Internet Engineering Task Force
Internet-Draft
Expires: July, 2004

Roland Bless
Xiaoming Fu
Robert Hancock
Seong-Ho Jeong
Cornelia Kappler
Sung-Hyuck Lee
Jukka Manner, Ed.
Paulo Mendes
Hannes Tschofenig
January, 2004

Mobility and Internet Signaling Protocols

<draft-manyfolks-signaling-protocol-mobility-00.txt>

Status of this Memo

This document is a submission to Next Steps in Signaling Working Group. Comments should be submitted to the nsis@ietf.org mailing list.

Distribution of this memo is unlimited.

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.

Abstract

IP mobility, which in its simplest form includes routing changes, can have major influence e.g. on the protocols designed within the NSIS Working Group. This draft is a first step in helping us to decide on

how the problems caused by mobility should be handled within Internet signaling protocols, and a stimulus to further security work.

Table of Contents

1. Introduction

The mobility of IP-based nodes incurs route change, usually at edge of the network. Route change may also be caused by reasons other than mobility, such as routing protocol adaptation in response to varying network conditions, or host multihoming. Normal IP mobility (i.e., Macro-mobility) also involves change of mobile node IP addresses. Since IP addresses are usually part of flow identifiers, change of IP addresses implies change of flow identifier.

Micro mobility usually does not cause change of the global IP addresses, but affects the routing paths within the local access

network. Some Local Mobility Management (LMM) mechanisms may change the IP address assigned to the mobile node within the access network, for example, mechanisms based on a hierarchy of mobility handling routers. Some protocols either use tunneling to forward packets towards the new location of the mobile node, or set and update per-host routing entries in the network, as for instance, ad-hoc routing protocols.

This draft addresses mobility-related considerations for NSIS. The goals of this draft are to analyze the effects of mobility on the NSIS Transport Layer Protocol (NTLP) and on the NSIS Signaling Layer Protocol (NSLP), and to make sure there are no initial design mistakes that break the protocols in mobile environments. The NTLP is an application independent protocol to transport service-related information between nodes in a network. Each specific service has its own NSLP protocol.

The goals of this draft are not to suggest a design for a separate mobility-specific NSIS protocol or to intentionally delay the current work. We expect that this study will actually speed-up the current work on the NSIS protocols. We do not intend to present specific implementation issues in this document, but rather propose how the NSIS protocols should be designed to work in a mobile environment.

A further goal of this draft is to give guidance to people proposing new NSLP protocols. The guidelines in this draft would help those people make sure their NSLP protocols truly work in wired and wireless/mobile environments. Moreover, a goal is to stimulate further discussions related to the security and authentication issues in a mobile environment making use of the NSIS protocols. We expect there to be a common (minimal) set of functions that the NTLP and NSLP need to support. Furthermore, we intend to capture any additional issues that would need specific precautions, e.g. in future NSLPs.

The discussion is divided into two parts. The first part discusses the very basic functionality needed within the NTLP and NSLP protocols. We expect these features to be from most parts already available within the NSIS protocols, or at least can be added with little effort.

The second part takes the discussions to more specific scenarios, including support for multihoming and inter-working issues with a number of mobility-related protocols. These functions would be looked

at once the first versions of the NSIS protocols are finished.

2. Terminology

For terminology related to wireless and mobile networking, we refer to [Seamoby-terms].

Session

A single application layer flow of information between end points that occur during the span of a single connection for which some network control state information is to be manipulated or monitored. IP mobility may cause the mapping between sessions and flows to change, and IP multihoming may mean there is more than one flow for a given session.

Session Identifier

The identifier used to relate signaling messages to a specific session.

Flow

A sequence of packets sent from a particular source to a particular (unicast or multicast) destination for which special handling is provided by the combination of header fields. Only unicast, unidirectional flows are considered in this document.

Flow Identifier

The identifier used to uniquely identify a particular data flow for which the specific service is requested from the network. All packets associated with the same flow will be assigned the same

flow identifier by the source. The flow identifier contains information about the flow which should receive a particular treatment, and it may consist of a combination of the typical 5-tuple or, for example, source IP address, destination IP address, and flow label in IPv6-based networks. See [I-D.ietf-nsis-fw] for more information.

Crossover Node (CRN)

A Crossover Node is a node that for a given function is a merging

point of two or more separate sets of state information, and not only a physical route splitting point. In the context of this draft, we can distinguish several logical (but not necessarily physically) different CRNs:

NTLP CRN, after a routing change, the node closest to the end host from which the NTLP state information towards the CN does not change.

NSLP CRN, after a routing change, the node closest to the end host from which the NSLP state information towards the CN does not change. The NSLP CRN may be different for different NSLPs.

NSIS CRN, either an NTLP CRN or an NSLP CRN.

Upstream CRN, after a handover, the node closest to the data receiver from which the state information towards the data sender does not change.

Downstream CRN, after a handover, the node closest to the data sender from which the state information towards the data receiver does not change.

Mobility CRN, node at which from the point of view of mobility mgmt old and new paths merge, e.g. MAPs in HMIPv6. Note in general: mobility CRN is may or may not be equal neither to NSLP CRN nor to NTLP CRN.

Routing CRN, node at which, using normal routing, old and new paths merge. In case of HMIP, mobility CRN is also routing CRN. However, in case of "normal" MIP with optimized routing, mobility mgmt doesn't know a CRN, whereas routing does. Depending on the location of nodes, the routing CRN may or may not be equal to the NSLP CRN or to NTLP CRN.

Path Update

The procedure for the re-establishment of NSIS state on the new path, the teardown of NSIS state on the old path, and the update of NSIS state on the common path due to route change or mobility. This is used to improve mobility handling for the affected flows.

Upstream Path Update: Path Update for the upstream signaling which is initiated by a signaling initiator on the common path

(e.g., a CN, a HA, or a GFA/MAP).

Downstream Path Update: Path Update for downstream signaling which is triggered by a signaling initiator on the new path (e.g., MN, mobile agent, or an AR).

Dead Peer Discovery (DPD)

The procedure for finding a dead NSIS peer due to a link or node failure and due to a mobile node moving away.

Downlink

The direction from the CN towards the mobile node.

Uplink

The direction from the mobile node towards the CN.

Receiver

The node in the network which is receiving the data packets in a flow.

Sender

The node in the network which is sending the data packets in a flow.

NSIS Transport-Layer Protocol (NTLP)

Description...

NSIS Signaling-Layer Protocol (NSLP)

Description

Downstream direction

Direction from a data source to the destination.

Upstream direction

Direction from data destination towards its source.

3. Framework

This section describes various mobility scenarios for the detailed discussions of mobility issues in NSIS signaling, using basic mobile IP (v4 and v6) handover as a starting point...

Our assumptions in this document and the framework are:

Session-ID is used to index state

Even if a mobile node has a permanent IP address (its home address), this cannot be used to index state in the network since the home address may not easily be visible to interior nodes. Other types of mobile nodes (e.g. using SIP or other application layer techniques) may not have permanent addresses at all. After a movement it obtains a new CoA, which is the basis for routing its data. If signaling-associated state is indexed based on some temporary data plane information, such as CoA, the state indexed by previous CoAs might be inaccessible for the signaling after most handover procedures.

Double state installation in the unchanged path should be avoided. This can only be done by establishing a relationship between the old and the new flow. This is essentially the same problem faced to tear down state in the old path.

Routing may be asymmetric

IP packets arriving to and leaving the MN may be routed differently. This may be due to the basic triangular routing of MIPv4, or due to the operation of an LMM protocol, or due to asymmetric routing caused by the basic operation of the IP routing protocols themselves.

The CN is not a mobile device

We may later add text to consider a mobile CN, too.

3.1.1. Global Mobility

Each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. In basic mobile IP, while situated away from its home, a mobile node is also associated with a care-of-address, which provides information about the mobile node's current location. After a mobile node registers its (primary) care-of address with a router on its home link, known as ``home agent'', data packets addressed to the mobile node's home address are routed to its care-of-address using IP-in-IP

encapsulation in the home agent (known as "mobile IP tunneling"). A mobile node can also register with the corresponding node, and then data packets addressed to the mobile node's home address are either IP-in-IP encapsulated (in mobile IPv4) or routed using routing header to its care-of-address (in mobile IPv6). Data packets sent by the mobile node to the corresponding node can be either directly routed to the destination's IP address, or routed through the home agent via a reverse tunnel (if reverse tunnel is used). Here, the following characteristics are important for support of NSIS signaling:

3.1.2. Other Global Mobility Approaches

Other approaches to provide global mobility work in the same generic way as MIP, but having a globally unique identifier for the MN, and then using an out-of-band signaling mechanism to provide the current location of the MN to the requesting party. For example, approaches based on DNS updates provide the globally unique DNS name, and the an up-to-date physical location of the MN. Another example is SIP, where users have globally unique identifiers, e.g. names, and the system is able to provide to the caller the current physical location of the called party.

3.1.3. Global Mobility Procedures

1) A flow associated with a mobile node, either sent or received by the mobile node, may continue to desire signaling services after a mobile IP handover. NSIS needs to be able to signal for such flows upon a mobile node's movements.

2) Either the sender or the receiver of a mobile node's flow can initialize an NSIS signaling. It is essential to require the NSIS signaling initiator to be authorized to initialize the signaling. Note that nodes within the network may also initiate NSIS signaling for the given session, for example, to handle route changes in the middle of the network, or to support seamless handovers.

3) The paths for a mobile node's outgoing traffic to the corresponding node and incoming traffic from the corresponding node may differ from each other.

4) Data traffic, in either direction between a mobile node and the corresponding node, can be routed directly, routed indirectly using a routing header, or IP-in-IP encapsulated, or the combination of them in different segments of the data transmission, depending on the mobility mode (route optimization or triangle routing; use reverse tunneling or not; mobile IPv4 or IPv6; whether LMM is used; etc.).

5) A mobile node's handover can be either intra-domain (inside one access network domain) or inter-domain (from one access network domain to another), which mainly concerns with topology information exchange, authorization and accounting issues. This is elaborated in Section 10.7 in [nsis-req].

6) A mobile node can support multiple care-of-addresses at one time, if it is connected to multiple access networks simultaneously. Although only one primary care-of-address will be used for routing traffic from the corresponding node to the mobile node, this multi-homing feature potentially can be used to enhance the NSIS signaling performance.

3.2. Local Mobility Management

Localized mobility management [lmm] mechanisms reduces the latency in mobility management signaling upon Care of Address change. These schemes, such as fast handover [fmip] and hierarchical mobile IPv6 [hmipv6], complicates the features identified in Section 3.1, for example, by associating new scoped care-of-addresses for a mobile node, and introducing one or more IP-in-IP encapsulated segment(s) in the path traversed by the communicating traffic. The additional CoA and IP-in-IP tunnels have implications for both the NTLP and NSLP. For example, NTLP needs to decide to perform tunnel handling when such tunnels exist in the same path that NTLP messages also traverse, while NSLP states may be updated according to the updated CoA in the localized domain. A discussion of these advanced characteristics is detailed in Section 11.

3.3. Context Transfer and Candidate Access Router Discovery

The NSIS protocol suite should be able to operate independently of Seamoby protocols such as Context Transfer Protocol (CTP) and

Candidate Access Router Discovery (CARD). Significant performance gains can be achieved if NSIS signaling can interact with such protocols.

When a mobile node has a choice of Candidate Access Routers (CARs) to perform handover, the Candidate Access Router Discovery (CARD) procedure can be used to identify those CARs along with their capabilities (for instance, QoS resource availability). In NSIS terminology, CARD can be used to find an appropriate NSLP-aware New AR (NAR), and the path leading to the node, where the NSLP state could be installed.

The Context Transfer Protocol (CTP) is used to transfer the NSLP state from the Previous AR (PAR) to NAR. When CTP is used, service-aware contexts could be quickly re-established in the NAR without requiring the MN to explicitly perform all protocol flows for those services from scratch. NSLP information, such as QoS state, can be transferred using CTP. However, this also concerns the NTLP because the context transfer for NSLP states takes place between PAR and NAR, which is vertical to the direction of normal NSIS signaling (which is between MN and CN).

With the help of CARD and CTP, NSIS signaling can quickly re-establish the NSLP state on the new path by reducing the state re-setup delay. However, making use of the CARD and CT protocols requires the ability from NTLP/NSLP to do (at that stage) off-path signaling on-behalf of the MN; this has implications on the authorization of signaling.

3.4. Examples of handovers

The discussions in this document focus on the effects a mobile end host can have on NSIS signaling. The fundamental concern in handover events is how the signaling can be localized in order to minimize the latency of setting up resources on a new path. Here one of the critical issues is the location and operation of a cross-over node. This section seeks to identify the various scenarios that emerge in different types of handovers and network setups.

There are several issues that affect how resources are set up in a

mobile environment:

1. Is the access router (AR) running
   a) NTLP,
   b) NTLP and the NSLP being signaled about, or
   c) neither of the two protocols?

2. Where are the
   a) NTLP CRN,
   b) NSLP CRN,
   c) mobility CRN, and
   d) routing CRN?

3. Does the interface at a given CRN routing towards the MN
   a) change after a handover, or
   b) remains the same?

4. Are the incoming and outgoing packets from the MN
   a) routed through the same routing path (symmetric), or
   b) through different paths (asymmetric routing)?

Note here, that the NSLP protocol can not be run without the transport part of the protocol, the NTLP.

Figure 3.1 presents possible setups in an access network that supports NTLP and NSLPs. It is assumed that all MNs in this draft are NSLP aware node. To provide examples of the issues mentioned above, consider the following scenarios:

- MN is connected to AR 1 and makes a handover to AR2. If incoming packets are arriving through Path A, the NTLP CRN is router a (RTRa) and its interface changes, and the NSLP CRN is RTRb and its interface does not change. If resources are set up for outgoing packets and the outgoing path changes to Path B due to the network routing table, resources must be set up on this new path, and possibly removed on the old Path A.

- MN is connected to AR 2 and makes a handover to AR 3. If packets were flowing, and will still continue to flow, on both directions, through Path B, the NTLP and NSLP nodes after the AR do not change, that is, the interface towards the MN at the NTLP and NSLP CRNs does not change.

- MN is connected to AR 8 and makes a handover to AR 7. If data was

flowing through Path F before the handover and now flows through Path E, the CRNs are not visible in the figure; RTRj is the first-hop NTLP and NSLP node from the MN and CRNs are somewhere else. If data flows through Path F before and after the handover, there is no NTLP or NSLP CRNs, as RTRm remains the first-hop NTLP and NSLP node.

