

Internet Engineering Task Force (IETF)
Request for Comments: 5847
Category: Standards Track
ISSN: 2070-1721

V. Devarapalli, Ed.
WiChorus
R. Koodli, Ed.
Cisco Systems
H. Lim
N. Kant
Stoke
S. Krishnan
J. Laganier
Qualcomm Inc.
June 2010

Heartbeat Mechanism for Proxy Mobile IPv6

Abstract

Proxy Mobile IPv6 (PMIPv6) is a network-based mobility management protocol. The mobility entities involved in the Proxy Mobile IPv6 protocol, the mobile access gateway (MAG) and the local mobility anchor (LMA), set up tunnels dynamically to manage mobility for a mobile node within the Proxy Mobile IPv6 domain. This document describes a heartbeat mechanism between the MAG and the LMA to detect failures, quickly inform peers in the event of a recovery from node failures, and allow a peer to take appropriate action.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5847>.

Copyright Notice

Copyright (c) 2010 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	3
2. Terminology	3
3. Heartbeat Mechanism	3
3.1. Failure Detection	4
3.2. Restart Detection	5
3.3. Heartbeat Message	6
3.4. Restart Counter Mobility Option	7
4. Exchanging Heartbeat Messages over an IPv4 Transport Network	8
5. Configuration Variables	8
6. Security Considerations	8
7. IANA Considerations	9
8. Acknowledgements	9
9. References	9
9.1. Normative References	9
9.2. Informative References	10

1. Introduction

Proxy Mobile IPv6 (PMIPv6) [RFC5213] enables network-based mobility for IPv6 hosts that do not implement any mobility protocols. The protocol is described in detail in [RFC5213]. In order to facilitate the network-based mobility, the PMIPv6 protocol defines a mobile access gateway (MAG), which acts as a proxy for the Mobile IPv6 [RFC3775] signaling, and the local mobility anchor (LMA), which acts similar to a home agent, anchoring a mobile node's sessions within a PMIPv6 domain. The LMA and the MAG establish a bidirectional tunnel for forwarding all data traffic belonging to the mobile nodes.

In a distributed environment such as a PMIPv6 domain consisting of LMAs and MAGs, it is necessary for the nodes to 1) have a consistent state about each other's reachability, and 2) quickly inform peers in the event of recovery from node failures. So, when the LMA restarts after a failure, the MAG should (quickly) learn about the restart so that it can take appropriate actions (such as releasing any resources). When there are no failures, a MAG should know about the LMA's reachability (and vice versa) so that the path can be assumed to be functioning.

This document specifies a heartbeat mechanism between the MAG and the LMA to detect the status of reachability between them. This document also specifies a mechanism to indicate node restarts; the mechanism could be used to quickly inform peers of such restarts. The Heartbeat message is a Mobility Header message (protocol type 135) that is periodically exchanged at a configurable threshold of time or sent unsolicited soon after a node restart. This document does not specify the specific actions (such as releasing resources) that a node takes as a response to processing the Heartbeat messages.

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 [RFC2119].

3. Heartbeat Mechanism

The MAG and the LMA exchange Heartbeat messages every HEARTBEAT_INTERVAL seconds to detect the current status of reachability between them. The MAG initiates the heartbeat exchange to test if the LMA is reachable by sending a Heartbeat Request message to the LMA. Each Heartbeat Request contains a sequence number that is incremented monotonically. The sequence number on the last Heartbeat Request message is always recorded by the MAG, and is used to match the corresponding Heartbeat Response. Similarly, the

LMA also initiates a heartbeat exchange with the MAG, by sending a Heartbeat Request message, to check if the MAG is reachable. The format of the Heartbeat message is described in Section 3.3.

A Heartbeat Request message can be sent only if the MAG has at least one proxy Binding Cache entry at the LMA for a mobile node attached to the MAG. If there are no proxy Binding Cache entries at the LMA for any of the mobile nodes attached to the MAG, then the Heartbeat message SHOULD NOT be sent. Similarly, the LMA SHOULD NOT send a Heartbeat Request message to a MAG if there is no active Binding Cache entry created by the MAG. A PMIPv6 node MUST respond to a Heartbeat Request message with a Heartbeat Response message, irrespective of whether there is an active Binding Cache entry.

