

MEXT Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 11, 2011

R. Wakikawa (Editor)
Toyota ITC USA.
May 10, 2011

Home Agent Reliability Protocol (HARP)
draft-ietf-mip6-hareliability-09.txt

Abstract

The home agent can be a single point of failure when Mobile IPv6 and its associated supporting protocols are 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 mechanisms of home agent failure detection, home agent state transfer, and home agent switching for home agent redundancy and reliability.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 11, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
1.1.	Terminology	4
1.2.	Problem Statement and Requirements	6
2.	Protocol Overview	8
3.	Home Agent Configuration	11
3.1.	Network Configuration	12
3.2.	Home Agent Address Configuration	13
4.	Home Agent Operations	13
4.1.	Home Agent List Management	13
4.2.	Detecting Home Agent Failure	14
4.3.	Processing the HARP Messages	15
4.3.1.	IP field and Security Descriptions of HARP message	15
4.3.2.	Processing Home Agent Hello (HA-HELLO)	16
4.3.3.	Processing Home Agent Switch Over (SWO-REQ/REP)	17
4.3.4.	Processing Home Agent Switch Back (SWB-REQ/REP)	19
4.4.	State Synchronization	20
4.4.1.	Binding Cache Information Management	21
4.4.2.	IP field and Security Descriptions of State Synchronization message	21
4.4.3.	Requesting State of Mobile Nodes (SS-REQ)	21
4.4.4.	Sending State Information (SS-REP)	22
4.4.5.	Synchronizing State (SS-REP and SS-ACK)	23
4.5.	Switching the Active Home Agent	24
4.6.	Consideration of Routing and Neighbor Discovery Protocol (VHARP)	25
4.7.	Interworking with VRRP	26
4.8.	Retransmissions and Rate Limiting	26
5.	Mobile Node Operation	27
5.1.	Home Agent Addresses Discovery	27
5.2.	IPsec/IKE Establishment to Home Agents	28
5.3.	Synchronizing State: K-bit treatment	28
5.4.	Receiving Home Agent Switch message	29
6.	Messages Format	29
6.1.	New Mobility Header Messages	29
6.1.1.	HARP Message Format	29

6.1.2.	State Synchronization Message Format	33
6.1.3.	Home Agent Rekey Message	35
6.2.	New Mobility Options	36
6.2.1.	Binding Cache Information Option	36
6.2.2.	State Synchronization Status Option	38
6.2.3.	AAA Information Option	39
7.	Security Considerations	40
8.	Protocol Constants	40
9.	Protocol Configuration Variables	40
10.	IANA Considerations	40
11.	Additional Authors	41
12.	Acknowledgements	42
13.	References	42
13.1.	Normative References	42
13.2.	Informative References	43
	Author's Address	44

1. Introduction

In Mobile IPv6 [[RFC-3775](#), [ID-3775bis](#)] and its derivative protocols like NEMO Basic Support [[RFC-3963](#)] and Dual Stack Mobile IPv6 [[RFC-5555](#)], if a home agent loses binding cache state or even network connectivity due to its failure, or some other reason, the result is 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 have uninterrupted service even when one home agent crashes or loses state. The Home Agent Reliability Protocol (HARP) is designed to manage standby home agents, and switch service from an active to a standby home agent in the case of an active home agent failure.

1.1. 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 both to a mobile host [[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:

Home Agent Reliability Protocol (HARP)

A protocol between Mobile IPv6 home agents that provides reliability by moving state and service from an active home agent to a standby, in the case of an active home agent failure. HARP can accomodate multiple home agents being placed on the same home link, or on different links, by grouping them into a redundant home agent set. One of home agents is selected as an active home agent. If the active home agent fails, a standby home agent can take over and become the active home agent. Since each home agent is assigned individual home agent addresses, a mobile node is aware of home agent failures and needs to register its binding to the new active home agent again.

Virtual Home Agent Reliability Protocol (VHARP)

A protocol between Mobile IPv6 home agents which provides reliability by cloning an active home agent. Unlike HARP, the standby home agents are an exact copy of the active home agent, including home agent IP address. It is similar to the virtual router concept of VRRP [[RFC-3768](#), [RFC-5798](#)] and HSRP [[RFC-2281](#)].

The VHARP operations are transparent to a mobile node.

Active Home Agent

A home agent that is currently serving the mobile nodes.

Standby Home Agent

A home agent which can 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.

Active Home Agent Address

An IPv6 address of the Active Home Agent.

Standby Home Agent Address

An IPv6 address of the Standby Home Agent.

Redundant Home Agent Set

A grouping which includes an active and standby home agent(s). The Group Identifier is used to identify a redundant home agent set. Operators need to configure a unique value per redundant home agent set.

Virtual Home Agent Address

A home agent address shared among home agents in a redundant home agent set. It is similar to virtual router address specified in VRRP [RFC-3768, [RFC-5798](#)]. The address is only activated on an active home agent.

Home Agent Preference

This preference value was 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. A home agent SHOULD NOT share the same preference value with other home agents. Meanwhile, operators can also define an independent value for the home agent reliability protocol. It is useful when operators want to assign different operational policies to the preference values of DHAAD and the

Home Agent Reliability Protocol.

New Messages

Home Agent Reliability Protocol (HARP) message defined in [Section 6.1.1](#):

SwitchOver Request (SWO-REQ)

SwitchOver Reply (SWO-REP)

SwitchBack Request (SWB-REQ)

SwitchBack Reply (SWB-REP)

Switch Complete (SW-COMP)

Home Agent HELLO (HA-HELLO)

State Synchronization (SS) message defined in [Section 6.1.2](#):

State Synchronization Request (SS-REQ)

State Synchronization Reply (SS-REP)

State Synchronization Reply-Ack (SS-ACK)

[1.2.](#) Problem Statement and Requirements

In Mobile IPv6 [RFC-3775, [RFC-4877](#)], a mobile node registers and establishes a binding with only one home agent. 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 fail-over 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 generic 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 agent(s). This includes the Binding Cache, AAA information, and other Mobile IPv6 and NEMO related information. Note that the Home Agent Reliability Protocol only exchanges state information for active 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 [RFC-4285], synchronizing the configurations of the Authentication protocol is out of scope in this document. Operators MAY correctly set the configuration information 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. (Note this is out of scope in this document.)

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, [RFC-5798](#)] or HSRP [[RFC-2281](#)], it can re-use that mechanism to detect home agent failure. In addition, periodic Hello messages are introduced in

this document to detect active an home agent's service availability.

Failure Notification

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

2. Protocol Overview

HARP works when one or more home agents are provisioned on a home link, or different links, and these are grouped into a redundant home agent set. One home agent is selected as the active home agent and receives binding updates from mobile nodes. According to [RFC- 3775, [RFC-4877](#)], an active home agent maintains not only binding cache information, but also IPsec/IKEv2 states per mobile node, because Mobile IPv6 relies on IPsec for securing the signaling, and optionally user plane traffic.

If the active home agent fails, all state information associated with a mobile node is lost. As a result, all mobile nodes served by the failed home agent will be disconnected. In HARP, other home agents, called standby home agents, exchange the required information with the active home agent. In case of a failure of the active home agent, HARP can let a standby home agent take over for the failed home agent with this information about the active mobile nodes.

