MULTIMOB Group J.C. Zuniga
INTERNET-DRAFT A. Rahman
Intended Status: Standards Track InterDigital Communications, LLC
Expires: January 2012 L.M. Contreras
C.J. Bernardos
Universidad Carlos III de Madrid
I. Soto
Universidad Politecnica de Madrid
July 11, 2011

Support Multicast Services Using Proxy Mobile IPv6
<draft-zuniga-multimob-smspmip-06.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 [RFC6224]. In this document, an enhancement is proposed to the base solution to use a multicast tree mobility anchor 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 [RFC6224]. In this document, an enhancement is proposed to the base solution to use a multicast tree mobility anchor (MTMA) 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 [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 in such a way that an MN using that service is not aware of 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.

This means 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 exclusively.

- 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 Solution

A PMIPv6 domain may handle data from both unicast and multicast sources. This document addresses an optimization of the base solution specified for multicast support in PMIPv6 domains [RFC6224] by introducing a complementary network entity, named multicast tree mobility anchor (MTMA), and defining the architecture and protocol flows derived from it. An MTMA can be used to serve as the mobility anchor for multicast traffic. 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.

This section describes how the MTMA works in scenarios of MN attachment and multicast mobility. We first concentrate on the case of both LMA and MTMA 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 MTMA1 is dedicated to multicast traffic. The tree mobility anchor MTMA1 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 MTMAs can be deployed by the operator (not shown in Figure 1).

Also in this architecture, all MAGs that are connected to the MTMA must support the MLD proxy [RFC4605] function. Specifically in Figure 1, each of the MAG1-MTMA1 and MAG2-MTMA1 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 MTMA1 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 (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, despite of that, this draft considers that every MN demanding multicast-only services is previously registered in a PMIPv6 unicast domain to get a unicast IP address. This registration can be required also for several purposes such as remote management, billing, etc.

```
                                   +--------------+
                                   |Content Source|
                                   +--------------+
                                          |
                                          |
         ***  ***  ***  ***      ***  ***  ***  ***
        *   **   **   **   *    *   **   **   **    *
       *                    *  *                     *
       *  Unicast Traffic   *  *  Multicast Traffic  *
       *                    *  *                     *
        *   **   **   **   *    *   **   **   **   *
         ***  ***  ***  ***      ***  ***  ***  ***
                 |                        |
                 |                        |
                 |                        |
              +-----+                 +------+
     Unicast  | LMA1|                 | MTMA1|     Multicast
      Anchor  +-----+                 +------+      Anchor
                  \\                    // ||
                   \\                  //  ||
                    \\                //   ||
                     \\              //    ||
                      \\            //     ||
                       \\          //      ||
                        \\        //       ||
                         \\      //        ||
                          \\    //         ||
                          +-----+       +-----+
                          | MAG1|       | MAG2|      MLD Proxy
                          +-----+       +-----+
                          |     |          |
                          |     |          |
                        {MN1} {MN2}      {MN3}
```

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

3.2 Deployment Scenarios

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.

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 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.

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 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 2 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 2. PMIPv6 domain with ratio N:1

The figure 2 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.

3.2.3 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 3 shows the scenario here described.

The figure 3 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.

```
            +----------------+       +----------------+
            |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 3. PMIPv6 domain with ratio 1:N

3.2.4 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 4 adapts the PMIPv6 domain with ratio N:1 scenario of figure 2 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 4. PMIPv6 domain with H-LMA

Figure 4 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.

3.3 Multicast Establishment

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

```
           MN1                   MAG1       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 5. MN Attachment and Multicast Service Establishment

In Figure 5, MAG1 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 MAG1 and MTMA. 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, MTMA (assuming that this is a new multicast group which MAG1 had not

previously subscribed to). Multicast traffic will then flow from MTMA 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 we only consider the tunnel of MAG2 with MTMA (for multicast traffic) and we assume that MN2 does not receive unicast traffic. Of course, if it was desired to support unicast traffic, this is served by a tunnel between MAG2 and LMA to transfer unicast traffic.

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

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 Proxy MLD) 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       |           |            |          |
            |            |           |            |          |
            |            |           |==Multicast Tunnel === |
            |            |           |            |          |
            |---------Rtr Sol------ >|            |          |
            |            |           |--- PBU --->|          |
            |            |           |            |          |
            |            |           |<-- PBA ----|          |
            |<-----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 MTMA 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 entity (i.e. LMA and MTMA) 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 LMA and MTMA.

3.5.3 MAG to MTMA attach requirements

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

3.5.4. Data structure stored by 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 (as 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.

3.6 Advantages

An advantage of the proposed MTMA 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 [RFC6224] 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 nodes (MTMAs), 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 specific multicast elements minimize the replication of multicast packets (the Tunnel Convergence problem), in certain scenarios, compared to [RFC6224]. Figures 7 and 8 illustrate this point visually. For this simple scenario, it can be observed that the multicast MTMA 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 MTMA 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-MTMA2 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.)

