MEXT (MIP6) Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 15, 2009

R. Wakikawa (Editor)
Toyota ITC

July 14, 2008

Home Agent Reliability Protocol
draft-ietf-mip6-hareliability-04.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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.

This Internet-Draft will expire on January 15, 2009.

Copyright Notice

Abstract

The home agent can be a single point of failure when Mobile IPv6 is operated in a system. It is critical to provide home agent reliability in the event of a home agent crashing or becoming unavailable. This would allow another home agent to take over and continue providing service to the mobile nodes. This document describes the problem scope briefly and provides a mechanism of home agent failure detection, home agent state transfer, and home agent switching for home agent redundancy and reliability.

Table of Contents

1. Introduction

In Mobile IPv6 [RFC-3775] and NEMO Basic Support [RFC-3963], if a home agent loses the binding cache state, due to failure or some other reason, it results in a loss of service for the mobile nodes. It is beneficial to provide high availability and redundancy for a home agent so that mobile nodes can avail of uninterrupted service even when one home agent crashes or loses state. The Home Agent Reliability protocol is designed to synchronize the Mobile IPv6 states between active and standby home agents as VRRP[RFC-3768] or HSRP [RFC-2281]. A home agent maintains not only binding cache but also IPsec and IKE related states per mobile node. Mobile IPv6 mandates IPsec encryption for signaling of home binding registration (BU and BA) and return routability (HoTI and HoT) as described in [RFC-3776]. However, IPsec states synchronization is out of scope in this document. The scope of Home Agent Reliability protocol is limited to the management of Mobile IPv6 related states. Thus, we define two different approaches such as Home Agent Virtual Switch and Home Agent Hard Switch depending on the capability of IPsec state synchronization. In both cases, a mobile node maintains only one home binding at any given time.

Home Agent Virtual Switch

The Home Agent Virtual Switch operation can be used only if IPsec state synchronization mechanism is available (outside of Home Agent Reliability Protocol). The IPsec state per mobile node MUST be shared between the active home agent and standby home agents in the background. The active and the standby home agents are addressed by the same home agent address, although only the active home agent is accessible with the home agent address. Each mobile node negotiates just one Security Association with the active home agent. In case there is a failure of the active home agent, the standby home agent takes over without the mobile node being aware of the change in the home agent.

Home Agent Hard Switch

In the Home Agent Hard Switch, IPsec/IKE states synchronization is not required. The home agents are addressed by different IP addresses. When an active home agent fails, a mobile node will receive a notification (Home Agent Switch message [RFC-5142]) from a standby home agent, and send a Binding Update to the standby home agent. In order for the mobile node to receive the Home Agent Switch message securely from the standby home agent, the mobile node needs to establish an IPsec SA with both the active and the standby home agents beforehand.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119].

In this document, the term mobile node refers to both a mobile node [RFC-3775] and a mobile router [RFC-3963].

Mobility related terms used in this document are defined in [RFC-3775] and [RFC-3753]. In addition or in replacement of these, the following terms are defined or redefined:

Active Home Agent

A home agent that is currently serving the mobile nodes.

Standby Home Agent

A home agent which will serve the mobile nodes when the active home agent fails.

Failed Home Agent

A home agent that is not available due to hardware or software failure, system maintenance, etc.

Redundant Home Agent Set

A group of active and standby home agent(s). The Group Identifier is used to identify a redundant home agent set.

Virtual Home Agent Address

A home agent address shared among home agents in a redundant home agent set and used only in the Home Agent Virtual Switch case. The address is only acitivated on an active home agent.

Home Agent Preference

This preference value is originally defined for Dynamic Home Agent Address Discovery (DHAAD) in RFC3775. This protocol re-uses this preference value for home agent selection when an active home agent has failed. However, an operator can also define an independent value used only for the home agent reliability protocol if the operator wants to have different preference values for DHAAD and the home agent reliability protocol. A home agent SHOULD NOT use the same preference value of other home agents.

3. Problem Statement and Requirements

In Mobile IPv6 [RFC-3775], a mobile node registers and establishes a binding with only one home agent. Thus the home agent represents the possibility of a single point of failure for Mobile IPv6. A home agent is responsible for multiple mobile nodes on its home link. The failure of the home agent may then result in the loss of connectivity for numerous mobile nodes located throughout the Internet. To overcome this problem, Mobile IPv6 allows deployment of multiple home agents on the home link so that upon the failure of a home agent, a mobile node can re-establish its connection through a new home agent. However, the base Mobile IPv6 specification does not address home agent failover and dynamic transfer of service from one home agent to another. This transfer of service from the failed home agent to a new active home agent requires coordination or pre-configuration among the home agents regarding security associations, transfer of mobile node bindings, and other service information for reliable Mobile IPv6 service in a deployment scenario.

For the home agent reliability solution, we define the following requirements:

Reliable Home agent service

> Multiple home agents are available for a home prefix and one of them actively serves the mobile nodes. A standby home agent takes over when the active home agent becomes unavailable. The transfer of the MN-HA association should be transparent to applications and should not take longer than the care-of-addresses update procedure described in Mobile IPv6 [RFC-3775].

Availability of a redundant home agent set

> Availability of an active home agent address and a standby home agent address at the bootstrapping period for the mobile node is assumed.

State Synchronization

> The information for mobile nodes must be able to be synchronized between an active home agent and standby home agents. This includes the Binding Cache, AAA information, other Mobile IPv6 and NEMO related information. Note that the Home Agent Reliability protocol exchanges only running states of mobile nodes. Therefore, we do not have any specific operation for synchronizing the configuration information. For instance, when Mobile IPv6 is operated with Authentication protocol, the synchronizing the configurations of the Authentication protocol is out of scope in

this document. Operators MAY correctly configure in multiple home agents.

Consideration of IPsec/IKE transfer

An active home agent maintains several IPsec and IKE states for mobile nodes. These states are synchronized within the redundant home agent set. The details are described in Section 10.

Secured Message Exchanges

The messages used between the home agents to transfer binding cache information MAY be authenticated and encrypted.

Failure Detection

Redundant home agents must actively check for possible failure of an active home agent. If a home agent supports an existing failure detection mechanism such as VRRP[RFC-3768] or HSRP [RFC-2281], it can re-use that mechanism to detect the home agent failure. On the other hand, periodic Hello messages are introduced to detect active home agent's service availability in this document.

Failure Notification

If necessary, a mobile node is notified about the active home agent failure by the standby home agent.

4. Protocol Overview

4.1. Home Agent Virtual Switch

A mobile node remains unaware about the change in the active home agent since the home agents have replicated all session state including IPsec/IKE/ESP states. IPsec/IKE/ESP state transfer is out of scope of this document.

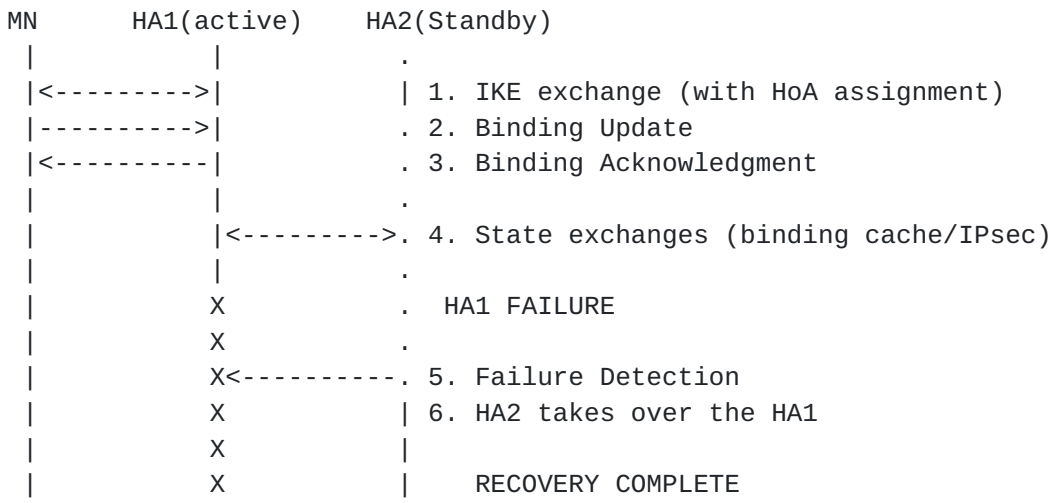

Figure 1: Overview of Home Agent Virtual Switch

The operations of the Home Agent Virtual Switch mode are shown in Figure 1. A mobile node first attempts the IKE exchange for Security Association (SA) setup and home address assignment (1). After binding registration is done (2, 3), the active home agent pushes all the states of its mobile nodes with a state synchronization message (4). The standby home agent(s) is not active unless it takes over from a failed home Agent.

When the active home agent's failure is detected (5), the standby home agent activates the virtual home agent address on its interface and takes over for the failed home agent. All the home agents in the redundant home agent set share a virtual home agent address and the routing will ensure only the active home agent will be reachable using that virtual home agent address. The standby home agent can serve all the mobile nodes for which the states are synchronized, without any further message exchange, because it has all the necessary information which it obtained from the failed home agent.