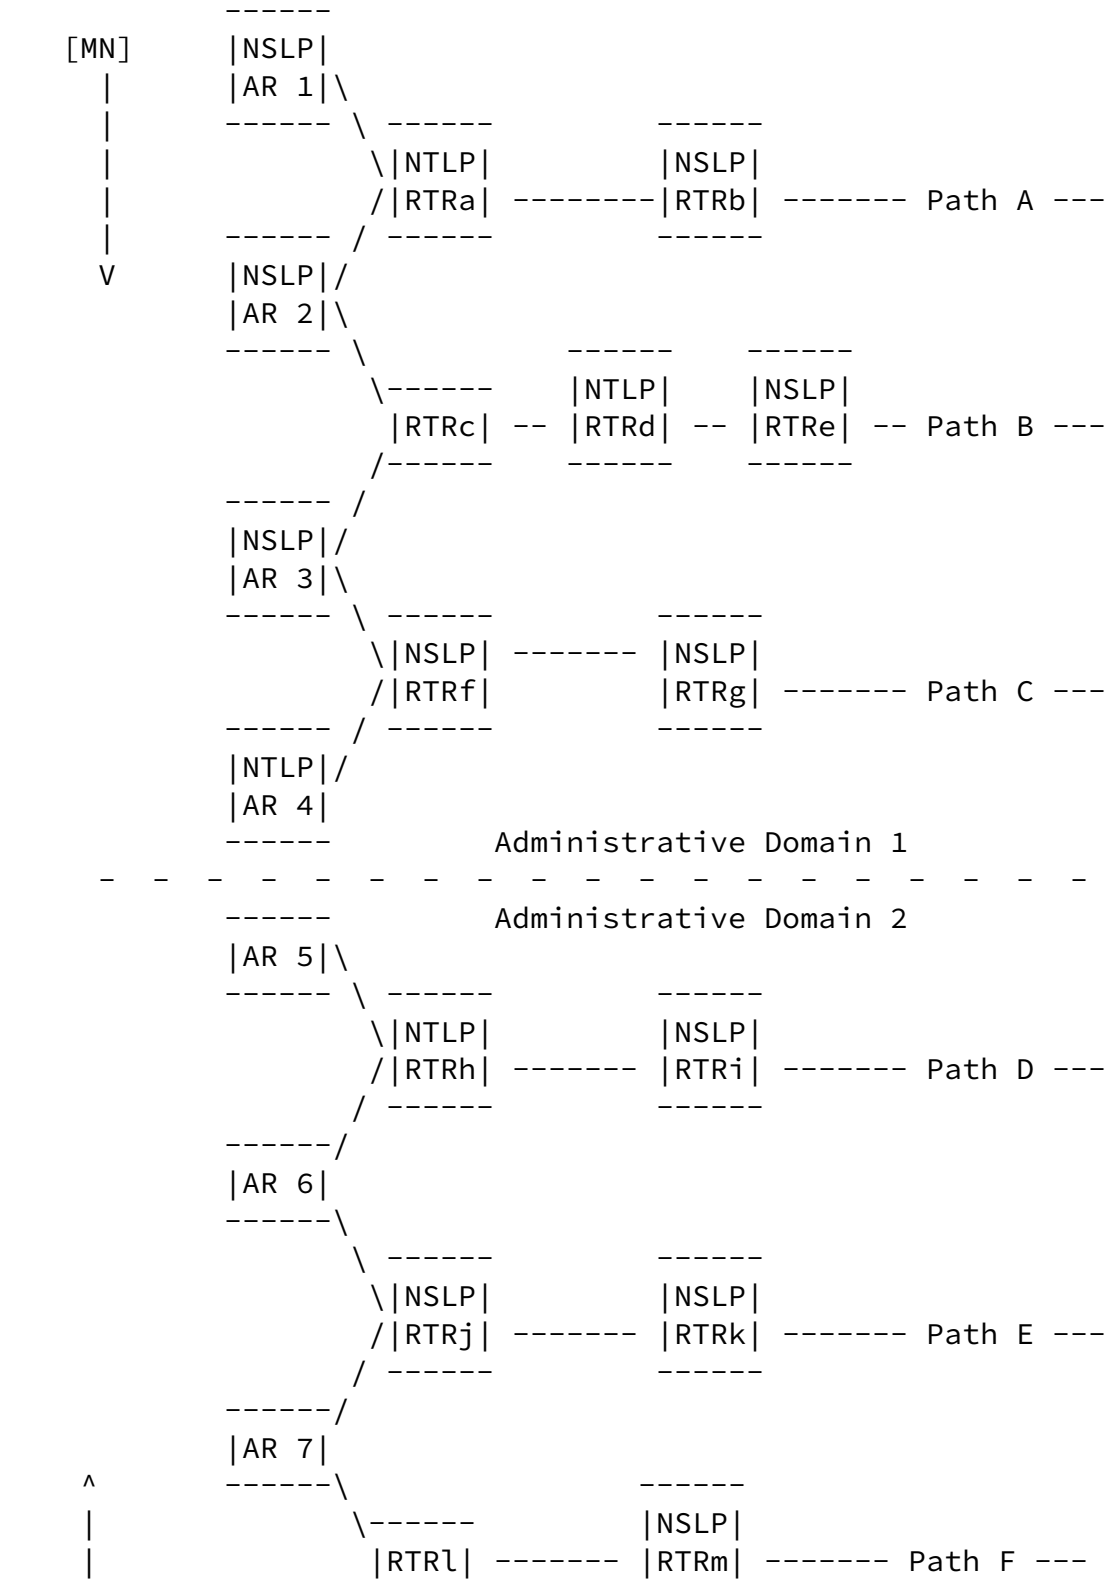
[MN]
|NSLP|
|AR 1|
|AR 2|
|AR 3|
|NTLP|
|AR 4|
Administrative Domain 1
Administrative Domain 2
|AR 5|
|AR 6|
|AR 7|
|NTLP|
|RTRa|
|NSLP|
|RTRb|
Path A
|RTRc|
|RTRd|
|RTRe|
Path B
|NSLP|
|RTRf|
|RTRg|
Path C
|RTRh|
|RTRi|
Path D
|RTRj|
|RTRk|
Path E
|RTRl|
|RTRm|
Path F

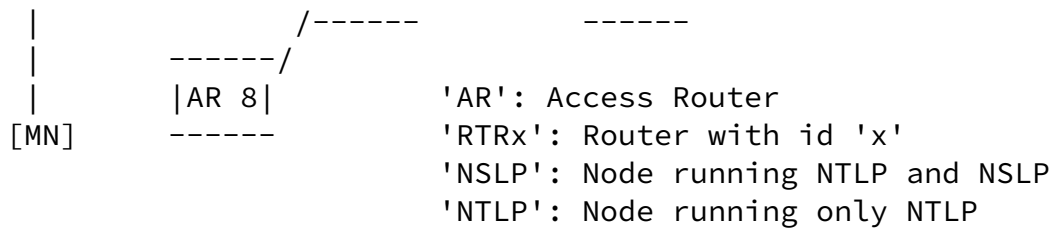

Figure 3.1: Examples of network architectures

3.5. Short Problem Statement

Based on the discussion in this section, we can highlight some critical issues that need to be solved to allow NSIS protocols to work in a mobile and wireless environment. There are a number of existing signaling protocols [nsis-analysis], however most of them do not support signaling in mobility scenarios, or the support is insufficient. RSVP [RFC2205], for example, relies on flow's fixed source and destination IP addresses, and ports information to identify signaling sessions and how signaling messages are routed. RSVP also lacks an overall consideration of mobile IP features.

An NSIS signaling protocol in mobility scenarios needs to consider the following issues:

o Change of route and possibly change of MN IP address

Topology changes entail the change of routes for data packets sent to or from the mobile node and may lead to a change of host IP addresses.

o Latency of route change

This change of routing and IP addresses is typically much faster than traditional route change (for example, those due to failure, adding or removal of nodes/links), which makes existing signaling protocols to be either unable to handle or at least create an additional overhead.

o Explicit routes

Moreover, MIPv6 adds the possibility to define explicit routers,

which creates further issues with routing paths.

o IP-in-IP encapsulation

The possible use of IP-in-IP encapsulation segments in the end-to-end path for routing traffic from the corresponding node to the mobile node (vice versa), associated with a route change of the same mobility behavior, makes signaling in mobility scenarios more complicated.

o Localization of state teardown

In the case of MN movement, only NSIS state between (former point of attachment of) MN and the NSIS CRN needs to be torn down: The NSIS state between CRN and CN may have to be updated e.g. to reflect a new Flow ID, but in any event it is still needed. Besides, tearing it down may necessitate re-authentication and re- authorization. Complication is added because the NSIS CRN is only known once the NSIS signaling along the new path is completed.

o Ping-pong type handover

In a ping-pong type handover, the MN returns to the previous AR after staying with the new AR only for a short while. On the one hand, NSIS should remove quickly in order to free resources. On the other hand, in the case of ping-pong type handover, state would need to be reestablished soon again, also adding overhead.

O Upstream local repair vs downstream local repair

Since upstream and downstream path need not be equal, upstream and downstream CRNs need not be equal, either. In fact, local repair needs to be handled independently for upstream and downstream flows, including, e.g. discovery of upstream and downstream CRN.

In summary, NSIS signaling needs to work with basic mobility which is an extension of general route (topology) change, and typically also includes IP address changes, supports mobile IP tunnels and multihoming. Path repair should be localized, and handled independently for upstream and downstream flows.

4. Cross-over Node Discovery and Path Update

This section discusses how to discover the crossover node (CRN) in general and the role of the CRN (especially in Path Update). This section also discusses how the NAT/FW NSLP affects the CRN discovery.

4.1. Background

Route change and main characteristics of mobile IP system described in Section 3.5 cause the session path to be changed, and therefore NSIS state needs to be re-established along the new path. In this case, the exiting protocols which do not support signaling in dynamic environment where route change or mobility event occurs can cause the following basic problems:

* Double reservation Problem

Since the state on the old path still remains as it is after re-establishing the state along the new path due to route change or mobility, the double reservation problem occurs. Although the existing state on the old path can be torn down by timeout of soft state, the refresh timer value in the core or wired network is quite long (e.g., 30 seconds in RSVP). As a result, in case of QoS-NSLP, the double reservation problem leads to the waste of resources and call blocking (especially in mobility scenarios). Therefore, the state along the old path should be removed immediately after state is set up on the new path.

* End-to-end signaling Problem

End-to-end signaling has to be considered as the issue-must-be-avoided in mobility scenarios, because it causes most problems related to service quality, signaling performance, and resource

availability. The route change caused by mobility, which complicates QoS signaling more, may result in the change of flow identifier. The change of flow identifier requires state update along the entire path to reflect the physical location of the MN: end-to-end signaling. This also incurs a long state setup delay, signaling overhead, and double reservation which affects network performance. Ultimately, the long state setup delay will particularly gives rise to the service blackout or degradation for multimedia application in the mobility environment.

To quickly re-establish NSIS state and improve scalability of NSIS signaling when route change and mobility event occurs, NSIS signaling

should be localized, and the localized signaling procedure is referred to as path update (see the terminology section). This is because, although NSLP messages are initiated by an MN or CN and sent to the opposite terminating point of the path, the path in the wireless access network usually changes only partially. Therefore, the NSLP/NTLP should limit the scope of signaling information to a local section of the signaling path.

The most appropriate node to do path update can be the CRN which is not a simple route splitting point, but an NSLP-aware merging point where the old and new session paths meet. To minimize the impact on seamless service, the CRN should be discovered by the NTLP as quickly as possible (see Section 4.2.2 for details), and afterward the involved NSLP should be triggered by the NTLP for necessary actions, for example, path update in case of QoS-NSLP.

4.2. Crossover Node Discovery

4.2.1. Types of Crossover Node

There can be various types of CRN according to normal route change, mobility management, signaling states, and flow direction. In the context of NSIS, this section mainly discusses the types of CRN according to NSIS signaling states and flow direction.

From the perspective of NSIS states (i.e., NSLP and NTLP sates), the types of CRN are basically divided as follows. First, from the NSLP's point of view, the CRN is a signaling application-aware node in the network where the signaling flows meet. Second, from the NTLP's point of view, it is a network node where more than one NTLP states (e.g., messaging association in [ntlp]). are stored. Although there can be various types of CRN according to state information, the CRN required for QoS-NSLP operation is NSLP CRN which has the corresponding signaling application information to perform the path update. Therefore, the CRN for the path update should be a logical NSLP-aware merging point rather than just a physical route splitting point. This implies that the CRN should not only be NTLP-aware but also NSLP-aware even if a CRN is detected by the NTLP. In this process, the NTLP should know the corresponding signaling application (e.g. QoS-NSLP) at the NSLP layer.

The path change of a specific NSLP session flows can be caused by either route change or mobility, so basically there are two different

types of merging point in the network according to the direction of signaling flows. The path change of downstream signaling flows may result in forming a downstream crossover node (DCRN) where the logical incoming interfaces start to converge, and the path change of upstream signaling flows may result in forming a upstream crossover node (UCRN) where the logical outgoing interface begins to diverge. Therefore, in general, the path change causes convergence or divergence of packets in data plane and/or in control plane (e.g., considering 3G/4G networks).

There are some differences between route change and mobility in forming DCRN and UCRN, and the asymmetric characteristics of routing adds complexity to the CRN discovery. When a generic route change occurs, the path change of signaling flows results in forming a chain of two CRNs, which is referred to as a divergence and convergence pair (see Section 4.4.1).

When a mobility event occurs, the asymmetric characteristics of routing between downstream and upstream directions can affect the location of the CRN. For instance, the handover of an MN (as an NI) will create a DCRN, and the handover of an MN (as an NR) will form a UCRN. However, the DCRN and the UCRN may be the same merging point in the network or may be different due to the asymmetric characteristics of routing although a CN is the same.

The CRN will be temporarily formed for path update, and how long the CRN will be involved in the path update depends on the period and method of re-establishing NSIS states in mobility scenarios. If sates, for example, are pre-established during handover to support multimedia applications seamlessly, candidate NEs can be ferreted out by interacting with Seamoby protocols. In this case, the candidate CRN(s) also is (are) discovered to localize the signaling to obtain performance gains in the network (see Section 4.2.2).

4.2.2. Determination of Crossover Node

A CRN can be discovered at both NTLP and NSLP layers. The CRN discovery at the NSLP layer can be done by NSLP signaling messages sent from the signaling initiator. For example, NSLP can realize it is a CRN by comparing the Source Identification Information (SII) contained in the incoming signaling message to that of previously stored in the node. However, in particular, the CRN would want to delete NTLP state when a particular NSLP is not supported there and NTLP state is not needed any more. Therefore, CRN discovery can be considered as an extension to the peer discovery at the NTLP level (e.g., using GIMPS query-response [ntlp]). In general, GIMPS message has message routing state information such as flow/session/signaling application identifier, so signaling application can be identified at the NTLP level. For example, in the connection mode of NTLP, when NTLP establishes messaging association between two adjacent peers,

the NTLP peers exchange message routing state information through

GIMPS query and response messages. Therefore, although CRN is discovered at the NTLP level, the discovered CRN is actually NSLP-aware node which has a involved signaling application.

There can also be two different approaches in CRN discovery according as whether the discovery is coupled with signaling message or not: Coupled approach and separated approach. In this case, the CRN discovery at each NSIS level depends on the used approach (see Section 6.1.4).

For CRN discovery, some session information such as the flow identifier and session identifier can be used. In addition, incoming/outgoing interfaces (e.g., Logical Interface Number: LIN) may also be used together with the session information. The CRN discovery can be further divided into UCRN discovery and DCRN discovery according to which node is a signaling initiator.

The session identifier in the GIMPS message is used to easily identify the involved session because it remains the same while the flow identifier may (or may not) change due to handover. The flow identifier is used to specify the relationship between the address information and the state re-establishment (e.g., QoS-NSLP state re-establishment). That is, the changed flow identifier indicates topological changes (i.e., old path, new path, and common path) and so the state re-establishment is required.

The logical interface number (LIN) can be used to establish or delete NSIS associations between peers. This identifier is also used to determine the CRN. NSIS entities may be able to use the interface number to locally distinguish each logical interface identifier between adjacent NTLP peers. Note that the LIN can be included in the NSIS message, but it can also be considered as an implementation issue.

In general, when a route change due to mobility occurs, CRN can be recognized by comparing the existing session information (e.g., the session and flow identifiers) with the session information included in the peer discovery message initiated by an NI (e.g., an MN or a CN) through a different LIN (e.g., an incoming/outgoing LIN). If the session identifier is still the same and the flow identifier and LIN has been changed, the current NSLP-aware node realizes it is the CRN. Note that the node which performs the CRN discovery should check whether the CRN has been discovered or not before realizing it is the

CRN.

Optionally, a mobility object can also be used to indicate that the MN has experienced a handover and a route change has occurred [Jeong01] [Lee01]. In this case, the NSIS protocol (or node) may need to interact with mobility protocols to detect the CRN immediately. For example, the CRN discovery may need to be triggered in parallel with the transmission of the binding update (BU) message (of MIP).

The mobility object may be defined in the NTLP message (e.g., GIMPS payload) to notify any mobility event explicitly, and it contains

various mobility-related fields such as handover_init field and mobility_event_counter field. The handover_init field can be used to explicitly inform that a handover is initiated for fast state re-establishment. The mobility_event_counter field can be used to detect the latest handover event to avoid confusion about where to send a confirmation message which indicates that the CRN has been found.

