

MULTIMOB Group
Internet-Draft
Expires: August 22, 2013

H. Asaeda
NICT
P. Seite
France Telecom
February 18, 2013

PMIPv6 Multicast Routing Optimization with PIM-SM
draft-asaeda-multimob-pmip6-ropt-with-pim-00

Abstract

This document describes IP multicast routing optimization with PIM-SM in Proxy Mobile IPv6 (PMIPv6) environment. The Mobile Access Gateway (MAG) and the Local Mobility Anchor (LMA) are the mobility entities defined in the PMIPv6 protocol and act as PIM-SM routers. The proposed protocol optimization addresses the tunnel convergence problem and cooperates with seamless handover mechanisms.

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 August 22, 2013.

Copyright Notice

Copyright (c) 2013 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. Conventions and Terminology	4
3. Overview	4
3.1. Multicast Communication in PMIPv6	4
3.2. Protocol Sequence for Multicast Channel Subscription	6
4. Multicast Tunnel (M-Tunnel)	7
4.1. Packet Encapsulation	7
4.2. M-Tunnels Connecting to Multiple PIM-SM Routers and ECMP Routing	8
5. Local Mobility Anchor Operation	9
6. Mobile Access Gateway Operation	10
7. Mobile Node Operation	10
8. Localized Multicast Routing	10
9. Smooth Handover	11
10. IANA Considerations	14
11. Security Considerations	14
12. Acknowledgements	14
13. References	14
13.1. Normative References	14
13.2. Informative References	15
Authors' Addresses	15

1. Introduction

Proxy Mobile IPv6 (PMIPv6) [1] enables network-based mobility for IPv6 mobile nodes (MNs) that do not implement any mobility protocols. The Local Mobility Anchor (LMA) is the topological anchor point to manages the mobile node's binding state. The Mobile Access Gateway (MAG) is an access router or gateway that manages the mobility-related signaling for an MN. An MN is attached to the Proxy Mobile IPv6 Domain (PMIPv6-Domain) that includes LMA and MAG(s), and is able to receive data coming from outside of the PMIPv6-Domain through LMA and MAG.

Network-based mobility support for unicast is addressed in [1], while multicast support in PMIPv6 is not discussed in it. Since LMA and MAG set up a bi-directional IPv6-in-IPv6 tunnel for each mobile node and forwards all mobile node's traffic according to [1], it highly wastes network resources when a large number of mobile nodes join/subscribe the same multicast sessions/channels, because independent data copies of the same multicast packet are delivered to the subscriber nodes in a unicast manner through MAG.

The base solution described in [12] provides options for deploying multicast listener functions in PMIPv6-Domains without modifying mobility and multicast protocol standards. However, in this specification, MAG MUST act as an MLD proxy [2] and hence MUST dedicate a tunnel link between LMA and MAG to an upstream interface for all multicast traffic. This limitation does not allow to use Protocol-Independent Multicast - Sparse Mode (PIM-SM) [3] native routing on MAG, and hence does not solve the tunnel convergence problem; MAG receives the same data from multiple LMAs when MAG attaches to them for mobile nodes and has subscribed the same multicast channel to them. It does not enable direct routing and does not optimize source mobility.

This document describes IP multicast routing optimization using PIM-SM in Proxy Mobile IPv6 (PMIPv6) environment. The Mobile Access Gateway (MAG) and the Local Mobility Anchor (LMA) are the mobility entities defined in the PMIPv6 protocol and act as PIM-SM routers. The proposed protocol optimization assumes that both LMA and MAG enable the Protocol-Independent Multicast - Sparse Mode (PIM-SM) multicast routing protocol [3]. The proposed optimization uses a dedicated GRE [4] tunnel for multicast, called M-Tunnel between MAG and PIM-SM router such as LMA. The proposed protocol optimization addresses the tunnel convergence problem and provides seamless handover. It can cooperate with localized routing and direct routing to deliver IP multicast packets for mobile nodes and source mobility. In this document, because multicast listener mobility is mainly focused on, the detail specification of source mobility is not

described.

This document does not require to change unicast communication methods or protocols defined in [1], and therefore both unicast and multicast communications for mobile nodes in PMIPv6-Domain are enabled if this extension is implemented.

