

Network Working Group
Request for Comments: 5142
Category: Standards Track

B. Haley
Hewlett-Packard
V. Devarapalli
Azaire Networks
H. Deng
China Mobile
J. Kempf
DoCoMo USA Labs
January 2008

Mobility Header Home Agent Switch Message

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies a new Mobility Header message type that can be used between a home agent and mobile node to signal to a mobile node that it should acquire a new home agent.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Scenarios	3
3.1. Overloaded	3
3.2. Load Balancing	3
3.3. Maintenance	3
3.4. Functional Load Balancing	3
3.5. Home Agent Renumbering	4
4. Home Agent Switch Message	4
5. Home Agent Operation	6
5.1. Sending Home Agent Switch Messages	6
5.2. Retransmissions	7
5.3. Mobile Node Errors	7
6. Mobile Node Operation	8
6.1. Receiving Home Agent Switch Messages	8
6.2. Selecting a Home Agent	9
7. Operational Considerations	9
8. Protocol Constants	10
9. IANA Considerations	10
10. Security Considerations	10
11. References	11
11.1. Normative References	11
11.2. Informative References	11
Acknowledgments	11

1. Introduction

RFC 3775 [RFC3775] contains no provision to allow a home agent to inform a mobile node that it needs to stop acting as the home agent for the mobile node. For example, a home agent may wish to handoff some of its mobile nodes to another home agent because it has become overloaded or it is going offline.

This protocol describes a signaling message, called the Home Agent Switch message, that can be used to send a handoff notification between a home agent and mobile node.

The Home Agent Switch message does not attempt to solve all general problems related to changing the home agent of a mobile node. In particular, this protocol does not attempt to solve:

- o The case where the Home Address of a mobile node must change in order to switch to a new home agent. This operation should be avoided using this message.

- o Determining when a home agent should actively move mobile nodes to another home agent. This decision should be made by a backend protocol, for example, as described in [hareliability].

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 [RFC2119].

3. Scenarios

Here are some example scenarios where a home agent signaling message would be useful.

3.1. Overloaded

There are a number of reasons a home agent might be considered overloaded. One might be that it is at, or near, its limit on the number of home bindings it is willing to accept. Another is that it has reached a pre-determined level of system resource usage -- memory, cpu cycles, etc. In either case, it would be desirable for a home agent to reduce the number of home bindings before a failure occurs.

3.2. Load Balancing

A home agent might know of other home agents that are not as heavily loaded as itself, learned through some other mechanism outside the scope of this document. An operator may wish to try and balance this load so that a failure would disrupt a smaller percentage of mobile nodes.

3.3. Maintenance

Most operators do periodic maintenance in order to maintain reliability. If a home agent is being shutdown for maintenance, it would be desirable to inform mobile nodes so they do not lose mobility service.

3.4. Functional Load Balancing

A Mobile IPv6 home agent provides mobile nodes with two basic services. It acts as a rendezvous server where correspondent nodes can find the current care-of address for the mobile node, and as an overlay router to tunnel traffic to/from the mobile node at its current care-of address.

A mobility service provider could have two sets of home agents to handle the two functions. The rendezvous function could be handled by a machine specialized for high-speed transaction processing, while the overlay router function could be handled by a machine with high data throughput.

A mobile node would start on the rendezvous server home agent and stay there if it does route optimization. However, if the original home agent detects that the mobile node is not doing route optimization, but instead reverse-tunneling traffic, it could redirect the mobile node to a home agent with better data throughput.

3.5. Home Agent Renumbering

Periodically, a mobility service provider may want to shut-down home agent services at a set of IPv6 addresses and bring service back up at a new set of addresses. Note that this may not involve anything as complex as IPv6 network renumbering [RFC4192]; it may just involve changing the addresses of the home agents. With a signaling message, the service provider could inform mobile nodes to look for a new home agent.

4. Home Agent Switch Message

The Home Agent Switch message is used by the home agent to signal to the mobile node that it needs to stop acting as the home agent for the mobile node, and that it should acquire a new home agent. Home Agent Switch messages are sent as described in Section 5.

The message described below follows the Mobility Header format specified in Section 6.1 of [RFC3775]:

[illegible]

The Home Agent Switch Message uses the MH Type value (12). When this value is indicated in the MH Type field, the format of the Message Data field in the Mobility Header is as follows:



of Addresses

An 8-bit unsigned integer indicating the number of IPv6 home agent addresses in the message. If set to zero, the mobile node MUST perform home agent discovery.

Reserved

An 8-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 Addresses

A list of alternate home agent addresses for the mobile node. The number of addresses present in the list is indicated by the "# of Addresses" field in the Home Agent Switch message.

Mobility Options

A 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 MUST follow the format specified in Section 6.2 of [RFC3775]. The receiver MUST ignore and skip any options that it does not understand.