This type of confirmation may be needed when the MN moves toward the second new AR immediately after it undergoes a handover to the first new AR from the old AR, because the CRN discovery message from the second new AR may arrive earlier than from the first new AR. The mobility object may also be defined in the NSLP in a similar way. In this case, there should be some relationship between the mobility objects of the NTLP and the NSLP.

If an MN is an NI when a route change due to mobility occurs, the MN begins to transmit signaling messages toward a CN in the downstream direction. In this case, an NSLP-aware node recognizes that the session paths converge, and then this node performs the comparison of session information checking the incoming LIN. After determining that the CRN has not been discovered yet, the NSLP-aware node realizes it is the DCRN.

When an MN (as an NR) undergoes handover, the UCRN can be determined by checking the outgoing LIN of signaling flow from a CN. In this case, the UCRN should be the first node where the signaling flow begins to diverge. Since UCRN is determined by whether outgoing path diverges or not, the UCRN discovery is more complex than the DCRN discovery. If NSIS operates with HMIPv6 and an MAP is an NSIS-aware node, the UCRN can be locally discovered in an access network by the method above. If the UCRN is discovered between the MN and the MAP, the path update can be actually localized for upstream flows. However, note that when interworking with HMIPv6, it is still an open

question how these nodes decide locally whether they are indeed the UCRN.

The CRN discovery may also be initiated during handover (i.e., before handover is completed). However, in this case, a more efficient mechanism is needed to find a candidate CRN. For example, after a mobility event is detected by the NTLP, the current AR may use CARD to transfer the context for fast QoS-NSLP state re-establishment. After the candidate AR is found, CTP can be used to transfer the context which includes the QoS-NSLP session information for fast QoS-NSLP state re-establishment. If an appropriate AR is found and the context transfer is completed, a candidate CRN can be discovered easily since the candidate CRN discovery is basically the same as above.

In some cases, however, it may not be always possible to use mobility-related protocols such as CT and CARD. In this case, the MN can initiate the CRN discovery only after it changes the point of attachment. To expedite the discovery process, it may be useful to transmit the peer discovery message (by the NTLP) and the first binding update message at the same time.

4.3. Path update

This section discusses possible procedures for path update according to the direction of signaling flows. As discussed in Section 4.1, the CRN can be a crucial point for path update, since the CRN is the NSLP-aware merging point of the old and new paths. From the perspective of path update, the CRN plays the role of initiating re-installation on the new path, teardown on the old path, and update of NSIS state on the common path.

In mobility scenarios, the flow identifier for NSIS signaling may Change. Since the flow identifier is used to identify the signaling state installed along the path, the procedures for path update should include state update along the entire path to reflect the topological change of the MN. The CRN discovery is different according to the direction of signaling flow in mobility scenarios, and the DCRN operates independently of the UCRN although DCRN and UCRN can be simultaneously ferreted in bi-directional state establishment. Therefore, the procedures for path update may differ according to the direction of signaling flows. For downstream signaling, path update is triggered by the MN (or mobile agent) or an AR, which is referred to as Downstream Path Update. For the upstream signaling, path update is initiated by a CN, a HA, or a GFA/MAP, which is referred to as

Upstream Path Update. In this case, each signaling initiator has to be authorized for secure signaling.

4.3.1. State Setup and Update

In both types of path update, NSIS protocol needs to interact with mobility signaling and the Seamoby protocols (during or posterior handover) to obtain performance gains through fast re-establishment of the NSIS states along the new path. In this case, NSIS needs to monitor for detecting the movement through several methods [nsis-fw]. After detecting a mobility event, the NSIS protocol can check resource availability on the new path (or new candidate path) through CARD or other mechanisms during handover in order to check the possibility of state re-establishment on the new path in advance.

In the downstream path update, if resource availability is assured, an MN initiates the NSIS signaling for state setup toward a CN along the new path. The DCRN discovery is implicitly done by this type of signaling initiated by the MN. In this case, the node where old and new logical session paths converge realizes that it is the DCRN, and afterward the DCRN sends a response message toward the MN to notify of NSLP state installed (e.g., in QoS-NSLP) or installs the NSLP states as responding the NSLP signaling initiated by the MN (e.g., as in RSVP). In the downstream path update, the sender-initiated approach (e.g., QoS-NSLP) leads to faster setup than the receiver-initiated approach as RSVP. And then, the DCRN sends a refresh message toward the signaling destination to update the changed flow identifier on the common path and also sends a teardown message toward the old AR to delete the NSIS states along the obsolete path.

In the case of upstream path update, the CN (or a HA/ a GFA/MAP) sends a refresh message toward the MN to perform path update. UCRN is discovered implicitly by the CN-initiated signaling along the shared path, and the node from which the common path begins to diverge into the old and new logical session paths realizes that it is the UCRN. In this case, the CN should be informed of the movement event using an NSIS signaling message sent by the MN or monitoring the mobility signaling. After the UCRN is determined as described in Section 4.2.2, it may send a refresh message to the MN along the new path while establishing the NSIS association between the updated peers, and afterward the UCRN may send a teardown message toward the old AR to delete the NSIS state on the obsolete path.

The state update in control plane on the shared/common path to reflect the changed flow identifier brings issues on the end-to-end signaling. Although the state update does not give rise to re-process AAA and admission control, it may lead to the signaling overhead. If NSIS protocol interacts with Hierarchical Mobile IPv6 scheme, the NSIS session only has the changed flow identifier between an MAP/GFA and an MN. However, whether the update of the flow identifier for the session can be considered only between an MN and an MAP to avoid end-to-end signaling is still an open issue.

One of the goals of path update is to avoid double reservations (in case of QoS signaling) on the shared path described in Section 4.1. The double reservation problem may be solved by establishing a signaling association using the unique session identifier. That is, NSLP state can be shared even if different flow identifiers changes. For example, QoS-NSLP state (for resource reservation) can be used by packets for either flow.

4.3.2. State Teardown

After re-establishment of the NSIS state along the new path, the state on the obsolete path should be quickly removed by path update mechanism to prevent the waste of resources due to double reservation (and resource allocation problem by call blocking) and to reduce the cost of using the resources in the access network as described in Section 4.1. Although the release of the existing state on the old path can be accomplished by timeout of soft state, the refresh timer value may be quite long and the maintenance of the NSIS state on the obsolete path may not be necessary. Therefore, the transmission of a teardown message is particularly preferred to the use of refresh timer to quickly delete the old state.

The CRN is an appropriate point to initiate the teardown toward the old AR after re-establishment of the state along the new path. In this case, the release of old state on the obsolete path can be accomplished by comparing outgoing LINs and through reverse routing using SII. This can prevent the teardown message from being forwarding toward along the common path. However, whether the teardown message can be sent toward the opposite direction to the state initiating node is still an open question. This also leads to

authorization problem because a node which does not initiate signaling for establishing the NSIS state can delete the state.

To avoid the waste of resources, the resources on the old path should be removed as soon as possible after re-establishing the state along the new path. However, this may not be appropriate for fast handover of a ping-pong type where an MN may return to the previous AR after staying at a new AR for a short while. When to delete the state along the obsolete path remains still an open issue.

If the old AR is the last node due to handover, its NSLP may trigger an error message to indicate that NSLP messages cannot be forwarded any further. This error message may cause the removal of the old states. However, although the error message is initiated, the state on the old path should not be deleted before re-establishing the state along the new path. This issue can be solved by using the handover_init field of mobility object mentioned in Section 4.2.2. When an MN, for example, detects a handover, the QoS-NSLP of the MN constitutes the MOBILITY object (the handover_init field) in the QoS-NSLP signaling message and send it to the current AR (the old AR), which prevent the current AR from initiating the error message indicating the dealt with more detail in Section 5.

4.4. Effect of generic routing changes

4.4.1. CRN discovery

In case of generic route change, the CRN can be a node which detects the change of a data flow in the network. When the downstream or upstream data flow begins to travel, a node can detect the route change by interacting with NSIS, routing protocol, and detection method based on network monitoring, data packet monitoring, or signaling message monitoring [nsis-fw]. The node detecting the route change starts to discover the next peer via the NTLP peer discovery message exchanges and continues to do the peer discovery until discovering a node which already has the involved NSLP states. Whenever a new peer is discovered, NSIS creates as an association with the previous peer using the LINs.

In case of downstream data flow, the first NSLP-aware node where the signaling flow starts to diverge can be considered as a diverging DCRN, and the first NSLP-aware node where the signaling flow begins to converge can be identified as a converging DCRN. As mentioned in Section 4.2.1, the route change of signaling flows results in forming a chain of divergence and convergence CRN pair in the network. For upstream signaling flow, the first NSLP-aware node where the signaling flow diverges can be considered as a diverging UCRN, and the first NSLP-aware node where the signaling flow converges can be identified as a converging UCRN. However, How an NSLP-aware node identified itself whether the first node which converges and diverges is CRN is still an open question.

4.4.2. Path Update

In generic route change, since the flow identifier does not change, state update along the common path is not performed. Therefore, state re-establishment along the new path and teardown along the old path are only carried out. There is also no difference between downstream signaling and upstream signaling compared to mobility scenarios because the diverging CRN should interact with the converging CRN for each signaling flow.

In downstream path update, the diverging and converging DCRN pair is discovered after route change as described in Section 4.4.1. In this case, the diverging DCRN initiates signaling to establish NSLP states on the new path toward the converging DCRN by sending the RESERVE message [QoS-NSLP]. Note that in the coupled approach, peer discovery is done simultaneously with state re-establishment (see Section 6.1.4), and so a diverging node and a converging node are implicitly identified as DCRN.

If each node between the diverging DCRN and the converging DCRN can not delete their NSLP state for itself (i.e., refresh timer), the converging DCRN can trigger the removal of the obsolete state by interworking with the diverging DCRN. Therefore, the converging DCRN begins to delete the NSIS states on the obsolete path in the reverse direction (e.g., toward the diverging DCRN) after installing state on the new path. In this case, the diverging DCRN should be able to identify that the teardown message (e.g., RESERVE message in [QoS-NSLP]) from the converging DCRN should not be delivered beyond the diverging DCRN. For this purpose, the teardown message may have a "Path Update (PU)" flag in its header field, or the destination address of the teardown message may be that of the diverging DCRN, and converging DCRN should know the reverse routing information to send the teardown message toward diverging DCRN (e.g., using SII in [QoS-NSLP]). For example, the diverging DCRN can prevent the teardown message from being forwarded toward a sender by discerning the í98PUí98 flag. However, whether the teardown message can be sent toward the opposite direction to the original state initiator is still an open question. This also leads to authorization problem because a node which does not initiate signaling for establishing the NSIS state may delete the state.

For the upstream path update, the divergence and convergence UCRN pair also follows the same procedure as above.

4.5. Open Issues

There are some open issues that should be discussed in the later version of this document, and they are summarized as follows.

- In the Interworking with HMIPv6, how can the nodes decide locally whether they are indeed the UCRN?

- Can the update of the flow identifier for the session when interworking with HMIPv6 be considered only between an MN

and an MAP to avoid end-to-end signaling?

- Can the teardown message be sent toward the opposite direction of the state initiator?

- When is the right time to delete the state along the obsolete path for fast handover of a ping-pong type?

- How can the crossover node be discovered in the specific multicasting/multihoming cases?

- How does the NAT/FW NSLP affect the CRN discovery?

5. Dead Peer Discovery (DPD)

5.1. Overview

A dead peer can occur either because a link or a network node, including the MN and CRN, failed, or because the mobile node moved away without informing NSLP/NTLP (it is recommended to link mobility- and nsis signaling such that this does not happen). Hence, DPD is the fall-back mechanism for dealing with mobility which is not currently hooked into the NSIS protocol suite.

The procedures for handling DPD should be the same no matter why a peer is dead, because an NE discovering a dead peer cannot judge the specific reason. That is, DPD due to a link or node failure, and DPD due to an MN moving away should trigger the same reaction. In any case, dead peers should be discovered as soon as possible to minimize service interruption. Subsequently, NSIS needs to find a different path interacting with the routing protocol. Thereby, NSIS needs to take into account the possibility that no path to the dead peer

exists. Once the new path is found, NSIS state needs to be set up along the new path, and NSIS state needs to be torn down along the old path. However, care must be taken to terminate teardown at the CRN since the NSIS state on the common path should not be deleted.

5.2. Failure Cases and Impact of Dead Peers

5.2.1. Failure Cases

Dead peers of interest in mobility scenarios include CRN, MN, and AR. In general, it is possible that only NSIS functions (i.e., NTLP/NSLP) of the node may fail, or the physical hardware.

As mentioned above, an MN may either fail or move. When it fails, it becomes a dead peer. When it moves, it either changes or retains its IP address (e.g., CoA). If it moves and changes its IP address without notifying NSLP/NTLP, it also becomes a dead peer. If it moves and keeps its IP address, we need to solve a rerouting problem rather

than a dead peer problem.

5.2.2. Impact of Dead Peers

The failure of a (potential) NSIS CRN may result in incomplete state re- establishment on the new path and incomplete teardown of the old path after handover. In this case, a new CRN should be discovered immediately by the CRN discovery mechanism described in Section 4.

The failure or movement of an MN may cause the 'invalid NR' problem [draft- lee-nsis-nslp-mobility-01.txt] where the NR is the MN. [the following text could be added for clarification: If the MN moves, an error message, e.g., can-not- be-forwarded-further, should be generated by the MN, since this message may prevent the teardown of NSIS state on the old path before NSIS state is re-established on the new path]. We may need to also consider the case where the MN is not the NR, but a router in the access network (possibly the AR) is proxying for the MN instead.

If the MN moves without changing its IP address, usually this is a micro mobility scenario. Two basic ways for handling micro mobility are currently used:

- By source node routing [HMIPv6] towards the MN, i.e. coding the new route explicitly in each packet. It is difficult to do nsis-signaling for such a scenario, except by also source-node routing signaling messages.

- By changing the so-called Regional CoA, which is not visible outside the micro mobility region. Packets destined to the MN are always addressed to the Mobility CRN. The Mobility CRN tunnels the packet to the MN [HMIPv4, HMIPv6]. Mobility in this case is not noticed by NSIS, because NSIS signaling is tunneled the same way as data packets. A separate NSIS state needs to be set up for the tunnel, and the NSIS state for the old tunnel needs to be torn down.

Micro mobility with unchanged IP address is also handled in ad-hoc routing protocols in which per-host routing entries are changed in the routing tables. Hence in this case, mobility results in rerouting, just as when an intermediate node or link fails.

If the MN moves with a changed IP address, the MN reappears somewhere else and tries to set up NSIS state along the new path. The requirements derived from this scenario contradict those derived from a true MN failure, where the MN does not reappear:

- In the case of MN movement, teardown of NSIS state should be terminated at the NSIS CRN (cf Sec. 3.5) However, the NSIS CRN is only known once the NSIS signaling along the new path is completed. Therefore, state along the new path needs to be established first, and only then the old state should be torn down. (See also discussion in Sec. 4).

In contrast, in case of MN failure, NSIS state should be removed along the entire path as quickly as possible. A CRN does not exist.

Recall it is impossible for the NE discovering a dead peer to distinguish these two cases. We therefore need to settle on a single mechanism for handling both.

The failure of an AR may make the interactions with Seamoby protocols (such as CARD and CT) impossible. In this case, the neighboring peer closest to the dead AR may need to interact with CARD and CT.

5.3. DPD procedures in NTLP

The procedures of how to do DPD should be handled by the NTLP. In fact, the DPD can be considered as an extension to the GIMPS peer discovery. The transmission of peer discovery messages may be separated from the transmission of regular signaling messages. It is also possible to combine both types of messages for efficiency in message delivery. For example, the detection of an NSIS peer and the establishment of an NSIS state can be performed using an NSIS message at the same time.

