MULTIMOB Working Group JC. Zuniga
Internet-Draft InterDigital
Intended status: Informational LM. Contreras
Expires: September 5, 2012 Telefonica I+D
CJ. Bernardos
UC3M
S. Jeon
Instituto de Telecomunicacoes
Y. Kim
Soongsil University
March 4, 2012

Multicast Mobility Routing Optimizations for Proxy Mobile IPv6
draft-ietf-multimob-pmipv6-ropt-00

Abstract

The MULTIMOB group has specified a base solution to support IP multicasting in a PMIPv6 domain [RFC6224]. In this document, some enhancements to the base solution are described. These enhancements include the use of a multicast tree mobility anchor as the topological anchor point for multicast traffic, as well as a direct routing option where the MAG can provide access to multicast content in the local network. These enhancements provide benefits such as reducing multicast traffic replication and supporting different PMIPv6 deployments scenarios.

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 September 5, 2012.

Copyright Notice

Copyright (c) 2012 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

Proxy Mobile IPv6 [RFC5213] is a network-based approach to solving the IP mobility problem. In a Proxy Mobile IPv6 (PMIPv6) domain, the Mobile Access Gateway (MAG) behaves as a proxy mobility agent in the network and does the mobility management on behalf of the Mobile Node (MN). The Local Mobility Anchor (LMA) is the home agent for the MN and the topological anchor point. PMIPv6 was originally designed for unicast traffic. However, a PMIPv6 domain may handle data from both unicast and multicast sources.

The Internet Group Management Protocol (IGMPv3) [RFC3376] is used by IPv4 hosts to report their IP multicast group memberships to neighboring multicast routers. Multicast Listener Discovery (MLDv2) [RFC3810] is used in a similar way by IPv6 routers to discover the presence of IPv6 multicast hosts. Also, the IGMP/MLD proxy [RFC4065] allows an intermediate (i.e. edge) node to appear as a multicast router to downstream hosts, and as a host to upstream multicast routers. IGMP and MLD related protocols however were not originally designed to address IP mobility of multicast listeners (i.e. IGMP and MLD protocols were originally designed for fixed networks).

The MULTIMOB group has specified a base solution to support IP multicast listener mobility in a PMIPv6 domain [RFC6224], which describes deployment options without modifying mobility and multicast protocol standards. The PMIPv6 allows a MAG to establish multiple PMIPv6 tunnels with different LMAs, e.g. up to one per MN. In the presence of multicast traffic, multiple instances of the same traffic can converge to the same MAG. Hence, when IP multicasting is applied into PMIPv6, it leads to redundant traffic at a MAG. This is the so-called "Tunnel Convergence problem".

To address this issue, a comprehensive solution is proposed in this document, consisting of two complementary enhancements: multicast anchor and direct routing. The former uses a multicast tree mobility anchor (MTMA) as the topological anchor point for remotely delivering multicast traffic, while the latter uses direct routing taking advantage of local multicast source availability, allowing a MAG to connect directly to a multicast router for simple access to local content. Neither of the schemes has any impact on the MN to support multicast listener mobility.

The MTMA details are described in section 3. Section 4 describes the direct routing technique. Section 5 describes the details about the dynamic selection at the MAG between direct routing (e.g. for local access) and MTMA (e.g. for remote access).

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 [RFC2119].

This document uses the terminology defined in [RFC5213], [RFC6275], and [RFC3810]. Specifically, the definition of PMIPv6 domain is reused from [RFC5213] and reproduced here for completeness.

Proxy Mobile IPv6 Domain (PMIPv6-Domain): Proxy Mobile IPv6 domain refers to the network where the mobility management of a mobile node is handled using the Proxy Mobile IPv6 protocol as defined in [RFC5213]. The Proxy Mobile IPv6 domain includes local mobility anchors and mobile access gateways between which security associations can be set up and authorization for sending Proxy Binding Updates on behalf of the mobile nodes can be ensured.

In this draft we refine such definition from the point of view of the kind of traffic served to the MN in the following way:

PMIPv6 unicast domain: PMIPv6 unicast domain refers to the network covered by one LMA for unicast service. This service allows MN mobility as it moves from one MAG to another associated to that LMA regarding its unicast traffic.

PMIPv6 multicast domain: PMIPv6 multicast domain refers to the network covered by one network element named MTMA (defined below) for multicast service in such a way that an MN using that service is not aware of mobility as it moves from one MAG to another.

Direct routing: it uses native multicast infrastructure for retrieving multicast data. For the operator having its own local content, this technique also includes the case that content source is directly connected to a MAG.

From the definitions above, it can stated that a PMIPv6 domain can have several PMIPv6 unicast domains and PMIPv6 multicast domains. Additionally, some other definitions are introduced, as follows.

MTMA or multicast tree mobility anchor: an entity working as topological anchor point for multicast traffic.

H-LMA or Hybrid-LMA: an entity dedicated to both unicast and multicast services, that is, it is able to work as both LMA and MTMA simultaneously.

3. Multicast Tree Mobility Anchor (MTMA)

An MTMA can be used to serve as the mobility anchor for multicast traffic. Typically, the MTMA will be used to get access to multicast content remotely.