The HEARTBEAT_INTERVAL SHOULD NOT be configured to a value less than 30 seconds. Deployments should be careful in setting the value for the HEARTBEAT_INTERVAL. Sending Heartbeat messages too often may become an overhead on the path between the MAG and the LMA. It could also create congestion in the network and negatively affect network performance. The HEARTBEAT_INTERVAL can be set to a much larger value on the MAG and the LMA, if required, to reduce the burden of sending periodic Heartbeat messages.

If the LMA or the MAG do not support the Heartbeat messages, they respond with a Binding Error message with status set to 2 (unrecognized mobility header (MH) type value) as described in [RFC3775]. When the Binding Error message with status set to 2 is received in response to a Heartbeat Request message, the initiating MAG or the LMA MUST NOT use Heartbeat messages with the other end again.

If a PMIPv6 node has detected that a peer PMIPv6 node has failed or restarted without retaining the PMIPv6 session state, it should mark the corresponding binding update list or binding cache entries as invalid. The PMIPv6 node may also take other actions, which are outside the scope of this document.

The detection of failure and restart events may be signaled to network operators by using asynchronous notifications. Future work may define such notifications in a Structure of Management Information Version 2 (SMIPv2) Management Information Base (MIB) module.

3.1. Failure Detection

A PMIPv6 node (MAG or LMA) matches every received Heartbeat Response to the Heartbeat Request sent using the sequence number. Before sending the next Heartbeat Request, it increments a local variable

MISSING_HEARTBEAT if it has not received a Heartbeat Response for the previous request. When this local variable MISSING_HEARTBEAT exceeds a configurable parameter MISSING_HEARTBEATS_ALLOWED, the PMIPv6 node concludes that the peer PMIPv6 node is not reachable. If a Heartbeat Response message is received, the MISSING_HEARTBEATS counter is reset.

3.2. Restart Detection

The section describes a mechanism for detecting failure recovery without session persistence. In the case that the LMA or the MAG crashes and reboots and loses all state with respect to the PMIPv6 sessions, it would be beneficial for the peer PMIPv6 node to discover the failure and the loss of session state and establish the sessions again.

Each PMIPv6 node (both the MAG and LMA) MUST maintain a monotonically increasing Restart Counter that is incremented every time the node reboots and loses PMIPv6 session state. The counter MUST NOT be incremented if the recovery happens without losing state for the PMIPv6 sessions active at the time of failure. This counter MUST be treated as state that is preserved across reboots. A PMIPv6 node includes a Restart Counter mobility option, described in Section 3.4, in a Heartbeat Response message to indicate the current value of the Restart Counter. Each PMIPv6 node MUST also store the Restart Counter for all the peer PMIPv6 nodes with which it currently has sessions. Stored Restart Counter values for peer PMIPv6 nodes do not need to be preserved across reboots.

The PMIPv6 node that receives the Heartbeat Response message compares the Restart Counter value with the previously received value. If the value is different, the receiving node assumes that the peer PMIPv6 node had crashed and recovered. If the Restart Counter value changes or if there was no previously stored value, the new value is stored by the receiving PMIPv6 node.

If a PMIPv6 node restarts and loses PMIPv6 session state, it SHOULD send an unsolicited Heartbeat Response message with an incremented Restart Counter to all the PMIPv6 nodes that had previously established PMIPv6 sessions. Note that this is possible only when the PMIPv6 node is capable of storing information about the peers across reboots. The unsolicited Heartbeat Response message allows the peer PMIPv6 nodes to quickly discover the restart. The sequence number field in the unsolicited Heartbeat Response is ignored and no response is necessary; the nodes will synchronize during the next request and response exchange.

'U'

Set to 1 in Unsolicited Heartbeat Response. Otherwise, set to 0.

'R'

A 1-bit flag that indicates whether the message is a request or a response. When the 'R' flag is set to 0, it indicates that the Heartbeat message is a request. When the 'R' flag is set to 1, it indicates that the Heartbeat message is a response.

Sequence Number

A 32-bit sequence number used for matching the request to the reply.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer that is a multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The receiver MUST ignore and skip any options that it does not understand. At the time of writing this document, the Restart Counter mobility option, described in Section 3.4, is the only valid option in this message.

