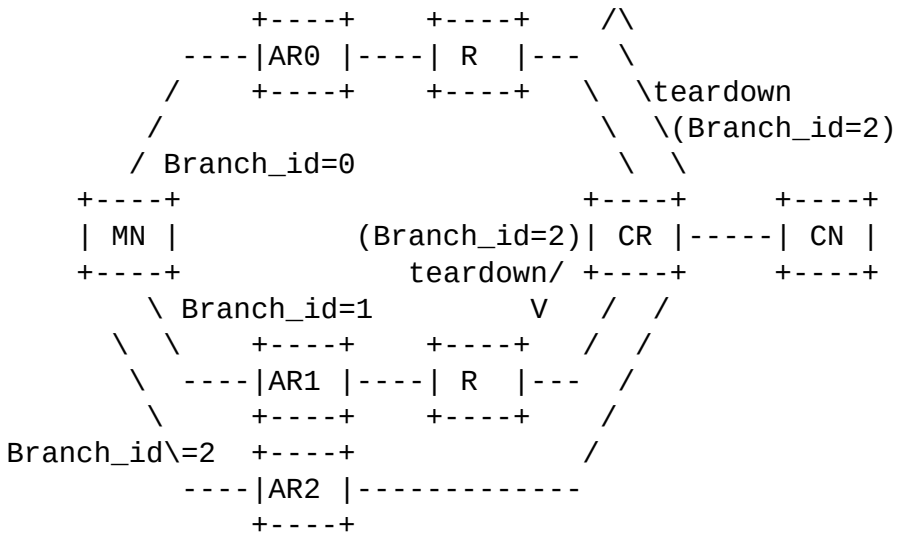

Figure 2: Local Repair Proposal

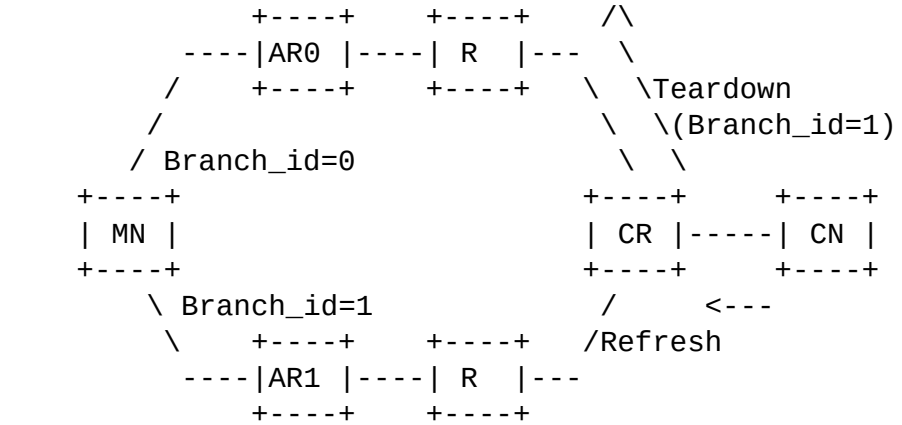

Figure 3: Local Repair case with CN as data sender

4.2 NSLP and NTLP/NSLP Interactions

The creation, update or removal of NTLP state also triggers associated NSLP entity(ies) to create, update or release related NSLP state accordingly. For example, before sending an NTLP refresh, the NTLP entity can trigger QoS NSLP and let the latter to request a Reserve refresh message toward the QoS NSLP target address. However, not all NSLP functionalities can be supported. For example, it may be impossible/unnecessary to request NSLP response upon the receipt of the NSLP teardown message in the NSLP responder, as the NTLP state along the path has been already released. In general, upon a NTLP notification, the NSLP should decide by its own which its next behavior is. (Issue 11)

4.3 Routing Interface

There are two approaches for NSIS to interact with routing (issue 12). One is API, upon which NSIS can query the mobility routing status. Another is to provide route change notification to NSIS whenever a mobility route change happens (e.g., binding cache changes). As the first approach typically introduces some delay in state management per handoff, we suggest the notification (active) way. This can also simplify the trigger design when the NTLP refresh timer is different from the timer for binding management.

The natural interface for NSIS to access routing interface would be in the NTLP.

4.4 Operation of Proposed Mobility Support Mechanisms

What we eventually need to support mobility in NSIS might be:

- a mobility module (to determine state index, decide which is the CR and do local repair) co-located with NTLP.
- a mobility tunnel-aware module co-located with NSIS peer discovery.
- a mobility object/field in NTLP header to indicate branch-id and/or scope.
- logical interface handles (LIH) in addition to PHOP and NHOP in the NTLP state.
- a mobility route change notification interface in mobility-aware nodes.

A pseudocode for proposed solution for mobility support in NSIS is described in Fig. 4.

5 Security Considerations

The introduction of the session identifier to tackle some mobility related issues has security implications which are described in [19]. Some questions raised in this context may deserve discussions within the NSIS working group.

