MULTIMOB Group   J.C. Zuniga
INTERNET-DRAFT   A. Rahman
Intended Status: Standards Track   InterDigital Communications, LLC
Expires: April 2011   L.M. Contreras
                                                            C.J. Bernardos
                                        Universidad Carlos III de Madrid
                                                                 I. Soto
                                       Universidad Politecnica de Madrid
                                                          October 25, 2010

Support Multicast Services Using Proxy Mobile IPv6
draft-zuniga-multimob-smspmip-04.txt

Status of this Memo

Copyright and License Notice

Abstract

The MULTIMOB group has specified a base solution to support IP multicasting in a PMIPv6 domain [I-D.draft-ietf-multimob-pmipv6-base-solution]. In this document, an enhancement is proposed to the base solution to use a dedicated multicast LMA as the topological anchor point for multicast traffic, while the MAG remains as an IGMP/MLD proxy. This enhancement provides benefits such as reducing multicast traffic replication and supporting different PMIPv6 deployments scenarios.

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.

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 [RFC4605] allows an intermediate (edge) node to appear as a multicast router to downstream hosts, and as a host to upstream multicast routers. IGMP and MLD related protocols 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 [I-D.draft-ietf-multimob-pmipv6-base-solution]. In this document, an enhancement is proposed to the base solution to use a dedicated multicast LMA (M-LMA) as the topological anchor point for multicast traffic, while the MAG remains as an IGMP/MLD proxy. This enhancement allows different PMIPv6 deployment scenarios. It also eliminates the so called "Tunnel Convergence problem" where the MAG may receive the same multicast packet from several LMAs. There are no impacts to the MN to support multicast listener mobility from this document.

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 [RFC2119].

This document uses the terminology defined in [RFC5213], [RFC3775], 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 this specification. 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.

Additionally, some definitions are introduced, as follows.

- U-LMA or Unicast-LMA: LMA entity dedicated to unicast service exclusively.

- M-LMA or Multicast-LMA: LMA entity dedicated to multicast service exclusively.

- H-LMA or Hybrid-LMA: LMA entity dedicated to both unicast and multicast services.

3 Solution

A PMIPv6 domain may handle data from both unicast and multicast sources. A dedicated multicast LMA can be used to serve as the mobility anchor for multicast traffic. Unicast traffic will go normally to the other LMAs in the PMIPv6 domain. This section describes how the multicast LMA works in scenarios of MN attachment and multicast mobility. We first concentrate on the case of both LMAs (multicast and unicast) defining a unique PMIPv6 domain, and then different deployment scenarios are presented.

3.1 Architecture

Figure 1 shows an example of a PMIPv6 domain supporting multicast mobility. LMA1 is dedicated to unicast traffic, and LMA2 is dedicated to multicast traffic. The multicast traffic LMA (LMA2) 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 (not shown in Figure 1) in a given PMIPv6 domain. Similarly, more than one multicast dedicated LMA can be deployed by the operator (not shown in Figure 1).

Also in this architecture, all MAGs that are connected to the multicast LMA must support the MLD proxy [RFC4605] function. Specifically in Figure 1, each of the MAG1-LMA2 and MAG2-LMA2 tunnel interfaces defines an MLD proxy domain. The MNs are considered to be on the downstream interface of the MLD proxy (in the MAG), and LMA2 is considered to be on the upstream interface (of the MAG) as per [RFC4605]. Note that 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.

As shown in Figure 1, MAG1 may connect to both unicast and multicast LMAs. 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  | LMA1|                 | LMA2 |     Multicast
      Anchor  +-----+                 +------+      Anchor
                  \\                    // ||
                   \\                  //  ||
                    \\                //   ||
                     \\              //    ||
                      \\            //     ||
                       \\          //      ||
                        \\        //       ||
                         \\      //        ||
                          \\    //         ||
                          +-----+       +-----+
                          | MAG1|       | MAG2|      MLD Proxy
                          +-----+       +-----+
                          |     |          |
                          |     |          |
                        {MN1} {MN2}      {MN3}
```

Figure 1. Architecture of Dedicated LMA as Multicast Anchor

3.2 Deployment Scenarios

From the network architecture point of view, there are a several options when considering the dedicated multicast LMA (M-LMA) approach. These options can be distinguished in terms of the number

of unicast and multicast LMAs present in a PMIPv6 domain and the service relationship that a set of MN get from them, in the form of a "U-LMA : M-LMA" ratio. According to that, it is possible to differentiate the following approaches:

- A set of MNs is served in a PMIPv6 domain by two LMAs, one for multicast service, the other one for unicast, in such a way that the ratio is 1:1.

- A set of MNs is served in a PMIPv6 domain by several LMAs, one for multicast service, while the rest for unicast, in such a way that the ratio is N:1.

- A set of MNs is served in a PMIPv6 domain by several LMAs, one for unicast, while the rest are devoted to multicast service, in such a way that the ratio is 1:N.

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

3.2.1 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 LMAs, one for unicast and the other one for multicast. All the MNs of the set are served by these two LMAs as they move in the PMIPv6 domain.

3.2.2 PMIPv6 domain with ratio N:1

This approach basically refers to the situation where a common set of MNs is served by a unique LMA 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 multicast) remains always the same as it moves in the PMIPv6 domain.

Figure 2 shows the scenario here described.