The MTMA connects to the MAG as described in [RFC6224] and it can reuse native PMIPv6 features such as tunnel establishment and security [RFC5213], heartbeat [RFC5847], etc. Unicast traffic will go normally to the LMAs in the PMIPv6 domain as described in [RFC5213].

This section describes how the MTMA works in scenarios of MN attachment and multicast mobility. It concentrates on the case of both LMA and MTMA defining a unique PMIPv6 domain. Some other different deployment scenarios are presented in Appendix A.

3.1. Overview

Figure 1 shows an example of a PMIPv6 domain supporting multicast mobility. The LMA is dedicated to unicast traffic, and the MTMA is dedicated to multicast traffic. The MTMA can be considered to be a form of upstream multicast router with tunnel interfaces allowing remote subscription for the MNs. Note that there can be multiple LMAs for unicast traffic in a given PMIPv6 domain (not shown in Figure 1 for simplicity). Similarly, more than one MTMA could be deployed by the operator (not shown in Figure 1).

As shown in Figure 1, MAG1 may connect to both unicast (LMAs) and multicast (MTMAs) entities. Thus, a given MN may simultaneously receive both unicast and multicast traffic. In Figure 1, MN1 and MN2 receive unicast traffic, multicast traffic, or both, whereas MN3 receives multicast traffic only.

```
                                   +--------------+
                                   |Content Source|
                                   +--------------+
                                          |
                                          |
         ***  ***  ***  ***      ***  ***  ***  ***
        *   **   **   **   *    *   **   **   **    *
       *                    *  *                     *
       *  Unicast Traffic   *  *  Multicast Traffic  *
       *                    *  *                     *
        *   **   **   **   *    *   **   **   **   *
         ***  ***  ***  **       ***  ***  ***  ***
                 |                       |
                 |                       |
                 |                       |
              +-----+                 +------+
     Unicast  | LMA |                 | MTMA |     Multicast
      Anchor  +-----+                 +------+      Anchor
                  \\                    // ||
                   \\                  //  ||
                    \\                //   ||
                     \\              //    ||
                      \\            //     ||
                       \\          //      ||
                        \\        //       ||
                         \\      //        ||
                          \\    //         ||
                          +-----+       +-----+      MLD Proxy
                          | MAG1|       | MAG2|
                          +-----+       +-----+
                          |     |          |
                          |     |          |
                        {MN1} {MN2}      {MN3}
```

Figure 1: Architecture of Multicast Tree Mobility Anchor (MTMA)

3.2. Operations of the Mobile Node

The MN operation is not impacted by the existence of an MTMA as anchor for the multicast traffic being subscribed. The MN will act according to the stated operations in [RFC5213] and [RFC6224].

This draft considers that every MN requesting multicast-only services is previously registered in a PMIPv6 unicast domain to get a unicast IP address. The registration can also be required also for several purposes such as remote management, billing, etc.

3.3. Operations of the Mobile Access Gateway

There are two main functionalities in the MAG when it is connected to an MTMA. One when the MAG incorporates MLD proxy functions as per [RFC4605]. The other case is when the MAG functions as a multicast router as per [RFC4601] or [RFC4607].

The following sections describe the MAG for both cases in more detail.

3.3.1. MAG as MLD Proxy

If the MAG has MLD proxy functionality only, once the MLD proxy instance is configured to obtain the multicast traffic remotely from the MTMA, the system behavior remains static.

In case of remote subscription, all MAGs that are connected to the MTMA must support the MLD proxy [RFC4605] function. Specifically in Figure 1, each of the MAG1-MTMA and MAG2-MTMA tunnel interfaces define an MLD proxy domain. The MNs are considered to be on the downstream interface of the MLD proxy (of the MAG), and the MTMA is considered to be on the upstream interface (of the MAG) as per [RFC4605]. Note that the MAG could also be an IGMP proxy. For brevity this document will refer primarily to MLD proxy, but all references to "MLD proxy" should be understood to also include "IGMP/MLD proxy" functionality.

3.3.1.1. Multicast Establishment

Figure 2 shows the procedure when MN1 attaches to MAG, and establishes associations with LMA (unicast) and MTMA (multicast).

```
MN1                  MAG        LMA       MTMA
|                (MLD Proxy) (Unicast) (Multicast)
MN attaches to MAG1   |          |          |
|                     |          |          |
|------Rtr Sol----- ->|          |          |
|                     |--PBU -- >|          |
|                     |          |          |
|                     |<-- PBA --|          |
|                     |          |          |
|                     |=Unicast= |          |
|                     |  Tunnel  |          |
|<-----Rtr Adv ------ |          |          |
|                     |          |          |
|< ------ Unicast Traffic------ >|          |
|                     |          |          |
|                     |==Multicast Tunnel ==|
|                     |          |          |
|<--MLD Query --------|          |          |
|                     |          |          |
MN requires           |          |          |
multicast services    |          |          |
|                     |          |          |
|---MLD Report (G) -->|          |          |
|                     |          |          |
|                     |---- Aggregated ---> |
|                     |    MLD Report (G)   |
|                     |          |          |
|                     |          |          |
|< --------- Multicast Traffic ----------- >|
|                     |          |          |
```

Figure 2: MN Attachment and Multicast Service Establishment for MTMA

In Figure 2, MAG first establishes the PMIPv6 tunnel with LMA for unicast traffic as defined in [RFC5213] after being triggered by the Router Solicitation message from MN1. Unicast traffic will then flow between MN1 and LMA.

