MULTIMOB Group S. Figueiredo
Internet Draft Universidade de Aveiro
Intended status: Informational S. Jeon
Expires: January 16, 2013 Instituto de Telecomunicacoes
R. L. Aguiar
Universidade de Aveiro
July 16, 2012

IP Multicast Use Case Analysis for PMIPv6.based Distributed Mobility Management

draft-sfigueiredo-multimob-use-case-dmm-02.txt

Status of this Memo

Copyright Notice

Abstract

Mobile networks are changing towards distributed mobility management, tackling inefficiencies in network management and packet routing. Identifying IP multicast use cases applicable into DMM would be meaningful before exploring solution spaces. This document describes use cases when IP multicast is applied on DMM environment using Base Solution approach specified in [RFC6224], and presents issues for each use case.

Table of Contents

1. Introduction

Centralized mobility management approach brings several drawbacks such as single point of failure, non-optimal routing and severe overloading on the anchor point. It is expected to be much severe as data traffic consumed by mobile devices increases.

In order to tackle these problems, distributed mobility management (DMM) is introduced, bringing the mobility anchor closer to the MN.

IP multicast provides an efficient method for distributing multimedia contents, but it was mainly designed for fixed networks. [RFC6224] specified the Base Solution applicable to network-based Proxy Mobile IPv6 (PMIPv6) protocol, based on centralized mobility management. IP multicast should be also specified in DMM, but the application of the Base Solution for multicast support in PMIPv6 standardized in IETF MULTIMOB WG needs to be identified first.

This document briefly describes use cases of IP multicast in a PMIPv6-based DMM environment, following DMM Requirements [DMMREQ], and introduces consequent problems. Both listener and sender cases are studied.

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], [RFC6275], and [RFC3810], and [RFC4601]. Also, new entities are defined relying on the PMIPv6 entities specified in [RFC5213]:

- Mobility Access Router (MAR): A router with the capability of acting both as a mobility anchor and as an access router, in a per flow basis.

- Previous Mobility Access Router (P-MAR): The MAR where the MN was attached to previously to the network-layer mobility process, and that may be acting as an anchor for one or multiple flows.

- New Mobility Access Gateway (N-MAR): The MAR to which the MN is currently attached, providing the access functionality and thus delivers all the flows destined to the MN.

- Multicast Listener Discovery Proxy (MLD-P): An entity providing MLD based forwarding following the operation defined in [RFC4605].

In the current document, only MLDv2-based signaling is considered, targeting IPv6 networks (REQ3 from [DMMREQ]).

3. Use Cases Description

3.1. Assumptions

This draft refers to requirements of DMM as a base reference architecture, with the major goal of showing multicast use cases [DMMREQ]. Following the recommendation of reusing the existing mobility protocols, the identified IP multicast model derives from PMIPv6 Base multicast mobility solution [RFC6224]. As such, MAG and LMA functionalities defined in [RFC5213] are assumed to be installed in a mobility access router (MAR), defined in this document. A MAR provides tunnel-based forwarding to provide a home network prefix (HNP)-based flow with the necessary IP session continuity whenever the MN moves to another MAR.

3.2. Multicast listener support

When a MN initially attaches to the P-MAR (as shown in Figure 1), it receives a HNP address which will be associated with communications started at that MAR. As the P-MAR detects the new logical link, it transmits a general MLD Query message - to which the MN will not yet reply, as it is not yet running any multicast session. The P-MAR then adds the downstream logical link to the MLD Proxy instance [RFC4605]. In this case, i.e. when users subscribe to multicast content only after associating with the MAR, the MLD Proxy will set its uplink to the multicast infrastructure. When the MN intends to start receiving the multicast session, it will send an unsolicited MLD Report, triggered by its application. On receiving the latter message, the MLD Proxy tries to join the multicast channel(s) by sending an aggregated MLD Report through the MLD Proxy upstream interface. Note that the same MLD Proxy instance will be assigned to all MNs which initiated their multicast subscriptions in the current MAR (i.e. the MNs having no multicast mobility session). When the joining procedure ends, multicast data is transmitted through the same interfaces, until reaching the MN.

```
                  +----------------+
                  |   Multicast    |
                  | Infrastructure |
                  +----------------+
                          *
                          * (S,G)
                          *
                     +-----------+               +-----------+
                     |  P-MAR    |---------------|  N-MAR    |
                     |           |***************|           |
                     | (MLD-P)   |---------------| (MLD-P)   |
                     +-----------+               +-----------+
                          *                            *
                          *                            *
                       +------+                     +------+
                       |  MN  |       ----->        |  MN  |
                       +------+                     +------+
```