3.4. Restart Counter Mobility Option

The following shows the message format for a new mobility option for carrying the Restart Counter value in the Heartbeat message. The Restart Counter mobility option is only valid in a Heartbeat Response message. It has an alignment requirement of 4n+2.




Figure 3: Restart Counter Mobility Option

Type

An 8-bit field that indicates that it is a Restart Counter mobility option. It MUST be set to 28.

Length

An 8-bit field that indicates the length of the option in octets excluding the Type and Length fields. It is set to 4.

Restart Counter

A 32-bit field that indicates the current Restart Counter value.

4. Exchanging Heartbeat Messages over an IPv4 Transport Network

In some deployments, the network between the MAG and the LMA may be IPv4-only and not capable of routing IPv6 packets. In this case, the Mobility Header containing the Heartbeat message is carried as specified in Section 4 of [RFC5844], i.e., the Mobility Header is part of the UDP payload inside an IPv4 packet (IPv4-UDP-MH).

5. Configuration Variables

The LMA and the MAG must allow the following variables to be configurable.

HEARTBEAT_INTERVAL

This variable is used to set the time interval in seconds between two consecutive Heartbeat Request messages. The default value is 60 seconds. It SHOULD NOT be set to less than 30 seconds or more than 3600 seconds.

MISSING_HEARTBEATS_ALLOWED

This variable indicates the maximum number of consecutive Heartbeat Request messages for which a PMIPv6 node did not receive a response before concluding that the peer PMIPv6 node is not reachable. The default value for this variable is 3.

6. Security Considerations

The Heartbeat messages are just used for checking reachability between the MAG and the LMA. They do not carry information that is useful for eavesdroppers on the path. Therefore, confidentiality protection is not required. Integrity protection using IPsec [RFC4301] for the Heartbeat messages MUST be supported on the MAG and the LMA. RFC 5213 [RFC5213] describes how to protect the Proxy Binding Update and Acknowledgement signaling messages with IPsec. The Heartbeat message defined in this specification is merely another subtype of the same Mobility Header protocol that is already being protected by IPsec. Therefore, protecting this additional message is

possible using the mechanisms and security policy models from these RFCs. The security policy database entries should use the new MH Type, the Heartbeat message, for the MH Type selector.

If dynamic key negotiation between the MAG and the LMA is required, Internet Key Exchange Protocol version 2 (IKEv2) [RFC4306] should be used.

7. IANA Considerations

The Heartbeat message defined in Section 3.3 must have the type value allocated from the same space as the 'MH Type' name space in the Mobility Header defined in RFC 3775 [RFC3775].

The Restart Counter mobility option defined in Section 3.4 must have the type value allocated from the same name space as the mobility options defined in RFC 3775 [RFC3775].

8. Acknowledgements

A heartbeat mechanism for a network-based mobility management protocol was first described in [NETLMM]. The authors would like to thank the members of a NETLMM design team that produced that document. The mechanism described in this document also derives from the path management mechanism described in [GTP].

We would like to thank Alessio Casati for first suggesting a fault handling mechanism for Proxy Mobile IPv6.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4306] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

9.2. Informative References

[NETLMM] Levkowitz, H., Ed., Giaretta, G., Leung, K., Liebsch, M., Roberts, P., Nishida, K., Yokota, H., and M. Parthasarathy, "The NetLMM Protocol", Work in Progress, October 2006.

[GTP] 3rd Generation Partnership Project, "3GPP Technical Specification 29.060 V7.6.0: "Technical Specification Group Core Network and Terminals; General Packet Radio Service (GPRS); GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface (Release 7)""", July 2007.

Authors' Addresses

Vijay Devarapalli (editor)
WiChorus
3950 North First Street
San Jose, CA 95134
USA

EEmail: vijay@wichorus.com

Rajeev Koodli (editor)
Cisco Systems
USA

EEmail: rkoodli@cisco.com

Heeseon Lim
Stoke
5403 Betsy Ross Drive
Santa Clara, CA 95054
USA

EEmail: hlim@stoke.com

Nishi Kant
Stoke
5403 Betsy Ross Drive
Santa Clara, CA 95054
USA

EEmail: nishi@stoke.com

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

EEmail: suresh.krishnan@ericsson.com

Julien Laganier
Qualcomm Incorporated
5775 Morehouse Drive
San Diego, CA 92121
USA

EEmail: julienl@qualcomm.com