2. Conventions and 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 [5].

The following terms used in this document are to be interpreted as defined in the Proxy Mobile IPv6 specification [1]; Mobile Access Gateway (MAG), Local Mobility Anchor (LMA), Mobile Node (MN), Proxy Mobile IPv6 Domain (PMIPv6-Domain), LMA Address (LMAA), Proxy Care-of Address (Proxy-CoA), Mobile Node's Home Network Prefix (MN-HNP), Mobile Node Identifier (MN-Identifier), Proxy Binding Update (PBU), and Proxy Binding Acknowledgement (PBA).

3. Overview

3.1. Multicast Communication in PMIPv6

Required components to enable IP multicast are multicast routing protocols and host-and-router communication protocols. This document assumes PIM-SM [3] as the multicast routing protocol and Multicast Listener Discovery (MLD) as the host-and-router communication protocol. This document allows mobile nodes to participate in Any-Source Multicast (ASM) and Source-Specific Multicast (SSM) [6]. However, in order to explicitly participate in SSM, mobile nodes MUST support either MLDv2 [7] or Lightweight-MLDv2 (LW-MLDv2) [8].

The architecture of a Proxy Mobile IPv6 domain is shown in Figure 1. LMA and MAG are the core functional entities in PMIPv6-Domain. The entire PMIPv6-Domain appears as a single link from the perspective of each mobile node.

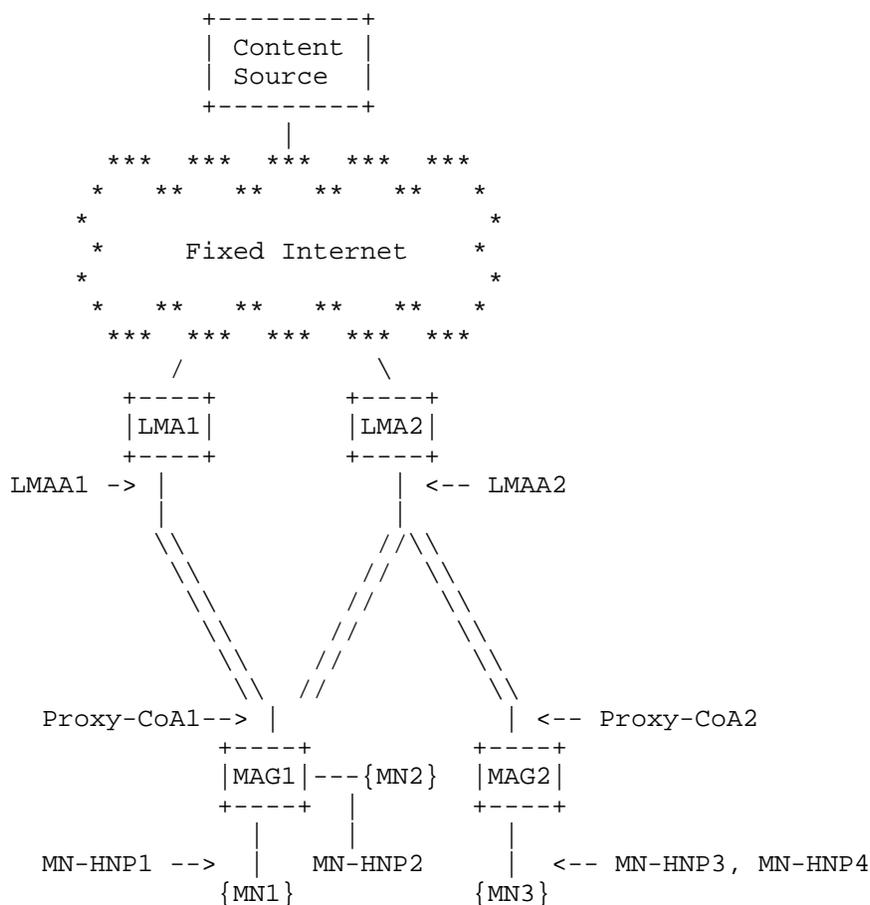


Figure 1: Proxy Mobile IPv6 Domain

When a mobile node wants to subscribe/unsubscribe a multicast channel, it sends MLD Report messages specifying sender and multicast addresses to the access link. The attached MAG detects this membership information and sends the PIM Join/Prune message to the corresponding LMA over a bi-directional GRE tunnel called M-Tunnel (described in Section 4) when the LMA is selected as the previous-hop router for the multicast channel, or sends the PIM Join/Prune message to the adjacent upstream multicast router for the multicast channel. When the LMA or the adjacent router receives the PIM Join/Prune message, it coordinates the corresponding multicast routing tree if necessary and starts forwarding the data.