6 Open Issues

- End-to-end addressing v.s. peer-to-peer addressing of NSIS signaling messages;
- Message format and routing interface definition;
- An Obsoleted_Branch_Removal flag for deciding whether to remove actively states in the obsoleted branch;
- Interworking with Context Transfer.

```
IF(a node detects a mobility route change)/*MN,CN,HA or FA/MAP/GFA/...*/
  flow_id = (nCoA, CN, ...);
  next_hop = discovery(dest); /* if the node is a tunnel entry or inside
     a tunnel, dest is the tunnel exit; otherwise dest is normal dest */
  IF(next_hop != old_next_hop) /* this is a new branch */
    branch_id = (branch_id + 1) % MAX_VAL;
    creat_state(&state,branch_id,flow_id,session_id,...);
  ELSE /*update the existing state with the new flow_id */
    update(&state(session_id), flow_id);
  ENDIF
  /* send a refresh, local repair if it is a new branch */
  mesg = new_msg (''refresh'', branch_id, flow_id, session_id,...);
  send_msg(&mesg);
ENDIF

/* Determine the cross-over router */
IF(a node gets a msg with t_flag == 0)
  IF(session_id(msg) == session_id(state))
    /* decide whether to update branch_id */
     IF(branch_id(msg) >= (branch_id(state)+1) % MAX_VAL)
       branch_oid = branch_id(state);
       update(&state(session_id), branch_id(msg));
     ENDIF
     /* then send a teardown mesg in the old path */
     mesg=new_msg(''teardown'', branch_oid, flow_id, session_id,...);
     send_msg(&mesg); /* send the teardown mesg along old path */
     forward_msg(&msg); /* forward refresh mesg on */
 ELSE /* it is not the cross-over router */
    IF(session_id(msg) does not exit in local state)
      create(&state(session_id, flow_id(&msg)));
    ELSE
      update(&state(session_id), flow_id(&msg));
    ENDIF
 forward_msg(&msg);
 ENDIF
ENDIF

 /* Determine the end of teardown msg */
IF(a node gets a teardown msg)
  IF(branch_id(&msg) > branch_id(&state))
     remove(&state);
     forward_msg(&msg);
  ENDIF /*else stop here, don't forward on the teardown msg*/
ENDIF
```

Figure 4: Pseudocode for proposed mobility support in NSIS

7 Acknowledgment

The authors would like to acknowledge discussions with Rene Soltwisch, Robert Hancock, Andrew McDonald, Hemant Chaskar and other members of the NSIS working group for numerous discussions on various aspects of this topic.

8 Authors' Addresses

Xiaoming Fu
Institute for Informatics
University of Goettingen
Lotzestr. 16-18
37083 Goettingen
Germany
EMail: fu@cs.uni-goettingen.de

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
EMail: hgs@cs.columbia.edu

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Hannes.Tschofenig@siemens.com

Appendix – Data Flows under Different IP Mobility Schemes

We summarize possible data flows under different IPv6 mobilitiy schemes as Figures 5-10 (note we are not discussing about mobility management messages flows; IPv4 is quite similar and not presented here).

They can be further summarized into two cases:

- Fully IPv6 normal routing, or
- A normal routing segment + one or more IP-in-IP tunnel(s)

The implications is we need to take special care to tunnel segment. If the NSIS NTLP follows similar concept in CASP, i.e., NTLP messages are addressed in a peer-to-peer way, the concern regarding tunnels becomes:

1) how to address&deliver discovery messages in tunnel entry points,

2) whether to signal into tunnels (if so, again how to address discovery messages).

Some other implications behind:

- For fully IPv6 normal routing, the cross-over router can be located anywhere between MN and CN;

- For routing with tunnel segment, the cross-over router can only be located somewhere between the HA (a tunnel entry) and the CN (for fast/hierarchical mobility, could be in a small segment between HA and CN). This further implicates, if we want to do local repair in such case, we need to signal into the tunnel; also, tunnel end points are therefore needed to support NSIS.

Figure 5: MIPv6: MN−− >CN, no reverse tunnel

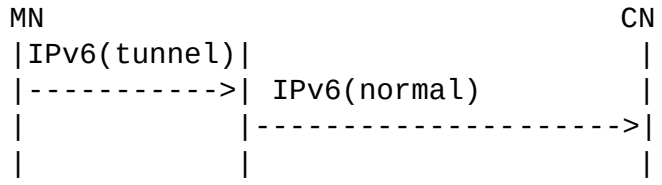

Figure 6: MIPv6: MN−− >CN, with reverse tunnel

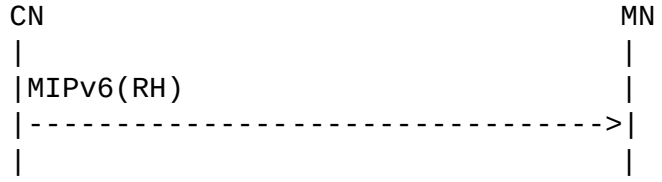

Figure 7: MIPv6: CN−− >MN, route optimization

References

[1] R. Hancock *et al.*, “Next steps in signaling: Framework.” Internet Draft, work in progress, Nov. 2002.

[2] H. Chaskar, “Requirements of a qos solution for mobile ip,” 2003. Work in progress.