```
            +----------------+       +----------------+
            |Content Source A|       |Content Source B|
            +----------------+       +----------------+
                   |                       |
                   |                       |
         ***  ***  ***  ***  ***  ***  ***  *** *** *** ***
        *   **   **   **   **  **   **   **   **   **  **  *
       *                                                    *
       *                 Fixed Internet                     *
       *        (Unicast & Multicast Traffic)               *
        *   **   **   **   **  **   **   **   **   **  **  *
         ***  ***  ***  *** *** ***  ***  ***  ***  ***  ***
           |                     |                      |
           |                     |                      |
           |                     |                      |
        +------+        +------------------+        +------+
        | LMA1 |        |       LMA2       |        | 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 2. PMIPv6 domain with ratio N:1

The figure 2 proposes an architecture where there are two LMAs, LMA1 and LMA3, acting as U-LMAs, while there is another one, the LMA2, working as dedicated M-LMA. The tunnels among MAGs and LMAs represented by lines ("||") indicate a tunnel transporting unicast traffic, while the tunnels depicted with circles ("o") show a tunnel transporting multicast traffic.

In the figure it can be observed that all the MNs are served by LMA2 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 U-LMA and M-LMA is permanently maintained.

3.2.3 PMIPv6 domain with ratio 1:N

This approach is related to an 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 LMAs for multicast service as they move in the PMIPv6 domain. Each particular MN association with the LMAs (unicast and multicast) remains always the same as it moves in the PMIPv6 domain.

Figure 3 shows the scenario here described.

```
            +----------------+       +----------------+
            |Content Source A|       |Content Source B|
            +----------------+       +----------------+
                   |                      |
                   |                      |
         ***  ***  ***  ***  ***  ***  ***  *** *** *** ***
        *   **   **   **   **  **   **   **   **   **  **  *
       *                                                    *
       *                 Fixed Internet                     *
       *        (Unicast & Multicast Traffic)               *
        *   **   **   **   **  **   **   **   **   **  **  *
         ***  ***  ***  *** *** ***  ***  ***  ***  ***  ***
           |                     |                      |
           |                     |                      |
           |                     |                      |
        +------+        +-----------------+          +------+
        | LMA1 |        |       LMA2      |          | 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    \\     oo
        +------+      +--------+     +--------+     +--------+
        | MAG1 |      |  MAG2  |     |  MAG3  |     |  MAG4  |
        +------+      +--------+     +--------+     +--------+
        |      |       |      |       |      |       |      |
        |      |       |      |       |      |       |      |
     {MN10}  {MN11}  {MN20} {MN21}  {MN30} {MN31} {MN40} {MN41}
```

Figure 3. PMIPv6 domain with ratio 1:N

The figure 3 proposes an architecture where the LMA2 is the unique U-LMA for a certain group of MNs, while there are two others LMAs, LMA1 and LMA3, act as M-LMAs for different subsets of MNs of the same group. Each M-LMA could be devoted to carry on a different content (for instance, LMA1 for source A and LMA3 for source B) or not. Looking at the picture, the subset formed by MN10, MN11, MN20 and MN21 is served by LMA1 for multicast. The rest of MNs are being served by LMA3 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 U-LMA and M-LMA is permanently maintained.

3.2.4 PMIPv6 domain with H-LMA

The H-LMA is defined as an LMA which simultaneously transports unicast and multicast service. In the context of the dedicated M-LMA solution, an H-LMA can play the role of M-LMA for an entire group of MNs in a PMIPv6 domain, while acting simultaneously as U-LMA for a subset of them. The figure 4 adapts the PMIPv6 domain with ratio 1:N scenario of figure 3 to the case where LMA2 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 |        |       LMA2      |          | 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 4. PMIPv6 domain with H-LMA

The figure 4 presents a PMIPv6 network where there are two pure unicast LMAs, LMA1 and LMA3, and a hybrid LMA, the LMA2. The LMA2 is a dedicated M-LMA 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 LMA2, but it is possible to distinguish three subsets from the unicast service perspective. 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 LMA2. 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 U-LMA and M-LMA is permanently maintained.