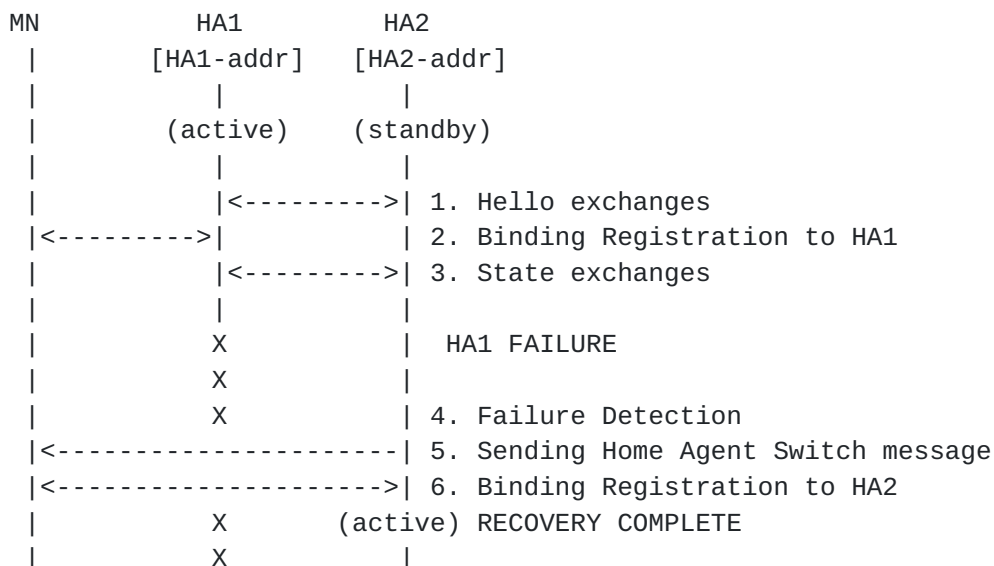


Figure 1: Overview of Home Agent Reliability Protocol (HARP)

Figure 1 shows an example of the HARP operations. HA1 and HA2 belong to the same redundant home agent set and are assigned with an individual IP address (HA1-addr and HA2-addr) at the home link. Each home agent can be seen as an individual home agent by mobile nodes. All the home agents periodically send a Hello message (named HA-HELLO) to exchange the home agent information, as well as monitor the state of the active home agent (1). The mobile node registers its binding with only the active home agent (2). The active home agent synchronizes the active mobile node information with the other standby home agents periodically (3).

HARP introduces the new HA-HELLO message for failure detection, but it may use any type of information to detect that failure. After detecting the failure of the active home agent (4), the standby home agent whose preference value is the highest takes over for the failed home agent. Once completed, the standby home agent sends a Home Agent Switch message to all the mobile nodes that were registered at the failed home agent (5). The standby home agent puts its own address in the Home Agent Address field of 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 (6). For protecting the Home Agent Switch message, the mobile node should have an IPsec Security Association (SA) with the standby home agent before the failover. The mobile node may pre-establish multiple IPsec SAs with all the home agents.

Although the active home agent manages IPsec/IKEv2 states per mobile node, HARP does not offer any recovery mechanism of these states by itself. IPsec/IKE state synchronization is out of scope in this document. If IPsec/IKEv2 state can be recovered from the active home agent on the standby home agent, HARP can be operated in a slightly different manner called Virtual-HARP (VHARP). Unlike HARP, the standby home agents are an exact copy of the active home agent. It is similar to the virtual router concept of VRRP [RFC-3768, [RFC-5798](#)] and HSRP [RFC- 2281]. Note that HARP is mandatory and VHARP is optional in this document. VHARP is shown in Figure 2.

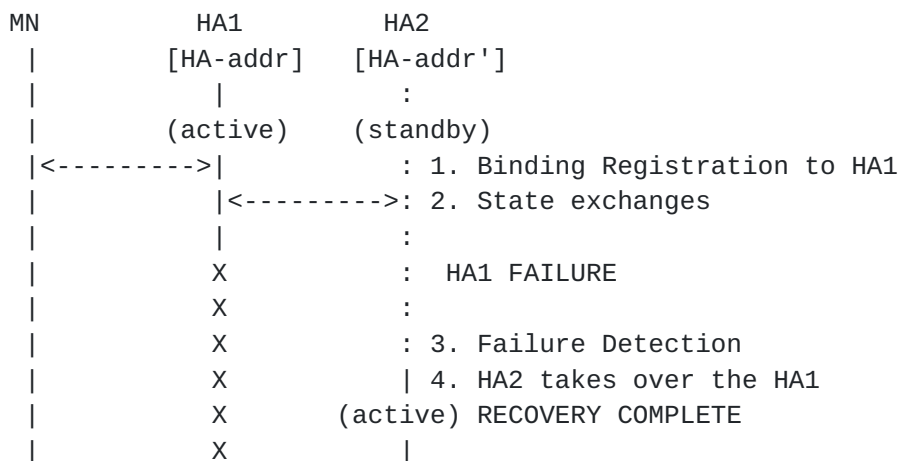


Figure 2: Overview of Virtual Home Agent Reliability Protocol (VHARP)

All the home agents (HA1 and HA2) in the redundant home agent set share a virtual home agent address (HA-addr), and routing ensures only the active home agent will be reachable using that virtual home agent address. After a mobile node's binding registration (1), the active home agent pushes the states of all of its mobile nodes to the other standby home agents (2). In VHARP, all the states of a mobile node need to be synchronized. For example, information such as the Binding Cache, and Authentication, Authorization, and Accounting (AAA) information.

After detecting the active home agent has failed (3), the standby home agent whose preference value is the highest takes over the failed home agent. The standby home agent activates the virtual home agent address on its interface attached to the home link. The virtual home agent address activation can be operated by VRRP. Since all the necessary states of mobile nodes have already been transferred to this standby home agent, the standby home agent can immediately start acting as the active home agent (4). Unlike HARP, the mobile node is not required to re-register its binding to a new active home agent. The mobile node may use the IKEv2 resumption mechanism [[RFC-5723](#)] to resume it's IPsec SA with the new active home agent.

This document offers a new management mechanism of active and standby home agents by using a new Mobility Header (MH) message called a HARP message as shown in Figure 3. This mechanism can be used in both HARP and VHARP. Each home agent exchanges its own home agent information with the other home agents in its redundant home agent set by a Home Agent HELLO message (HA-HELLO) (1). The HA-HELLO message can also be used to monitor the availability of the active home agent.

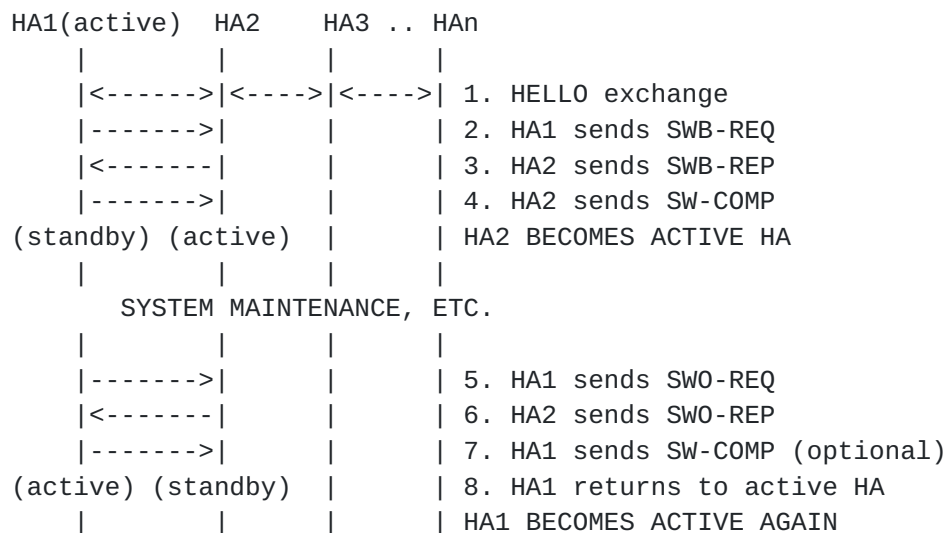


Figure 3: Home Agent Management

In some scenarios, the active home agent may need to stop serving mobile nodes for system maintenance. This specification enables manual intervention for home agent management. As shown in Figure 3, the active home agent (HA1) sends a SwitchBack Request message (SWB-REQ) to a standby home agent (HA2) (2). HA2 will acknowledge the message by sending a SwitchBack Reply message (SWB-REP) to HA1 (3). In HARP operation, it could take a long time to complete home agent failover since all mobile nodes must re-register with the new home agent. During this failover operation, HA1 may continue serving the mobile nodes until the switch-over is completed. When HA2 decides the switch-over has completed, it MAY send an optional message, SW-COMP, to HA1 (4). As soon as HA2 sends the SW-COMP, it becomes the active home agent. HA1 becomes a standby home agent when it receives SW-COMP. If SW-COMP is not used, HA2 and HA1 change their status appropriately.