When the MAG detects mobile node's handover, it can proceed the seamless handover procedures. Since both PMIPv6 and multicast

protocols (i.e., MLD and PIM-SM) do not have functions for multicast context transfer in their original protocol specifications, the external functions or protocols should be used for handover. One of the possible ways is the use of "mobile node's Policy Profile", as it could include "multicast channel information", which expresses mobile node's subscribing multicast channel list, as well as the mandatory fields of the Policy Profile specified in [1]. Mobile node's Policy Profile is provided by "policy store" whose definition is the same as of [1].

3.2. Protocol Sequence for Multicast Channel Subscription

A mobile node sends unsolicited MLD Report messages including source and multicast addresses when it subscribes a multicast channel. When the MAG operating as a PIM-SM router receives MLD Report messages from attached mobile nodes, it sends PIM Join messages to its neighboring routers (Figure 2) to join the multicast delivery tree (if not joined). When the upstream router for the requested channel is LMA (or a non-adjacent PIM router), the MAG sends the corresponding PIM Join messages to the LMA (or the non-adjacent PIM router) using M-Tunnel (see Section 4). When the upstream router for the requested channel is a directly attached adjacent router, the MAG sends the corresponding PIM Join messages to the adjacent upstream router natively. The LMA or such routers then join the multicast delivery tree and forward the packets to the downstream MAG.

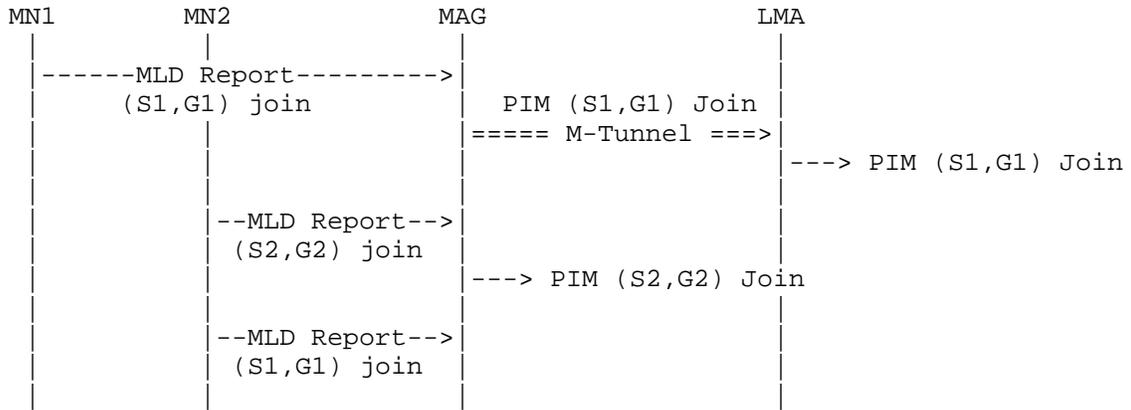


Figure 2: MLD Report and PIM Messages Transmission

PIM-SM relies on Multicast Routing Information Base (MRIB) to select a single upstream interface for each multicast channel by the Reverse Path Forwarding (RPF) algorithm. After all potential upstream interfaces including M-Tunnel are recognized in MRIB (as described in Section 4), the MAG enabling PIM-SM can select either M-Tunnel

interface or physical interface for a multicast channel by the Reverse Path Forwarding (RPF) algorithm. The proposed architecture for PMIPv6 with PIM-SM therefore does not cause the tunnel convergence problem, and hence duplicate packets are not forwarded to the MAG.

4. Multicast Tunnel (M-Tunnel)

4.1. Packet Encapsulation

M-Tunnel is a bi-directional GRE tunnel [4] dedicated for PIM messages and IP multicast data transmissions. The tunnel end-point of M-Tunnel is a MAG that is a PIM-SM capable router. Another tunnel end-point is also a PIM-SM capable router. The typical use case of M-Tunnel is to establish a bi-directional tunnel link between LMA and MAG; therefore LMA shall be another tunnel end-point. M-Tunnel can be established in a bootstrap phase of MAG (without detecting a multicast channel subscription request from a mobile node) and kept while the MAG enables PIM routing functions to forward multicast packets. An M-Tunnel is not set up per mobile node basis, but per MAG basis; it can be shared with mobile nodes attached to the MAG as seen in Figure 3.