Additional results can be found in [ERCIM], where both solutions are compared by simulation under realistic traffic conditions. It can be shown that, for multicast traffic, the number of channels that a node (LMA in the base solution, MTMA in the proposed multicast architecture) has to serve does not decrease linearly with the reduction of the number of MNs associated to that node. The key factor is the set of channels subscribed by the MNs. In fact, as the number of MNs increases in the PMIPv6 domain, we have less advantage for having several nodes serving multicast, as each of them will probably manage all the multicast channels (or at least the popular ones) anyway.

```
                                   +--------------+
                                   |Content Source|
                                   +--------------+
                                          |
                                          |
                                        +---+     Packet destined
                                        | 1 |   for Multicast group "G"
                                        +---+
                                          |
         ***  ***  ***  ***      ***  ***  ***  ***
        *   **   **   **   *    *   **   **   **     *
       *                    *  *                     *
       *  Unicast Traffic   *  *  Multicast Traffic  *
       *                    *  *                     *
        *   **   **   **   *    *   **   **   **    *
         ***  ***  ***  ***      ***  ***  ***  ***
                 |                       |
                 |                     +---+
                 |                     | 2 |
                 |                     +---+
                 |                       |
              +-----+                +------+
     Unicast  | LMA1|                | MTMA2|     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 the MTMA architecture

```
                       +--------------+
                       |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 |
                    +---+        +---+  +---+
                        \\       //       ||
                         \\     //        ||
                          \\   //         ||
                           \\ //          ||
                          +-----+       +-----+
                          | MAG1|       | MAG2|      MLD Proxy
                          +-----+       +-----+
                          |     |          |
                        +---+ +---+      +---+
                        | 7 | | 8 |      | 9 |
                        +---+ +---+      +---+
                          |     |          |         All MNs in same
                          |     |          |       multicast group "G"
                        {MN1} {MN2}      {MN3}
```

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

4 Consideration of MAG as multicast router in the tunnel interface to MTMA

In the architecture described before, all MAGs that are connected to the MTMA are considered to act as MLD proxies. This follows the MAG characterization provided in [RFC6224]. However, interesting advantages can be derived from the fact of converting the MAG node in a multicast router in the tunnel interface towards the MTMA, that is, in implementing PIM protocol ([RFC4601], [RFC4607]) in the tunnel interface, in case the MAG connects to more than one MTMA in the PMIPv6 domain.

This could be the case, for instance, in which a PMIPv6 domain provides access to MNs of different home networks, each home network using a distinct MTMA to provide multicast service in the PMIPv6 domain. With the MAG working as a multicast router in the tunnel interface, in a source-specific multicast scenario [RFC4607], the MAG could send the PIM request to the corresponding MTMA based on the multicast source address.

Another possible scenario for connecting more than one MTMA to a MAG could be the case of a home network using different MTMAs to serve different content over the same PMIPv6 domain for scalability reasons, or as a way to provide backup in case of MTMA failure.

5 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].

6 IANA Considerations

This document makes no request of IANA.

7 References

7.1 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.

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

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

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

7.2 Informative References

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

[ERCIM] L.M. Contreras, C.J. Bernardos, I. Soto, "On the efficiency of a dedicated LMA for multicast traffic distribution in PMIPv6 domains", 5th ERCIM Workshop in eMobility, Vilanova i la Geltru, Spain, June 2011.

[ETSI] ETSI TS 102 034, "Digital Video Broadcasting (DVB); Transport of MPEG-2 TS Based DVB Services over IP Based Networks", v1.4.1, August, 2009.

Appendix A. Overhead analysis of the proposed MTMA architecture.

This appendix provides an analysis of the overhead introduced by the proposed multicast architecture. In this solution an MTMA is used to serve the multicast traffic to the MNs. The MAGs in the PMIPv6 domain

are connected to the MTMA through a tunnel which is used to deliver the multicast flows subscribed by the MNs attached to the MAG.

A very common way for video delivery over IP networks is the transport of MPEG-2 Transport Streams (TS) encapsulated in RTP/UDP/IP datagrams, as described in [ETSI].

An MPEG-2 transport stream is a packet of 188 bytes. So, an Ethernet frame with 1500 bytes of payload can carry a maximum of up to 7 MPEG-2 TS packets.

When encapsulating those 7 MPEG-2 TS packets in RTP/UDP/IP datagrams we are forming a datagram of length 7*188 (MPEG-2 TS) + 12 (RTP) + 8 (UDP) + 40 (IPv6) = 1376 bytes.

In the proposed multicast architecture, such datagram should be transported over the tunnel existing between a MAG and the MTMA. That tunnel implies an IP-in-IP encapsulation, that is, an additional 40 byte length header should be added to the datagram. In this situation, the overhead caused by the MTMA approach can be calculated as 40 / (40 + 1376) = 2,8%.

This results in a minimal overhead derived from the use of the tunnel between MTMA and MAG.

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: contreras.uc3m@gmail.com

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

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