The Binding Refresh Advice mobility option defined in Section 6.2.4 of [RFC3775] is valid for the Home Agent Switch message.

If no home agent addresses and no options are present in this message, no padding is necessary and the Header Len field in the Mobility Header will be set to zero.

5. Home Agent Operation

5.1. Sending Home Agent Switch Messages

When sending a Home Agent Switch message, the sending node constructs the packet as it would any other Mobility Header, except:

- o The MH Type field MUST be set to (12).
- o If alternative home agent addresses are known, the sending home agent SHOULD include them in the list of suggested alternate home agents. The home agent addresses field should be constructed as described in Section 10.5.1 of [RFC3775], which will randomize addresses of the same preference in the list.
- o The "# of Addresses" field MUST be filled-in corresponding to the number of home agent addresses included in the message. If no addresses are present, the field MUST be set to zero, forcing the mobile node to perform home agent discovery by some other means.
- o If the home agent is able to continue offering services to the mobile node for some period of time, it MAY include a Binding Refresh Advice mobility option indicating the time (in units of 4 seconds) until the binding will be deleted.

The Home Agent Switch message MUST use the home agent to mobile node IPsec ESP (Encapsulating Security Payload) authentication SA (Security Association) for integrity protection.

A home agent SHOULD send a Home Agent Switch message when a known period of unavailability is pending so the mobile node has sufficient time to find another suitable home agent.

The sending node does not need to be the current home agent for the mobile node, for example as described in [hareliability], but it MUST have a security association with the mobile node so the message is not rejected. In this case, the Home Agent Switch message SHOULD only contain the address of the home agent sending the message in the Home Agent Addresses field, which implies that the mobile node should switch to using the sender as its new home agent.

5.2. Retransmissions

If the home agent does not receive a response from the mobile node -- either a Binding Update message to delete its home binding if it is the current home agent, or a Binding Update message to create a home binding if it is not the current home agent -- then it SHOULD retransmit the message until a response is received. The initial value for the retransmission timer is INITIAL-HA-SWITCH-TIMEOUT.

The retransmissions by the home agent MUST use an exponential back-off mechanism, in which the timeout period is doubled upon each retransmission, until either the home agent gets a response from the mobile node to delete its binding, or the timeout period reaches the value MAX-HA-SWITCH-TIMEOUT. The home agent MAY continue to send these messages at this slower rate indefinitely.

If the home agent included a Binding Refresh Advice mobility option, then it SHOULD delay any retransmissions until at least one half of the time period has expired, or INITIAL-HA-SWITCH-TIMEOUT, whichever value is less.

5.3. Mobile Node Errors

If a mobile node does not understand how to process a Home Agent Switch message, it will send a Binding Error message as described in Section 6.1.

If a mobile node is unreachable, in other words, it still has a home binding with the home agent after reaching the timeout period of MAX-HA-SWITCH-TIMEOUT, the home agent SHOULD NOT make any conclusions about its status.

In either case, the home agent SHOULD attempt to continue providing services until the lifetime of the binding expires.

Attempts by the mobile node to extend the binding lifetime with a Binding Update message SHOULD be rejected, and a Binding Acknowledgement SHOULD be returned with status value 129 (Administratively prohibited) as specified in Section 6.1.8 of [RFC3775].

6. Mobile Node Operation

6.1. Receiving Home Agent Switch Messages

Upon receiving a Home Agent Switch message, the Mobility Header MUST be verified as specified in [RFC3775], specifically:

- o The Checksum, MH type, Payload Proto, and Header Len fields MUST meet the requirements of Section 9.2 of [RFC3775].
- o The packet MUST be covered by the home agent to mobile node IPsec ESP authentication SA for integrity protection.

If the packet is dropped due to the above tests, the receiving node MUST follow the processing rules as Section 9.2 of [RFC3775] defines. For example, it MUST send a Binding Error message with the Status field set to 2 (unrecognized MH Type value) if it does not support the message type.

Upon receipt of a Home Agent Switch message, the mobile node MUST stop using its current home agent for services and MUST delete its home binding by sending a Binding Update message as described in Section 11.7.1 of [RFC3775]. This acts as an acknowledgement of the Home Agent Switch message. Alternately, if the sender of the message is not the current home agent, sending a Binding Update message to create a home binding will act as an acknowledgement of the Home Agent Switch message. Retransmissions of Binding Update messages MUST use the procedures described in Section 11.8 of [RFC3775].

If a Binding Refresh Advice mobility option is present, the mobile node MAY delay the deletion of its home binding and continue to use its current home agent until the calculated time period has expired.

If the Home Agent Switch message contains a list of alternate home agent addresses, the mobile node SHOULD select a new home agent as described in Section 6.2, and establish the necessary IPsec security associations with the new home agent by whatever means required as part of the mobile node/home agent bootstrapping protocol for the home agent's mobility service provider. If no alternate home agent addresses are included in the list, the mobile node MUST first perform home agent discovery.