3.3 Multicast Establishment

Figure 5 shows the procedure when MN1 attaches to MAG1, and establishes associations with LMA1 (unicast) and LMA2 (multicast).

```
              MN1                   MAG1       LMA1       LMA2
               |                (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 5. MN Attachment and Multicast Service Establishment

In Figure 5, MAG1 first establishes the PMIPv6 tunnel with LMA1 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 LMA1.

For multicast traffic, a multicast tunnel may have been pre-configured between MAG1 and the multicast LMA (LMA2). Or the multicast tunnel 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 MAG1. MAG1 acting as a MLD Proxy as defined in [RFC4605] will then send an Aggregated MLD Report to the multicast anchor, LMA2

(assuming that this is a new multicast group which MAG1 had not previously subscribed to). Multicast traffic will then flow from LMA2 towards MN1.

3.4 Multicast Mobility

Figure 6 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 MAG2 is connected only to LMA2 (multicast) and does not receive unicast traffic. Of course, if it was desired to support unicast traffic, the architecture will easily allow MAG2 to also connect to LMA1 to support unicast traffic.

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 of a new multicast group (from MAG2's point of view), this will then result in an Aggregated MLD Report being sent to LMA2 from MAG2. This message will be sent through a pre-established (or dynamically established) multicast tunnel between MAG2 and LMA2.

When MN2 detaches, MAG1 may keep the multicast tunnel with the multicast LMA2 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 Proxy MLD) signaling will handle a large part of the multicast mobility management for the MN.

```
          MN2          MAG1       MAG2         LMA1       LMA2
           |        (MLD Proxy) (MLD Proxy)  (Unicast)(Multicast)
           |            |           |            |          |
         MN Attached    |           |            |          |
          To MAG1       |           |            |          |
           |            |           |            |          |
           |            |========= Multicast Tunnel ======= |
           |            |           |            |          |
         MN Detaches    |           |            |          |
          From MAG1     |           |            |          |
           |            |           |            |          |
           |            |           |            |          |
         MN Attaches    |           |            |          |
          To MAG2       |           |            |          |
           |            |           |            |          |
           |            |           |==Multicast Tunnel === |
           |            |           |            |          |
           |---------Rtr Sol------ >|            |          |
           |            |           |            |          |
           |<-----Rtr Adv --------- |            |          |
           |            |           |            |          |
           |            |           |            |          |
           |<---------MLD Query---- |            |          |
           |            |           |            |          |
           |---MLD Report (G) ----> |            |          |
           |            |           |            |          |
           |            |           |---- Aggregated -----> |
           |            |           |    MLD Report (G)     |
           |            |           |            |          |

           |            |           |            |          |
           |< --------- Multicast Traffic ---------------- >|
           |            |           |            |          |
           |            |           |            |          |
```

Figure 6. Multicast Mobility Signaling

3.5 PMIPv6 enhancements

This section describes the enhancements to the Proxy Mobile IPv6 [RFC5213] protocol required to support the M-LMA architecture.

3.5.1 New 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 LMA (i.e. U-LMA and M-LMA) may be serving the mobile node.

3.5.2 Policy Profile Information with Multicast Parameters

A given mobile node's policy profile information must be updated to be able to store the IPv6 addresses of both the U-LMA and M-LMA.

3.5.3 MAG to M-LMA attach requirements

The MAG procedures must be updated to be able to handle simultaneous attach for a given mobile node to both the U-LMA and M-LMA. For example, packets coming from a given mobile node must be screened to determine if it should be sent to the U-LMA or to the M-LMA.

3.6 Advantages

An advantage of the proposed dedicated multicast LMA (M-LMA) architecture is that it allows a PMIPv6 domain to closely follow a simple multicast tree topology for Proxy MLD forwarding (cf., sections 1.1 and 1.2 of [RFC4605]). In contrast, the combined unicast/multicast LMA as proposed in [I-D.draft-ietf-multimob-pmipv6-base-solution] will be a more complex set of trees.

Another advantage of the proposed dedicated multicast solution is that it allows a gradual network upgrade of a PMIPv6 domain to support multicast functionality. This is because the operator does not have to upgrade all the LMAs in the network to support multicast functionality. Only certain LMAs, dedicated to multicast support, will have to be upgraded to support the new multicast functionality. Also, multiple deployment scenarios are supported as required by the operator for expected traffic distributions.

A final advantage is that a dedicated multicast LMA minimizes replication of multicast packets (the Tunnel Convergence problem), in certain scenarios, compared to [I-D.draft-ietf-multimob-pmipv6-base-solution]. Figures 7 and 8 illustrate this point visually. For this simple scenario, it can be observed that the dedicated multicast LMA topology (Figure 7) generates 6 packets for one input multicast packet. In comparison, the combined unicast/multicast LMA topology (Figure 8) generates 8 packets for one input multicast packet.

In general, it can be seen that the extra multiplication of packets in the combined unicast/multicast LMA topology will be proportional to the number of LMAs, and the number of MNs (in a given MAG)

associated to different LMAs, for a given multicast group. The packet multiplication problem aggravates as more MNs associated to different LMAs receive the same multicast traffic when attached to the same MAG. Hence, the dedicated multicast architecture significantly decreases the network capacity requirements in this scenario.

(Note that in Figure 7, it is assumed that MN1 and MN2 are associated with MAG1-LMA1, and MN3 is associated with MAG2-LMA2 for multicast traffic. In Figure 8, it is assumed that MN1 is associated with MAG1-LMA1, MN2 is associated with MAG1-LMA2, and MN3 is associated with MAG2-LMA2 for multicast traffic. In both Figures 7 and 8, it is assumed that the packets are transmitted point to point on the last hop wireless link.)