After maintenance is complete and HA1 is back online, HA1 sends a SwitchOver Request (SWO-REQ) to HA2 in order to become the active home agent again (5). HA2 acknowledges it by sending a SwitchOver Reply (SWO-REP) back to HA1 (6). HA1 now starts the home agent failover operation. After the switch-over is complete, HA1 sends a SW-COMP to HA2 (7). Then, HA1 becomes the active home agent and HA2 becomes a standby home agent (8).

3. Home Agent Configuration

3.1. Network Configuration

HARP supports two different configurations for standby home agents. Standby home agents can be placed on the same home link or on different links. Figure 4 depicts the configuration where home agents serving the same home network are located on the same link as defined in [RFC-3775].

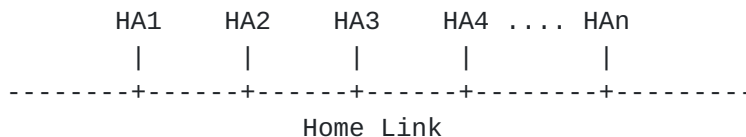


Figure 4: Local Recovery Configuration

Figure 5 illustrates when standby home agents are located on different links (illustrated as Recovery Link in Figure 5). Most large operators have a very stringent requirement on network availability even in the worst type of disaster or outage. This configuration can achieve home agent recovery even if the entire home link fails. This is called geographic redundancy, and is a well-known configuration for Telecommunications operators. In Figure 5, home agents (HA1-HA4) are placed in geographically separated regions (region-1 and region-2). If region-1 suffers down time for any reason, all the sessions will be seamlessly taken over by the nodes in region-2. Note that HA3 and HA4 cannot receive packets meant for the home network until the route on the Routers is changed. The routing must also be updated to direct the packets meant for the home link to the recovery link.

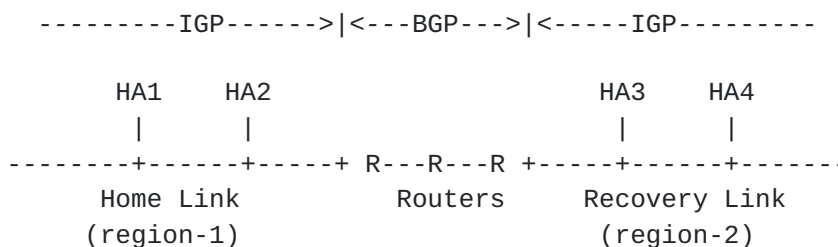


Figure 5: Global Recovery Configuration

3.2. Home Agent Address Configuration

In HARP, each home agent obtains its individual IPv6 address from its serving home prefix. In VHARP, all the home agents use a virtual home agent address generated from the home prefix.

In addition, each home agent running VHARP needs to obtain its individual IPv6 address from its attached link. This IPv6 address is used only for VHARP operations between home agents and is not revealed to mobile nodes for binding registration.

All the home agents MUST join the ALL_HA_MULTICAST_ADDR. In VHARP, each home agent joins the multicast group with its individual IPv6 address, but not with the virtual home agent address. This multicast address can be used to exchange HA-HELLO messages among the home agents. Alternatively, if a remote home recovery link is defined, each home agent unicasts the HARP messages to home agents configured at the remote recovery link.

4. Home Agent Operations

4.1. Home Agent List Management

In Mobile IPv6, each home agent periodically sends router advertisements with the Home Address Information option [[RFC-3775](#)]. HARP introduces a HARP HA-HELLO message to replace the router advertisement. There are several reasons to use HA-HELLO messages instead of a Router Advertisement such as:

- o A HA-HELLO message can be sent beyond the link, while a router advertisement cannot. In case of geographic redundancy, Router Advertisements cannot be sent to the recovery link unless the home link and the recovery link are virtually connected, for example, by L2TP.
- o A HA-HELLO message is defined to manage additional information, such as Group ID and Active/Standby Status of the home agents in the home agent list.
- o A HA-HELLO message is exchanged only between home agents, while a Router Advertisement is also processed by mobile nodes attached to a home link. A HA-HELLO does not introduce any burden to the mobile nodes even if it is frequently sent on the home link.

When a HA-HELLO is used to exchange home agent information, each home agent SHOULD NOT process the Home Agent Information option carried by a Router Advertisement. Router Advertisements are only processed by

mobile nodes. Operators may define different configuration values to the parameters of the home agent information for a HA-HELLO and a Router Advertisement.

This document requires additional information to be added the conceptual Home Agents list defined in [\[RFC-3775\]](#). The additional information is learned through HA-HELLO message exchange.

- o Group ID of a redundant home agent set. It is learned through the Group ID field of the HA-HELLO.
- o HA-HELLO Interval. This value is locally configured at every home agent by operators, and is learned through the Hello Interval field of the HA-HELLO.
- o Individual home agent addresses used in the VHARP operation. This information is only required when VHARP is used in addition to the virtual home agent address. It is learned through the Source Address of the HA-HELLO message.
- o VHARP capability. This information is learned through the V flag of the HA-HELLO message.
- o Current mode (HARP or VHARP). This information is learned through the M flag of the HA-HELLO message.
- o Active status. This information is learned through the A flag of the HA-HELLO message.

4.2. Detecting Home Agent Failure

Active and standby home agents can monitor each other in several ways. One method is to reuse other failure detection mechanisms defined in VRRP [\[RFC-3768, RFC-5798\]](#) and HSRP [\[RFC-2281\]](#). However, VRRP and HSRP are not sufficient since they cannot detect the case where the system is running, but the Mobile IPv6 stack is not operational. Failure events used in HARP/VHARP are listed below.

Loss of HA-HELLO

HARP/VHARP is an extension to Mobile IPv6, and can monitor the availability of Mobile IPv6 services on each home agent by periodically sending a HA-HELLO as a heart-beat. This HA-HELLO can be exchanged frequently enough to detect a failure without any additional overhead to mobile nodes attached to the home link. In the event that a standby home agent does not receive any HA-HELLOs from its peer for a configurable duration of time, the standby home agent assumes the peer home agent has failed. Details of the

Hello message are described in [Section 4.3.2](#).

Monitored Server Failure by the Active Home Agent

There may be a number of critical servers, such as an AAA, in the network that are essential for ongoing Mobile IPv6 sessions at the home agent. Operators can have a policy in place for which 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 agent and the recovery operation should be started.