There are cases where an NE does not deliver signaling messages successfully to its NSIS peer along the signaling path, for example, when an NF (or NR) was disconnected from the network due to one of the failures described above, causing a change of signaling path in the network. Such dead peers which are no longer reachable should be detected. Some possible DPD procedures are described below.

A peer discovery message can be periodically transmitted to the neighboring peer (e.g., responding node in [GIMPS]), and the responding node can send a response message. To determine if the peer is alive, the use of a timer may be helpful. For example, the response message may need to be received by the sender (e.g., querying node in [GIMPS]) of the peer discovery message before the timer expires. Otherwise, the responding node can be considered dead.

It is important to check the validity of the peer discovery messages for security protection. For example, it may be necessary to determine if the peer discovery message has been received from the authorized peer. Cookies such as query-cookie and response-cookie [GIMPS] may be useful for this purpose.

According to the [GIMPS], the NTLP itself does not provide for teardown of NTLP state because, as opposed to NSLP state, it is not very expensive. NTLP instead relies on time-out. Upon DPD, NTLP informs the local NSLPs about it, and may even send a notification to other NTLP peers upstream to inform other NSLPs which it does not support locally. It is an open question when to stop propagating this information [GIMPS], which is not specific to mobility.

Local NSLPs (e.g., QoS-NSLP) could either initiate a teardown of the corresponding NSLP state upstream, i.e., in the direction opposite to

the dead peer (possibly accelerated expiration as described in [GIMPS] if the node is not authorized to do this) or send a notification upstream which might result in the NI to take action.

Actually in [QoS-NSLP] it is not fixed yet what must happen. A dead peer may lead to rerouting, or sometimes, as sometimes in the case of a dead NR, no new route being discovered. Rerouting may be noticed by more than one NE using one of the detection mechanisms described in [GIMPS]. Furthermore, more than one NE may be able to reroute around the situation (see Fig. 8 in [GIMPS]). The NE closest to the flow sender should become the NSLP CRN.

Finding the new path and establishing state can be done as described in Section 4. The relative timing of state teardown and re-establishment is still an open question as discussed in Section 5.2.2.

6. Case Examples

The movement of end-hosts leads to changes in the data path due to the change of their point of attachment in the network. This results in the original data path between a sender and a receiver to be divided into three paths, all of which intersect at a CRN: the unchanged path from the CN until the crossover router, the new path from the crossover router until the new location of the MN and the old path from the crossover router until the old location of the MN.

Due to rerouting of data packets after handovers, signaling-associated states need to be updated or removed. This concerns with which information is needed for indexing states and where and when a creation, update or removal of these states is required. If signaling-associated state is indexed based on flow-Ids, the state indexed by the flow-ID referred to the old path might be inaccessible for the signaling after most handover procedures. Hence, it is assumed that signaling-associated state is indexed by session-IDs.

This section provides concrete examples of the signaling done in a handover situation.

6.1. NSIS in the hard-handover case

This example is called hard-handover, or break-before-make handover, in which the NSIS signaling, required to update the MN path, happens only after the MN is attached to the new access point.

To update the path between the CN and the MN, state needs to be installed in the new path, released from the old one and updated in the unchanged path. The NSIS signaling required for this operation may be triggered by the mobile node, mobility agent(s), or by the access router at which the mobile node is attached to. In any situation, the CRN is the starting point or finish point of the NSIS signaling messages.

Since we have to assume that the MN can be a sender as well as a receiver, the first difficulty to find a CRN is posed by the asymmetric characteristic of routing, as mentioned in Section 4.2. Due to routing asymmetry there is no reason for the UCRN and the DCRN to be the same in the, even for the same correspondent node. Therefore the NSIS signaling sequence required to update the path depends on the role of the MN, which is the one detecting the attachment to a new network.

Figures 6.1 and 6.2 illustrate the signaling needed to update the MN path in the upstream and downstream direction when the MN is a sender or a receiver. In both figures, the "r", "s", and "u" indicate NSIS messages to remove state in the old path, set state in the new path and update state in the unchanged path. The "t" in Figure 2 represents the triggering message that the MN sends to the CN. This triggering message can be for instance a mobility-binding message.

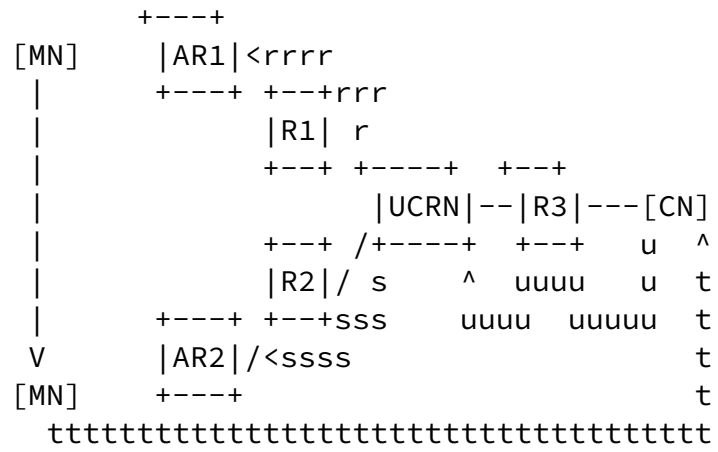

Figure 6.1: MN as a receiver

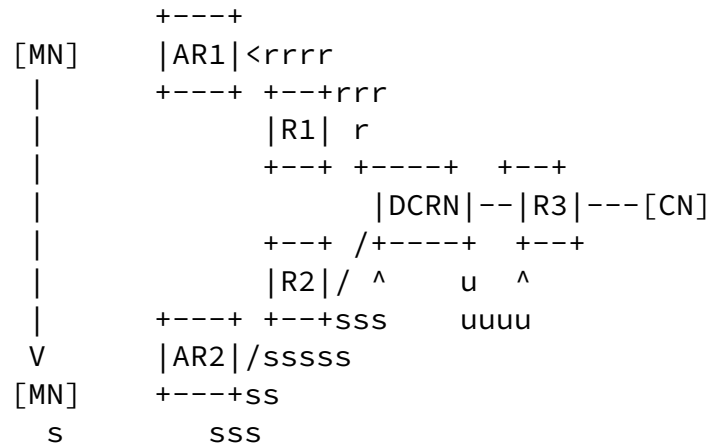

sssssssss

Figure 6.2: MN as a sender

When the MN is the sender, it is assumed that it is the MN that starts signaling over the new path to discover the CRN. Also, it can be considered that it is the CRN that starts removing state in the old path and updating state in the unchanged path.

When the MN is the receiver, a general assumption would be that it is

the CN that starts signaling over the unchanged path to discover the CRN. The CN can start signaling for instance after receiving a binding update message or after detecting a change in its binding entry. However, if local mobility management mechanisms are used for mobility, it can be also possible that mobility agents such as the handover manager or the mobility anchor point can start the update of the MN path. Also, it can be considered that it is the CRN that starts signaling to update the old and new path.

A path update as the one used by RSVP [RFC2205], where the update may be triggered by a local route change, is impossible or at least difficult in any of the two described situations. The reason is that, the node experiencing a route change can only be the MN, CN, HA, or other mobility agents, which are not necessarily the CRN. Moreover, it is the CoA that changes and not the route to the old CoA.

6.1.1. Signaling on the Unchanged Path

The network control state on the unchanged path must be updated to reflect new flow identification, if the flow-ID contains a CoA. This leads to the problem of requiring end-to-end signaling, which should be avoided to decrease the control load overhead. However, it should be possible to avoid AAA and admission control processing.

6.1.2. Signaling on the New Path

Updating state in a new path may be conditioned by the session ownership and the availability of resources. In the latter case, when the network is overloaded, it is preferable to keep state belonging to previously established flows while blocking new requests. Therefore, signaling to update mobile sessions should have priority over local requests for resources.

6.1.3. Signaling on the Old Path

The signaling to release resources over the old path should be done as soon as possible to avoid wasting resources. Typically, the changed path is located inside an access network, where resources are relatively expensive, thus it might be inefficient to wait for typical soft-state timeouts. However, immediately releasing resources along the old path might cause problems. In case of a ping-pong type of movement, the immediately release of state in the old path can have a performance impact higher than the cost of keeping that state. This higher performance impact happens, because resources along the old path might be reused after a very short time period. This means that the MN may return to the previous access network shortly after leaving it, which brings some problems about deciding when to release state in the old path.

6.1.4. Interaction between NTLP and NSLP Signaling

When the MN is the sender, there is a tight relationship between signaling to discover the CRN and signaling to update the new path. When the MN is the receiver, signaling to discover the CRN is tightly related with signaling to update the unchanged path. This means that in each one of these situations, signaling to update the new path or the unchanged path can be done simultaneously with the signaling needed to discover the CRN (coupled approach) or can be done after the CRN discovery process (uncoupled approach).

Although all NEs in a new path have to be discovered by the NTLP peer discovery mechanism, the interaction between NTLP and NSLP signaling to discover the CRN (NTLP and NSLP) and to update state depends upon the used approach. It is assumed that, even though NSIS messages follow standard IP routing, different NSLPs can have different NSLP CRNs and that the NTLP CRN can be NSLP-unaware. However, it is assumed that each NSLP CRN is a NTLP node.

In the coupled approach, both types of CRNs (NTLP and NSLPs) are discovered by using the same signaling message. In this case the NTLP message transports different NSLP objects, and in the discovered NTLP nodes the appropriate NSLP state will be updated/set (NTLP trigger the corresponding NSLP, if any). One advantage of the coupled

approach is that it encompasses few update time.

In the uncoupled approach, a NTLP message is sent to discover the NTLP CRN and after that a different NSLP message is sent to discover the NSLP CRN for each type of NSLP. This approach utilizes additional NSLP-based signaling, and NTLP is essentially used as a transparent underlying mechanism to transport NSLP messages. Therefore, NTLP should be able to notify 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. One advantage of the uncoupled approach is to allow the NSLP signaling to be timely independent from the NTLP one, allowing a more flexible management of the different NSLPs.

6.1.5. Routing of NTLP messages

As stated in the NSIS framework document [nsis-fw], there are two ways to address a signaling message being transmitted between NEs:

o Peer-to-peer, where the message is addressed to a neighboring NE that is known to be closer to the destination NE.

o End-to-end, where the message is addressed to the flow destination directly, and intercepted by an intervening NE.

The peer-to-peer signaling is called connection mode in the definition of GIMPS [ntlp], while the end-to-end signaling is called datagram mode in the same document.

Each one of these types of messages is necessary for some aspects of the NTLP operation. In particular, initial discovery of the next peer in a new path will usually require end-to-end addressing, whereas reverse routing, signaling on the old path and on the unchanged path can be done based on peer-peer addressing. The latter case is possible since NEs in such paths should have already messaging associations. The used mode is not visible to the NSLP, and the information needed in each case is available from the flow-ID or locally stored as NTLP state.

End-to-end routing of signaling messages at NTLP level can be based on the flow-ID, while the peer-to-peer signaling can be based on the session-ID. In a flow-ID-based approach, NTLP has to rely on a mapping between certain fields of the flow-ID, e.g., destination IP

address and additional IP header information, and local IP routing table. In a session-ID-based approach, NTLP can route the signaling messages based on a mapping between the session-ID and the local NTLP-level state, such as the address of the next/previous NE. If there is no existing state, a next-peer discovery has to be performed to create such a state.

As the association of different flow-IDs to a single session-ID is a problem common to many signaling applications, the association between both identifiers might be done at the NTLP. However, this association could also be done at the NSLP layer, if the method used to perform such association is specific to each application. In either case, it is assumed that the session-ID should be visible within the NTLP, allowing it to perform an enhanced forwarding control for packets belonging to that session.

Another three related identifiers, namely the message identifier(message-ID) introduced in RFC 2961 [RFC2961], the branch identifier (branch-ID) suggested in CASP [I-D.schulzrinne-nsis-casp][fu03], and the Reservation Sequence Number (RSN) proposed in QoS NSLP [I-D.ietf-nsis-qos-nslp], have been also discussed as potential mechanisms useful for mobility support in NSIS. All of three indicate the order in which corresponding signaling messages are processed by the corresponding signaling entities (RSVP, CASP-NTLP and QoS-NSLP daemons, respectively) and try to address the out-of-order problems of signaling messages.

Message-ID, together with Epoch object in RFC 2961, concerns with signaling messages between peering neighbors, where the out-of-order problem can come from retransmission/refresh. It was not designed for mobility support specifically. As an extension to message-ID concept, the branch ID can be used for detecting out-of-order signaling messages along different branches each of which can consist of multiple hops. It can be useful to avoid explicit teardown messages from being forwarded on the unchanged path. Different from the branch ID, the RSN is meaningful in a QoS-NSLP node for protecting out-of-order problems in each branch, which can consist of multiple QoS NEs.

6.2. Example of Signaling of an Anticipated Handover

The term "seamless mobility" is often referred to mean that the MN is able to keep an ongoing session seamlessly (without experiencing

perceivable service interruption or performance penalty) during and after moving from one access network to another. Measures to achieve seamless mobility include soft handover and anticipated handover. The former requires the MN to keep the old path, while data is received over the new path. This approach is only possible if the MN is multihomed. The present section discusses fast state installation by using anticipated handovers, in which the MN signals the new path while still connected to the old one.

With an anticipated handover, state in the new path can be set in advance, which means before the MN gets any layer 3 connectivity to the new access router. Anticipated handovers require the discovery of candidate access points or access routers (CARD may help, cf. section 3.3), and the ability of the MN to trigger the signaling to set the new path over the old access router. However, the new path up to/from the CRN must be also figured out which is especially not that easy in case the MN is acting as receiver. An anticipated handover checks resource availability along a potential new path before an MN actually changes its point of attachment. Therefore, if there are not enough resources available along the new path an unsuccessful handover (or period of QoS degradation) can be avoided and the mobile node can stay connected to its current point of attachment (if possible).

On the one hand anticipated handover offers mainly two advantages:

- reducing seamless handover latency, because most signaling to set resource in the new path is carried out in advance

- avoiding unsuccessful handovers or unnecessary periods of QoS degradation

The first point may be especially important in inter-domain handover scenarios where signaling procedures will take longer. On the other hand, the higher resource consumption may be considered as disadvantage if resources along the new path could be reserved successfully, but the old path is still used (dual reservation after split or merge CRN). However, as already mentioned in section 6.1.3 an indication for handover completion (thereby triggering release of the obsolete path) may help to keep this period as short as possible.

To perform an anticipated handover, MNs do not have to be multihomed. However, anticipated handovers may involve some kind of NSIS proxy [ntlp] on the new access network to signal on the new path on behalf of the MN. If we assume that the end-to-edge communication is done between the MN and its access router, some study is required to determine how to signal between the MN currently access router and the NSIS proxy in the new access network, e.g., how to discover the most suitable NSIS proxy, and to establish a communication between access networks. The latter issue involves out-of-path signaling.

Moreover, in some anticipated signaling scenarios, NSIS signaling cannot be triggered by the mobility protocol, which required some study about other possible triggers, such as:

- Cross-layer triggering. For instance, the layer-2 mechanism can give some information about a possible movement.

- Context-awareness triggering. For instance, information about a lower traffic load in some neighbor access networks can trigger the establishment of state in a new path.

7. Multihoming-related Issues

[This will include discussions about how multihoming affect NSIS signaling.]

[JM: How would SCTP and multiple associations between communicating parties actually work with NTLP/NSLP?]

8. Interactions with Mobility Signaling

This section discusses the interactions between NTLP/NSLP signaling and various mobility-related protocol, including MIPv4, MIPv6, various LMM protocols, CTP, and CARD.

8.1. Mobility Management Protocols

Basically, NSIS signaling is assumed to be performed after mobility protocols take place, i.e., completing a handover. This involves a number of issues:

- whether NTLP or NSLP should have an interface with the mobility protocols.

- which information needs to be obtained from mobility protocols. As routing of NSIS messages is handled by the NTLP, NTLP should interact with mobility protocol for the IP address change, start/completion of a binding process and mobile IP tunnel information.