Figure 1 Multicasting architecture using distributed mobility management

When the MN moves from P-MAR, the N-MAR is required to establish a tunnel for IP session continuity of the flows being sent towards and from the MN's HNP assigned by the P-MAR. This implies that N-MAR has an appropriate method to know the P-MAR. Multiple ideas are supposed to be made at the solution stage of DMM WG, therefore it is out of scope of this document. Following the operation of the MLD Proxy [RFC4605], after the bidirectional tunnel establishment, the following process takes place. First, the N-MAR sends a General MLD Query, and verifies whether the MN is admissible for multicast sessions. Then, the MLD Proxy at the N-MAR adds the downstream interface corresponding to the MN, and configures the upstream interface towards the MN's P-MAR. This is simple and applicable as a network-based multicast DMM approach. However, this approach leads to a couple of relevant issues.

3.2.1.1. Duplicated Traffic

One of the problems is traffic duplication. This is a result of the tunnel convergence problem occurring in [RFC6224]. As shown in Figure 2, MN1 and MN3, which moved from MAR1 and MAR3, respectively, are currently located at the MAR2. Through their respective tunnels, they receive multicast packets of the same channel through different anchoring MARs. This causes duplicated traffic, converging to the MAR2, with the magnitude of replicated traffic, which may be much bigger than that of PMIPv6 when we assume that the number of MARs in future DMM domains is expected to be much larger than that of LMAs at core level within a PMIPv6 domain.

```
          +----------------+
          | Multicast Tree | *
          +----------------+     *
                  *        *         *
                  *          *           *
                  *            *             *
            (S,G) *       (S,G)  *               *  (S,G)
                  *                *                 *
             +----------+  (-->)  +----------+ (<--)   +----------+
             |   MAR1   |---------|   MAR2   |---------|   MAR3   |
             |          |*********|          |*********|          |
             | (MLD-P)  |---------| (MLD-P)  |---------|  (MLD-P) |
             +----------+  Tun.1  +----------+  Tun.2  +----------+
                                    *   *   *
                                   *    *    *
                                  *     *     *
              +---+  move      +---+  +---+ +---+   move   +---+
              |MN1|  --->      |MN1|  |MN2| |MN3|   <---   |MN3|
              +---+            +---+  +---+ +---+          +---+

                (<--/-->) : direction of the multicast packet flow
```

Figure 2 Data replication

3.2.1.2. Non-optimal routing

Another issue is non-optimal routing (Figure 3). If we consider a significantly large domain, multicast packets MAY traverse a long distance, depending on the setup direction of the upstream interface of MLD Proxy instances. The issue is more obvious if we assume all MARs are connected to the multicast infrastructure.

```
            +----------------+
            |   Multicast    |
            | Infrastructure |
            +----------------+
                     *
                     * (S,G)
                     *
                +-----------+                 +-----------+
                |  P-MAR    |------      ------|  N-MAR    |
                |           |****** ... ******|           |
                |(MLD-P)    |------      ------| (MLD-P)   |
                +-----------+                 +-----------+
                     *                              *
                     *                              *
                  +-------+                     +-------+
                  |  MN   |       ----->        |  MN   |
                  +-------+                     +-------+
```

Figure 3 Non-optimal routing problem

3.3. Multicast sender support

Mobile multicast sender using PMIPv6 base solution is being defined in [SENDER]. Basically, MLD Proxies provide the ability for MAGs to forward the multicast traffic to the MN’s corresponding LMA.