[4.3](#). Processing the HARP Messages

[4.3.1](#). IP field and Security Descriptions of HARP message

The HARP message format is defined in [Section 6.1.1](#). If a HARP message is unicast, the destination address **MUST** be one of the Home Agents in the same Redundant Home Agent set. If it is a HA-HELLO message, designated by setting the type field to 4, it can be multicast. The destination address **MUST** be set to the ALL_HA_MULTICAST_ADDR address. The source address **MUST** be set to the sender's home agent address. Note that in VHARP, the virtual home agent address **SHOULD NOT** be set to the source or destination address. Instead, the IP address of the interface the packet is being sent from **SHOULD** be used.

If a HARP message is unicast, it **SHOULD** be secured by IPsec ESP. If a HA-HELLO message is multicast, multicast extensions to IPsec [RFC-5374] **SHOULD** be applied. If all the home agents are placed in a secure transport network to exchange a HARP message, authentication and encryption **MAY** be omitted. Which security verification is used depends on operational policy. If security verification fails for a received HA-HELLO, the HA-HELLO **MUST** be discarded.

The following operations **MUST** be performed when transmitting a HARP message:

- o The incremented latest Sequence Number **MUST** be set in the Sequence Number field. 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).

- o The sender's Group ID MUST be set in the Group ID field.
- o The V-flag MUST be set if the sender is capable of VHARP.
- o The M-flag MUST be unset if the sender is operating in HARP mode.
- o The M-flag MUST be set if the sender is operating in VHARP mode.
- o The A-flag MUST be set if the sender is the active home agent.

The following functions MUST be performed when a HARP message is received:

- o The Group ID in the HARP message MUST match the receiver's Group ID.
- o The source address of the HARP message MUST belong to a home agent in the receiver's redundant home agent set.
- o The M-flag MUST match the receiver's operating mode.
- o The Sequence Number value in the HARP message MUST be larger than the last received Sequence Number value. When the sequence number rollover occurs, the sequence number value in the HA-HELLO MUST be zero.

If any one of the above checks fails, the receiver SHOULD discard the HARP message.

4.3.2. Processing Home Agent Hello (HA-HELLO)

4.3.2.1. Sending HA-Hello Messages

Each home agent MUST send a HA-HELLO in the following cases:

- o UNSOLICITED: Each home agent SHOULD periodically send a HA-HELLO. The time interval is configured locally at each home agent.
- o UNSOLICITED: When a home agent detects its local information has changed it should immediately send a HA-HELLO.
- o SOLICITED: When a home agent receives a HA-HELLO with the R-flag set, the HA-HELLO can be sent to the destination home agent.

A home agent can solicit a HA-HELLO from a particular home agent(s) in the same redundant home agent set by unicasting or multicasting a HA-HELLO with the R-flag set. Soliciting a HA-HELLO happens when:

- o A new home agent boots up. The new home agent SHOULD solicit HA-Hello messages by multicasting a HA-Hello message with the R-flag set to 1.
- o If a HA-HELLO has not been received after the specified Hello Interval, a HA-HELLO MAY be solicited to the home agent.
- o A home agent entry in the redundant home agent set is about to be removed due to home agent lifetime expiration. A HA-HELLO SHOULD be solicited from the home agent whose lifetime is soon expired.

In addition to [Section 4.3.1](#), the following operations MUST be performed when transmitting a HA-HELLO.

- o The Type field MUST be set to 4.
- o The R-flag MUST be set if the sender is soliciting a HA-HELLO from the other home agent(s).
- o The appropriate home agent configuration values MUST be copied to the Home Agent Preference, the Home Agent Lifetime, and Hello Interval fields.

[4.3.2.2](#). Receiving Hello Messages

The receiver MUST perform the verification of the HA-HELLO described in [Section 4.3.1](#). After the verification, the receiver copies the values stored in the HA-HELLO message to the corresponding home agent list entry according to [Section 4.1](#).

If the home agent lifetime field in the HA-HELLO is set to 0, the receiver MUST remove the sending home agent 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 4.3.2.1](#).

[4.3.3](#). Processing Home Agent Switch Over (SWO-REQ/REP)

When a standby home agent decides to become an active home agent, the standby home agent sends a SwitchOver Request (SWO-REQ) to the current active home agent in the following way:

- o It MUST be unicast to only the current active home agent.
- o It MUST be sent from a standby home agent. The active home agent MUST NOT generate this message.

When an active home agent receives a SWO-REQ, it MUST do the following verifications and operations, in addition to what is described in [Section 4.3.1](#)

- o If the receiver of the SWO-REQ is not an active home agent, it MUST send a SWO-REP with the Status field set to 130 (Not active home agent).
- o If the sending home agent does not belong to the same redundant home agent set, a SWO-REP message MUST be sent to the sender with the Status field set to 132 (Not in same redundant home agent set).
- o If there are any other reasons that the receiver cannot accept the SWO-REP, the active home agent MUST reply a SWO-REP with the Status field set to 129 (Administratively prohibited).
- o Otherwise, the active home agent MUST become a standby home agent and reply with a SWO-REP message with the Status field set to 0 (Success).

When a standby home agent receives a SWO-REP, it MUST do the following verifications and operations, in addition to what is described in [Section 4.3.1](#):

- o If the receiver is an active home agent, the SWO-REP MUST be discarded.
- o If the standby home agent receives an unsolicited SWO-REP which is not in reply to an SWO-REQ it has sent, it MUST ignore the SWO-REP.
- o Otherwise, if the Status field of the SWO-REP is 0 (Success), the standby home agent (the receiver of the SWO-REP) immediately becomes an active home agent.
- o 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.

4.3.4. Processing Home Agent Switch Back (SWB-REQ/REP)

When an active home agent decides to become a standby home agent, it sends a SWB-REQ to one of the standby home agents. The reason for the active home agent sending this message might be due to an administrative intervention, or an event like Monitored Server Failure by the active home agent, or due to a Routing Peer/Link Failure. The following operations MUST be performed when SWB-REQ is sent:

- o It MUST be unicast to only one of the standby home agents in the same redundant home agent set.
- o It MUST be sent from an active home agent. A standby home Agent MUST NOT generate this message.

When a home agent receives a SWB-REQ message, it verifies the message as follows:

- o If the sending home agent of the SWB-REQ is not an active home agent, a SWB-REP MUST be sent in which the Status field is set to 130 (Not active home agent).
- o If the sending home agent does not belong to the same redundant home agent set, a SWB-REP MUST be sent in which the Status field is set to 132 (Not in same redundant home agent set).
- o Otherwise, the receiving home agent MUST send a SWB-REP with the Status field set to 0 (Success).
- o After sending the SWB-REP, the standby home agent MUST NOT become an active home agent immediately. This is because the active home agent is still active until it receives the SWB-REP acknowledging the SWB-REQ it sent. The standby home agent SHOULD change to active after LINK_TRAVERSAL_TIME. The default value of LINK_TRAVERSAL_TIME is defined in [Section 9](#).

When a home agent receives a SWB-REP message, it verifies the message as follows:

- o If the standby home agent received an unsolicited SWB-REP not in reply to its own SWB-REQ, it SHOULD ignore the SWB-REP.
- o If the Status field of the SWB-REP is 0 (Success), the active home agent should immediately become a standby home agent. The sending home agent of SWB-REP becomes an active home agent after LINK_TRAVERSAL_TIME.

- o If the value in the Status field is greater than 128, the receiver of the SWB-REP (active home agent) cannot become a standby home agent and MUST continue to be an active home agent.