4.2. Home Agent Hard Switch

The overview of the Home Agent Hard Switch is shown in Figure 2. This mode is not transparent to the mobile node when the active home agent failure occurs.

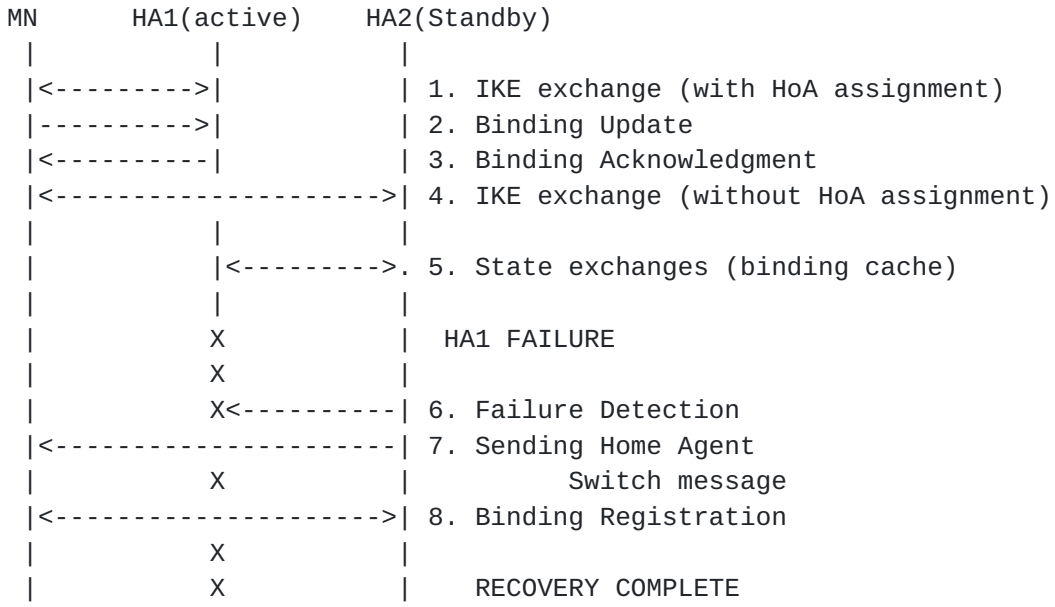

Figure 2: Overview of Home Agent Hard Switch

The mobile node establishes IPsec/IKE state with all the home agents in the redundant home agent set beforehand (1 and 4), however it registers its binding only with the active home agent (2 and 3). When an active home agent fails, a standby home agent uses a pre-existing IPsec SA to notify the mobile node about the failure by securely sending a Home Agent Switch message. In order to discover home agent addresses, two different mechanisms are defined, as described in Section 9.4.1. The active home agent synchronizes the required states of the mobile nodes, such as Binding Cache and AAA information, with other standby home agents periodically (5). The mobile node MUST NOT request a home address(es) assignment through the IKE exchange to the standby home agent when it establishes an SA with it (4).

When the standby home agent detects the failure of the active home agent (6), it sends a Home Agent Switch message to all the mobile nodes that were registered with the failed home agent (7). The Home Agent Switch message must be encrypted by a pre-established IPsec SA. After the switch message, the mobile node MUST send a binding update to the new active home agent in order to update the Mobile IPv6

tunnel end points (8).

4.3. Home Agent Management

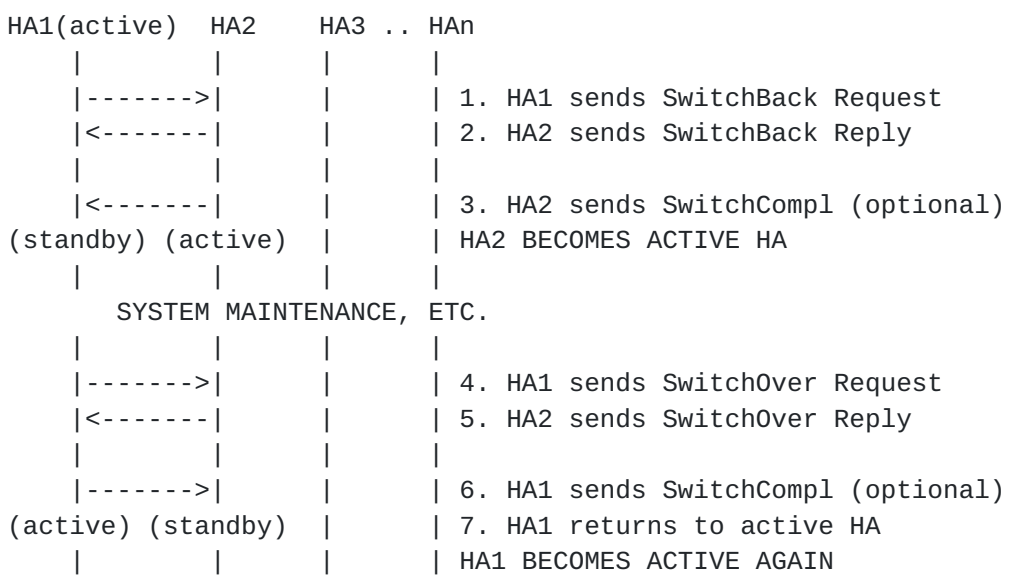

Figure 3: Manual Home Agent Change

In some scenarios the active home agent may need to stop serving mobile nodes for system maintenance. This specification provides for a manual home agent switch by using SwitchBack Request and Reply messages. As shown in Figure 3, the active home agent (HA1) sends a SwitchBack Request message to a standby home agent (HA2). As soon as HA2 receives the message, it becomes the active home agent. HA2 will acknowledge the message by sending a SwitchBack Reply message to HA1. HA1 becomes a standby home agent when it receives the SwitchBack Reply. After the downtime, HA1 sends a SwitchOver Request to HA2 in order to become the active home agent again.

The SwitchCompl message is used only in the Home Agent Hard Switch. As shown in Section 9, it takes certain time to complete the home agent switch. Thus, the old active home agent continues serving the received packets for the mobile nodes during the switch process. As soon as the new home agent completes the recovery, it sends SwitchCompl message to the previous active home agent. SwitchCompl is an optional operation in this specification.

5. Messages

5.1. New Mobility Header Messages

5.1.1. State Synchronization Message

This message is used to exchange state corresponding to a particular mobile node(s). It MUST be unicasted and MUST be authenticated by IPsec ESP. This message has the MH Type value TBD.

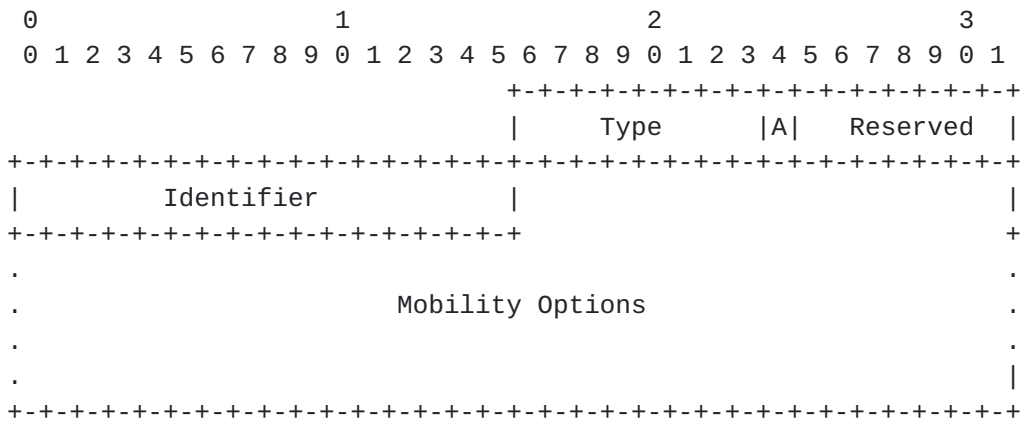

Figure 4: State Synchronization Message

Type

8-bit unsigned integer. It can be assigned one of the following values:

0: Request

This message is called State Synchronization Request (SS-REQ) and used to solicit the active state corresponding to a particular mobile node.

1: Reply

The message is called State Synchronization Reply (SS-REP) and is used between the home agents in the redundant home agent set to exchange binding cache information and any other information related to providing mobility service to the mobile nodes either periodically or in response to a SS-REQ.

2: Reply-Ack

The message is called State Synchronization Reply-Ack (SS-ACK) and is used to acknowledge to the SS-REP. This message is optional and is specially used when the links between home agents are not reliable.

Ack flag

This flag is valid only for SS-REP. If the sender requires explicit acknowledgment by SS-ACK, it MUST set this flag.

Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Identifier

A 16-bit identifier to aid in matching state synchronization message. The identifier should never be set to 0. It should always be more than 1.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The encoding and format of defined options are described in [RFC-3775]. The receiver MUST ignore and skip any options which it does not understand. This message requires at least one mobility option, therefore, there is no default length for this message.

One of the following options is mandatory in SS-REQ message. Multiple same options can be stored in the same SS-REQ message, (ex. two IP address options for two mobile nodes):

* IP Address Option (Sub-type: Home Address) defined in [RFC-5268]. If a home agent wants the Binding Cache information for a particular mobile node, it includes the mobile node's home address in an IPv6 Address Option. If a home agent want to solicit all the active mobile nodes' states, it can include the unspecified address (0::0) in an IPv6 address option.