- whether to use traditional routing/NSIS interface to trigger NSIS signaling, or use the start/completion information of binding

processes which changes the characteristics of flows (e.g., in CN, MN or HA, as explained below). If NSIS relies on (typically) after seconds of routing change detection of a routing/NSIS interface mechanism to obtain route change and tunnel change information, it can be less processing-intensive, but the time for reacting and behaving in mobility scenarios is a concern. Using start/completion information of binding processes allows faster state recovery and removal, but its disadvantage is that fast or ping-pong movements may result in considerable signaling overhead and possible errors.

- how to coordinate several entities which involved with changing of flow characteristics (e.g., the binding processed in MN, CN, HA, ...).

- how to coordinate the mobility binding update interval and NSIS signaling interval. Mobility bindings can take place periodically even for the MN with the same point of attachment. This concerns with the signaling service latency (e.g., installation of packet classifier, etc.).

- how to disseminate the mobility information within the two-layer architecture. In general, NTLP is needed to be involved with mobility anyway, for example to route NSIS messages along the new path, or to transport explicit release signaling messages along old path. Therefore, it is reasonable to assume NTLP should be able to notify NSLP to update state (by 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.

The following Figures 8.1 (a)-(f) illustrates the characteristics of flows sent between the corresponding node and the mobile node under different mobility scenarios (MIPv6, FMIPv6, HMIPv6; the cases of flows sent from the mobile node to the corresponding node for FMIPv6 and HMIPv6 are similar and omitted here).

To summarize, an ideal interface between NSIS signaling and mobility protocols would be that whenever a mobility protocol changes a characteristics in any place for the flows, NSIS signaling should be able to react accordingly as soon as possible. However, as identified above, an overall coordination/synchronization needs further study.

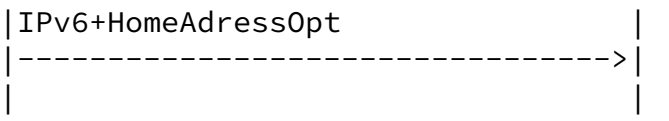

(a) MIPv6: MN-->CN, no reverse tunnel

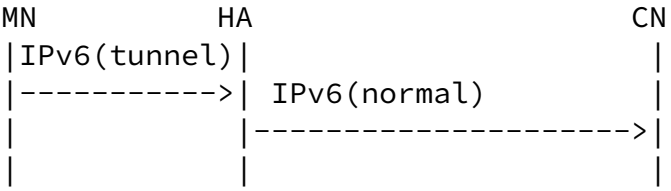

(b) MIPv6: MN-->CN, with reverse tunnel

Fig. 8.1: Implications for flows under different mobility scenarios

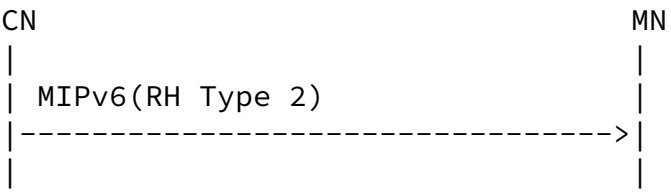

(c) MIPv6: CN-->MN, route optimization

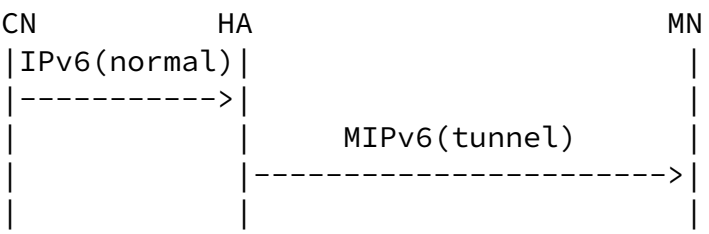

(d) MIPv6: CN-->MN, no route optimization

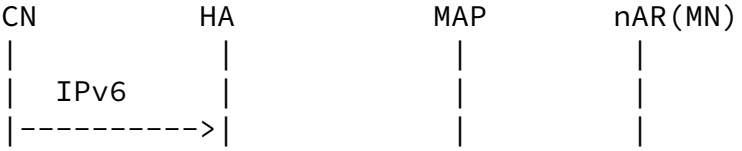

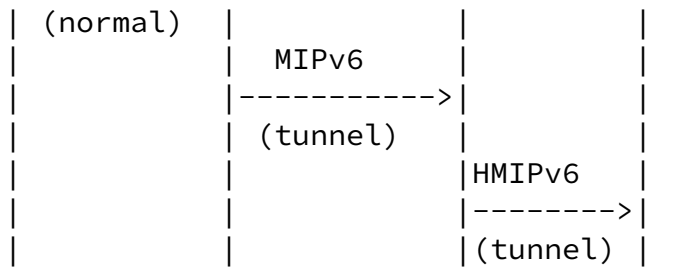

(e) HMIPv6: CN-->MN, no route optimization

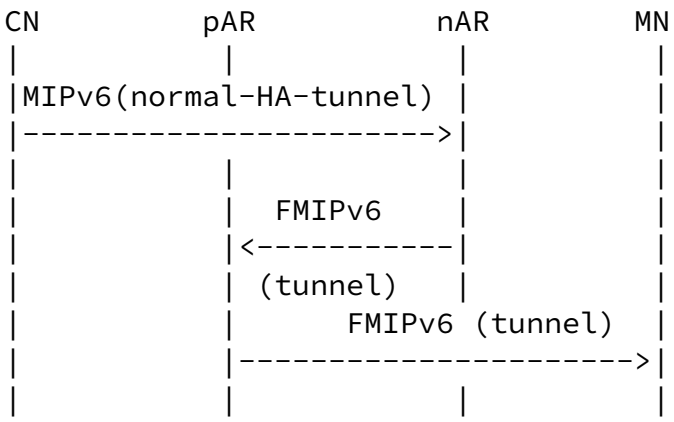

(f) FMIPv6: CN-->MN, no route optimization

Fig. 8.1: Implications for flows under different mobility scenarios

8.2. Interactions with Seamoby Protocols

Although the NSIS protocol suite operates in a path-coupled way, the interactions between NSIS and Seamoby protocols have an effect on parts that are out of the signaling path. In the context of handovers between old and new access routers, there can be performance optimization issues in the following two areas: selection the optimal access router to handover to and transfer of state information between the old and new access routers to avoid having to regenerate it in the new access router after handover. The Seamoby working group is developing protocols solutions for these functions (CARD and CTP respectively), but a discussion of the way in which these functions interact with NSIS signaling is necessary.

As mentioned in Section 3.3, significant performance gains can be achieved if NSIS signaling can interact with such protocols although they can operate independently. In this case, some questions arise: which scenario these protocols can be used in, or what the mode of interaction should be: pre-establishment and re-establishment approaches or passive triggering mode where NSIS is triggered by CARD/CTP and active triggering mode where NSIS triggers these protocols.

In general, CARD is required to identify candidate ARs (CARs) for handover and find capabilities of these CARs prior to the initiation of the IP-layer handover. CTP is used to quickly re-establish context transfer-candidate services without processing for those services from scratch and additionally to provide an interoperable solution that supports various layer 2 radio access technologies. However, in the pre-establishment using NSIS, CARD/CTP can interact with the NSIS protocol suite to ferret NSIS-aware candidate AR where an MN will move and establish NSIS state before the completion of handover. That is, the interaction between NSIS and CARD/CTP prevents resources from being excessively pre-reserved. In this approach, for the fast setup of NSIS state, path update along the candidate NEs may also be achieved simultaneously with (or through) discovering a candidate CRN.

When a handover, for example, is initiated, the current AR receiving the movement detection information (e.g., 'RtSolPr' message in FMIPv6 [FMIP] if NSIS also interacts with this mobility protocol) from an MN may interact with the CARD to ferret an appropriate (NSIS-aware) access router (or a few candidate access routers (CARs) may be found). In this process, the NTLP of the current AR should be able to recognize whether the CAR is an NSIS-aware node after sending the 'capability reply' message (of CARD). The QoS-NSLP of the AR may need to be interaction with the CTP to transfer the QoS-NSLP state information to the newly discovered NSIS-aware candidate AR. After receiving the context, the NTLP of the candidate AR may be able to begin to trigger the discovery of a candidate CRN using the QoS-NSLP state information in the coupled approach or separated approach. QoS and context transfer issues have already been considered already some time ago in [I-D.thomas-nsis-rsvp-analysis]. More recently [CTP-Interop] and [Lee01] present ways and some open issues for

interoperation between NSIS and CTP in both predictive mode and non-predictive mode.

In the re-establishment approach, CARD can be used to only check

admission control status on the new path before handover is completed, and CTP can be used to transfer NSIS state (e.g., QoS-NSLP state) to a CAR to quickly re-establish the state along the new path after handover. The main objective of interaction in this approach is to reduce state setup delay and packet losses due to handover.

In case of passive triggering mode, CARD/CTP may use NSIS signaling to check the admission control status of CAR and pre-establish NSIS state on the candidate path and discover a candidate CRN in the pre-establishment approach. In this case, resource availability on the path between the AR and CN should also be discovered using NSIS signaling. A possible example is that some entity in a candidate AR can trigger NSIS using resource and reservation information from the current AR to find out about how much resources would be available on the new path.

In the context of NSIS, the NSIS protocol generally can trigger the CARD/CTP to transfer its own state information from the current AR to CAR: active triggering mode where the NSIS protocol should monitor the operation of these protocols. Note that the NSIS protocol, in the first place, should interact with mobility protocols (i.e., usually with FMIPv6), or be coupled with movement detection mechanisms to timely initiate the CARD/CTP in both reservation approaches.

If NSIS does not consider interworking with CARD/CTP or it is not possible to use these protocols, NSIS protocol in itself may be able to discover the CAR as an extension of NTLP peer discovery mechanism in the separated approach, and to check whether resources on the candidate path is available or not before the completion of handover. However, this also makes NSIS protocol perform path-decoupled signaling, and whether these functions in NSIS can be implicitly developed is an open issue.

9. Additional issues

This section highlights some important issues not discussed earlier in this draft.

9.1. Both End-Hosts are Mobile

Considerations about signaling between two mobile devices. Until now, we are assuming a non-mobile corresponding node. Problems can show up if both devices start to signal at the same time.

9.2. Uni- and Bi-directional State Establishment

It should be possible to support unidirectional NSIS state establishment in both sender- and receiver-oriented modes. For example, in case of QoS-NSLP, the MN (as a sender) can initiate a reservation setup for its outgoing flows in the sender-initiated mode. With the receiver-initiated approach, the MN (as a sender) requests the receiver to make a reservation, thus allowing the receiver to initiate a reservation for the flow. After handover of the MN (as a sender) to a new AR, the state re-establishment should be performed in the similar way.

In addition to the unidirectional NSIS state establishment above, bidirectional state establishment can also be supported. In the basic case, bidirectional NSIS signaling can simply use a separate instance of the same signaling mechanism in each direction. Although the bidirectional data flows have the same end points, the paths in the two directions do not need to be the same. Therefore, the CRN of the downstream path may be different from that of the upstream path in mobility scenarios. As a matter of course, the Session ID in the downstream reservation should be different from that of the upstream reservation. If the routes (i.e., upstream and downstream paths) are symmetric, an NSIS single signaling message can be used to install state in both directions. If the routes are asymmetric, an NSIS signaling message from the originator (e.g., MN or CN) can trigger an independent signaling message from the responder.

9.3. State Management

The main objective of NSIS is to manage state information along the path taken by a data flow. For state management, the NSIS protocol suite normally use a soft-state approach to manage state in NEs where the state created by the NSIS message has to be periodically refreshed.

At the NTLP layer, the state is maintained through the exchange of GIMPS query/response messages between adjacent peers [ntlp]. In this case, the peer relationship is maintained using a timer which implies how long the association between the peers can be considered valid. That is, if it has not been refreshed until the timer expires (e.g., after 30 seconds as a default value), the peer relationship is removed. The management of state (i.e., routing state and messaging association) can be controlled in this way.

At the NSLP layer, the peer-to-peer refresh messages can also be used for state management. In case of QoS-NSLP, states should be set up and maintained for the reservation of desired resources. In this context, the operation of QoS-NSLP is similar to that of RSVP [RFC 2205]. An example of state management at the QoS-NSLP layer is as follows. Upon receiving a RESERVE message, an NE (specifically the QoS-NSLP) sets up state for QoS reservation. This state will be deleted unless it is refreshed by a RESERVE message before the refresh timer expires. The peer-to-peer based refreshment allows the

QoS-NSLP to appropriately select the refresh time by considering the current network environment. For example, it may set the refresh timer value in the mobile/wireless (access) network to a smaller value than that in the core (wired) network [QoS-NSLP]. Note that, however, unlike the QoS-NSLP, the refresh time of NTLP state doesn't need to be adjusted according to the type of the network from the perspective of resource utilization.

In case of QoS-NSLP, the main objective of the adjustment of the refresh time is to minimize the waste of resources due to double reservation. Setting the refresh time in the access network differently from that of the backbone network can be done by manual configuration or an adaptive technique. A possible example of such adaptive techniques is to use a field, e.g., 'REFRESH' field of the mobility object (or Refresh object). The 'REFRESH' field may consist of 'M' bit for indicating the type of the network (e.g., the mobility-supporting access network) and 'PRE' bit for fast QoS re-establishment (e.g., pre-reservation). The refresh timer value of pre-reservation state should be maintained for a short period of time.

In mobile and wireless networks, the QoS-NSLP (rather than the NTLP) should be able to set the refresh timer value depending on the part of the network (e.g., an access network or backbone network) or the reservation style (e.g., pre-establishment or re-establishment). For example, in case of pre-reservation, upon receiving the mobility object during handover, the QoS-NSLP of the NE which is supposedly involved in the QoS signaling can set the 'PRE' bit of the outgoing QoS-NSLP message. In this case, if the refresh timer value of 'PRE' bit is set to a little higher value than the estimated handover latency, the MN can be provided with seamless QoS service using the pre-reserved resources, and the resources which are pre-reserved but unused will be timely released after handover. Note that after handover, QoS-NSLP should restore the original refresh timer value in

order to avoid the overhead due to the frequent transmissions of the refresh message (e.g., 'PRE' bit is reset to null and 'M' bit is kept to be 1). Note that, however, procedure for computing the refresh time is not part of the NSIS protocol. Thus, how to set the refresh timer value of the 'M' and 'PRE' bits according to mobility scenarios is also an implementation issue.

9.4. State establishment in Network Mobility

The network mobility (NEMO) Working Group is focusing on managing the mobility of a mobile network (e.g., a leaf network) which changes its point of network attachment as a unit through one or more mobile routers (MRs). In this case, the leaf network consists of one or more MNs and/or fixed hosts, and it may include multiple heterogeneous network interfaces. An MR basically has a Home Agent (HA) and bi-directional tunneling between the MR and HA to preserve session continuity while the MR moves into other point of network attachment. The MR as a single node obtains a CoA as in the MIP mechanism, which allows nesting of mobile networks. However, the nested mobile

networks cause the pinball routing problem because flows of each mobile network may transit multiple HAs through multiple bi-directional tunnelings. A mobile network can also have multihoming-related issues through either a single MR which has multiple interfaces to the network, or multiple MRs which attach the mobile networks to the network.

The solutions in the NEMO WG will support preservation of route aggregation in the network when flows of MNs (and/or fixed hosts) in a mobile network are sent to the same CN. In this case, aggregate state installation, e.g., for aggregate reservation (or group reservation), should be considered to guarantee resources along the aggregated route between the MR of the mobile network and the CN. This aggregate state installation issue also requires careful consideration in view of mobility related issues in NSIS. To deal with aggregate state installation in network mobility, issues such as multihoming and pinball routing problem caused by the nested mobile network and various scenarios for network mobility should be considered. However, it is recommended that such issues be handled in liaison with NEMO WG which is still at its early stage in developing solutions for the route optimization problem. Therefore, it is premature to specify details on the aggregate state installation in this draft.