For multicast traffic, a multicast tunnel may have been pre-configured between MAG and MTMA, or may be dynamically established when the first MN appears at the MAG.

MN1 sends the MLD report message (when required by its upper layer applications) as defined in [RFC3810] in response to an MLD Query from MAG. The MAG, acting as a MLD Proxy defined in [RFC4605], will then send an Aggregated MLD Report to the multicast anchor, MTMA (assuming that this is a new multicast group which the MAG had not previously subscribed to). Multicast traffic will then flow from the

MTMA towards MN1.

3.3.1.2. Multicast Mobility

Figure 3 illustrates the mobility scenario for multicast traffic. Specifically, MN2 with ongoing multicast subscription moves from MAG1 to MAG2. Note that, for simplicity, in this scenario we only consider the tunnel of MAG2 with MTMA (for multicast traffic) and we does not show any unicast traffic. Of course, if it was desired to support unicast traffic, it would be served by a tunnel between MAG2 and LMA.

According to the baseline solution signaling method described in [RFC6224], after MN2 mobility, MAG2 acting in its role of MLD proxy will send an MLD Query to the newly observed MN on its downlink. Assuming that the subsequent MLD Report from MN2 requests membership for a new multicast group (from MAG2's point of view), this will then result in an Aggregated MLD Report being sent to the MTMA from MAG2. This message will be sent through a multicast tunnel between MAG2 and MTMA (pre-established or dynamically established) .

When MN2 detaches, MAG1 may keep the multicast tunnel with the multicast MTMA if there are still other MNs using the multicast tunnel. Even if there are no MNs currently on the multicast tunnel, MAG1 may decide to keep the multicast tunnel for potential future use.

As discussed above, existing MLD (and MLD proxy) signaling will handle a large part of the multicast mobility management for the MN.

```
           MN2          MAG1       MAG2         LMA       MTMA
           |        (MLD Proxy) (MLD Proxy)  (Unicast)(Multicast)
           |            |           |            |          |
         MN Attached    |           |            |          |
          to MAG1       |           |            |          |
           |            |           |            |          |
           |            |========= Multicast Tunnel ======= |
           |            |           |            |          |
         MN Detaches    |           |            |          |
          from MAG1     |           |            |          |
           |            |           |            |          |
           |            |           |            |          |
         MN Attaches    |           |            |          |
          to MAG2       |           |            |          |
           |            |           |            |          |
           |---------Rtr Sol------ >|            |          |
           |            |           |--- PBU -- >|          |
           |            |           |            |          |
           |            |           |<-- PBA ----|          |
           |            |           |            |          |
           |<-----Rtr Adv --------- |            |          |
           |            |           |            |          |
           |            |           |==Multicast Tunnel === |
           |            |           |            |          |
           |<---------MLD Query---- |            |          |
           |            |           |            |          |
           |---MLD Report (G) ----> |            |          |
           |            |           |            |          |
           |            |           |---- Aggregated -----> |
           |            |           |    MLD Report (G)     |
           |            |           |            |          |
           |            |           |            |          |
           |< --------- Multicast Traffic ---------------- >|
           |            |           |            |          |
           |            |           |            |          |
```

Figure 3: Multicast Mobility Signaling for MTMA

3.3.2. MAG as Multicast Router

If the MAG is a multicast router, the system behavior when operating with remote subscription is as described before, considering that a multicast routing protocol is running between the MAG and the MTMA on the tunnel interface. Even once the MAG has decided to obtain the multicast traffic remotely based for instance on routing information and/or network management criteria, this decision can be dynamically changed if such criteria changes. This behavior is further described in section Section 6.2.

3.4. Operations of the Multicast Tree Mobility Anchor

The MTMA provides connectivity to the multicast infrastructure out of the PMIPv6 domain. The MTMA itself could either act as an additional MLD proxy (only in the case where all the connected MAGs act also as MLD proxies), reporting to a further node an aggregated view of the subscriptions in a PMIPv6 multicast domain; or it can act as a designated multicast router for all the MAGs in a PMIPv6 multicast domain. The MTMA will then request the multicast content on behalf of the MAGs (and MNs behind them). In addition, the MTMA will create and maintain the corresponding multicast forwarding states per each tunnel interface towards the MAGs. Whatever the role played, when the MAGs act as MLD proxy, the MTMA becomes the MLD querier of the MLD proxy instance located in each MAG.

4. Direct Routing

Direct routing uses native multicast infrastructure, allowing a MAG to directly connect a multicast router in the PMIPv6 domain. A MAG can act as a MLD proxy or multicast router for redirecting multicast packets.

The main purpose of direct routing is to provide optimal routing for local content. As a consequence, it alleviates the MTMA of the channel management and data delivery of locally available content. Unicast traffic will go normally to the LMAs in the PMIPv6 domain.

This section describes how the direct routing works in scenarios of MN attachment and multicast mobility.

4.1. Overview

Figure 4 shows the architecture for the local routing case using native multicasting infrastructure [I-D.deng-multimob-pmip6-requirement].

The LMA is dedicated to unicast traffic, and the multicast traffic is obtained from an upstream multicast router present in the PMIPv6 domain. Note that there can be multiple LMAs for unicast traffic (not shown in Figure 1) in a given PMIPv6 domain.