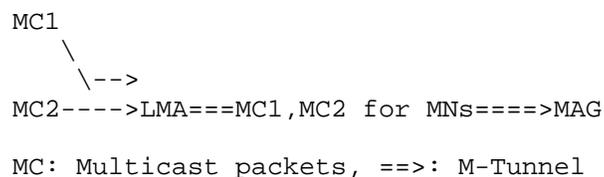


Figure 3: Multicast packet forwarding through M-Tunnel

In order for the PIM routing protocol to use an M-Tunnel for multicast forwarding, an M-Tunnel interface must be recognized by the PIM routing protocol as the upstream multicast interface for MAG. It is done by the configuration of static multicast routes, such as "ip mroute 0.0.0.0 0.0.0.0 gre0" or "ip mroute 1.1.1.0 255.255.255.0 gre0". By such configuration, MAG inserts the multicast forwarding paths using the M-Tunnel into its MRIB. MAG then selects the M-Tunnel interface as the corresponding RPF interface, and forwards the PIM Join/Prune messages over the M-Tunnel. If operators want to select other interface, e.g. a physical interface, as the upstream multicast interface for some specific source prefixes, e.g. sources inside the PMIPv6-Domain, they can *additionally* configure the specific multicast routes with longer prefixes. This configuration will be used for direct routing. Then the MAG selects as the appropriate upstream router according to the MRIB entry. Note that

the case having multiple M-Tunnels configured on MAG is described in Section 4.2.

The format of the tunneled multicast packet forwarded from LMA to MAG is shown below. "S" and "G" are the same notation used for (S,G) multicast channel.

```

IPv6 header (src= LMAA, dst= Proxy-CoA) /* Outer Header */
  GRE header                               /* Encapsulation Header */
    IPv6 header (src= S, dst= G)           /* Inner Header */
      Upper layer protocols                 /* Packet Content */

```

Figure 4: Multicast packet format tunneled from LMA to MAG

When a PIM message is sent from MAG to LMA, the src and dst addresses of the outer tunnel header will be replaced to Proxy-CoA and LMAA, respectively. To convey a PIM message, the src address of the inner packet header is changed to either LMA's or MAG's link-local address. The dst address of the packet header is assigned based on the PIM's condition (see [3]).

In order to establish M-Tunnel, LMA and MAG need to negotiate GRE encapsulation and GRE keys for M-Tunnel. The GRE Key option to be used for the negotiation of GRE tunnel encapsulation mode and exchange of the uplink and downlink GRE keys is defined in [9]. It is also possible to use the static fixed GRE keys for M-Tunnel.

4.2. M-Tunnels Connecting to Multiple PIM-SM Routers and ECMP Routing

There can be multiple LMAs in a PMIPv6-Domain each serving a different group of mobile nodes. In that case, a MAG will connect to multiple LMAs with different M-Tunnels having different GRE keys. For example, in Figure 5, MAG1 establishes two M-Tunnels with LMA1 and LMA2, and MAG2 establishes one M-Tunnel with LMA2.

A MAG that has multiple M-Tunnels, such as MAG1 in Figure 5, must decide a single upstream M-Tunnel interface for an RP or a source address or prefix. There are two ways to decide a single upstream M-Tunnel for a MAG. One is only with static MRIB configuration by operation. For example, operators can configure each M-Tunnel interface as the RPF interface for specific source address(es) or prefix(es) one by one. Each M-Tunnel interface is then inserted into the MAG's MRIB and used for different source address(es) or prefix(es).

The other way to select a single upstream M-Tunnel interface is with PIM ECMP [13]. A MAG enabling PIM routing functions selects a path by the ECMP algorithm as described in [13]. The PIM ECMP function

chooses the PIM neighbor with the highest IP address or the best hash value over the destination and source addresses. The algorithm choosing a single interface is based on an operator's decision. When operators decide to use PIM ECMP to select a single upstream M-Tunnel from multiple M-Tunnels, both the MAG and the tunnel end-point PIM-SM routers (e.g., LMAs) MUST enable PIM ECMP.