* Mobile Network Prefix Option. If a home agent wants to know the forwarding state setup for a particular Mobile Network Prefix, it includes a Mobile Network Prefix Option as defined in [RFC-3963].

* Vendor Specific Mobility Option. If a home agent wants vendor specific information, it can solicit with this option as defined in [RFC-5094].

One of the following options is mandatory in SS-REP:

* Binding Cache Information Option

* AAA Information Option

* Vendor Specific Mobility Option

5.1.2. Home Agent Control Message

This message is used to control the status of a home agent to either active or standby. This message MUST be unicasted between home agents and MUST be authenticated and encrypted by IPsec ESP. The Home Agent Control message has the MH Type value TBD. If no options are present in this message, no padding is necessary and the Header Len field will be set to 1.

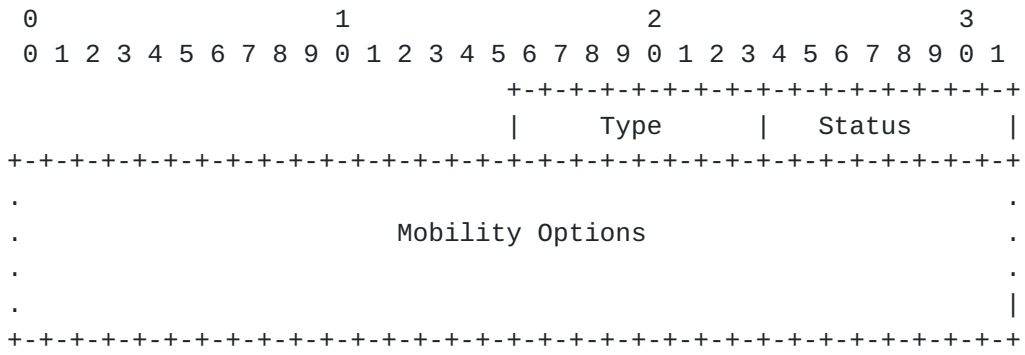

Figure 5: Home Agent Control Message

Type

8-bit unsigned integer. It can be assigned one of the following values:

0: SwitchOver Request (SO-REQ)

It is unicasted by a standby home agent that desires to become the active home agent. The receiver of the message MUST transition to standby state as soon as the message is received and validated successfully.

1: SwitchOver Reply (SO-REP)

It is used to acknowledge the receipt of the corresponding SO-REQ.

2: SwitchBack Request (SB-REQ)

It is unicasted by an active home agent that desires to become a standby home agent. The receiver of this message SHOULD transition to active state as soon as the message is received and validated successfully.

3: SwitchBack Reply (SB-REP)

It is used to acknowledge the receipt of the corresponding SB-REQ.

4: Switch Complete (SW-COMP)

This message is used to indicate the completion of switch over, (i.e. sending home agent switch messages and receiving binding update messages from all the served mobile nodes).

Status

8-bit unsigned integer indicating the disposition of a SO-REQ or SB-REQ. This field is only valid in SO-REP and SB-REP. The following Status values are defined:

0: Success

128: Reason unspecified

129: Administratively prohibited

130: Not active home agent (The receiver of SO-REQ is not the active home agent)

131: Not standby home agent (The receiver of SB-REQ is already the active home agent)

132: Not in same redundant home agent set

Mobility Options

No options are defined in this specification

5.1.3. Home Agent Hello Message

The Home Agent Hello (HA-HELLO) message MUST be either unicasted or multicasted to carry home agent information among the redundant home agent set. The HA-Hello message is defined for two purpose: 1) an alive check and 2) home agent information exchange. A HA-HELLO SHOULD be authenticated and encrypted by IPsec ESP when it is unicasted. If a HA-Hello message is multicasted, IPsec ESP cannot be applied. In this case the redundant home agent set should be located in a secure network. Alternatively, all the home agents MUST have a secure channel with each other. The HA-Hello has the MH Type value TBD. If no options are present in this message, 0 octets of padding are necessary and the Header Len field will be set to 2.

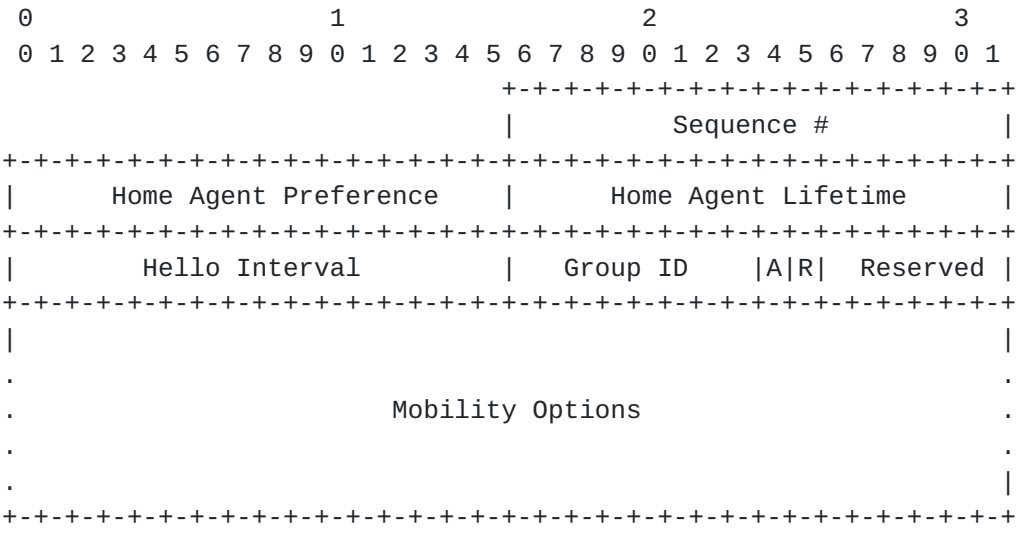

Figure 6: Home Agent Hello Message

Sequence #

16-bit unsigned integer. The Sequence number of the HA-Hello message can be used to verify whether this Hello message is the latest one or not.

Home Agent Preference

16-bit unsigned integer. The preference for the home agent sending the HA-Hello message. This preference is the same as the Home Agent Preference value of the Home Agent Information option as defined in [RFC-3775]. However, operators MAY use a different

preference value for this operation.

Home Agent Lifetime

16-bit unsigned integer. The lifetime for the home agent sending the HA-Hello message. This lifetime is the same as the Home Agent Lifetime value of the Home Agent Information option as defined in [RFC-3775].

Hello Interval

16-bit unsigned integer. The interval for the home agent sending this Hello message.

Group Identifier

8-bit unsigned integer. This value is used to identify a particular redundant home agent set.

A flag

Active Home Agent flag. If this flag is set, the sender of this HA-Hello message is an active home agent.

R flag

HA-HELLO requesting flag. If this flag is set, the receiver of this HA-Hello message must send back a HA-Hello message to the sender.

Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Mobility Options

No valid options are defined in this specification.

5.1.4. Home Agent Switch Message

This message is defined in Section 9.4.3. The Home Agent Reliability protocol extends this message for the Home Agent Hard Switch.

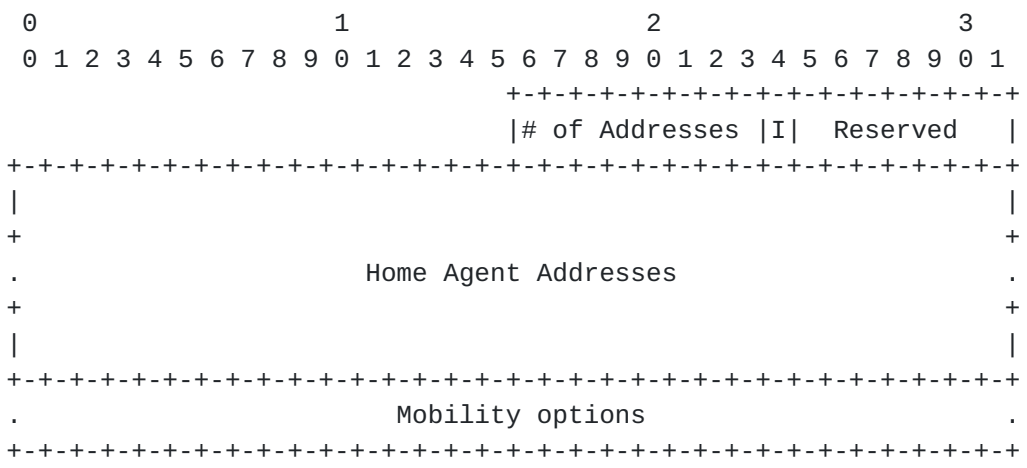

Figure 7: Home Agent Switch Message

IPsec Re-key (I)

The IPsec re-key (I) bit is set to indicate that the mobile node SHOULD start an IPsec re-key with the home agent specified in the Home Agent Addresses field. This flag is used when a failed home agent recovers and needs to re-establish IPsec SA/IKE state with a mobile node. When this flag is set, the mobile node does not switch the home agent, but only re-key the SA.