As shown in Figure 4, a MAG may connect to both unicast (LMA) and multicast (MR) nodes. Thus, a given MN may simultaneously receive both unicast and multicast traffic.

As seen in Figure 4, each MAG has a direct connection (i.e., not using the tunnel interface) with a multicast router. To facilitate

IGMP/MLD signaling and multicast packets forwarding, a MLD proxy function defined in [RFC4605], or multicast routing function SHOULD be placed on the MAG.

```
                           Multicast Tree
                                  :
                                  :         || - PMIPv6 Tunnel
       +----------+        +----------+     |  - Multicast Data Path
       |   LMA    |        |    MR    |
       +----------+        +----------+
            ||  \\           /      |
            ||   \\         /       |
            ||    \\       /        |
            ||     \\     /         |
            ||      \\   /          |
            ||       \\ /           |
            ||        \\            |
            ||        /\\           |
            ||       /  \\          |
            ||      /    \\         |
            ||     /      \\        |
            ||    /        \\       |
       +----------+       +----------+
       |  P-MAG   |       |  N-MAG   |    MLD Proxy or Multicast Router
       +----------+       +----------+
            :                   :
        +------+           +------+
        |  MN  |   ----->  |  MN  |
        +------+           +------+
```

Figure 4: Architecture for direct routing based PMIPv6 multicasting

4.2. Operations of the Mobile Node

The MN operation is not impacted by the direct routing option. The MN will act according to the stated operations in [RFC5213] and [RFC6224].

This draft considers that every MN requesting multicast-only services is previously registered in a PMIPv6 unicast domain to get a unicast IP address. This registration can also be required for several purposes such as remote management, billing, etc.

4.3. Operations of the Mobile Access Gateway

There are two main functionalities in the MAG when it supports direct routing. One is the when the MAG incorporates MLD proxy functions as per [RFC4605]. The other case is when the MAG functions as a multicast router as per [RFC4601] or [RFC4607].

The following sections describe the MAG for both cases in more detail.

4.3.1. MAG as MLD Proxy

In case the MAG only incorporates MLD proxy functionality, for every one of the MLD proxy instances invoked in the MAG it is necessary to define at configuration time the upstream interface from where the multicast traffic will be received. This decision requires to define whether the multicast subscription by an MLD proxy instance for all the multicast channels will be local (if the upstream interface points to a multicast router internal to the PMIPv6 domain) or remote (in case of the upstream interface is the bi-directional tunnel towards the LMA, for the architecture in [RFC6224], or the MTMA, for the multicast listener optimization described in this document).

4.3.1.1. Multicast Establishment

If the MAG has MLD proxy functionality only, once the MLD proxy instance is configured to obtain the multicast traffic locally, the system behavior remains static.

In Figure 5, the MAG first establishes the PMIPv6 tunnel with LMA for unicast traffic as defined in [RFC5213] after being triggered by the Router Solicitation message from the MN. Unicast traffic will then flow between the MN and LMA.

For multicast traffic, it is assumed that the upstream interface of the MLD proxy instance has been configured pointing to a multicast router internal to the PMIPv6 domain (or towards an additional MLD proxy node in the domain), for all the multicast channels (which, in consequence, have to be local). There should be direct connectivity between the MAG and the local multicast router (or additional MLD proxy).

Upon detecting node attachment from an incoming interface, the MAG adds each downstream interface to the MLD Proxy instance with upstream link to a MR according to the standard MLD proxy operations and sends an MLD Query message towards the MN. The MN sends the MLD report message (when required by its upper layer applications) as defined in [RFC3810] in response to an MLD Query from MAG.Upon receiving the MLD Report message from each incoming interface, the MAG checks the MLD Proxy instance associated with the downstream interface and then the MLD Report messages will be aggregated and forwarded to the upstream link associated with the MR (assuming that this is a new multicast group which the MAG had not previously subscribed to). Multicast traffic will then flow from the local multicast router towards the MN.

```
     MN                     MAG           LMA           MR
      |                 (MLD Proxy)    (Unicast)    (Multicast)
 MN attaches to MAG1          |              |              |
      |                       |              |              |
      |------Rtr Sol----- ->|              |              |
      |                       |---- PBU --->|              |
      |                       |              |              |
      |                       |<--- PBA ----|              |
      |                       |              |              |
      |                       |== Unicast ==|              |
      |                       |    Tunnel    |              |
      |<---- Rtr Adv -------|              |              |
      |                       |              |              |
      |<------- Unicast Traffic --------->|              |
      |                       |              |              |
      |                       |              |              |
      |<--- MLD Query ------|<------ MLD Query ---------|
      |                       |              |              |
  MN requires                 |              |              |
  multicast services          |              |              |
      |                       |              |              |
      |-- MLD Report (G) -->|              |              |
      |                       |              |              |
      |                       |------- Aggregated ------->|
      |                       |       MLD Report (G)       |
      |                       |              |              |
      |                       |              |              |
      |<------------ Multicast Traffic ---------------->|
      |                       |              |              |
```