4.4. State Synchronization

The State Synchronization (SS) message format is defined in [Section 6.1.2](#). It can carry multiple aspects of the state information associated with a mobile node by setting mobility options in the Mobility Options field. The following list shows examples of the mobility options which can be specified in the state synchronization message:

- o IPv6 Home Address (Binding Cache Option)
- o Binding Cache Information (Binding Cache Option)
- o NEMO Mobile Network Prefix (Mobile Network Prefix Option [RFC-3963])
- o IPv4 Care-of Address (IPv4 Care-of Address Option [[RFC-5555](#)])
- o IPv4 Home Address (IPv4 Home Address Option [[RFC-5555](#)])
- o Binding Identifier (Binding Identifier Option [[RFC-5648](#)])
- o AAA states (AAA Information Option)
- o Miscellaneous states (Vendor Specific Mobility Option [[RFC-5094](#)])

When a home agent needs to send the state of multiple mobile nodes in a single state synchronization message (SS-REQ or 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 related to the mobile node, if necessary.

In HARP, since a mobile node will re-register to the new active home agent after a home agent failover, it is not necessary for the standby home agents to synchronize all the mobile nodes' state information. The standby home agents only need to collect the home address information of all the mobile nodes served by the active home agent. The information is used to send Home Agent Switch messages to all the mobile nodes when a home agent failure occurs.

In the case of VHARP, home agent fail-over is accomplished without the mobile nodes having to perform re-registration. Therefore, standby home agents need to copy the complete state information of each mobile node registered with the active home agent.

4.4.1. Binding Cache Information Management

In HARP, each standby home agent learns the partial binding cache information such as a pair of a home address and a mobile node's registering home agent address.

In VHARP, a standby home agent ideally copies the received binding cache information and other mobile node's information into the appropriate database so that it can act as an active home agent as soon as it takes over the failed home agent.

4.4.2. IP field and Security Descriptions of State Synchronization message

A state synchronization message is unicast. The destination address MUST be one of the home agents in the same Redundant Home Agent set. The source address MUST be set to the sender's home agent address. Note, that in VHARP, the virtual home agent address MUST NOT be set to the source address, the IP address of the interface the packet is being sent from SHOULD be used.

The message SHOULD be secured by IPsec ESP. If all the home agents are placed in a secure transport network to exchange the state synchronization message, authentication and encryption MAY be omitted. If security verification fails for a received state synchronization message, the message MUST be discarded. The choice of security mechanism used depends on the operational model of the network.

4.4.3. Requesting State of Mobile Nodes (SS-REQ)

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:

- o It MUST set the Type field to 0 (Request).
- o It MUST set a random value in the Identifier field that does not coincide with any other currently pending Requests.
- o It MUST include a Binding Cache Information option(s) in which the Home Address field is set to the target home address. Other fields of the Binding Cache Information option can be omitted.
- o If the request is for the state of all the mobile nodes registered at the destination home agent for the SS-REQ message, it MUST set the Home Address field of the Binding Cache Information option to the unspecified address (::).

- o If the sender is requesting information about multiple mobile nodes, it MUST include multiple binding cache information options in a single SS-REQ. The sender SHOULD NOT send multiple SS-REQs per mobile node.
- o It MUST send a SS-REQ to the active home agent of the target mobile node(s).

When a home agent receives a SS-REQ, it MUST perform the verification described in [Section 4.4.2](#) and the following:

- o If the receiver does not have binding cache information for the target mobile node(s) specified in the received Binding Cache Information option(s), it MUST ignore the SS-REQ and MUST NOT reply with a SS-REQ.
- o Otherwise, the receiver MUST reply with a SS-REP, including all the state information of the target mobile node(s).

4.4.4. Sending State Information (SS-REP)

An SS-REP message(s) SHOULD be sent when:

1. The active home agent receives an SS-REQ.
2. The active home agent creates or deletes a binding cache entry for a particular mobile node.

The active home agent MAY additionally send an SS-REP message in the following cases:

1. The active home agent updates the state information for all sessions that have changed since the last update in a periodic interval.
2. Often in VHARP, the active home agent MAY update a binding cache entry for a particular mobile node whenever the binding cache entry is updated. If an active home agent sends an SS-REP message whenever the local state information changes, such as a binding cache change, the number of the SS-REP messages can be quite large.

The following rules must be applied when the active home agent constructs a SS-REP:

- o It MUST copy the Identifier field of the SS-REQ to the same field of the SS-REP, if the SS-REP is sent in response to the SS-REQ.

- o It MUST set the Identifier field to zero (0) if the SS-REP is sent without solicitation (no SS-REQ).
- o It MUST include the required mobility options in the SS-REP.
 - * In HARP, a partial Binding Cache Information option (the Home Address Field only) MUST be included in the SS-REP.
 - * In VHARP, a full Binding Cache Information option, and other required options shown in [Section 6.1.2](#), MUST be included in the SS-REP.
- o It MUST include the state of all the active mobile nodes registered at the active home agent in the SS-REP when the unspecified address is found in the Home Address mobility option carried with the SS-REQ. The message may be fragmented depending on the total size needed to carry all states.

4.4.5. Synchronizing State (SS-REP and SS-ACK)

When a home agent receives a SS-REP, it MUST take the following operations:

- o If no options are carried in the SS-REP, the home agent MUST ignore the SS-REP.
- o If the sender of SS-REP is not in the same global home agent set, the home agent MUST reject the SS-REP and MUST send SS-ACK with the Status Synchronization Status option in which status value is set to [130: Not in same global home agent set]
- o The receiver MUST record the IPv6 address of the sender as the active home agent of the mobile node in its local binding cache.
- o The receiver MUST update its binding cache, and all other necessary information, in its database(s).
- o If the A-flag is set in the SS-REP, the receiver MUST reply with an SS-ACK.

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

4.5. Switching the Active Home Agent

In HARP, the standby home agent which is going to be active MUST send a Home Agent Switch message [[RFC-5142](#)] to all the mobile nodes that were being served by the failed home agent. The following rules MUST be applied when transmitting a Home Agent Switch message.

- o MUST use IPsec ESP to protect the Home Agent Switch message.
- o MUST set the address of the standby home agent address who is the sender of this Home Agent Switch message in the Home Agent Address field of the Home Agent Switch message [[RFC-5142](#)].

If there are a large number of mobile nodes served by the failed home agent, the overhead of sending Home Agent Switch messages is high. This overhead cannot be avoided if the active home agent suddenly stopped serving mobile nodes due to an unexpected reason (crash, network trouble, etc). However, if this switch-over is an administrative operation (maintenance, etc), the previous active home agent may continue serving the mobile nodes until the switch-over is complete. Until the mobile node sends a binding update to the new active home agent, it still sends packets to the previous home agent.

When the new active home agent completes the switch-over, it SHOULD send a SW-COMP to the previous active home agent. 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 be shutdown or detached from the network safely.

In VHARP, after detecting the active home agent has failed, the standby home agent whose preference value is the highest MUST take over for the failed home agent. The standby home agent MUST activate the virtual home agent address and its virtual MAC address. A virtual MAC address as introduced in [[RFC-3768](#), [RFC-5798](#)] SHOULD be used in VHARP. If VHARP is run with VRRP and HSRP as described in [Section 4.7](#), the virtual home agent address can be treated as a virtual router address in VRRP and HSRP. Therefore, VRRP and HSRP can automatically activate the virtual home agent address on the standby home agent after their election mechanism has completed. 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.

When the failed home agent is restarted and wants to become the active home agent again, it MUST re-establish an IPsec SA with each mobile node, as all the mobile nodes will have purged their IPsec SA with the home agent when the failure occurred. Otherwise, it cannot

be 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 Rekey message to all the mobile nodes served 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 detail of the Home Agent Rekey message is described in [Section 6.1.3](#). The mobile node will re-key the SA by using The IKEv2 resumption mechanism [[RFC-5723](#)]. Alternatively, the mobile node MAY start a new IKE session with the recovered home agent.