10. Guidelines for Designers of new NSLPs

This section presents issues that must be taken into account when designing a new NSLP for a mobile and wireless environment. The main issues are:

- IP addresses of the communicating nodes can change during the lifetime of the session,
- The bandwidth of the last-hop link may be limited and vary drastically,
- Routing may be asymmetric,
- There may be tunnels that hide the original IP packet header,

- an NE before establishing new state should check whether is already has state with the same session ID but a different flow ID. If this is the case it needs to find out whether it is CRN and act accordingly (details to be spelled out).
- path repair should be localized
- procedures for handling DPD should be the same independent of whether a peer is truly dead or just changed its IP address because of handover.

<Add here a discussion of what those actually mean for the designer, e.g. answers to questions like "so what if the link has low bandwidth, what do I care?".

11. Summary of Split of functionality

<Summary of what functions should go where, NTLP or NSLP>

12. Security Considerations

This section describes authorization issues for mobility scenarios in NSIS. It tries to raise additional questions beyond those discussed in [SID].

For the discussion of various authorization problems we assume that initial authorization is strongly coupled to authorization handling in subsequent message interactions. Making this assumption, as we see

in the subsequent text, has some implication to the signaling message behavior. It is certainly possible that the entities who grant the initial reservation and those who subsequently cause modifications are not the same entities.

Please also note that NSIS does not mandate a single model for the initial authorization step (such as receiver has to provide authorization as in RSVP). Hence it is necessary to consider more combinations. As argued in [NSIS-AAA] it is necessary to consider cases where the sender, the receiver or both are authorizing a reservation. The reader should keep in mind that these signaling message exchanges are not only applicable for QoS reservations but also for other NSLP applications such as NAT/Firewall signaling. The concept of sender- and receiver-initiated reservations is very vague in case of NAT/Firewall signaling. There a concept of delayed authorization is suggested with requires both, the sender and the receiver, to authorize packet and NAT binding establishment.

Subsequently, we will consider the case where the mobile node acts as a data sender followed by a discussion of the CN as a data sender. Each scenario is separated into more individual scenarios.

12.1. MN as data sender

Figure 12.1 describes a scenario with the MN as a data sender which moves from one point of attachment to another. The two flows are merged at the DCOR. The path between the DCOR and the CN is referred as the shared segment. Along the shared segment it might be necessary to update the flow identifier but not the NSLP state itself.

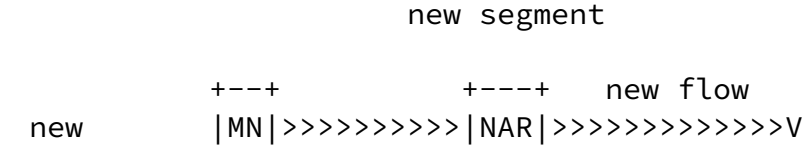

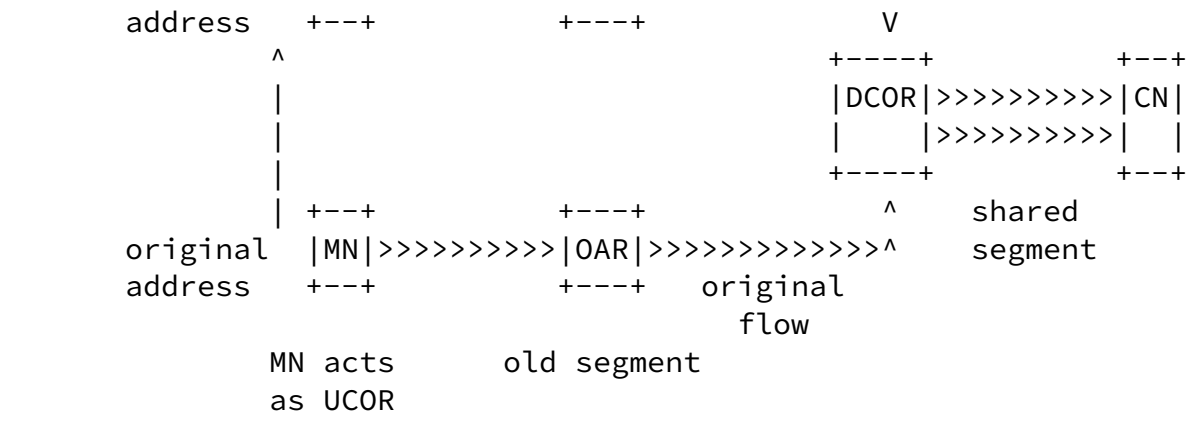

Figure 12.1: MN as data sender

We will start with the initial flow setup triggered by the MN and also authorized by the MN.

12.1.1. MN is authorizing entity

This scenario considers the initial flow setup executed by the MN whereby the MN provides authorization for the initial flow setup. The initial setup might be used to create state for subsequent authorization actions by the MN. It is obvious that the authorization for the NSLP application (e.g., QoS NSLP) has to provided. Depending on the underlying authorization model it might be either peer-to-peer or end-to-middle. This authorization decision can possibly be treated independent of the authorization issues discussed in this section

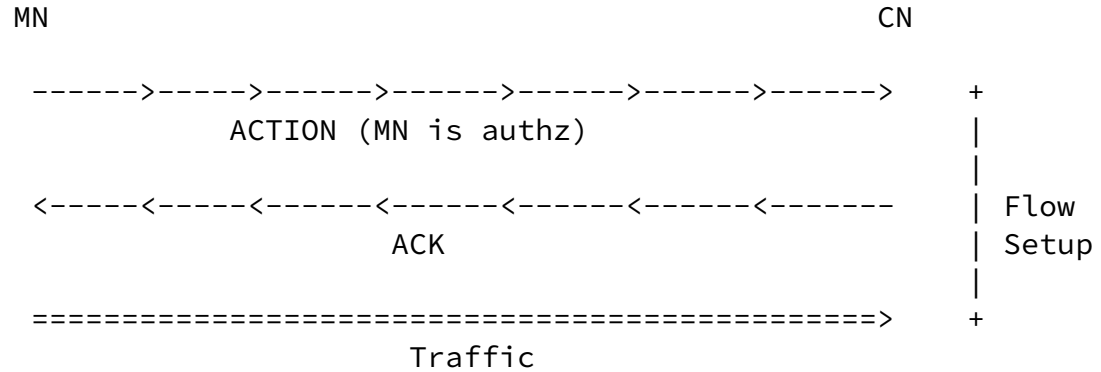

Figure 12.2: Initial MN authorized flow setup

The following questions seem to be interesting:

- Should the MN indicate that it is the authorizing entity for subsequent actions to all entities along the path?

- What information should be used for this purpose?

- Who should add this information? Should the visited network of the MN add something to the signaling message during the initial flow setup?

- How do other entities along the path learn this information?

Scenario: The MN authorizes DCOR

The movement of the mobile node after the initial flow setup requires authorization. Various session ownership authorization issues are illustrated in the [SID] draft itself. We will not repeat these issues in this document again.

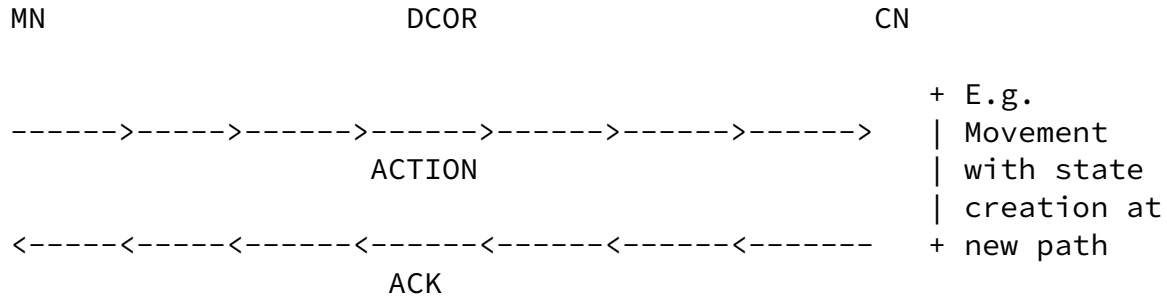

Figure 12.3: MN authorizes DCOR

The following questions are of interest:

- Why should DCOR execute something on behalf of the MN? (i.e., why should it trust the MN and what information can the DCOR use for verification?) As an example, the DCOR might delete state along the old segment.

- Should the DCOR alone be able to start signaling (the DCOR might be a designed node in some mobility protocols (e.g., MAP)) since it is the node which has more information that other nodes based on the mobility signaling protocols?

- How should other nodes between the MN and the DCOR and between the DCOR and the CN know that the DCOR is now acting on behalf of the MN?

Scenario: The CN triggers action

CN wants to tear-down flow or it wants to trigger an action in the network

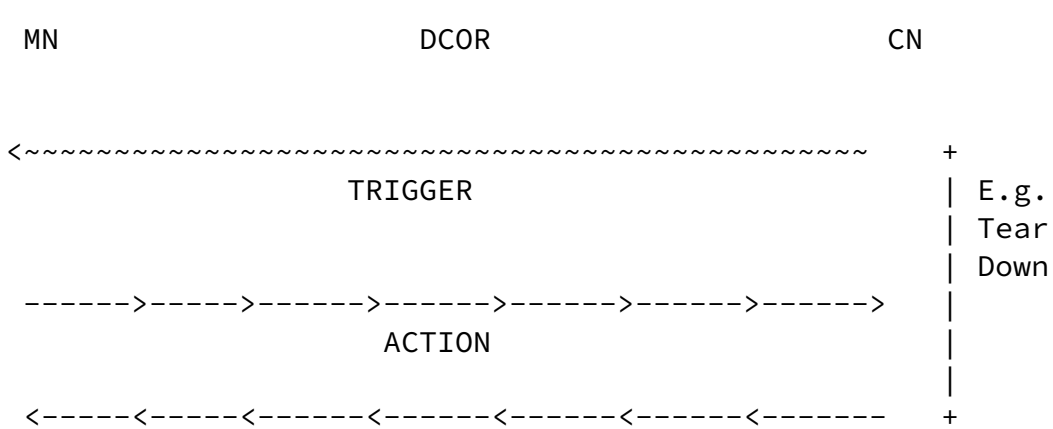

Figure 12.4: CN triggers action

Questions:

- Why should the MN trust the trigger?

- Is it possible to specify the security properties of the trigger message in more detail? (If this is an NSIS message then we could argue that hop-by-hop trust does not always replace end-to-end security.)

- (see for example the discussion in <draft-tschofenig-nsis-aaa-issues-01.txt> with regard to an indicator which entity to charge for the reservation.

- Should the CN restrict the actions of the MN (e.g., delete, update, create)? On the shared segment it might, for example, be possible to restrict the allowed action to a flow identifier update.

12.1.2. CN is authorizing entity

This scenario is quite similar to the CN triggering in Section 1.1.1. Two slightly different protocol variations will be considered. Authorizing some actions in the reverse data flow direction is more difficult as it can easily be seen in Figure 12.5.

Variant 1: CN asks MN to trigger action (on behalf of the CN)

In Figure 12.5 the CN authorizes the MN to start signaling after, for example, a movement. After receiving the trigger message (and some authorization information) the mobile node starts signaling along the new segment and automatically discovers the DCOR. The message travels along the shared segment to the CN and updates the flow identifier (if necessary). The MN might additionally allow the DCOR to delete the reservation along the old segment.

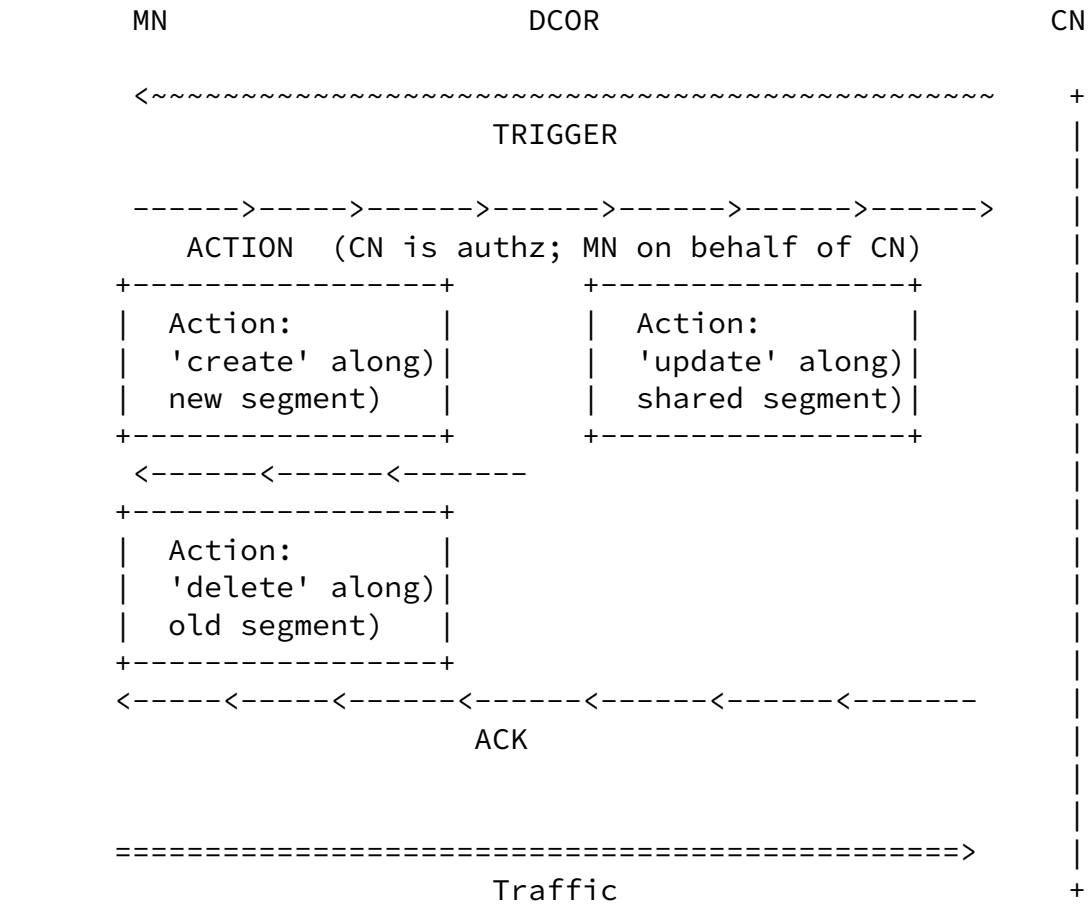

Figure 12.5: CN asks MN to trigger an action (on behalf of the CN)

Questions:

- How should the "delegation" mechanism work such that intermediate nodes believe the MN that it is acting on behalf of the CN?

- Is it possible to carry this information with the Trigger message from the CN and the MN?

Variant 2: CN uses installed state to route message backward

As a second variant the CN uses NSIS installed state to route a signaling message backward along the path. In some rare cases the DCOR node might be known already. In this case it is possible to stop the update process along the shared segment and to possibly mark installed state along the old segment for deletion. When the MN receives the message it again has to install state along the new segment towards the DCOR. The mobile node might also trigger the deletion of resources along the old segment together with this state creation (pessimistic delete). The optimistic delete operation is certainly more error prone.

As it can easily be seen from the description many assumptions have to be made in order for this signaling exchange to work. Hence, as a generic solution this approach is certainly not suggested. We included this scenario for completeness.