Figure 5: Multicast service establishment for direct routing

4.3.1.2. Multicast mobility

```
         MN           P-MAG       N-MAG        LMA        MR
          |             |           |           |          |
          |             |           |           |          |
          |<------------|<-- Multicast Data----------------|
          |             |       .   |           |          |
          |             |       .   |           |          |
          |             |       .   |           |          |
       Link         Handover        |           |          |
    Disconnected    Detection       |           |          |
          |             |           |           |          |
          |             |           |           |          |
          |             |    MN Attachment      |          |
          |             |           |           |          |
          |             |           |           |          |
          |------- Rtr. Sol. ------>|           |          |
          |             |           |           |          |
          |<------ MLD Query -------|           |          |
          |             |           |           |          |
          |------- MLD Report ----->|           |          |
          |             |           |    Aggregated        |
          |             |           |---- MLD Report ----->|
          |             |           |           |          |
          |             |           |           |          |
          |<------------------------|<-- Multicast Data ---|
          |             |           |           |          |
          |             |           |           |          |
          |             |           |--- PBU -->|          |
          |             |           |           |          |
          |             |           |<-- PBA ---|          |
          |             |           |           |          |
```

Figure 6: Multicast mobility signaling for direct routing

Figure 6 shows the handover operation procedure in the local direct routing architecture. When an MN hands off to the next MAG (N-MAG) from the previous MAG (P-MAG), the N-MAG detects the newly arrived attached MN and performs binding update procedure by exchanging PBU/PBA signaling messages with LMA. At the same time, a MLD Proxy instance detecting the new MN transmits an MLD query message to the MN. After receiving the MLD query message, the MN sends an MLD report message that includes the multicast group information. The N-MAG then sends an aggregated MLD report message to the upstream link associated with the MR. In the direct routing case, an upstream interface of MLD Proxy instance is decided towards certain multicast router based on the operator’s configuration or multicast routing, as compared to the base solution defined in [RFC6224] where it is determined for each MN based on the Proxy Binding Update List. When the N-MAG receives the multicast packets from the MR, it then simply forwards them without tunnel encapsulation. The N-MAG updates the MN's location information to the LMA by exchanging PBU/PBA signaling messages.

4.3.2. MAG as multicast router

If the MAG behaves as a multicast router, the MAG then implements a multicast routing protocol. This allows the MAG to make decisions about from where to receive the traffic of any multicast channel, based on routing information and/or network management criteria. The selected incoming interface for receiving multicast traffic will be then the one matching such criteria, and it could drive to either a local or remote subscription. Some situations are introduced in the next section.

If the MAG is a multicast router, the system behavior when operating with local subscription is as before, but extending the role of the MAG to be a multicast router, and running a multicast routing protocol among the MAG and local multicast router serving the multicast traffic. Once the MAG decides to obtain the multicast traffic locally based in routing information and/or network management criteria, this can be dynamically changed if such criteria change.

5. Functions and Requirements

A set of new functions and structures are needed in PMIPv6 to allow the use of the solution described in this document. The following sub-sections describe these required extensions.

5.1. Extension to the Binding Update List in MAG

The Binding Update List in the MAG must be updated to be able to handle the fact that more than one entity (i.e. LMA and MTMA) may be serving the mobile node.

5.2. Extension of the Policy Profile Information including multicast related parameters

A given mobile node's policy profile information must be updated to be able to store the IPv6 addresses of both the LMA and MTMA, for the remote subscription case.

Additionally, when the MAG act as multicast router in the local subscription case it is required to keep registration of the IP address for the rendez-vous point in the PMIPv6 domain, when PIM-SM is used. When using PIM-SSM, the IP addresses of the local multicast sources have to be also registered.

5.3. Data Structure in MTMA

The MTMA does not directly interact with the MNs attached to any of the MAGs. The MTMA only manages the multicast groups subscribed per MAG on behalf of the MNs attached to it. Having this in mind, the relevant information to be stored in the MTMA should be the tunnel interface identifier (tunnel-if-id) of the bi-directional tunnel for multicast between the MTMA and every MAG (e.g. similar to what it is stated in [RFC5213] for the unicast case), the IP addresses of the multicast group delivered per tunnel to each of the MAGs, and the IP addresses of the sources injecting the multicast traffic per tunnel to the multicast domain defined by the MTMA.

6. Dynamic Selection Support

As mentioned above, the MAG as multicast router provides some flexibility for choosing local versus remote multicast subscription. With this approach IP multicast traffic can selectively be received from the home, visited or local domains, and the selection of traffic can be based on operator policies. Considering PIM as the multicast routing protocol running on the MAG, it is possible to find out two situations where such dynamic selection can occur, according to the PIM flavor on place. For all the scenarios below we consider a certain multicast flow being injected by two different sources, one local to the PMIPv6 domain and one remote through the home network, by using an MTMA.

6.1. Use Cases

The MAG has different options to subscribe to a multicast group, such as:

- Via the tunnel with the LMA unicast [RFC6224]

- Via the tunnel with the MTMA (as described in Section 3)

- Via local subscription/routing (as described in Section 4)

Also, the content can be located in different places. For instance, the content might be locally available (e.g. TV channels offered in the visited domain), or the content might be remote (e.g. TV channels offered in the home domain). In case the content is available remotely at the home network it is preferred to subscribe via the MTMA tunnel to home. However, if the content is available locally, it is preferred to subscribe at the MAG (local break point) instead of via the home network. The MAG may therefore have to choose which approach needs to be taken to subscribe to a particular content requested by a particular MN.