Reserved

The reserve field is reduced from 8 to 7 bits

5.2. New Mobility Options

5.2.1. IP address Option

This option is already defined in the Fast Handovers for Mobile IPv6 (FMIP) specification [RFC-5268]. This document introduces new Sub-Type values for home agent address and Home Address.

Option-Code

* 4: Home Address

5.2.2. Binding Cache Information Option

The binding cache information option has an alignment requirement of 8n+2. The Binding Cache Information option is only valid in a State Synchronization message. Its format is as follows:

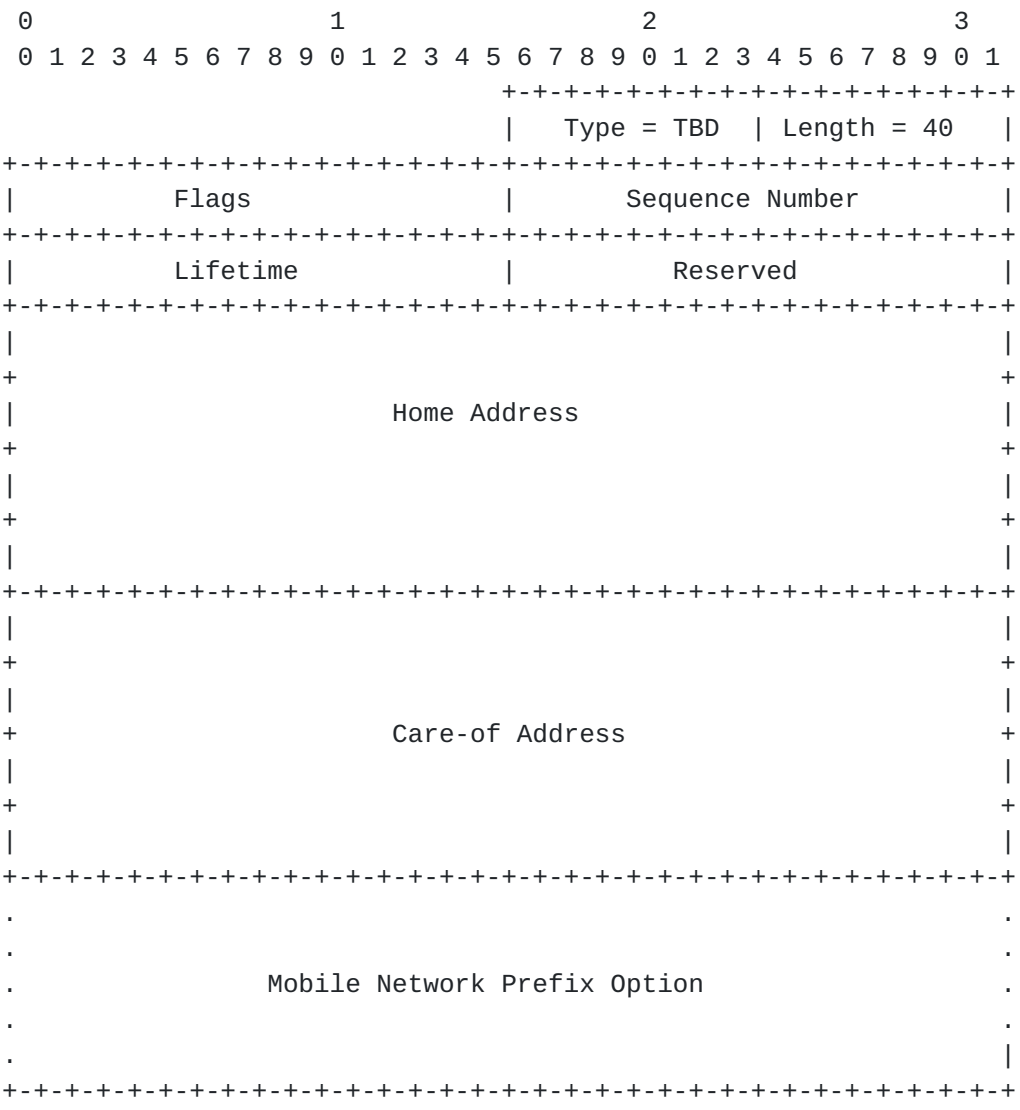

Figure 8: Binding Cache Information Option

The fields of Home Address, Care-of Address, Flags, Sequence Number, and Lifetime are copied from the registered binding of a particular mobile node or mobile router. The 8-bit Reserved field MUST be set to zero. If the R-flag is set to indicate this binding cache entry is for a mobile router, then this option will be immediately followed by one or more Mobile Network Prefix options.

5.2.3. AAA Information Option

This option is used to carry the AAA state of the mobile node's Mobile IPv6 sessions. The AAA state information can be carried in RADIUS or Diameter AVP formats including the user and session info. This information option is only valid in a State Synchronization

message.

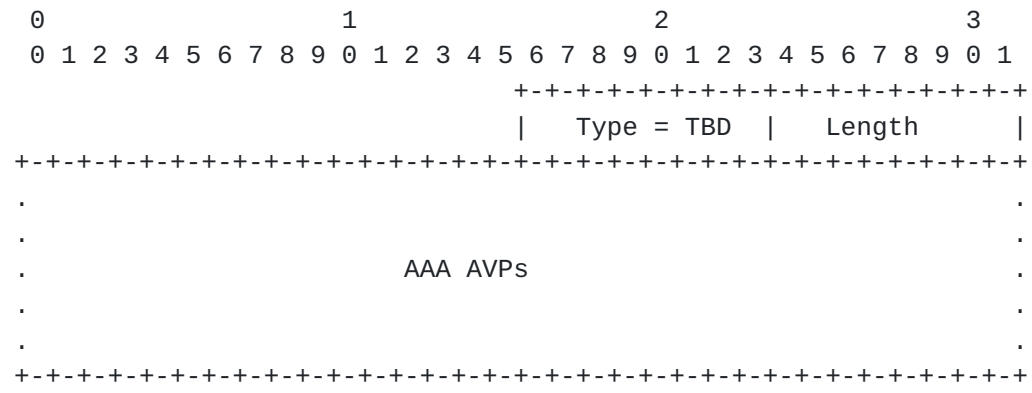

Figure 9: Vendor Specific Inforamtion Option

Type

8-bit Type value. The value is TBD.

Length

8-bit length value.

AAA AVPs

Series of TLV encoded AAA AVPs (including vendor specific AVPs) carrying AAA-related information for each Mobile IPv6 and IPsec/ IKE session.

6. Home Agent Configuration

6.1. Network Configuration

The Home Agent Reliability protocol supports two different configurations for standby home agents. Standby home agents can be placed on the same home link or on a different link.

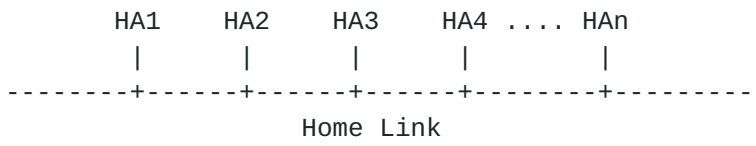

Figure 10: Local Recovery Configuration

Figure 10 depicts the configuration where home agents serving the same home network are located on the same link. For example, HA2, HA3 and HA4 are standby home agents of HA1. This is the same as what Mobile IPv6 defines for home agent configuration.

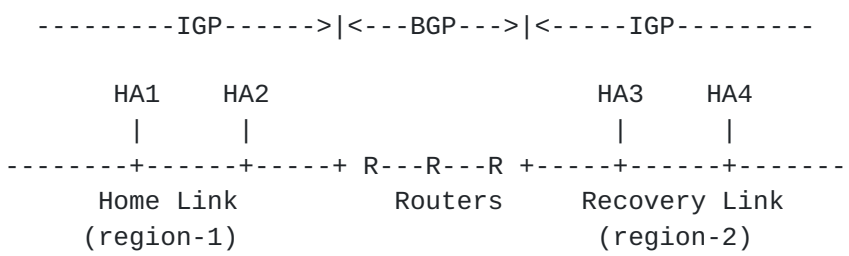

Figure 11: Global Recovery Configuration

Figure 11 illustrates when standby home agents are located on a different link (illustrated as Recovery Link in Figure 11). Most large operators have a very stringent requirement on network availability even in the worst type of disaster or outage. For example, HAs in region-1 are backed up by HAs in region-2. These two regions are geographically separated. If region-1 suffers a downtime due to any reason, all the sessions will be seamlessly taken over by the nodes in region-2. This is called geographic redundancy. This is a well-known configuration for Telecommunications operators. It can achieve home agent recovery even if the entire home link fails. In Figure 11, HA3 and HA4 are standby home agents of HA1 and HA2. In this case, HA3 and HA4 cannot receive packets meant for the home network until the route on the Routers is changed. The routing must

be also updated to direct the packets meant for the home link to the recovery link.

6.2. Address Configuration for Virtual Switch

Each standby home agent obtains its individual IPv6 address from its attached link. This IPv6 address is used by the home agent reliability protocol to exchange information with the associated home agents. The link between home agents should be secured.