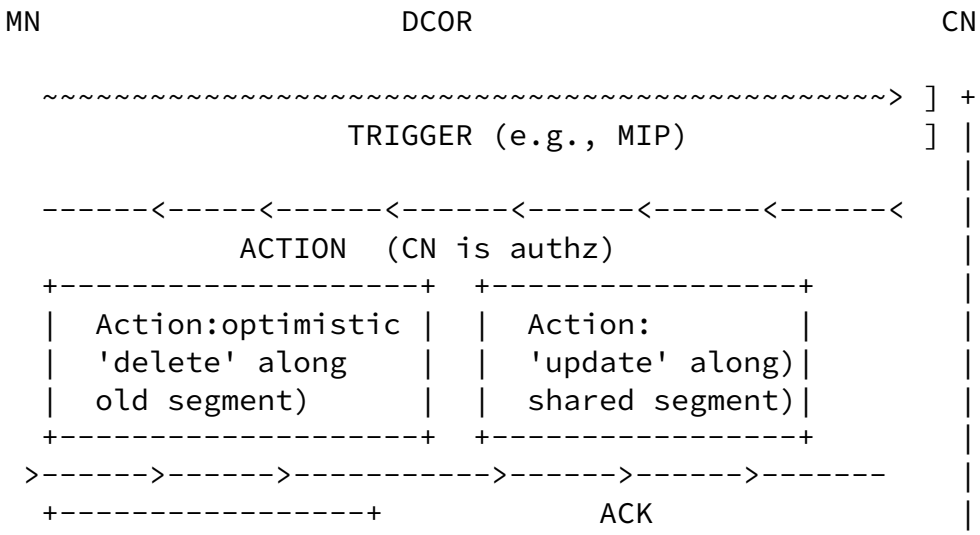

```
      |  Action:          |                              | Action
      |  'create' along)|                              |
      |  new segment)    |                              |
      +------------------+                              |
      <------<------<-------                            |
      +--------------------+                            |
      | Action:pessimistic|                             |
      | 'delete' along)    |                            |
      | old segment)       |                            |
      +--------------------+                            |
                                                        |
     ===============================================>   |
                         Traffic                        +
```

Figure 12.6: CN uses installed state to route message backward

Questions:

- Similarly as before the security properties of the trigger message need to be evaluated.

- It is not always possible to route the message backward from the CN to the MN:

  a) state at the new path might not be established (hence the signaling message cannot travel backward)
  b) the signaling message might not reach the MN via the old segment.

In the multi-homing case where the mobile node can be reached via more than two paths it is possible to execute this exchange. The same might be true for some local repair cases.

- The messages triggered by the MN (namely create state along the new segment and the pessimistic 'delete along the old segment) still need to be executed on behalf of the CN. Compared to the first variant there might be some room for optimization since the first message was transmitted by the CN.

12.1.3. MN and CN are authorized

If we argue that the authorization at the NSLP layer is somehow tight to the authorization issues discussed in this section (i.e.,

authorization for certain protocol actions) then we also have to consider the case where both nodes (MN and CN) have to contribute their part to the authorization decision. This situation appears for example in the NAT/Firewall signaling case but also in the area of QoS reservation where both parties might need to share costs.

If both end hosts are authorized then some signaling message exchanges are less difficult since the trigger message does not need to include authorization information. Some problems, however, do not disappear such as the session ownership problem and additional problems might be caused by certain solution approaches. Since this section does not discuss solutions the reader is referred to the [sid] draft which lists a few strawman proposals for addressing the session ownership problem.

12.2. CN as data sender

In this section we consider the scenarios where the CN acts as a data sender. Except for some minor issues we believe that the requirements caused by the signaling exchanges are mostly discussed in Section 1.3 already.

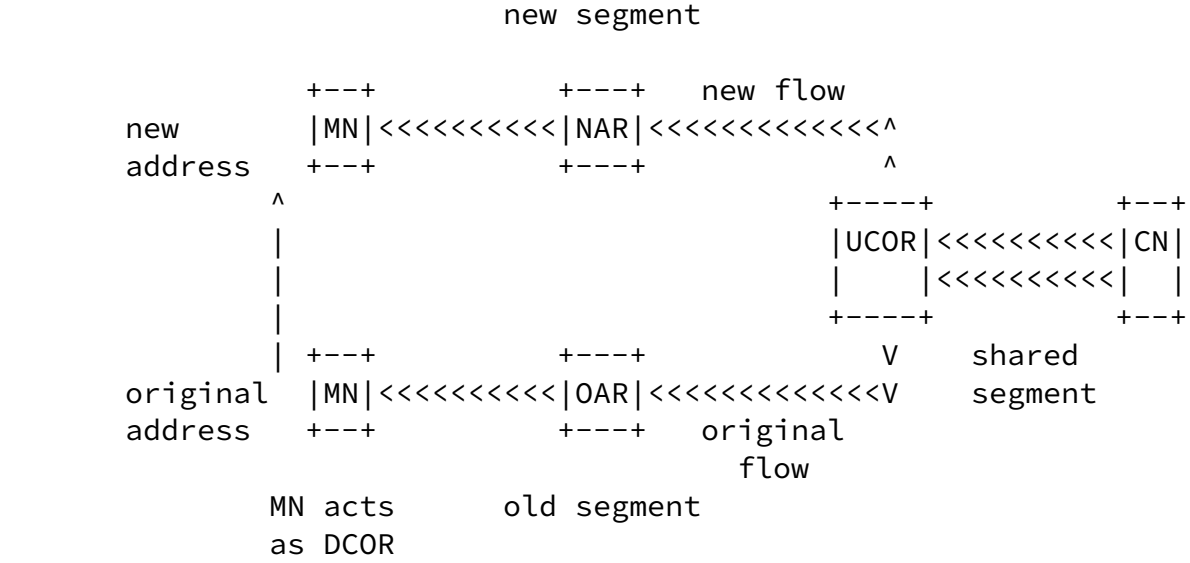

Figure 12.7: CN as data sender

Figure 12.7 shows the corresponding node as a data sender with the mobile node moving from OAR to NAR. As a consequence the CN is triggered to update its QoS reservation along the path. The shared segment only requires an update of the flow identifier (except if the home address is used) but authorization is unaffected since no change in authorization or security relationship is required. Then the signaling message hits the UCOR where the old segment and the new segment form a branch. If necessary then the CN authorized the deletion of the old reservation along the old segment and a new

reservation needs to be created along the new path. In some sense this procedure is similar to a route change. It is even possible that the MN is still reachable via both paths.

12.2.1. CN is authorizing entity

This scenario is similar to the one described in Section 1.3.1. No additional problems arise with a scenario where the CN is both data sender and also authorized. In Figure 12.8 the CN authorizes the UCOR to delete the old segment and to establish a new reservation along the new segment. Furthermore, at the shared segment only an update of the flow identifier might be necessary.

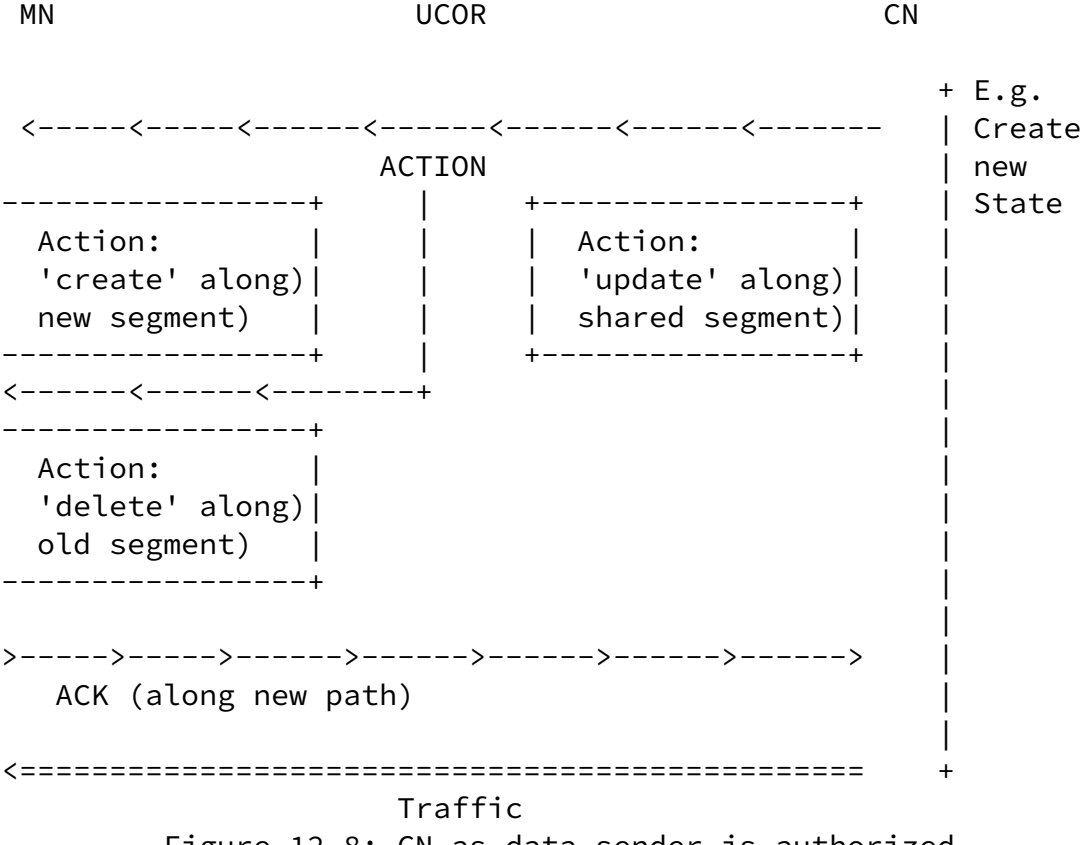

Figure 12.8: CN as data sender is authorized

Since the mobile node first detects the route change. A trigger to the CN allows the CN to quickly react on the route change. There are two variants:

1) The MN sends a trigger to CN and CN starts signaling flow as shown in Figure 12.8.

2) The MN routes the message back along the reverse path using the previously established state along the old route. This mechanism only works if the MN is able to send messages along the old path. As a generic mechanism this is not suggested.

3) An intermediate node act on its own. This might be possible that the UCOR is an entity which participates in the mobility signaling (e.g., MAP) exchange. Depending on the message exchange it needs to be studied whether the signaling message provides sufficient

authorization to trigger the NSIS exchange.

12.2.2. MN is authorizing entity

In this scenario we consider the case where the CN is the data sender but the MN authorizes actions. The considerations are similar to those elaborated in Section 1.3.2 where the MN is the data sender but the CN is the authorizing entity.

12.3. Multi-homing Scenarios

Multi-homing scenarios have the property that the more than one path belongs to a signaling session. In Figure 12.9 the MN uses two interfaces to route NSIS message towards the CN. Both individual sessions are tight together with the same session identifier. The MN needs to indicate that both reservations need to be kept alive (and the DCOR should not delete a reservation). At the shared segment only a single reservation is stored.

From an authorization point of view the session ownership issues is applicable since the DCOR needs to merge the two reservation into a single one along the shared segment.

12.3.1. MN as data sender

This section shows the multi-homing scenario with the MN as a data sender.

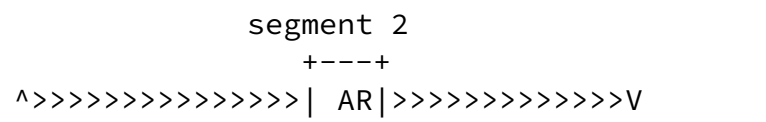

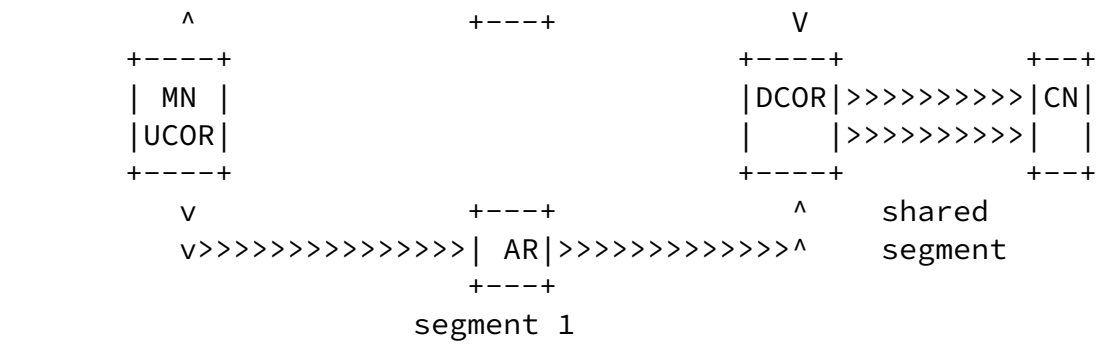

Figure 12.9: Multi-homed MN as data sender

If the MN is the authorizing entity then the session ownership problem needs to be solved. Without solving this type of authorization problem it is possible for an adversary to "join" the reservation at the shared segment. Furthermore, it is an open issue whether reservation merging is allowed only for cases where one flow identifier is used at different interfaces or even with different

flow identifiers.

If the CN is the authorizing entity then, again, some message needs to be sent from the CN to the MN to trigger the exchange or to route the request backward along the established path. The MN is reachable via the two paths.

Mobility handling might be smoother since it is possible that only one interface experiences an IP address change with all the previously discussed implications.

12.3.2. CN as data sender

This section shows the multi-homing scenario with the CN as a data sender. The scenario is simpler (for the CN authorizing case) than the one described in Section 1.5.1 since the signaling message along the shared segment travels the previously established path. This scenario is similar to the route change scenario. At the mobile node itself the two paths merge which again leads to a session ownership problem. How should the MN know whether a signaling message with the same session identifier appearing at a new interface belongs to the indicated session authorized by the CN.

If the MN is the authorizing entity then again communication between the end hosts is required as a trigger. Reverse routing might, in some cases, also be possible.

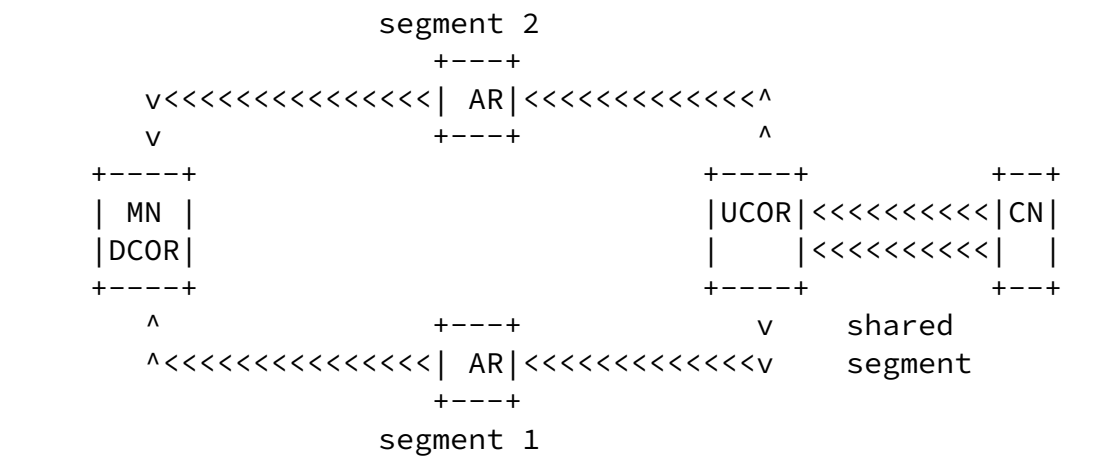

Figure 12.10: Multi-homed CN as data sender

12.4. Context Transfer

Figure 12.9 is taken from [CTP-Interop] and reused to illustrate authorization issues.

Intra-domain mobility with the help of the context transfer protocol can help to move established state information between different access nodes within the same administrative domain including security

associations, QoS parameters (QoS NSLP state), NTLP state and even authorization information. An authorization for a QoS reservation granted along one path through the access network might also valid at a different access router or even at a different path within the same administrative domain. Discussions in the EAP working group, however, reveal that this might not always be the case. However, if we extend the scheme from intra-domain context transfer to inter-domain context transfer then we might encounter some interesting authorization problems. Note that these issues do not only address authorization of QoS resources, but are more applicable to network access authentication and authorization in general. Network access authentication and authorization would not necessarily be executed again after attaching to the new domain. Instead, a trust relationship is established between the new and the old

administrative domain.