- If the IP address of the source injecting a certain multicast group is local (scope: local domain), the MAG should get access to it via local subscription (or routing, if the MAG is a multicast router).

- If IP address of the source injecting a certain multicast group is global (or the scope is broader than the local domain), the MAG may have to decide among the different available options (i.e. RFC6224, Local Routing, or MTMA). This can be achieved through some static or dynamic configuration at the MAG.

6.2. Any Source Multicast Scenario

This situation applies for both PIM-SM and BIDIR PIM variants. In this case, once the MAG receives the MLD report from the MN requesting the multicast channel in the form (*,G), the MAG could decide what multicast flow subscribes to (either the local or the remote one).

The subscription can be statically pre-configured or dynamically configured based on some rule. For instance, static configuration can be made per MN (user), such as "multicast traffic from user X should always go through the home (i.e., via the tunnel with the MTMA/LMA-as-per-RFC6224), while traffic from user Y should go via local subscription". Also, configuration profiles can also be more complex and include considerations on types of traffic or IP flows, such as "traffic of type A from user X should always go through the home, traffic of type B from user X should be subscribed locally" using routing information and/or network management criteria. Similarly, routing information can be received dynamically. For example, at user's registration time PBU/PBA signaling can be used to carry the profile information similar to what is described in [I-D.ietf-netext-pmipv6-sipto-option]. Also, routing information can be exchanged dynamically when the multicast group subscription is made.

In case of using PIM-SM, another scenario is possible. PIM-SM allows switching from a multicast shared-tree to a source-specific tree to optimize the path for traffic delivery. The location of the rendezvous point and the multicast source can either be in the PMIPv6 domain or the home network, so the optimization could be from local subscription to remote subscription or vice versa. The possibility of switching to a source-based tree, and the time for doing so is implementation-dependent, and this could be triggered immediately (e.g. after reception of the first multicast packet) or after some time, or may not even switch at all.

6.3. Source Specific Multicast Scenario

This situation applies for PIM-SSM. Then, in a source-specific multicast scenario [RFC4607], the MAG would send the PIM request to the corresponding interface based on the multicast source address indicated on the (S,G) subscription requested by the MN in the MLD Report, using the routing information.

7. IANA Considerations

TBD.

8. Security Considerations

This draft discusses the operations of existing protocols without modifications. It does not introduce new security threats beyond the current security considerations of PMIPv6 [RFC5213], MLD [RFC3810], IGMP [RFC3376] and IGMP/MLD Proxying [RFC4605].

9. Authors

Additional co-authors of this document are:

Akbar Rahman

InterDigital Communications, LLC

E-mail: akbar.rahman@interdigital.com

Ignacio Soto

Universidad Politecnica de Madrid

E-mail: isoto@dit.upm.es

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC4065] Kempf, J., "Instructions for Seamoby and Experimental Mobility Protocol IANA Allocations", RFC 4065, July 2005.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC4605] 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.

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

[RFC5847] Devarapalli, V., Koodli, R., Lim, H., Kant, N., Krishnan, S., and J. Laganier, "Heartbeat Mechanism for Proxy Mobile IPv6", RFC 5847, June 2010.

[RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

10.2. Informative References

[I-D.deng-multimob-pmip6-requirement]
Deng, H., Chen, G., Schmidt, T., Seite, P., and P. Yang, "Multicast Support Requirements for Proxy Mobile IPv6", draft-deng-multimob-pmip6-requirement-02 (work in progress), July 2009.

[I-D.ietf-netext-pmipv6-sipto-option]
Gundavelli, G., Zhou, X., Korhonen, J., and R. Koodli, "IPv4 Traffic Offload Selector Option for Proxy Mobile

IPv6", draft-ietf-netext-pmipv6-sipto-option-04 (work in progress), February 2012.

[RFC6224] Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

Appendix A. MTMA Deployment Use Cases

From the network architecture point of view, there are several options when considering the multicast tree mobility anchor (MTMA) approach. These options can be distinguished in terms of the number of LMAs and MTMAs present in a PMIPv6 domain and the service relationship that a set of MNs gets from them, in the form of a "LMA : MTMA" ratio. According to that, it is possible to differentiate the following approaches:

A set of MNs is served in a PMIPv6 domain by two entities, one MTMA for multicast service, and one LMA for unicast, in such a way that the ratio is 1:1 (one common PMIPv6 unicast and multicast domain).

A set of MNs is served in a PMIPv6 domain by several entities, one MTMA for multicast service, while the others (LMAs) for unicast, in such a way that the ratio is N:1 (N PMIPv6 unicast domains coexist with a unique multicast domain).

A set of MNs is served in a PMIPv6 domain by several entities, one LMA for unicast, while the others (MTMAs) are devoted to multicast service, in such a way that the ratio is 1:N (one single PMIPv6 unicast domain coexists with multiple multicast domains).

Scenarios with an N:M ratio are considered to be a combination of the previous ones.

A.1. PMIPv6 domain with ratio 1:1

This approach basically refers to the architecture presented in Figure 1. Within this approach, a common set of MNs is served by a couple of entities, one LMA for unicast and one MTMA for multicast. All the MNs of the set are served by these two elements as they move in the PMIPv6 domain.