The virtual home agent's IPv6 address which is known by the mobile node is shared with the standby home agents. When a home agent fails, the standby home agent activates the IPv6 address of the failed home agent and becomes the active home agent. The standby home agent should not activate the IPv6 address until it knows the active home agent is no longer reachable at the address, otherwise address duplication will occur. To guarantee transparency of the home agent virtual switch to mobile nodes located on the home link, the neighbor cache of the home agent IP address MUST be carefully operated. See Section 8.1 in detail.

6.3. Address Configuration for Hard Switch

Each standby home agent obtains its individual IPv6 address from its attached link. This IPv6 address is used by the home agent reliability protocol to exchange information with the associated home agents. The link between home agents should be secured.

Each home agent configures itself with a different IPv6 address from the same home prefix. This IPv6 address can be used for the Home Agent Reliability protocol if the standby home agents are located at the same link of the active home agent (Figure 10). In case of Figure 11, the router must carefully route packets to the standby home agents as described in Section 8.1. Once a mobile node registers its binding with the active home agent, it may solicit an IPsec/IKE exchange with standby home agents. These packets must be routed to the recovery link. This can be achieved by installing host routes for the standby home agents on the recovery link of the router.

7. Home Agent Common Operation

7.1. Home Agent List Management

In Mobile IPv6 [RFC-3775], each home agent periodically sends router advertisements with the Home Address Information option [RFC-3775] when there are multiple home agents present on a link. This information is managed in a home agent list. For the Home Agent Reliability Protocol, HA-HELLO messages are used to manage the home agent list. There are several reasons to use HA-HELLO message instead of Router Advertisement such as:

1. In the Home Agent Virtual Switch mode, if the standby home agents send unsolicited Router Advertisements to the home link, the mobile nodes attached to the home link are aware of the presence of standby home agents. However, the standby home agents must be hidden until the active home agent fails. HA-Hello messages are exchanged only between home agents.

2. As shown in Section 6.1, standby home agents are not always configured at the same link. In this case, Router Advertisements cannot be sent to the recovery link unless the home link and the recovery link are connected (ex. L2TP). HA-HELLO can be sent beyond the link.

3. The Home Agent Reliability protocol defines to manage additional information such as Group ID and Active/Standby Status of the home agents in the home agent list.

In Mobile IPv6, Router Advertisement are to carry the home agent information to both mobile nodes on the home link and the home agents. On the other hand, in the Home Agent Reliability protocol, HA-Hello is to exchange the information among the home agents and the Router Advertisement is used to notify the information to the mobile nodes. The home agents SHOULD NOT process the Home Agent Information option carried by Router Advertisement if HA-HELLO is available. Operators can define different values to the parameters of the home agent information for HA-HELLO and Router Advertisement. The management operation of the home agent list is unchanged and defined in [RFC-3775].

7.2. Detecting Home Agent Failure

The active and standby home agents can monitor each other in several ways. One method is to reuse other failure detection mechanisms defined in VRRP[RFC-3768] and HSRP [RFC-2281] (see Section 7.6). However, VRRP and HSRP cannot detect the case where the system is running but the Mobile IPv6 stack is not operational. In the Home

Agent Reliability protocol, a new message called HA-HELLO is periodically exchanged in the redundant home agent set as a heart-beat. If HA-HELLO is implemented as a part of Mobile IPv6 stack, it can detect the home agent failure (Mobile IPv6 stack failure). This HA-HELLO can also be exchanged frequently enough to detect the failure promptly and does not introduce any processing overhead to the mobile node attached to the home link.

Failure events used in the Home Agent Reliability protocol are listed below.

Loss of HA-HELLO

In the event that a standby home agent does not receive any HA-HELLO from its peer which is currently the active home agent for a configurable duration, the standby home agent assumes the active home agent's failure. Any home agents can also request the home agent information of the other home agent in the same redundant home agent set by sending HA-HELLO with R-flag set. If HA-HELLO is not replied from the target home agent within a configurable amount of time, that home agent peer is considered to have failed. The detail of the Hello message is described in Section 7.3.

Monitored Server Failure by the Active Home Agent

There may be number of critical servers such as AAA server in the network that are essential for the ongoing Mobile IPv6 sessions at the home agent. Operators can have a policy in place that the active home agent is treated as a failed home agent upon detecting that the link to such servers has failed.

Routing Peer/Link Failure

Operators may require the home agent to detect its next-hop routing peer failure. If the next-hop routing failure is fatal in nature, or due to some other routing policies, the active home agent is treated as a failed home gent and the recovering operation should be started.

7.3. Processing Hello Messages

HA-HELLO MUST be either unicasted or multicasted. A new multicast address (all_homeagent_multicast_address) will be assigned by the IANA. When all the home agents in a redundant home agent set are configured on a same home link, they MUST join the all_homeagent_multicast_address. On the other hand, if a home recovery link is separately defined as described in Figure 11, each home agent SHOULD unicast HA-HELLO.

7.3.1. Requesting Hello Message

A home agent can solicit HA-HELLO to a particular home agent in the same redundant home agent set by unicasting HA-HELLO with the R-flag set. The sender MUST fill the fields of the HA-HELLO with its home agent information. If a home agent needs to request HA-HELLO to all the home agents, it sends the HA-HELLO with R-flag set to the all_homeagent_multicast_address. Requesting HA-HELLO SHOULD be operated when:

1. A new home agent needs to collect the information of all the other home agents in the same redundant home agent set. The HA-HELLO with R-flag set is multicasted to all_homeagent_multicast_address.

2. A home agent entry in the redundant home agent list is about to be removed due to home agent lifetime expiration. The HA-HELLO with R-flag set is unicasted to the home agent whose lifetime is soon expired.

3. HA-HELLO has not been received during the specified hello interval. The HA-HELLO with R-flag set is unicasted to the target home agent.

7.3.2. Sending Hello Message

Each home agent periodically sends HA-HELLO for the home agent's failure detection. The interval time is configured at each home agent. Each home agent MUST also send a HA-HELLO in following case:

1. when a home agent receives a HA-HELLO with the R-flag set

2. When a home agent detects its local information such as home agent preference, home agent lifetime, and registration status change.

3. When a new home agent boots up, it SHOULD solicit Hello messages by multicasting a Hello message with the R-flag set in parallel with sending its own Hello message.

When a home agent sends HA-HELLO, the following rule MUST be applied.

- Whenever a home agent generates HA-HELLO, it MUST increment in the Sequence Number. The Sequence Number SHOULD be initialized to zero for the first Hello message. To accomplish sequence number rollover, if the sequence number has already been assigned to be the largest possible number representable as a 16-bit unsigned integer, then when it is incremented it will then have a value of

zero (0).

- It MUST also specify its own Group ID in HA-HELLO.
- If a home agent is the active home agent, it MUST set the A-flag in HA-HELLO.
- In the Home Agent Hard Switch mode, the source IPv6 address of HA-HELLO MUST be the home agent address.
- In the Home Agent Virtual Switch mode, the home agent local address MUST be used.

7.3.3. Receiving Hello Message

When a home agent receives HA-HELLO, it SHOULD verify the HA-HELLO as follows:

- If the HA-HELLO is not protected by IPsec ESP, it SHOULD be discarded. Note that, only if the HA-HELLO is sent on a dedicated link between the home agents, IPsec protection might not be always required. This depends on the operational policy.
- If HA-HELLO is sent from non global IPv6 address, it MUST be discarded.
- If the source IPv6 address of HA-HELLO is not belong to one of the home agents in the redundant home agent set, the HA-HELLO MUST be ignored.
- If the Group ID field of the received HA-HELLO and the receiver's Group ID is different, HA-HELLO MUST be discarded. HA-HELLO MUST NOT be sent to home agents whose Group ID is different from the sender.
- If the Sequence Number value in the HA-HELLO is equal to or less than the last received Sequence Number value stored in the home agent list entry, the HA-HELLO MUST be discarded.
- HA-HELLO satisfying all of above tests MUST be processed by receiver. The receiver copies home agent information in HA-HELLO to the corresponding home agent list entry. The home agent address of the sender is retrieved from the Source Address field of the IPv6 header of the HA-HELLO.
- If the home agent lifetime field in the HA-HELLO is set to 0, the receiver removes the sender from the home agents list.

- If the R-flag is set in the received HA-HELLO, the receiver MUST send a new HA-HELLO to the originator as described in Section 7.3.2.

7.4. Processing State Synchronization Messages

It is necessary for standby home agents to synchronize the state information of each mobile node registered with the active home agent. In the Home Agent Hard Switch mode, it is not necessary for the home agents to synchronize the complete binding cache information. The standby home agent needs the mapping information of the active home agent and the mobile node. The information is used to send the Home Agent Switch messages to all the mobile node served by the failed home agent.