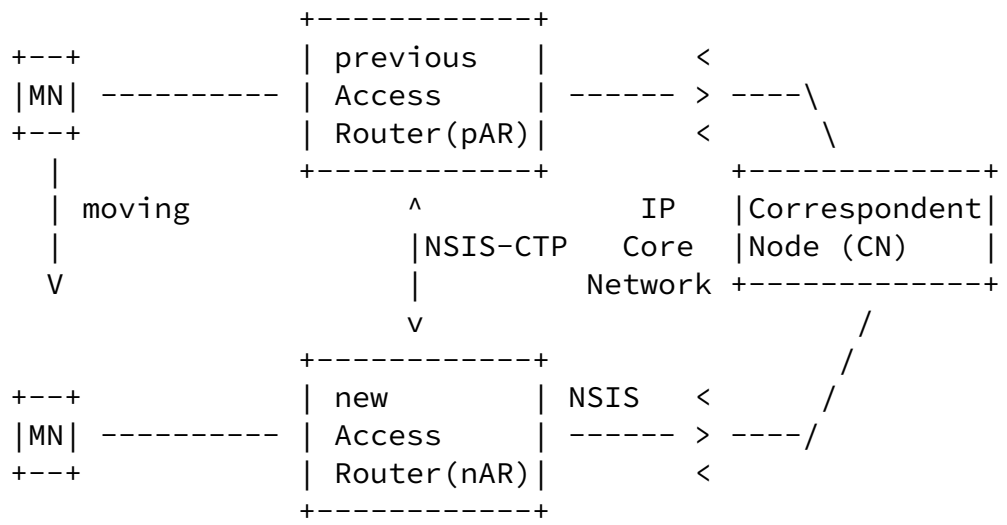

Figure 12.11: Context Transfer

For the signaling message exchange in the predictive and non-predictive CT mode it can be seen that the assumption is made that the MN is the authorizing entity. A communication with the CN does not take place and would certainly increase the complexity of the protocol exchange. The authorization properties of the Context Transfer procedure (and some micro-mobility schemes) need to be studied in more detail to see their implications for security.

The context transfer procedure seems to provide a simple solution for some session ownership problems (in case that the MN is authorized).

12.5. Proxy Scenario

The proxy scenarios refers to those cases where one of the end hosts or even both end hosts are not NSIS aware. Two security implications need to be studied.

First, there is an authorization issue with regard to the NSLP application. For QoS signaling the end host (and the user) has to authorize the QoS reservation since the reservation might require the user is charged for it. Since the end host is not NSIS aware some

other mechanism or protocol needs to be available with provides this functionality. For NAT/Firewall signaling delayed authorization assures that both end hosts authorize the packet filter creation at their local networks (particularly in case of missing trust

relationship between intermediate networks).

Second, the authorization issues which relate to the session ownership problem also need to be studied. Since the session ownership issues are a related to the signaling participating nodes and not to the users or the true end points we think that it does not lead to complications. This is, however, only true if we assume that authorization at the NSLP and authorization decisions for the signaling message handling is decoupled.

12.6. Implications for the costs of a QoS reservation

It is obvious that mobility support within NSIS raises security issues. A number of mobility scenarios with impacts on security are discussed in Section 7 of [NSIS-AAA]. Even if the signaling message exchange is restarted from scratch (i.e., using a new flow-ID), security handling within NSIS is affected. This type of processing is, however, mostly not a topic for this draft.

12.6.1. Missing Cost Control

A large number of service providers (e.g., wireless LAN hot spots) with complex roaming agreements create a non-transparent cost structure. In a traditional subscription-based scenario, users are subscribed to their home network and use this trust relationship to dynamically establish a financial settlement between the visited network and the home network. Additionally, security associations are dynamically established as part of this procedure. This is the typical AAA deployment scenario. In these scenarios users do not learn the identity of the access network as part of a regular authentication and key exchange protocol message exchange. The identity of the access network is possibly never revealed (in a secure fashion). The user is therefore only authenticated to the home network (and hopefully vice versa). While issuing a QoS reservation request to the next NSIS peer (for example in the access network), the end host is typically unaware of the cost of such an end-to-end QoS reservation. Without knowing the costs it is not possible to reject a too expensive QoS reservation.

Currently there is no standardized protocol available which allows users to communicate cost limits, to request cost information for network resources or to learn already accumulated costs for a particular reservation.

Especially in mobility environments - where an end host is likely to have access to different networks within a short time period - cost control is even more complicated.

Some protocol proposals try to merge existing mobility protocols with QoS signaling (i.e., a form of in-band signaling). Thereby the access network is queried (towards the crossover router or the MAP) for the possibility of making a QoS reservation (without actually making the reservation itself). Without a query mechanism, a user cannot take reservation costs into account when choosing between different access networks (or different access routers). Hence, the user might be able to refuse a more expensive service provider. The ability to allow a user to choose between different providers might be required - not only because of the availability of different access technologies (e.g., IEEE 802.1x, Bluetooth, UTRAN) and different service quality offered, but also for cost reasons.

Although real-time notifications of QoS reservation costs (cost control) to the user are out of the scope of NSIS, some interaction might be required.

12.6.2. Implications for Price Determination

The problem of determining the price of a QoS reservation has been mentioned in [NSIS-AAA] and closely relates to integrating the end host into the process of authorization. Even if the end host is aware of the price of a QoS reservation during reservation setup the price might change for a number of reasons:

o First, mobility in general causes a different path to be chosen and might therefore require a new price determination. End host mobility is visible to the end host itself, therefore the appropriate actions can be triggered by the end host to always determine the correct price.

o Route changes somewhere along the path, e.g., mobility in NEMO networks or even mobility in ad-hoc networks, create more problems, since the route change might not be visible for the end host. If price determination is based on the number of networks traversed and intermediate nodes which contribute to the total price of a QoS reservation, then a periodic price query is necessary. Discussions at the NEMO mailing list already point to this problem [Nemo-ML]. If the price of QoS reservation is associated with the authorization itself, then a periodic re-authorization based on the change of prices or on the accumulated costs is necessary.

12.7. Conclusion

This security considerations section illustrates the importance of authorization for NSIS in a mobility environment. Performance is important in mobility environments. Proper security handling accounts for a high percentage of the total performance. It is important to consider this aspect in the analysis of mobility proposals.

From the scenarios we can observe the following issues:

- Signaling in the direction of the data path is simpler than in the opposite direction.

- There are many similarities between the scenarios in Section X (MN as data sender) and Section Y (CN as data sender) particularly if we include multi-homing scenarios.

- Most issues are related to authorization problems that appear after the initial flow setup took place. In this case we speak about the following problem: "Is an entity allowed to perform the indicated action." Only a few problems are related to authorization problems which already appear during the initial signaling message exchange.

- If the data sender triggers the signaling message exchange and also provides authorization then the complexity can be kept fairly low.

- NSLP authorization decisions should be treated separately from authorization decisions which affect the signaling message exchange.

In the [SID] draft we tried to raise the question of a possible security goal. We list a few variations of this goal:

Version 1:

An off-path adversary MUST NOT be able to inject messages which are then accepted by NSIS nodes along the path. An NSIS node which once was along the data path is not treated as an adversary.

Version 2:

Only end points MUST be able to create messages and intermediate NSIS nodes MUST be able to verify them.

Version 3:

Only the session creator MUST be able to create messages which are then successfully verifiable by intermediate NSIS nodes.

Based on the first version of the goal it might be necessary to differentiate between NSIS nodes which are currently part of the signaling chain (i.e., nodes which are currently part of the NSIS signaling message processing) and NSIS nodes which previously were part of this chain.

Furthermore, it might be useful to differentiate between different messages:

- Refresh, delete or update messages: I assume that these messages are not idempotent and hence previous state along some nodes has to be established already. To provide security for the 'different peer' case it might still be required to provide some security for the session identifier. This issue is, however, not as dangerous as the threat described in the SID draft.

- Create messages: This message is particularly dangerous since it requires (in case of a sender-initiated message) no previous state. The threat description in the SID draft is immediately applicable. A receiver-initiated signaling message is, from a session identifier point of view, better since previously created state can be used. This might provide some security, although not too much considering the limited capabilities of the responder to truly provide some additional authorization capabilities (due to missing end-to-end security protection or in case of signaling proxies).

- Query alike messages: These message types require little protection since they do little harm to the state. They still might allow an adversary to gain information about the reserved resources.

- Error messages: These messages are also sensitive but are typically returned after a request was submitted. This is, however, not true for asynchronous error messages. Still some state has to be created to allow routing along the established path.

Currently, this analysis does not consider the different message types.

As a final conclusion we must state that more discussion is necessary to address security and mobility handling in an appropriate way.

Particularly, the expected NSIS signaling behavior must be described. The improvements due to mobility functionality within NSIS must be compared with the increased complexity. Careful analysis and performance evaluations are necessary.

13. Contributors

This draft initially written by Roland Bless, Xiaoming Fu, Robert Hancock, Seong-Ho Jeong, Cornelia Kappler, Sung-Hyuck Lee, Jukka Manner, Paulo Mendes, and Hannes Tschofenig.

14. Acknowledgments

This draft is based on four earlier drafts by (in alphabetical order) Jongho Bang, Roland Bless, Xiaoming Fu, Robert Hancock, Seong-Ho Jeong, Cornelia Kappler, Sung-Hyuck Lee, Byoung-Joon Lee, Jukka Manner, Paulo Mendes, Henning Schulzrinne, Charles Q. Shen, and Hannes Tschofenig.

15. Informative References

[mipv4] Perkins, C., "IP Mobility Support for IPv4," RFC 3344, Aug. 2002.

[mipv6] Johnson, D., Perkins, C. and J. Arkko, Mobility Support in IPv6, draft-ietf-mobileip-ipv6-24 (work in progress), June 2003.

[lmm] Carl Williams, "Localized Mobility Management Requirements", draft-ietf-mipshop-lmm-requirements-00 (work in progress), Oct 2003.

[fmip] Koodli, R., "Fast Handovers for Mobile IPv6", draft-ietf-mipshop-fast-mipv6-00 (work in progress), Oct 2003.

[hmipv6] Soliman, H., Castelluccia, C., Malki, K. and L. Bellier, "Hierarchical Mobile IPv6 mobility management (HMIPv6)", draft-ietf-mipshop-hmipv6-00 (work in progress), Oct. 2003.

[nsis-fw] Hancock, R. and et al., "Next Steps in Signaling: Framework", draft-ietf-nsis-fw-04 (work in progress), Sept 2003.

[ntlp] Schulzrinne, H. and R. Hancock, "GIMPS: General Internet

Messaging Protocol for Signaling", draft-ietf-nsis-ntlp-00 (work in progress), Oct. 2003.

[nsis-req] Brunner, M., "Requirements for Signaling Protocols", draft-ietf-nsis-req-09 (work in progress), August 2003.

[nsis-analysis] Manner, J., Fu, X. and P. Pan, "Analysis of Existing Quality of Service Signaling Protocols", draft-ietf-nsis-signalling-analysis-03 (work in progress), Oct 2003.

[RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, Sep 1997.

[RFC2746] Terzis, A., Wroclawski, J. and L. Zhang, "RSVP Operation Over IP Tunnels", RFC 2746, January 2000.

[RFC3583] Chaskar, H., "Requirements of a Quality of Service (QoS) Solution for Mobile IP", RFC 3583, September 2003.

[manner02] Manner, J., Lopez, A., Mihailovic, A., Velayos, H., Hepworth, E. and Y. Khouaja, "Evaluation of mobility and QoS interaction", Computer Networks vol.38, no.2, pp.137-163, February 2002.

[Seamoby-terms] Jukka Manner, Markku Kojo (editors) Mobility Related Terminology Internet Draft (work in progress), November, 2003.

[RFC2961] Berger, L., Gan, D., Swallow, G., Pan, P., Tommasi, F. and S. Molendini, "RSVP Refresh Overhead Reduction Extensions", RFC 2961, April 2001.

[I-D.schulzrinne-nsis-casp] Schulzrinne, H. and et al., "CASP - Cross-Application Signaling Protocol", draft-schulzrinne-nsis-casp-01 (work in progress), March 2003.

[fu03] Fu, X., Schulzrinne, H. and H. Tschofenig, "Mobility Support for Next-Generation Internet Signaling Protocols", Proceedings of IEEE Vehicular Technology Conference 2003-Fall, October 2003.

[I-D.ietf-nsis-qos-nslp] Van den Bosch, S. and et al., "NSLP for Quality-of-Service signaling", draft-iet f-nsis-qos-nslp-00 (work in progress), September 2003.

[I-D.westphal-nsis-qos-mobileip] Westphal, C. and H. Chaskar, "QoS Signaling Framework for Mobile IP", draft-westphal-nsis-qos-mobileip-00 (work in progress), June 2002.

[SID] Tschofenig, H. et al.: "Security Implications of the Session Identifier", <draft-tschofenig-nsis-sid-00.txt>, (work in progress), June 2003.

[CTP-Interop] "A Framework for Interoperation between NSIS and CTP", <draft-soltwisch-nsis-ctp-interop-00.txt>, (work in progress), October, 2003.

[NSIS-AAA] Tschofenig, H., et al.: "NSIS Authentication, Authorization and Accounting Issues", <draft-tschofenig-nsis-aaa-issues-01.txt>, (work in progress), March 2003.

[NSIS-Authz] Tschofenig, H., et al.: "QoS NSLP Authorization Issues", <draft-tschofenig-nsis-qos-authz-issues-00.txt>, (work in progress), June 2003.

[Nemo-ML] Alper, Y., "[nemo] AAA and NEMO", discussion in the IETF Nemo Mailing List (available at: http://www.nal.motlabs.com/pipermail/nemo/2003-February/000271.html), February 2003.

[Lee01] S.-H. Lee, and et al., "Mobility Functions in the QoS NSLP", Internet Draft, Work in progress, October 2003.

[Jeong01] S.-H. Jeong, and et al., "Mobility Functions in the NTLP", Internet Draft, Work in progress, October 2003.

[I-D.thomas-nsis-rsvp-analysis] Thomas, M., "Analysis of Mobile IP and RSVP Interactions", draft-thomas-nsis-rsvp-analysis-00 (work in progress), November 2002.

16. Author's Addresses

Questions about this document may be directed to:

Roland Bless
Institute of Telematics, Universitaet Karlsruhe (TH)
Zirkel 2
76128 Karlsruhe
Germany
EMail: bless@tm.uka.de
URI: http://www.tm.uka.de/~bless/

Xiaoming Fu
University of Goettingen
Telematics Group
Lotzestr. 16-18
Goettingen 37083
Germany
E-Mail: fu@cs.uni-goettingen.de

Robert Hancock
Siemens/Roke Manor Research Ltd
Romsey, Hants, SO51 0ZN
United Kingdom
Voice: +44-1794-833601
Fax: +44-1794-833434
E-Mail: robert.hancock@roke.co.uk

Seong-Ho Jeong
Hankuk University of FS
89 Wangsan, Mohyun
Yongin-si, Gyeonggi-do 449-791
Korea
Phone: +82 31 330 4642
E-Mail: shjeong@hufs.ac.kr

Cornelia Kappler
Siemens AG
Siemensdamm 62
Berlin 13627
Germany
EMail: cornelia.kappler@siemens.com

Sung-Hyuck Lee
SAMSUNG Advanced Institute of Technology
San 14-1, Nongseo-ri, Giheung-eup
Yongin-si, Gyeonggi-do 449-712
KOREA
Voice: +82-31-280-9585
Fax: +82-31-280-9569
E-Mail: starsu.lee@samsung.com

Jukka Manner
Department of Computer Science
University of Helsinki
P.O. Box 26 (Teollisuuskatu 23)

FIN-00014 HELSINKI
Finland
Voice: +358-9-191-44210
Fax: +358-9-191-44441
E-Mail: jmanner@cs.helsinki.fi

Paulo Mendes
DoCoMo Communications Laboratories Europe GmbH
Landsberger Str. 312
80687 Munich, Germany

Voice: +49-89-56824-226
Fax: +49-89-56824-300
E-mail: mendes@docomolab-euro.com

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

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.