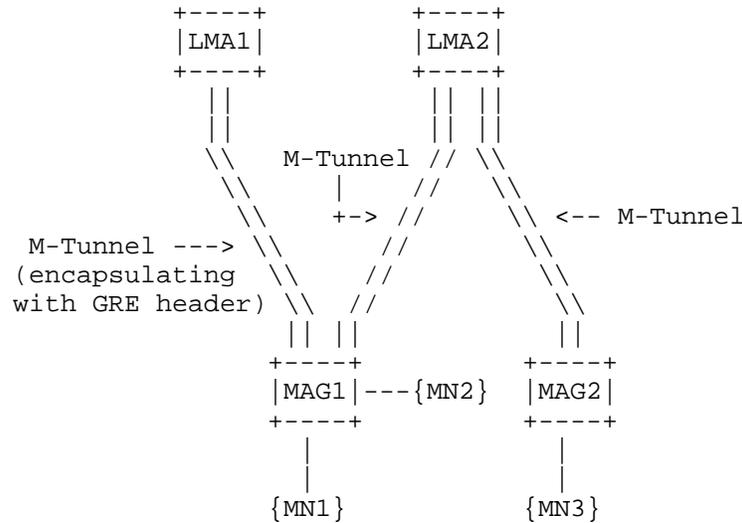


Figure 5: M-Tunnels established between LMA and MAG

5. Local Mobility Anchor Operation

The LMA is responsible for maintaining the mobile node's reachability state and is the topological anchor point for the mobile node's home network prefix(es). This document assumes that the LMA is capable of forwarding multicast packets to the MAG by enabling the PIM-SM multicast routing protocol [3]. The LMA acting as a PIM-SM multicast router may serve MAGs as downstream routers for some multicast channels when a mobile node is a multicast data receiver (or as upstream routers when a mobile node is a multicast data sender). The downstream (or upstream) MAG is connected to the LMA through the M-Tunnel for multicast communication.

When the LMA sets up the multicast state and joins the group as the MAG's upstream router, the multicast packets are tunneled to the MAG that requested to receive the corresponding multicast session. The MAG then forwards the packets to the mobile node according to the multicast listener state maintained in the MAG. [1] supports only point-to-point access link types for MAG and mobile node connection;

hence a mobile node and the MAG are the only two nodes on an access link, where the link is assumed to be multicast capable.

6. Mobile Access Gateway Operation

The MAG performs the mobility management on behalf of a mobile node. This document assumes that the MAG is PIM-SM capable and forwards multicast packets to the corresponding mobile nodes attached to MAG by enabling the PIM-SM multicast routing protocol. In addition, the MAG must maintain multicast membership status for the attached mobile nodes at the edge and forwards the multicast data to the member mobile nodes. This condition requires MAG to support MLDv2 [7] or LW-MLDv2 [8], as well.

When mobile nodes subscribe multicast channel(s), they send MLD Report messages with their link-local address to the MAG, and the MAG sends the corresponding PIM Join messages to the upstream router if the MAG has no multicast state for the requested channel(s). The upstream router is selected by the Reverse Path Forwarding (RPF) lookup algorithm, and that is either the LMA or an adjacent multicast router attached to the same link. If the LMA is the upstream router for the channel(s) for the MAG, the MAG encapsulates PIM Join messages using the M-Tunnel.

The optimal multicast routing path may not include the LMA, especially in localized routing as described in Section 6.10.3 of [1] and [10]. The localized routing option is designed to support node-to-node communication within PMIPv6-Domain where a local content source exists. Details are described in Section 8.

7. Mobile Node Operation

Mobile nodes attached to the MAG can behave as regular receiver hosts. A mobile node sends MLD report messages to the MAG when it wants to subscribe and unsubscribe IP multicast channels.

In order to subscribe/unsubscribe multicast channel(s) by unsolicited report messages and inform current membership state by solicited report messages, mobile nodes MUST support either MLDv1 [7], MLDv2 [7], or LW-MLDv2 [8], and SHOULD support MLDv2 or LW-MLDv2.

8. Localized Multicast Routing

Localized routing defined in [10] allows mobile nodes attached to the same or different MAGs to directly exchange unicast traffic by using

localized forwarding or a direct tunnel between the MAGs. Localized routing must be initiated both MAG and LMA. Localized routing is not persistent, and is initiated by two signaling messages, Localized Routing Initiation (LRI) and Local Routing Acknowledgment (LRA), sent by LMA or MAG.