7.4.1. Requesting State of a Particular Mobile Node(s)

When a home agent needs the state information for a particular mobile node or a subset of mobile nodes, it sends a SS-REQ message constructed as follows:

- It MUST set the Type field to 0 (Request).
- It MUST set a random value in the Identifier field that does not coincide with any other currently pending Requests.
- It MUST include an IP address mobility option(s) which subtype is set to the home address if the target is mobile node(s).
- It MUST include a Mobile Network Prefix mobility option(s) for mobile router(s).
- It MUST set the unspecified address (0::0) in the Home Address mobility option if it solicits the state of all the mobile nodes and mobile routers registering at the receiver of SS-REQ (i.e.destination of SS-REQ).
- In the Home Agent Virtual Switch, the sender of the SS-REQ MUST be a home agent local address of one of the home agents in the same redundant home agent set.
- In the Home Agent Hard Switch, the sender of the SS-REQ MUST be a home agent address of one of the home agents in the same redundant home agent set.
- The destination of the SS-REQ MUST be the active home agent for the requesting home address or mobile network prefix. The standby home agent MUST NOT reply the SS-RREP to the sender.

When a home agent receives the SS-REQ, it MUST verify if SS-REQ is constructed with the above rules. If SS-REQ satisfy all the above tests, the receiver of the SS-REQ MUST reply SS-REP including the state information of the requested mobile node(s) and/or mobile network prefix(es) as described in Section Section 7.4.2.

7.4.2. Synchronizing State

State synchronization messages SHOULD be sent when:

1. The active home agent receives SS-REQ.

2. The active home agent creates a binding cache entry for a particular mobile node.

3. The active home agent deletes a binding cache entry for a particular mobile node.

The active home agent MAY additionally send state synchronization message in following cases:

1. The active home agent update the state information for all sessions that changed since the last update in a periodic interval

2. Only for the Home Agent Virtual Switch mode, the active home agent updates a binding cache entry for a particular mobile node whenever the binding cache entry is updated. In the Home Agent Hard Switch mode, standby home agents only need the mapping information of a home address of the mobile node/router and the home agent address of the active home agent to which the mobile node/router is currently registering. This mapping is used to send a Home Agent Switch message.

If an active home agent sends a State Synchronization message whenever the local state information changes, such as a binding cache change, the number of the State Synchronization messages sent can be quite large.

All the state information of the requested mobile nodes is stored in the SS-REP. Following rules must be applied when the active home constructs SS-REP.

- If the SS-REP is sent in response to the SS-REQ, the active home agent MUST copy the Identifier field of the State Synchronization request message to the Identifier field in the SS-REP. Otherwise, it MUST set the Identifier field to 0.

- When the active home agent stores the state of multiple mobile nodes in a SS-REP, a Binding Cache Information option is used as a separator. For each mobile node, a Binding Cache Information option is placed first, followed by any other options such as AAA option. When the next Binding Cache Information option is reached in the State Synchronization message, it indicates the information of a different mobile node.

- If the unspecified address is found in the Home Address mobility option carried with the SS-REQ, the active home agent MUST return the state of all the active mobile nodes and mobile routers by the SS-REP. The IP fragmentation can be occurred depending on the total size of all the states.

- A SS-REP MUST be authenticated and encrypted by IPsec ESP.

- The destination and source home agents MUST belong to the same redundant home agent set.

- In the Home Agent Hard Switch, the IPv6 source address MUST be set to the home agent address of the sender.

- In the Home Agent Virtual Switch, the IPv6 source address MUST be set to the home agent local address of the sender.

When a home agent receives a SS-REP, it MUST verify whether the SS-REP is constructed with the above rules or not. If the SS-REP does not satisfy all the rules above, it is discarded. Otherwise, the following operation must be taken.

- The receiver of SS-REP MUST update its binding cache and all other necessary information such as AAA and vendor specific information in the particular database.

- In the Home Agent Hard Switch mode, the receiver MUST record the IPv6 address of the sender as the active home agent of the mobile node.

7.4.3. Reliable Transmission by Explicit Acknowledgement

Signaling messages of the Home Agent Reliability protocol are not guaranteed reliable transmission due to the Mobility Header use. This is not always critical, because the link between home agents is carefully managed as stable and reliable. However, operators may need more explicit notification to confirm the message exchanges between home agents. This specification provides an optional acknowledgment to SS-REP messages.

If an active home agent requires an acknowledgment of SS-REP, it set the Ack flag in the SS-REP. The receiver of such SS-REP will send back a SS-ACK. The receiver MUST copy the Identifier value received in the SS-REP into SS-ACK in order to match the SS-REP and SS-ACK.

7.5. Processing Home Agent Control Messages

7.5.1. Standby Home Agent becomes an Active Home Agent

When a standby home agent decides to become an active home agent, the standby home agent sends a SO-REQ to the active home agent. This message MUST be unicasted to the active home agent and MUST be encrypted and authenticated by IPsec ESP. The active home Agent MUST NOT generate this message.

When an active home agent receives a SO-REQ, it MUST operate the following verification and operations:

- If the SO-REQ is not protected by IPsec, it MUST be discarded.
- If the receiver of the SO-REQ is not an active home agent, it MUST send a SO-REP with the Status field set to 130 (Not active home agent).
- If the sender home agent does not belong to the same redundant home agent set, a SO-REP message MUST be sent to the sender with the Status field set to 132 (Not in same redundant home agent set).
- If the receiver is an active home agent, there is case where the active home agent cannot be standby home agent. In such case, the active home agent can reply a SO-REP with the Status field set to 129 (Administratively prohibited).
- Otherwise, the active home agent MUST become a standby home agent and reply with a SO-REP message with the Status field set to 0 (Success).

The SO-REP MUST be also protected by IPsec ESP. Otherwise, the message MUST be silently discarded. If the receiver of SO-REP did not send a SO-REQ message (i.e. unexpected SO-REP), the message MUST be ignored. If the Status field of the SwitchOver Reply message is 0 (Success), the receiving standby home agent immediately becomes an active home agent as described in Section 8.2. If the value in the Status field is greater than 128 an error has occurred. In this case, the receiver MUST NOT attempt to be an active home agent.

7.5.2. Active Home Agent becomes in-active

When an active home agent decides to become a standby home agent, it sends a SB-REQ to one of standby home agent. The reason for the active home agent to send this message can be administrative intervention, and events like Monitored Server Failure by the active home agent or Routing Peer/Link Failure. This message MUST be unicasted to one of the standby home agents and MUST be encrypted and authenticated by IPsec ESP. A standby home agent MUST NOT generate this message.

When a home agent receives a SwitchBack Request message, it first verifies the message.

- If the SwitchBack Request message is not protected by IPsec ESP, it MUST be discarded.

- If the sender home agent of the SB-REQ is not an active home agent, the receiver MUST reply a SB-REP with the Status field is set to 130 (Not active home agent).

- If the sending home agent does not belong to the same redundant home agent set, a SB-REP MUST be sent in which the Status field set to 132 (Not in same redundant home agent set).

- Otherwise, the receiving home agent MUST send a SB-REP with the Status field is set to 0 (Success).

After sending the SwitchBack reply, it MUST NOT become an active home agent immediately. This is because the active home agent is still active until it receives the SB-REP which is acknowledging the SB-REQ. The standby home agent SHOULD change to active at least after LINK_TRAVERSAL_TIME.

If a home agent receives a SB-REP, it MUST be protected by IPsec ESP, otherwise the message MUST be silently discarded. If the receiving home agent did not send a SB-REQ matched to the received SB-REP, the message MUST be silently discarded. If the Status field of the SB-REP is 0 (Success), the active home agent immediately becomes a standby home agent. The sender home agent of SB-REP becomes active home agent as described in Section 8.2. If the value in the Status field is greater than 128, the receiver of SB-REP (active home agent) cannot become a standby home agent and MUST continue to be an active home agent.

7.6. Interworking with VRRP

VRRP and HSRP specify an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. This operation is similar to the Home Agent Virtual Switch operation. For example, the VRRP router controlling the IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail over in the forwarding responsibility should the Master become unavailable. Although VRRP is used to guarantee home agent address reachability, it cannot be used for state synchronization and explicit switching of Master and Backup. Thus, the Home Agent Reliability protocol cannot be replaced by VRRP. This section explains how VRRP can interwork with the Home Agent Reliability protocol.

When VRRP is available, VRRP can replace the Hello message described in Section 5.1.3. However, some of information is missed by using VRRP. After receiving a VRRP message, each home agent SHOULD process the message and store the information as if it receives Home Agent Hello messages Section 7.3.3. The Home Agents SHOULD still perform binding cache synchronization as described in Section 7.4 and SHOULD support the Home Agent Switch message as described in Section 9.2.

In addition to this, VRRP is useful only if all home agents are located on the same link. If the home agents are topologically separated, the Home Agent Reliability protocol MUST be used.