A.2. PMIPv6 domain with ratio N:1

This approach basically refers to the situation where a common set of MNs is served by a unique MTMA for multicast service, but simultaneously there are subsets from that group of MNs which are served by distinct LMAs for unicast service as they move in the PMIPv6 domain. Each particular MN association with the LMAs (unicast) and MTMA (multicast) remains always the same as it moves in the PMIPv6 domain.

Figure 7 shows the scenario here described.

```
            +----------------+       +----------------+
            |Content Source A|       |Content Source B|
            +----------------+       +----------------+
                   |                      |
                   |                      |
         ***  ***  ***  ***  ***  ***  ***  *** *** *** ***
        *   **   **   **   **  **   **   **   **   **  **  *
       *                                                    *
       *                 Fixed Internet                     *
       *        (Unicast & Multicast Traffic)               *
        *   **   **   **   **  **   **   **   **   **  **  *
         ***  ***  ***  *** *** ***  ***  ***  ***  ***  ***
           |                     |                      |
           |                     |                      |
           |                     |                      |
        +------+        +-----------------+          +------+
        | LMA1 |        |       MTMA2     |          | LMA3 |
        +------+        +-----------------+          +------+
          || \\        oo    oo      oo   oo          //  ||
          ||  \\      oo     oo      oo    oo        //   ||
          ||   \\    oo      oo      oo     oo      //    ||
          ||    \\  oo       oo      oo      oo    //     ||
          ||     \\oo        oo      oo       oo  //      ||
          ||      \\         oo      oo        oo//       ||
          ||     oo\\        oo      oo         //        ||
          ||    oo  \\       oo      oo        //oo       ||
          ||   oo    \\      oo      oo       //  oo      ||
          ||  oo      \\     oo      oo      //    oo     ||
        +------+      +--------+     +--------+     +--------+
        | MAG1 |      |  MAG2  |     |  MAG3  |     |  MAG4  |
        +------+      +--------+     +--------+     +--------+
        |      |       |      |       |      |       |      |
        |      |       |      |       |      |       |      |
     {MN10}  {MN11}  {MN20} {MN21}  {MN30} {MN31} {MN40} {MN41}
```

Figure 7: PMIPv6 domain with ratio N:1

The Figure 7 proposes an architecture where there are two entities acting as LMAs, LMA1 and LMA3, while there is another one, named MTMA2, working as multicast tree mobility anchor. LMA1 and LMA3 constitute two distinct unicast domains, whereas MTMA2 forms a single multicast domain. The tunnels among MAGs and LMAs represented by lines ("||") indicate a tunnel transporting unicast traffic, while the tunnels among MAGs and MTMA2 depicted with circles ("o") show a tunnel transporting multicast traffic.

In the figure it can be observed that all the MNs are served by MTMA2 for the incoming multicast traffic from sources A or B. However, there are different subsets regarding unicast traffic which maintain distinct associations within the PMIPv6 domain. For instance, the subset formed by MN10, MN11, MN20 and MN21 is served by LMA1 for unicast, and the rest of MNs are being served by LMA3. For the scenario described above, the association between each MN and the corresponding LMA and MTMA is permanently maintained.

A.3. PMIPv6 domain with ratio 1:N

This approach is related to a scenario where a common group of MNs is served by a unique LMA for unicast service, but simultaneously there are subsets from that group of MNs which are served by distinct MTMAs for multicast service as they move in the PMIPv6 domain. Each particular MN association with the LMA and MTMAs (unicast and multicast respectively) remains always the same as it moves in the PMIPv6 domain.

Figure 8 shows the scenario here described.

```
            +----------------+       +----------------+
            |Content Source A|       |Content Source B|
            +----------------+       +----------------+
                   |                      |
                   |                      |
         ***  ***  ***  ***  ***  ***  ***  *** *** *** ***
        *   **   **   **   **  **   **   **   **   **  **  *
       *                                                    *
       *                 Fixed Internet                     *
       *        (Unicast & Multicast Traffic)               *
        *   **   **   **   **  **   **   **   **   **  **  *
         ***  ***  ***  *** *** ***  ***  ***  ***  ***  ***
           |                     |                      |
           |                     |                      |
           |                     |                      |
        +------+        +-----------------+          +------+
        | MTMA1|        |       LMA2      |          | MTMA3|
        +------+        +-----------------+          +------+
          oo oo        //    ||      ||   \\          oo  oo
          oo  oo      //     ||      ||    \\        oo   oo
          oo   oo    //      ||      ||     \\      oo    oo
          oo    oo  //       ||      ||      \\    oo     oo
          oo     oo//        ||      ||       \\  oo      oo
          oo      oo         ||      ||        \\oo       oo
          oo     //oo        ||      ||         \\        oo
          oo    //  oo       ||      ||        oo\\       oo
          oo   //    oo      ||      ||       oo  \\      oo
          oo  //      oo     ||      ||      oo    \\     oo
        +------+      +--------+     +--------+     +--------+
        | MAG1 |      |  MAG2  |     |  MAG3  |     |  MAG4  |
        +------+      +--------+     +--------+     +--------+
        |      |       |      |       |      |       |      |
        |      |       |      |       |      |       |      |
     {MN10}  {MN11}  {MN20} {MN21}  {MN30} {MN31} {MN40} {MN41}
```