[3] M. Thomas, “Analysis of mobile ip and rsvp interactions,” Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[4] J. Manner *et al.*, “Localized RSVP.” Internet Draft, work in progress, May 2002.

[5] C. Perkins, “Ip mobility support for ipv4,” 2002. Work in progress.

[6] D. Johnson, C. Perkins, and J. Arkko, “Mobility support in IPv6,” Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[7] S. Bradner, “Key words for use in RFCs to indicate requirement levels,” RFC 2119, Internet Engineering Task Force, Mar. 1997.

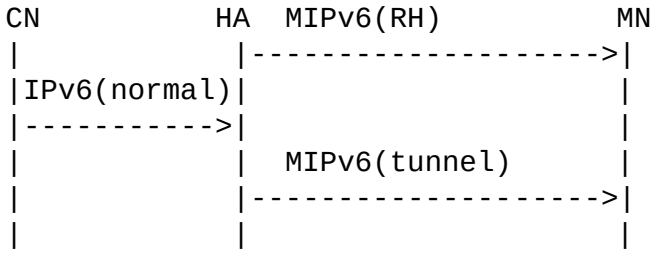

Figure 8: MIPv6: CN−− >MN, no route optimization

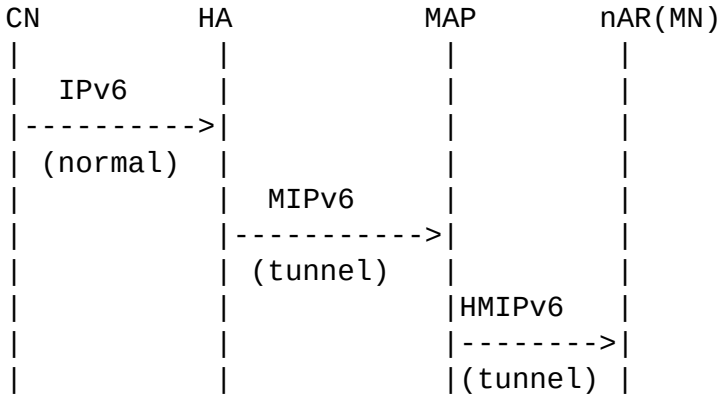

Figure 9: HMIPv6: CN−− >MN, no route optimization

[8] C. Westphal and H. Chaskar, “Qos signaling framework for mobile IP,” Internet Draft, Internet Engineering Task Force, June 2002. Work in progress.

[9] C. Shen *et al.*, “Mobility extensions to RSVP in an RSVP-mobile IPv6 framework,” Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[10] H. Schuzrinne, H. Tschofenig, X. Fu, J. Eisl, and R. Hancock, “Casp – cross-application signaling protocol.”

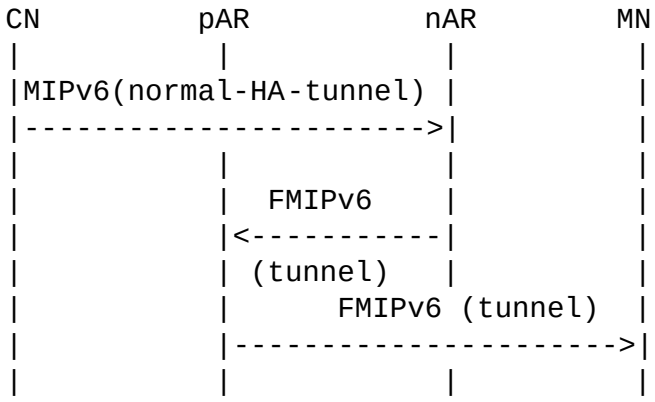

Figure 10: FMIPv6: CN−− >MN, no route optimization

Internet draft, work in progress, Sept. 2002.

[11] M. Brunner, “Requirements for QoS signaling protocols.” Internet Draft, work in progress, July 2002.

[12] C. Williams, “Localized mobility management requirements,” Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[13] H. Soliman, C. Castelluccia, K. Malki, and L. Bellier, “Hierarchical MIPv6 mobility management (HMIPv6),” Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[14] E. Gustafsson, A. Jonsson, and C. Perkins, “Mobile IPv4 regional registration,” Internet Draft, Internet Engineering Task Force, Mar. 2002. Work in progress.

[15] G. Dommety, A. Yegin, C. Perkins, G. Tsirtsis, K. Malki, and M. Khalil, “Fast handovers for mobile IPv6,” Internet Draft, Internet Engineering Task Force, Mar. 2002. Work in progress.

[16] K. Malki *et al.*, “Low latency handoffs in mobile IPv4,” Internet Draft, Internet Engineering Task Force, July 2002. Work in progress.

[17] G. Montenegro and Ed, “Reverse tunneling for mobile IP, revised,” RFC 3024, Internet Engineering Task Force, Jan. 2001.

[18] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource reservation protocol (rsvp) – version 1 functional specification.” RFC 2205, Sept. 1997.

[19] H. Tschofenig, H. Schulzrinne, *et al.*, “Security implications of the session identifier,” internet draft, Internet Engieering Task Force, June 2003. work in progress.