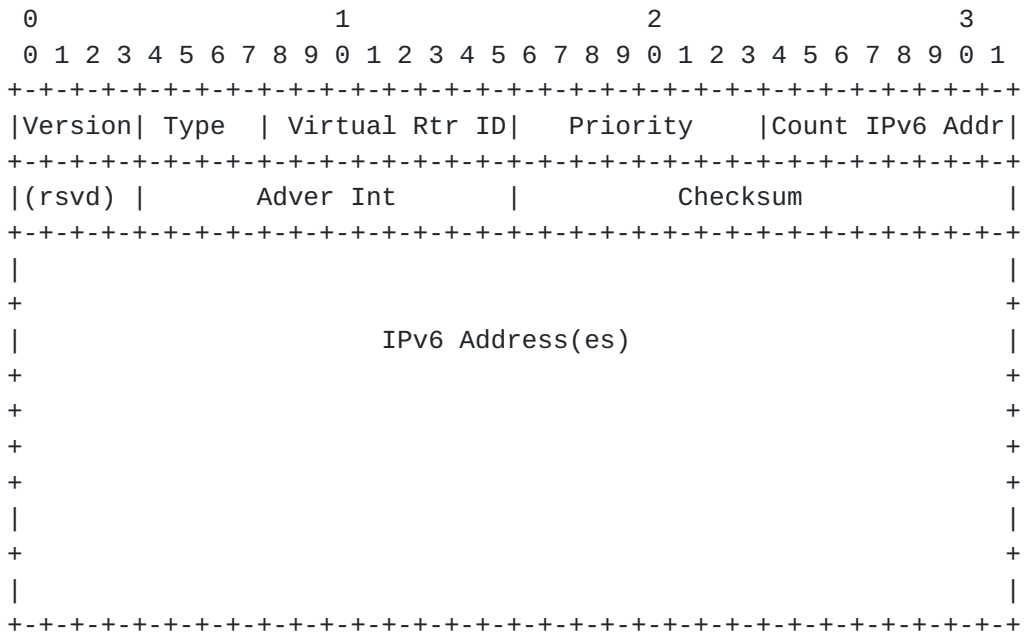

Figure 12: VRRP Packet Format

The message format of VRRP is described in Figure 12. Each field is mapped as follows:

Virtual Rtr ID

Group ID is stored in the Virtual Rtr ID field.

Priority

Home Agent Preference is stored in the Priority field. Note that VRRP only has 8 bits for the Priority field. Therefore, values larger than 255 MUST NOT be assigned to the preference value.

Count IPv6 IPv6 Addr

This field MUST be always be 1.

Advert Int

This field MUST be mapped to the Hello Interval field of the Home Agent Hello message, though it only has 12 bytes.

IPv6 address

A home agent address is stored in this field.

Home Agent Lifetime, Sequence Number and Flags field are missing in the VRRP packet format. Therefore, operators SHOULD use the same statically configured value for Home Agent Lifetime. Each home agent does not check freshness of received VRRP message because of no sequence number. If VRRP is used, a home agent cannot determine the active home agent from the VRRP message due to lack of A flag, and cannot request a VRRP advertisement to other home agents.

7.7. Retransmissions and Rate Limiting

Standby and active home agents are responsible for retransmissions and rate limiting of a SS-REQ, SO-REQ, SB-REQ messages for which they expect a response. The home agent MUST determine a value for the initial transmission timer:

- If the home agent sends a SS-REQ message, it SHOULD use an initial retransmission interval of INITIAL_STATE_SYNC_REQ_TIMER.

- If a standby home agent sends a SO-REQ message, it SHOULD use an initial retransmission interval of INITIAL_SWICHOVER_REQ_TIMER.

- If an active home agent sends a SB-REQ message, it SHOULD use an initial retransmission interval of INITIAL_SWICHBACK_REQ_TIMER .

If the sending home agent fails to receive a valid matching response within the selected initial retransmission interval, it SHOULD retransmit the message until a response is received. All of the above constants are specified in Section 11.

The retransmission MUST use an exponential backoff process as described in [RFC-3775] until either the home agent receives a response, or the timeout period reaches the value MAC_HARELIABILITY_TIMEOUT (16sec). The home agent SHOULD use a separate back-off process for different message types and different destinations. The rate limiting of Mobility Header messages is the same as one in [RFC-3775]. A home agent MUST NOT send Mobility Header Messages to a particular home agent more than MAX_UPDATE_RATE (3) times a second, which is specified in [RFC-3775].

8. Home Agent Virtual Switch

8.1. Consideration of Routing and Neighbor Discovery Protocol

This section gives a brief explanation of how a home agent interacts with routing and Neighbor Discovery Protocol (NDP) when the Home Agent Virtual Switch mode is used.

When a standby home agent becomes active in the Home Agent Virtual Switch mode, it MUST start to advertise the home agent address and the home prefix of the home addresses serviced by the redundant home agent set into the routing infrastructure. This operation is normally done using a route selector such as BGP or an OSPF modifier. For example, we can use the AS_Path prepend operation for BGP, and the Metric field in OSPF for the route selection. When each home agent participates in OSPF routing, each home agent should be configured with the appropriate metric matched to the home agent preference value. When the active home agent fails, OSPF detects the failure and can dynamically switch the route to the standby home Agent based on the OSPF cost value. If this cost conflicts with the home agent preference value due to configuration errors, the routers on the home link may not route packets to the desired standby home agent. In order to change the OSPF cost correctly and dynamically, The operator takes other existing approaches. For example, most of router vendors have a private MIB to set the cost via SNMP, though this is a vendor-specific function.

When an active home agent activates a home agent address, it SHOULD use a virtual MAC address as introduced in [RFC-3768]. When the active home agent is changed, the neighbor cache of the active home agent is not necessarily updated on mobile nodes located on the home link. Otherwise, the new home agent MUST update the neighbor cache entry for the home agent address on all the mobile nodes located on the home link. In addition, Mobile IPv6 uses proxy NDP to intercept packets meant for mobile nodes which are away from the home link. However, it is unnecessary for the new active home agent to overwrite the existing proxy neighbor entries of the mobile nodes.

8.2. Home Agent Recovery

After detecting the active home agent has failed, the standby home agent whose preference value is the highest MUST take over the failed home agent. The standby home agent MUST activate the virtual home agent address. If a virtual MAC address as introduced in [RFC-3768] is used, the standby home agent MUST start using that virtual MAC address as well. Since all the necessary state has already been transferred to this standby home agent before the active home agent failed, it can immediately start acting as the active home agent.

9. Home Agent Hard Switch

9.1. Home Agent Recovery

After detecting the active home agent has failed, the standby home agent whose preference value is the highest MUST take over the failed home agent. The standby home agent MUST send a Home Agent Switch message to all the mobile nodes that were registered at the failed home agent as described in Section 9.2, using the pre-established IPsec SA. The standby Home Agent MUST set its own address in the Home Agent Address field in the Home Agent Switch message so that it will receive the binding update from the mobile node as an acknowledgment of the sent Home Agent Switch message. The home agent switch-over is complete when it receives binding updates from all the mobile nodes. It is important to remark that sending Home Agent Switch messages to all the mobile nodes at once may bring non-negligible overhead to the home agent.

This overhead cannot be avoided if the active home agent suddenly stop serving mobile node because of unexpected reasons (crash, network trouble, etc). However, if this switch over is operated under the administrative operation (maintenance, etc), the previous active home agent may continue serving the mobile nodes until the switch over is completed. Until the mobile node sends a binding update to the new active home agent, it still sends the packet to the previous home agent in the Home Agent Hard Switch. Therefore, the new active home agent can notify the completion of switch-over to the previous active home agent by using Home Agent Control message as described in Section 9.3. As soon as this message is received, the previous active home agent can be shutdown or detached from the network safely.

9.2. Sending Home Agent Switch Messages

The standby home agent which is going to be active MUST send a Home Agent Switch message as defined in [RFC-5142] to all the mobile nodes that were being served by the failed home agent. The Home Agent Switch message must be securely sent to the mobile node by using IPsec ESP. The standby home agent MUST include only its own home agent address in the Home Agent Switch message. If there are a large number of mobile nodes served by the failed home agent, the overhead sending Home Agent Switch messages is high. Until a mobile node receives this Home Agent Switch messages, the mobile node's communication is discontinued. Therefore, until the standby home agent completes sending the Home Agent Switch message to all the mobile nodes and receives Binding Updates from all the mobile node, the failed home agent SHOULD serve mobile nodes if possible. This is the case when the active home agent is replaced by administrative

operation with the Home Agent Control messages as described in Section 9.3.

When a failed home agent recovers, it MUST re-establish an IPsec SA with each mobile node served by its redundant home agent set. Otherwise, it cannot be either a standby or active home agent for the mobile nodes. Therefore, as soon as the active home agent detects the recovery of the failed home agent, it sends a Home Agent Switch message with the I-flag set to all the mobile nodes serving by other home agents in the same redundant home agent set, and includes the recovered home agent address in the Home Agent Addresses field. The mobile node will re-key the SA, but it will not change the home agent by this home agent switch message which I-flag is set.

9.3. Notification of Home Agent Switch Completion

If the new active home agent completes the switch-over as described in Section 8.2, it SHOULD send a SW-COMP to the previous active home agent in the Home Agent Hard Switch case. Until the previous home agent receives this message, it SHOULD continue serving any mobile nodes that are registered with it. Once the previous home agent receives the SW-COMP message, it can stop home agent services.

9.4. Mobile Node Operation

9.4.1. Home Agent Addresses Discovery

In the Home Agent Hard Switch mode, a mobile node authenticates itself to two or more home agents and creates IPsec SAs with them during bootstrapping. When the active home agent fails, another home agent can use the pre-existing SA to notify the mobile node about the failure by sending a Home Agent Switch message.

In order to discover multiple home agent addresses, two different mechanisms are defined in the bootstrapping solution in the split scenario [RFC-5026]. One is DNS lookup by home agent Name, the other is DNS lookup by Service Name. DHCPv6 can also be used in the integrated scenario [ID-BOOTINT] to provide home agent provisioning to mobile nodes.