Figure 8: PMIPv6 domain with ratio 1:N

The Figure 8 proposes an architecture where the LMA2 is the unique LMA for a certain group of MNs, while there are two others entities, MTMA1 and MTMA3, acting as MTMAs for different subsets of MNs of the same group. MTMA1 and MTMA3 constitute two distinct multicast domains, whereas LMA2 forms a single unicast domain. Each MTMA could be devoted to carry on a different content (for instance, MTMA1 for source A and MTMA3 for source B) or not. Looking at the picture, the subset formed by MN10, MN11, MN20 and MN21 is served by MTMA1 for multicast. The rest of MNs are being served by MTMA3 also for multicast. Finally, all of them are served by LMA2 for unicast. For the scenario described above, the association between each MN and the corresponding LMA and MTMA is permanently maintained.

A.4. PMIPv6 domain with H-LMA

The H-LMA is defined as an entity which simultaneously transports unicast and multicast service, that is, it simultaneously works as LMA and MTMA. In the context of the MTMA solution, an H-LMA can play the role of MTMA for an entire group of MNs in a PMIPv6 domain, while acting simultaneously as LMA for a subset of them. The figure 9 adapts the PMIPv6 domain with ratio N:1 scenario of figure 7 to the case where MTMA2 is an H-LMA, which serves multicast traffic to all the MNs in the picture, and simultaneously, it is able to serve unicast traffic to the subset formed by MN30, MN40 and MN41.

```
            +----------------+       +----------------+
            |Content Source A|       |Content Source B|
            +----------------+       +----------------+
                   |                       |
                   |                       |
         ***  ***  ***  ***  ***  ***  ***  *** *** *** ***
        *   **   **   **   **  **   **   **   **   **  **  *
       *                                                    *
       *                   Fixed Internet                   *
       *         (Unicast & Multicast Traffic)              *
        *   **   **   **   **  **   **   **   **   **  **  *
         ***  ***  ***  *** *** ***  ***  ***  ***  ***  ***
           |                     |                       |
           |                     |                       |
           |                     |                       |
        +-------+       +-----------------+         +-------+
        | LMA1  |       |      H-LMA      |         | LMA3  |
        +-------+       +-----------------+         +-------+
          || \\        oo    db      db   oo          //  ||
          ||  \\      oo     db      db    oo        //   ||
          ||   \\    oo      db      db     oo      //    ||
          ||    \\  oo       db      db      oo    //     ||
          ||     \\oo        db      db       oo  //      ||
          ||      \\         db      db        oo//       ||
          ||     oo\\        db      db         //        ||
          ||    oo  \\       db      db        //oo       ||
          ||   oo    \\      db      db       //  oo      ||
          ||  oo      \\     db      db      //    oo     ||
        +-------+     +---------+    +---------+    +---------+
        | MAG1  |     |  MAG2   |    |  MAG3   |    |  MAG4   |
        +-------+     +---------+    +---------+    +---------+
        |      |       |      |       |      |       |      |
        |      |       |      |       |      |       |      |
     {MN10}  {MN11}  {MN20} {MN21}  {MN30} {MN31} {MN40} {MN41}
```

Figure 9: PMIPv6 domain with H-LMA

Figure 8 presents a PMIPv6 network where there are two pure unicast LMAs, LMA1 and LMA3, and a hybrid LMA, labeled as H-LMA in the figure. The H-LMA is an MTMA from the perspective of MAG1 and MAG4. The tunnels among MAGs and LMAs represented by lines ("||") indicate a tunnel transporting exclusively unicast traffic, the tunnels depicted with circles ("o") show a tunnel transporting exclusively multicast traffic, and the tunnels with mixed lines and circles ("db") describe a tunnel transporting both types of traffic simultaneously.

All of the MNs in the figure receive the multicast traffic from H-LMA (one single multicast domain), but it is possible to distinguish three subsets from the unicast service perspective (that is, three unicast domains). The first subset is the one formed by MN10, MN11 and MN 20, which receives unicast traffic from LMA1. A second subset is the one formed by MN21 and MN30, which receives unicast traffic from H-LMA. And finally, a third subset is built on MN31, MN40 and MN41, which receives unicast traffic from LMA3. For the scenario described above, the association between each MN and the corresponding LMA and H-LMA is permanently maintained.

Authors' Addresses

Juan Carlos Zuniga
InterDigital Communications, LLC
1000 Sherbrooke Street West, 10th floor
Montreal, Quebec H3A 3G4
Canada

Email: JuanCarlos.Zuniga@InterDigital.com
URI: http://www.InterDigital.com/

Luis M. Contreras
Telefonica I+D
Don Ramon de la Cruz, 82-84
Madrid 28006
Spain

Email: lmcm@tid.es

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/

Seil Jeon
Instituto de Telecomunicacoes
Campus Universitario de Santiago
Aveiro 3810-193
Portugal

Email: seiljeon@av.it.pt
URI: https://atnog.av.it.pt/~sjeon/

Younghan Kim
Soongsil University
Sangdo-dong, Dongjak-gu
Seoul 511
Republic of Korea

Email: yhkim@dcn.ssu.ac.kr
URI: http://dcnlab.ssu.ac.kr/