[4.6. Consideration of Routing and Neighbor Discovery Protocol \(VHARP\)](#)

This section gives a brief explanation of how a home agent interacts with routing and Neighbor Discovery when VHARP is used.

When a standby home agent becomes active in VHARP, 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 creates 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 would use its existing approaches. For example, most router vendors have a private MIB to set the OSPF 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](#), [RFC-5798](#)]. 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 Neighbour Discovery 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.

4.7. 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 VHARP. 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 failover 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 HARP/VHARP.

When VRRP is available, VRRP can replace the Hello message described in [Section 6.1.1](#). However, some information is missing by using just VRRP. After receiving a VRRP message, each home agent SHOULD process the message and store the information as if it had received a Home Agent Hello message, as described in [Section 4.3.2.2](#). The message format of VRRP can be found in [Section 5.1 of \[RFC-5798\]](#). Each field is mapped as follows:

- o Virtual Rtr ID: Group ID is stored in the Virtual Rtr ID field.
- o 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.
- o Count IPv6 IPv6 Addr: This field MUST be always be 1.
- o Max Advert Int: This field MUST be mapped to the Hello Interval field of the Home Agent Hello message, though it only has 12 bytes.
- o IPv6 address: A home agent address is stored in this field.

Home Agent Lifetime, Sequence Number and Flags fields are not present in the VRRP packet format. Therefore, operators SHOULD use the same statically configured value for Home Agent Lifetime. Each home agent does not check the freshness of received VRRP message because there is no sequence number.

4.8. Retransmissions and Rate Limiting

Home agents are responsible for retransmissions and rate limiting of SS-REQ, SWO-REQ, SWB-REQ messages for which they expect a response.

The home agent MUST determine a value for the initial transmission timer:

- o If the home agent sends a SS-REQ message, it SHOULD use an initial retransmission interval of INITIAL_STATE_SYNC_REQ_TIMER.
- o If a home agent sends a SWO-REQ or SWB-REQ message, it SHOULD use an initial retransmission interval of INITIAL_SWITCH_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 8](#).

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. 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 [\[RFC-3775\]](#).

5. Mobile Node Operation

This section describes the operations of a mobile node only when HARP is used. None of the operations in this section are required with VHARP.

5.1. Home Agent Addresses Discovery

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 agent 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.

5.2. IPsec/IKE Establishment to Home Agents

In this document, a mobile node needs to manage an IPsec SA with a home agent(s). The following mechanisms can be used to manage the IPsec SA(s) with a home agent(s):

- o IKEv1/v2 running per home agent (HARP) to establish multiple IPsec SAs for home agents.
- o The IKEv2 resumption mechanism [[RFC-5723](#)] to update an IPsec SA with the new home agent (VHARP)

If an IPsec/IKEv2 state synchronization mechanism is available in Virtual Private Network (VPN) products, none of above is required for the VHARP operation. The IPsec SAs per mobile node are seamlessly copied among multiple home agents.

The mobile node MUST follow the standard IKEv2 exchange in the bootstrapping solution of the split scenario [[RFC-5026](#)]. If multiple IKEv2 operations are run per home agent, the mobile node MUST NOT attempt the home address assignment to standby home agents.

5.3. Synchronizing State: K-bit treatment

When a mobile node moves and the care-of address changes, it can use the Key Management Mobility Capability (K) bit in the Binding Update in order to update the peer endpoint of the key management protocol, for example, the IKE Security Association.

If an active home agent receives a Binding Update with the K-bit set, it MUST process the Binding Update as specified in [[RFC-3775](#)]. In addition, the active home agent MUST notify the other standby home

agents of the care-of address change. To do so, it MUST send a State Synchronization Reply message, including a Binding Cache Information option, to all the other standby home agents. The flags of the Binding Update MUST be copied to the flags field of the Binding Cache Information option. The standby home agents will update the peer endpoint of the key management protocol upon detecting the K-bit it set in the Flag field of the Binding Cache Information option.

If the K-bit is not set in the Binding Update, an active home agent needs to rerun the key management protocol. The active home agent MUST send State Synchronization Reply messages, including Binding Cache Information options, to all the other standby home agents. The flags of the Binding Update MUST be copied to the flags field of the Binding Cache Information option. The standby home agents that receive the State Synchronization Reply message will detect the care-of address change and rerun the key management protocol.

5.4. Receiving Home Agent Switch message

A mobile node must follow the verification and operations specified in [[RFC-5142](#)] when it receives a Home Agent Switch message.

The Home Agent Switch message MUST be securely exchanged between a mobile node and a home agent by using IPsec ESP.

When the mobile node receives a Home Agent Switch message, if the message contains the IPv6 address of a standby home agent, it MUST select the standby home agent as its active home agent and MUST send a new Binding Update message to it.

The standby home agent address in the Home Agent Switch message MUST be equal to the sender of the Home Agent Switch message. If the IPv6 address stored in the Home Agent address field is different from the sender's source IPv6 address, the mobile node MUST send a binding update to the sender and MUST NOT use the IPv6 address in the Home Agent Switch message.

6. Messages Format

6.1. New Mobility Header Messages

6.1.1. HARP Message Format

The HARP message has the type field to identify different roles. The HARP message has the MH Type value TBD.

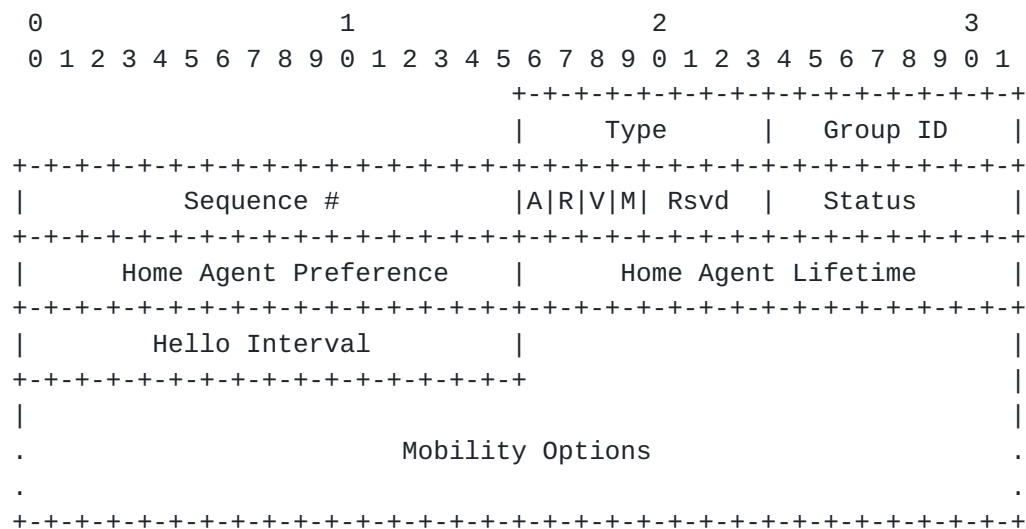


Figure 6: Home Agent Hello Message

Type

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

0: SwitchOver Request (SWO-REQ)

Unicast 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 (SWO-REP)

Used to acknowledge the receipt of the corresponding SWO-REQ.

2: SwitchBack Request (SWB-REQ)

Unicast 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 (SWB-REP)

Used to acknowledge the receipt of the corresponding SWB-REQ.

4: Switch Complete (SW-COMP)

Used to indicate the completion of a switch-over, (i.e. sending Home Agent Switch messages, and receiving binding update messages from all the served mobile nodes).