6.2. Selecting a Home Agent

In most cases, the home agent addresses in the Home Agent Switch message will be of other home agents on the home link of the mobile node (the computed prefix is the same). In this case, the mobile node SHOULD select a new home agent from the addresses as they are ordered in the list. If the first address in the list is unable to provide service, then the subsequent addresses in the list should be tried in-order.

In the case that the home agent addresses in the Home Agent Switch message are not all home agents on the home link of the mobile node (the computed prefix is different), the mobile node SHOULD select one with its home network prefix first, if available, followed by home agents with other prefixes. Choosing a home agent with a different prefix might require a change of the home address for the mobile node, which could cause a loss of connectivity for any connections using the current home address.

7. Operational Considerations

This document does not specify how an operator might use the Home Agent Switch message in its network. However, the following requirements are placed on its usage:

- o The use of this message needs to take into account possible signaling overhead, congestion, load from the mechanism itself, and the resulting registration to another home agent. A home agent may provide service for thousands, if not millions, of mobile nodes. Careless application of the Home Agent Switch message may cause the new home agent, or some other parts of the network, to suffer. As a result, it is REQUIRED that applications of this message either employ a feedback loop between resources of the new home agent and the sending of additional Home Agent Switch messages, or apply a maximum rate at which mobile nodes can be informed of the switch that is far below the designated capacity of new registrations that the set of home agents can process. If no other information is available, this maximum rate should default to MAX-HA-SWITCH-TRANSMIT-RATE.
- o In general, switching the home agent of a mobile node should only be done when absolutely necessary, since it might cause a service disruption if the switch to a new home agent fails, the new home agent is itself under an overload condition, or the network connection of the new home agent is congested.

Similarly, the path characteristics via the new home agent may be different, which may cause temporary difficulties for end-to-end transport layer operation.

- o If this message is being used for load-balancing between a set of home agents, they should all be configured with the same set of prefixes so a home agent switch does not require a change of the home address for a mobile node. That operation is NOT RECOMMENDED and should be avoided.

8. Protocol Constants

INITIAL-HA-SWITCH-TIMEOUT	5 seconds
MAX-HA-SWITCH-TIMEOUT	20 seconds
MAX-HA-SWITCH-TRANSMIT-RATE	1 per second

9. IANA Considerations

IANA has assigned a new Mobility Header type for the following new message described in Section 4:

(12) Home Agent Switch message

10. Security Considerations

As with other messages in [RFC3775], the Home Agent Switch message MUST use the home agent to mobile node ESP encryption SA for confidentiality protection, and MUST use the home agent to mobile node ESP authentication SA for integrity protection.

The Home Agent Switch message MAY use the IPsec ESP SA in place for Binding Updates and Acknowledgements, as specified in Section 5.1 of [RFC3775], in order to reduce the number of configured security associations. This also gives the message authenticity protection.

Some operators may not want to reveal the list of home agents to on-path listeners. In such a case, the Home Agent Switch message should use the home agent to mobile node IPsec ESP encryption SA for confidentiality protection.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

11.2. Informative References

- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [hareliability] Wakikawa, R., Ed., "Home Agent Reliability Protocol", Work in Progress, November 2007.

Acknowledgments

We would like to thank the authors of a number of previous documents that contributed content to this RFC:

- o Ryuji Wakikawa, Pascal Thubert, and Vijay Devarapalli, "Inter Home Agents Protocol Specification", March 2006.
- o Hui Deng, Brian Haley, Xiaodong Duan, Rong Zhang, and Kai Zhang, "Load Balance for Distributed Home Agents in Mobile IPv6", October 2004.
- o James Kempf, "Extension to RFC 3775 for Alerting the Mobile Node to Home Agent Unavailability", October 2005.
- o Brian Haley and Sri Gundavelli, "Mobility Header Signaling Message", September 2007.

Thanks also to Kilian Weniger, Jixing Liu, Alexandru Petrescu, Jouni Korhonen, and Wolfgang Fritsche for their review and feedback.

Author's Addresses

Brian Haley
Hewlett-Packard Company
110 Spitbrook Road
Nashua, NH 03062, USA
EMail: brian.haley@hp.com

Vijay Devarapalli
Azaire Networks
3121 Jay Street
Santa Clara, CA 95054 USA
EMail: vijay.devarapalli@azairenet.com

James Kempf
DoCoMo USA Labs
181 Metro Drive
Suite 300
San Jose, CA 95110 USA
EMail: kempf@docomolabs-usa.com

Hui Deng
China Mobile
53A, Xibianmennei Ave.
Xuanwu District
Beijing 100053
China
EMail: denghui@chinamobile.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