In the split scenario, a mobile node can use DNS lookup by Service Name to discover the home agents, as defined in [RFC-5026]. For example, if home agent reliability is required by a mobile node, DNS lookup by Service Name method is recommended for the mobile node to discover multiple home agents addresses. Therefore, mobile nodes will query the DNS SRV records with a service name of mip6 and protocol name of ipv6. The DNS SRV records includes multiple home agent addresses and different preference values and weights. The

mobile node SHOULD choose two or more home agents from the home agents list according to their preference value. Then the mobile node should authenticate itself to these home agents via an IKEv2 exchange.

In the integrated scenario, a mobile node can use DHCPv6 to get home agent provisioning from an MSP or ASP, as already defined in [ID-BOOTINT]. The only requirement is that the DHCPv6 response must include multiple home agents' information in order to support home agent reliability.

9.4.2. IKE/IPsec pre-establishment to Home Agents

After a mobile node obtains multiple home agent addresses, it needs to trigger multiple IKE exchanges with the multiple home agents selected from the home agent list. Since both IKEv1 and IKEv2 can be used to bootstrap Mobile IPv6, this solution does not introduce any new operations to co-operate with IKEv1 or IKEv2. It should initiate IKE for home agents as soon as home registration is complete.

The mobile node MUST follow the standard IKEv2 exchange in the bootstrapping solution of the split scenario [RFC-5026]. Home Address configuration maybe also be included, if necessary, for the first IKE exchange. After its Home Address is assigned or approved by the first home agent, mobile node SHOULD register itself with the second home agent with IKE using the same Home Address. Therefore, no home address configuration should be used in the second IKEv2 procedure. Note that the mobile node only sends a Binding Update message to the first home agent.

9.4.3. Receiving Home Agent Switch message

A mobile node must follow the operation specified in [RFC-5142] when it receives a Home Agent Switch message.

If the I-flag is set in the received Home Agent Switch message, the mobile node MUST re-key the SA with the home agent addresses stored in the Home Agent Addresses field. The mobile node MUST NOT change its active home agent when the I-flag is set. If the home agent address is not known from the bootstrapping described in Section 9.4.1, the mobile node MUST NOT start an IKE session with the unknown home agent. Instead, it SHOULD re-start home agent discovery again to update its home agent address information.

When the mobile node receives a Home Agent Switch message without I-flag set, and if the message contains the IPv6 address of a standby home agent, it MUST select the standby home agent in the switch message as the active home agent and send a new Binding Update

message to it. Note that the standby home agent address in the Home Agent Switch MUST be equal to the sender of the Home Agent Switch message. The standby Home agent expects the Binding Update as an acknowledgment of the Home Agent Switch message. The mobile node already has a pre-established SA with the standby home agents and should use that SA to send the Binding Update. If the address stored in the Home agent address field is different from the sender, the mobile node MUST send a binding update to the sender.

10. Security Considerations

Since Mobile IPv6 operation requires ESP in transport mode between the mobile node and the home agent, we will discuss the ESP field synchronization issues between the mobile node and the redundant set of home agents. This synchronization is required only for Home Agent Virtual Switch mode. Most of fields should be synchronized based on [RFC-4301]. The ESP header has the following fields:

SPI

This field identifies the SAD at the receiver.

The mobile node negotiates only one IPsec SA. Hence, the SPI value will remain unchanged upon home agent failover.

Sequence Number

This field is used for "anti-replay" feature of ESP. The transmitter must include this monotonically increasing number. The receiver may process the sequence number based on local policy.

The mobile node and the redundant home agent set will have the same set of sequence numbers for transmit and receive. Hence, synchronization of the sequence number field is mandatory in this mode of operation.

The SA1, SA2, SA3, SA4 could be synchronized between the home agents as these messages are not sent continuously. Moreover for the Binding Update case, if the mobile node is in the middle of sending a Binding Update to an active home agent for a binding refresh, and the active home agent is not available at that moment, the mobile node will not get any response from the active home agent. After a standby home agent becomes active, the mobile node will retry and it will receive the Binding Update from the mobile node with a sequence number that is +n from its last known sequence number for SA1. For the Binding Acknowledgment case (SA2), the standby home agent SHOULD add a random number to the last known sequence number over and above the replay window to ensure that the packet passes the replay check at the mobile node. The same applies to HoTi and HoT messages with SA3 and SA4. Note that this windowing of the sequence numbers for Mobile IPv6 signaling is only needed to cover the corner cases when Binding Update or HoTi is in-flight and the active home agent fails.

The technique explained above should work for user data packets if ESP is used to encrypt user data traffic as well. The actual switchover time and the routing infrastructure convergence time is the only latency that the user may perceive.

Initialization Vector

Since the Initialization Vector will be delivered in each exchange between a mobile node and home agent, this field is not necessarily synchronized between home agents.

Others

Other fields should be synchronized based on RFC4301 [RFC-4301]

In the Home Agent Hard Switch mode, the standby home agent needs to send a Home Agent Switch message using IPsec encryption. Since the mobile node has pre-established an IPsec SA with both the active and standby home agents, the standby home agent can send the message to the mobile node with the pre-established IPsec SA.

11. Protocol Constants

INITIAL_STATE_SYNC_REQ_TIMER: 3sec

INITIAL_SWICHOVER_REQ_TIMER: 1sec

INITIAL_SWICHBACK_REQ_TIMER 1sec

LINK_TRAVERSAL_TIME 150msec

12. IANA Considerations

- The values for following mobility header message MUST be assigned by IANA.
  - State Synchronization Message
  - Home Agent Control Message
  - Home Agent Hello Message
  - Home Agent Switch Message
- The values for following mobility options MUST be assigned by IANA.
  1. Binding Cache Information Option
  2. AAA Information Option
- New Option Code for the IP address option defined in [RFC-5268]

13. Additional Authors

This document is a result of discussions in the Mobile IPv6 Home Agent Reliability Design Team. The members of the design team that are listed below are authors that have contributed to this document:

Samita Chakrabarti

samita.chakrabarti@azairenet.com

Kuntal Chowdhury

kchowdhury@starentnetworks.com

Hui Deng

denghui@chinamobile.com

Vijay Devarapalli

vijay.devarapalli@azairenet.com

Sri Gundavelli

sgundave@cisco.com

Brian Haley

brian.haley@hp.com

Behcet Sarikaya

bsarikaya@huawei.com

Ryuji Wakikawa

ryuji@sfc.wide.ad.jp

14. Acknowledgements

This document includes a lot of text from [ID-LOCALHAHA] and [ID-HAHA]. Therefore the authors of these two documents are acknowledged. We would also like to thank the authors of the home agent reliability problem statement [ID-PS-HARELIABILITY] for describing the problem succinctly and Alice Qin for her work on the hello protocol.

15. References

15.1. Normative References

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

[RFC-3775] Johnson, D., Perkins, C., Arkko, J., "Mobility Support in IPv6", RFC 3775, June 2004.

[RFC-3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.

[RFC-5094] Devarapalli, V., "Mobile IPv6 Vendor Specific Option", RFC 5094, October 2007.

[RFC-5142] Haley, B., "Mobility Header Home Agent Switch Message", RFC-5142, November 2007.

[RFC-5026] Giaretta, G., "Mobile IPv6 bootstrapping in split scenario", RFC 5026, October 2007.

[ID-BOOTINT] Chowdhury, K. and A. Yegin, "MIP6-bootstrapping via DHCPv6 for the Integrated Scenario", draft-ietf-mip6-bootstrapping-integrated-dhc-06 (work in progress), April 2008.

15.2. Informative References

[RFC-3768] Hinden, R., "Virtual Router Redundancy Protocol (VRRP)", RFC 3768, April 2004.

[RFC-2281] Li, T., Cole, B., Morton, P., and D. Li, "Cisco Hot Standby Router Protocol (HSRP)", RFC 2281, March 1998.

[RFC-3776] Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", RFC 3776, June 2004.

[RFC-4301] Kent, S. and Atkinson, R., "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC-3753] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.

[RFC-5268] Koodli, R., "Mobile IPv6 Fast Handovers", RFC 5268, June

2008.

[ID-HAHA] Wakikawa, R., "Inter Home Agents Protocol Specification", draft-wakikawa-mip6-nemo-haha-spec-01 (expired), March 2006.

[ID-LOCALHAHA] Devarapalli, V., "Local HA to HA protocol", draft-devarapalli-mip6-nemo-local-haha-01 (expired), March 2006.

[ID-PS-HARELIABILITY] Faizan, J., "Problem Statement: Home Agent Reliability", draft-jfaizan-mipv6-ha-reliability-01 (expired), February 2004.

Appendix A. Change Log From Previous Versions

Changes from draft-ietf-mip6-hareliability-03

o Only Editorial Update and No Technical Change

Author's Address

Ryuji Wakikawa
Toyota ITC
6-6-20 Akasaka, Minato-ku
Tokyo 107-0052
Japan

Phone: +81-3-5561-8276
Fax: +81-3-5561-8292
Email: ryuji@jp.toyota-itc.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).