```
                                   +--------------+
                                   |Content Source|
                                   +--------------+
                                          |
                                          |
                                        +---+     Packet destined
                                        | 1 |   for Multicast group "G"
                                        +---+
                                          |
         ***  ***  ***  ***      ***  ***  ***  ***
        *   **   **   **   *    *   **   **   **     *
       *                    *  *                     *
       *  Unicast Traffic   *  *  Multicast Traffic  *
       *                    *  *                     *
        *   **   **   **   *    *   **   **   **   *
         ***  ***  ***  ***      ***  ***  ***  ***
                 |                       |
                 |                     +---+
                 |                     | 2 |
                 |                     +---+
                 |                       |
              +-----+                 +------+
     Unicast  | LMA1|                 | LMA2 |     Multicast
      Anchor  +-----+                 +------+      Anchor
                 \\                     //||
                  \\                   // ||
                   \\                 //  ||
                    \\               //   ||
                     \\          +---+  +---+
                      \\         | 3 |  | 4 |
                       \\        +---+  +---+
                        \\       //       ||
                         \\     //        ||
                          \\   //         ||
                           \\ //          ||
                          +-----+       +-----+
                          | MAG1|       | MAG2|      MLD Proxy
                          +-----+       +-----+
                          |     |          |
                        +---+ +---+      +---+
                        | 5 | | 6 |      | 7 |
                        +---+ +---+      +---+
                          |     |          |         All MNs in same
                          |     |          |       multicast group "G"
                        {MN1} {MN2}      {MN3}
```

Figure 7. Packet Flow in a Dedicated Multicast LMA

```
                       +--------------+
                       |Content Source|
                       +--------------+
                              |
                              |
                            +---+      Packet destined
                            | 1 |    for Multicast group "G"
                            +---+
                              |
         ***  ***  ***  ***  ***  ***  ***  *** ***
        *   **   **   **   **  **   **   **   **   *
       *                                            *
       *                 Fixed Internet             *
       *        (Unicast & Multicast Traffic)       *
        *   **   **   **   **  **   **   **   **   *
         ***  ***  ***  ***      *** ***  ***  ***
                 |                       |
               +---+                   +---+
               | 2 |                   | 3 |
               +---+                   +---+
                 |                       |
              +-----+                 +------+     Combined
              | LMA1|                 | LMA2 |      Unicast/Multicast
              +-----+                 +------+       Anchor
                 \\                   //  ||
                  \\                 //   ||
                   \\               //    ||
                    \\             //     ||
                    +---+        +---+  +---+
                    | 4 |        | 5 |  | 6 |
                    +---+        +---+  +---+
                        \\       //       ||
                         \\     //        ||
                          \\   //         ||
                           \\ //          ||
                          +------+      +------+     MLD Proxy
                          | MAG1|       | MAG2|
                          +------+      +------+
                          |     |          |
                        +---+ +---+      +---+
                        | 7 | | 8 |      | 9 |
                        +---+ +---+      +---+
                          |     |          |         All MNs in same
                          |     |          |       multicast group "G"
                        {MN1} {MN2}      {MN3}
```

Figure 8. Packet Flow in a Combined Unicast/Multicast LMA

4 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].

5 IANA Considerations

This document makes no request of IANA.

6 References

6.1 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.

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

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

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

[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.

6.2 Informative References

[I-D.draft-ietf-multimob-pmipv6-base-solution] Schmidt, T.C., Waehlisch, M., and S.Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-05 (work in progress), July 28, 2010.

Author's Addresses

Juan Carlos Zuniga
InterDigital Communications, LLC
Email: JuanCarlos.Zuniga@InterDigital.com

Akbar Rahman
InterDigital Communications, LLC
Email: Akbar.Rahman@InterDigital.com

Luis M. Contreras
Universidad Carlos III de Madrid
Email: luisc@it.uc3m.es

Carlos J. Bernardos
Universidad Carlos III de Madrid
Email: cjbc@it.uc3m.es

Ignacio Soto
Universidad Politecnica de Madrid
Email: isoto@dit.upm.es