4: Home Agent HELLO (HA-HELLO)

Used to carry home agent information among the redundant home agent set. MUST be either unicast or multicast. The HA-Hello message is defined for two purpose: 1) an alive check and 2) home agent information exchange.

Group Identifier

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

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.

(A)ctive flag

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

(R)equest 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.

(V)HARP capability flag

VHARP capability Flag. If a home agent is capable of IPsec/IKE state synchronization, it MUST set this flag.

(M)ode flag

A home agent MUST set this flag only when VHARP is used in the current operation. If the flag is unset, the home agent currently operates HARP. (HARP:0, VHARP:1)

Reserved

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

Status

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

0: Success

128: Reason unspecified

129: Administratively prohibited

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

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

132: Not in same redundant home agent set

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.

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 specification does not define any options valid for the HARP message.

6.1.2. State Synchronization Message Format

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

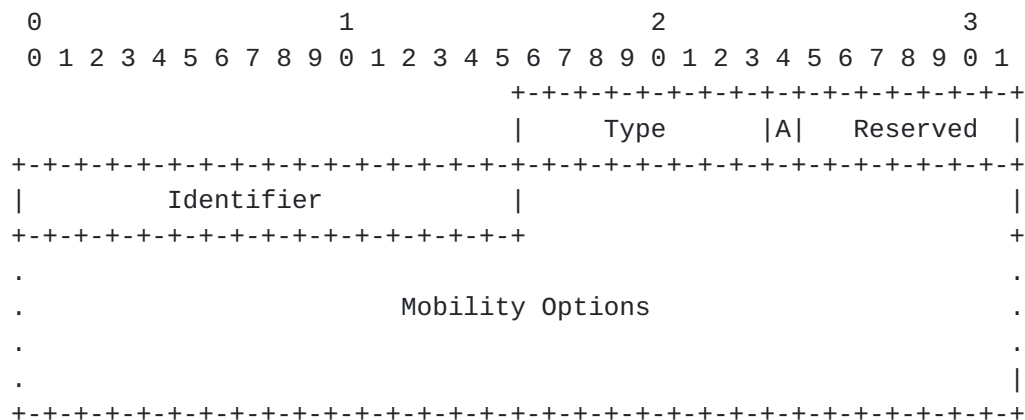


Figure 7: State Synchronization Message

Type

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

0: State Synchronization Request (SS-REQ)

Used to solicit the active state corresponding to a particular mobile node.

1: State Synchronization Reply (SS-REP)

Used between the home agents in the redundant home agent set to exchange binding cache and any other information related to providing mobility service to the mobile nodes. Sent either periodically or in response to a SS-REQ.

2: State Synchronization Reply-Ack (SS-ACK)

This message is optional and is used only when the links between home agents are not reliable.

(A)ck flag

This flag is valid only for SS-REP. If the sender requires explicit acknowledgment by an 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 messages. 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.

Binding Cache Information Option is mandatory in the SS-REQ message. Multiple options can be stored in the same SS-REQ message. A home agent includes the mobile node's home address in the Binding Cache Information option. If a home agent wants to solicit all the active mobile nodes' states, it can include the unspecified address (::) in an IPv6 address option.

Binding Cache Information option is mandatory in SS-REP. SS-REP can carry many mobility options. The following options are just examples.

- * AAA Information Option
- * Vendor Specific Mobility Option [[RFC-5094](#)]
- * Mobile Network Prefix Option [[RFC-3963](#)]
- * IPv4 Care-of Address Option [[RFC-5555](#)]
- * IPv4 Home Address Option [[RFC-5555](#)]
- * Binding Identifier Option [[RFC-5648](#)]

6.1.3. Home Agent Rekey Message

This message is used to indicate that the mobile node SHOULD start an IPsec re-key with the home agent specified in the Home Agent Addresses field. This message is used when a failed home agent recovers and needs to re-establish IPsec SA/IKE state with a mobile node. This message MUST be unicast to a mobile node by the active home agent and MUST be authenticated and encrypted by IPsec ESP. The Home Agent Rekey 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 2.

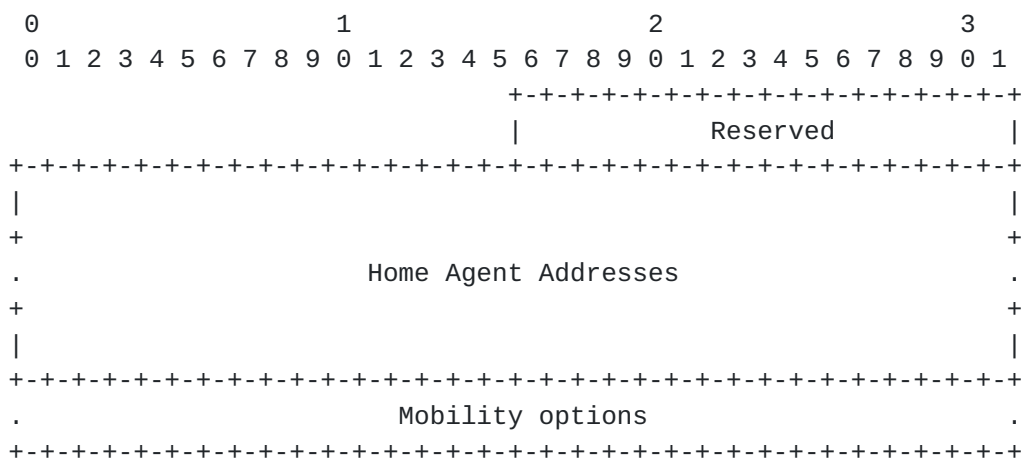


Figure 8: Home Agent Rekey Message

Reserved

A 16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Home Agent Address

The receiver of this message MUST re-key the security association with the specified home agent.

When a mobile node receives a Home Agent Rekey message, it MUST verify the message as follows:

- o The message MUST be sent from the receiver's active home agent. Otherwise, the message MUST be discarded.
- o The message MUST be protected by IPsec ESP. Otherwise, the message MUST be discarded.
- o The message SHOULD contain one of the standby home agent's addresses. If the home agent address is not known from the bootstrapping described in [Section 5.1](#), the mobile node MUST NOT start an IKE re-key session with the unknown home agent. Instead, it SHOULD re-start home agent discovery to update its home agent address information.

If all the above verifications are satisfied, the mobile node MUST re-key the SA with the home agent addresses stored in the Home Agent Addresses field.

[6.2.](#) New Mobility Options

[6.2.1.](#) Binding Cache Information Option

The Binding Cache Information option is used to carry Binding Cache Information of each mobile node. This option is only valid in a State Synchronization message. Its format is as follows:

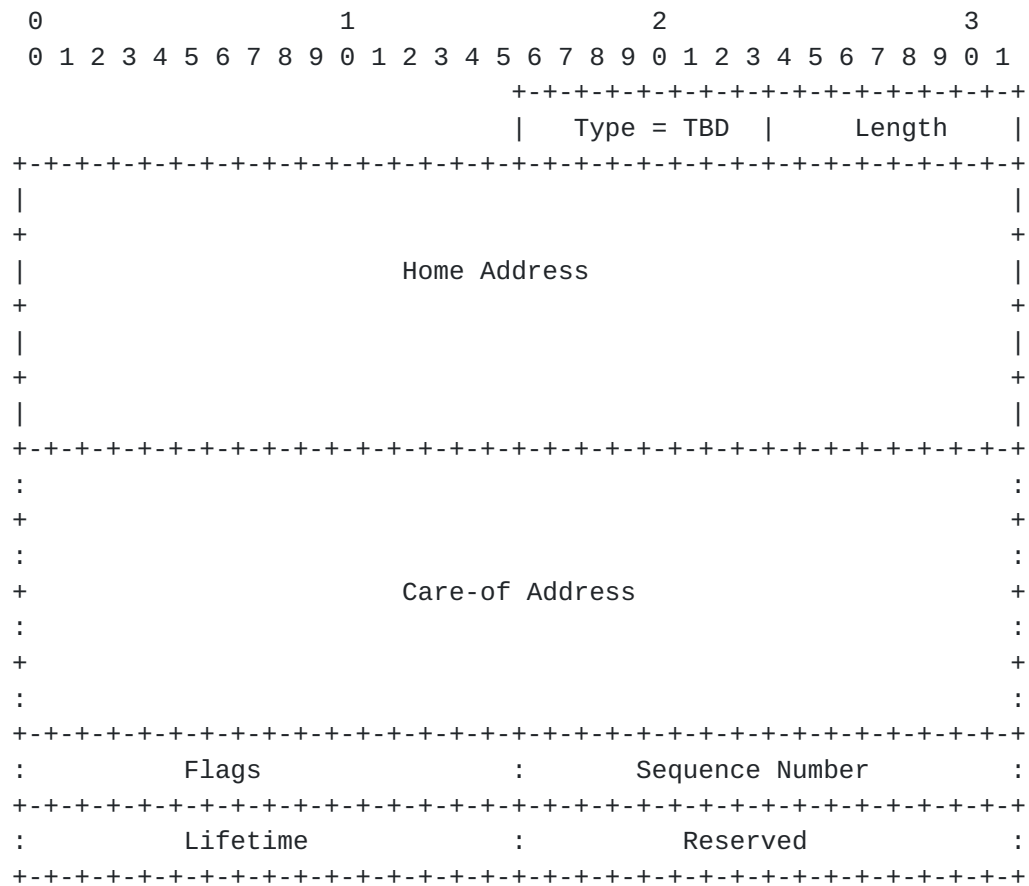


Figure 9: Binding Cache Information Option

Length

8-bit unsigned integer, representing the length in octets of the mobility option, not including the Option Type and Option Length fields. There are two valid length values, 16 and 40, depending on the number of fields in use. The alignment requirement is either $8n+6$ (length 16) or $8n+2$ (length 40).

Home Address

The Home Address of a mobile node.

Care-of Address

Flags

Sequence Number

Lifetime

Optional values only used in VHARP, in which case the corresponding value from the binding cache database of the active home agent is copied into each field.

Reserved

A 16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

6.2.2. State Synchronization Status Option

The State Synchronization Status option is used to carry the status value of an SS-ACK for a received SS-REP. In [ID-HAHA], SS-ACK is mandatory in response of an SS-REP to update global binding registration status.



Figure 10: State Synchronization Status Option

Status

8-bit unsigned integer indicating the status of the SS-REP.

* 0: Success

- * 128: Reason unspecified
- * 129: Malformed SS-REP
- * 130: Not in same global home agent set

Reserved

A 24-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Home Address

Corresponding home address of the mobile node.

6.2.3. 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.

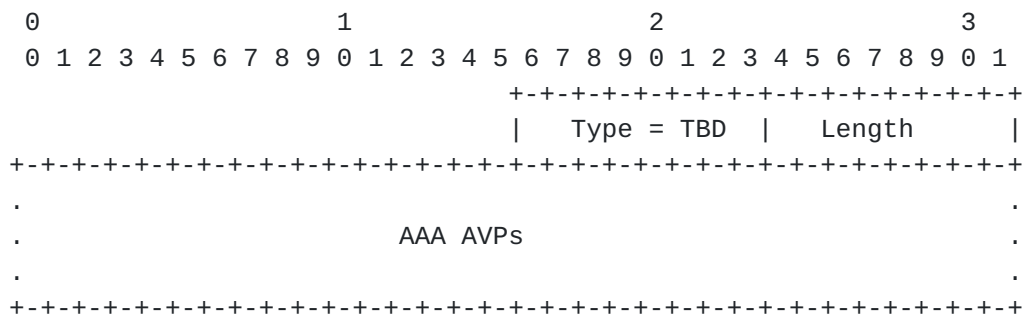


Figure 11: AAA Information Option

Length

8-bit unsigned integer, representing the length in octets of the mobility option, not including the Option Type and Option Length fields.

AAA AVPs

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

7. Security Considerations

All the messages newly defined in this document SHOULD be secured by IPsec ESP. When a HA-HELLO message is multicast, the multicast extensions to IPsec [[RFC-5374](#)] is used. In some operational scenarios, home agents are located deep in the core network and securely managed. If there is a secure transport network between home agents, some of security mechanism can be disabled, depending on administrative policy.

A Home Agent Switch message is reused for signaling between a home agent and a mobile node in HARP. It is protected by IPsec ESP as defined in [[RFC-5142](#)].

When an active home agent fails, mobile nodes using that home agent need to change their home agent to one of standby home agents. The mobile node needs to update or establish the IPsec SA with the new home agent as described in [Section 5.2](#). Existing mechanisms [[RFC5723](#)] are applied to this operation.

8. Protocol Constants

INITIAL_STATE_SYNC_REQ_TIMER: 3sec

INITIAL_SWITCH_REQ_TIMER: 1sec

MAC_HARELIABILITY_TIMEOUT 16sec

ALL_HA_MULTICAST_ADDR: TBD

9. Protocol Configuration Variables

LINK_TRAVERSAL_TIME: default 150msec

10. IANA Considerations

The following Extension Types MUST be assigned by IANA:

- o Home Agent Reliability Protocol (HARP) Message

- o State Synchronization (SS) Message
- o Binding Cache Information Option
- o AAA Information Option
- o A new link-local multicast address (ALL_HA_MULTICAST_ADDR) for all home agents will be assigned by the IANA.

11. 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

behcet.sarikaya@huawei.com

Ryuji Wakikawa

ryuji.wakikawa@gmail.com

12. 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.

13. References

13.1. 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.

[ID-3775bis] Perkins, C., Johnson, D., Arkko, J., "Mobility Support in IPv6", [draft-ietf-mext-rfc3775bis-10.txt](#), October 2010.

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

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

[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-5374] B. Weis, G. GrossD. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", [RFC 5374](#), November 2008

[RFC-5555] Soliman, H. et al, "Mobile IPv6 support for dual stack Hosts and Routers (DSMIPv6)", [RFC-5555](#), June 2009.

[RFC-5648] Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T.,

and K. Nagami, "Multiple Care-of Addresses Registration", [RFC 5648](#), October 2009.

[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.

[13.2.](#) Informative References

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

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

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

[RFC-4285] A. Patel, K. Leung, M. Khalil, H. Akhtar, K. Chowdhury, "Authentication Protocol for Mobile IPv6", [RFC 4285](#), January 2006

[RFC-4877] V. Devarapalli, F. Dupont, "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture", [RFC 4877](#), April 2007.

[RFC-5273] Y. Sheffer, H. Tschofenig, "Internet Key Exchange Protocol Version 2 (IKEv2) Session Resumption", [RFC 5273](#), January 2010.

[RFC-5798] S. Nadas, "Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6", [RFC 5798](#) (soon?), December 2009.

[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-mip6-ha-reliability-01](#) (expired), February 2004.

Author's Address

Ryuji Wakikawa
TOYOTA InfoTechnology Center, U.S.A., Inc.
465 Bernardo Avenue
Mountain View, CA 94043
USA

Email: ryuji.wakikawa@gmail.com