To support localized multicast routing with PIM-SM capable LMA and MAG, both LMA and MAG MUST include the routes organized by the localized routing procedure specified in [10] into their MRIBs. The exact mechanism to do this is not specified in this document and is left open for implementations and specific deployments.

To support localized routing for the case that a source node and a receiver node are attached to different MAGs but the same LMA (as seen in Section 6 of [10]), these MAGs must use the same tunneling mechanism for the data traffic tunneled between them. M-Tunnel defined in this document corresponds to the concept; these MAGs establish M-Tunnel and enable localized multicast routing.

9. Smooth Handover

The MAG is responsible for detecting the mobile node's movements to and from the access link and for initiating binding registrations to the mobile node's LMA. In PMIPv6, it does not require for mobile nodes to initiate to re-subscribe multicast channels, and the MAG keeps multicast channel subscription status for mobile nodes even if they move to a different MAG (i.e., n-MAG) in PMIPv6-Domain.

The MAG needs to join the multicast delivery tree when an attached mobile node subscribes a multicast channel. When the mobile node changes the network, it seamlessly receives multicast data from the new MAG according to the multicast channel information stored in the "MN's Policy Profile" or by some handover mechanisms such as [14] and [15]. Whether the MN's Policy Profile or a handover mechanism mobile operators use depend on their policy or implementation.

Here, a handover procedure using the MN's Policy Profile is described as an example. When the multicast channel information subscribed by mobile nodes is maintained in "MN's Policy Profile" stored in a policy store [1], the MAG can use the channel information to provide seamless handover. The procedures are described as follows and illustrated in Figure 6;

1. Figure 6 shows the examples that a mobile node has received multicast data from an upstream multicast router via p-MAG (*1) and from LMA via p-MAG (*2).

2. Whenever the mobile node moves a new network and attaches to n-MAG, the n-MAG obtains the MN-Identifier (MN-ID) and learns multicast channel information described in Mobile Node's Policy Profile associated to this MN-Identifier. Describing the method how the n-MAG identifies the p-MAG is out of scope of this document, while using the same mechanism described in [16] would be one of the possible methods.
3. If there are multicast channels the mobile node has subscribed but the n-MAG has not yet subscribed, n-MAG joins the corresponding multicast channels by sending the PIM Join message to its upstream router. If the upstream router is the LMA, the PIM messages are encapsulated and transmitted over the M-Tunnel (*4); otherwise the PIM messages are sent natively to the adjacent upstream router (*3).
4. The multicast data is forwarded from the LMA through the M-Tunnel between the LMA and n-MAG (*4) or from the adjacent upstream router (*3).

10. IANA Considerations

This document has no actions for IANA.

11. Security Considerations

TBD.

12. Acknowledgements

Many of the specifications described in this document are discussed and provided by the multimob mailing-list. Also, extensive comments were received from Sergio Figueiredo, Dirk von Hugo, and Stig Venaas.

13. References

13.1. Normative References

- [1] Gundavelli, S, Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [2] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [3] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [4] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [5] Bradner, S., "Key words for use in RFCs to indicate requirement levels", RFC 2119, March 1997.
- [6] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.
- [7] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [8] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790,

February 2010.

- [9] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", RFC 5845, June 2010.
- [10] Krishnan, S., Koodli, R., Loureiro, P., Wu, Q., and A. Dutta, "Localized Routing for Proxy Mobile IPv6", RFC 6705, September 2012.
- [11] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

13.2. Informative References

- [12] Schmidt, T., Waehlich, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.
- [13] Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, November 2000.
- [14] Contreras, LM., Bernardos, CJ., and I. Soto, "PMIPv6 multicast handover optimization by the Subscription Information Acquisition through the LMA (SIAL)", draft-ietf-multimob-handover-optimization-01.txt (work in progress), December 2012.
- [15] von Hugo, D. and H. Asaeda, "Context Transfer Protocol Extension for Multicast", draft-vonhugo-multimob-cxtp-extension-03.txt (work in progress), February 2013.
- [16] Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", RFC 5949, September 2010.
- [17] Asaeda, H., "IGMP/MLD-Based Explicit Membership Tracking Function for Multicast Routers", draft-ietf-pim-explicit-tracking-04.txt (work in progress), January 2013.

Authors' Addresses

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

Email: asaeda@nict.go.jp

Pierrick Seite
France Telecom
4, rue du Clos Courtel
BP 91226, Cesson-Sevigne 35512
France

Email: pierrick.seite@orange-ftgroup.com