To allow the sender to deliver multicast content to the multicast tree, the MLD Proxy should configure its upstream interface towards a multicast router [PM-HOME]. Depending on the network topology, it may also be configured towards a MLD Proxy placed on a neighbor MAR. On the multicast source’s mobility (Figure 4), an identical operation to the listener mobility case is expected from the MLD Proxy behavior. In this case, the source uploads multicast traffic through one of MLD Proxy’s downstream interfaces, and the traffic is forwarded through the uplink interface towards the P-MAR.

```
                      +------+         +-----------------+
                      |  RP  |---------|   Multicast     |
                      +------+         | Infrastructure |
                          *            +-----------------+
                          * (S,G)               |
                          *                     |
                     +-----------+          +-----------+
                     |  P-MAR    |----------|  N-MAR    |
                     |           |**********|           |
                     | (MLD-P)   |----------| (MLD-P)   |
                     +-----------+          +-----------+
                          *                      *
                          *                      *
                       +------+              +------+
                       |   S  |   ---->      |  S   |
                       +------+              +------+
```

Figure 4 Multicast sender mobility

3.3.1.1. Triangular routing

When a source moves to N-MAR from P-MAR, multicast data will be sent through the mobility tunnel between N-MAR and P-MAR (Figure 5). If a listener (L1) attaches to the same MAR (N-MAR), it will receive the multicast data through multicast infrastructure, following the regular configuration of MLD Proxy. Hence, the multicast data is routed non-optimally between the source and the listener, going from N-MAR to P-MAR, to the multicast routing tree, and then back to N-MAR again before reaching the listener.

```
                 +------+         +----------------+
                 |  RP  |*********|   Multicast    |
                 +------+         | Infrastructure |
                     *            +----------------+
                     * (S,G)             *
                     *                   *
                +-----------+      +-----------+
                |   P-MAR   |-------|   N-MAR   |
                |           |*******|           |
                | (MLD-P)   |-------| (MLD-P)   |
                +-----------+      +-----------+
                     *                *     *
                     *                *     *
                  +------+  move +------+ +-----+
                  |  S   |  ---> |   S  | |  L1 |
                  +------+       +------+ +-----+
```

Figure 5 Triangular routing after source mobility

A similar problem occurs in the opposite process, i.e. if a multicast source starts transmitting multicast content at a MAR, and a listener moves to the same MAR while receiving the source’s content (Figure 6).

```
                 +------+         +----------------+
                 |  RP  |*********|   Multicast    |
                 +------+         | Infrastructure |
                     *            +----------------+
                     * (S,G)             *
                     *                   *
                +-----------+      +-----------+
                |   N-MAR   |-------|   P-MAR   |
                |           |*******|           |
                | (MLD-P)   |-------| (MLD-P)   |
                +-----------+      +-----------+
                  *       *              *
                  *       *              *
              +------+  +----+  move  +----+
              |  S   |  | L1 |  <---  | L1 |
              +------+  +----+        +----+
```

Figure 6 Triangular routing after listener mobility

When the source and the listener are within a same MAR (MAR2), if both the source and listener try to send the session and receive it, respectively, the traffic will be optimally sent to the listener without going through native multicast infrastructure. As the traffic reaches the MLD Proxy via the downstream interface to which the source is attached, it will be sent through the downstream interface to which the listener sent the MLD Report. However, if the source and the listener move to different MARs, the traffic will traverse the following non-optimal path, even though they share a common anchor:

```
  Source -> MAR1 -> MAR2 -> Multicast Tree -> MAR2 -> MAR3
```

This problem is depicted in Figure 7.

```
                          +----------------+
                          | Multicast Tree |
                          +----------------+
                                 *  *
                                 *  *
                                 *  *
                                 *  *
                                 *  *
        +----------+  (-->)  +----------+  (-->)  +----------+
        |   MAR1   |---------|   MAR2   |---------|   MAR3   |
        |          |*********|          |*********|          |
        | (MLD-P)  |---------| (MLD-P)  |---------|  (MLD-P) |
        +----------+  Tun.1  +----------+  Tun.2  +----------+
            *                  *       *              *
            *                 *         *             *
            *                *           *            *
          +---+    move    +---+       +---+  move  +---+
          | S |    <---    | S |       | L |  -->   | L |
          +---+            +---+       +---+        +---+

             (<--/-->) : direction of the multicast packet flow
```

Figure 7 Multicast traffic non-optimal routing due to both mobile sender and listener

3.3.1.2. Non-distributed anchoring

REQ1 from [DMMREQ] refers that "DMM MUST enable a distributed deployment of mobility management of IP sessions so that the traffic can be routed in an optimal manner without traversing centrally deployed mobility anchors". When a MN subscribes to a new multicast session with existing multicast mobility session, the Aggregated MLD Report containing all the MN’s multicast subscriptions will be sent from the current MLD Proxy through the same uplink interface, i.e. towards a single multicast mobility anchor. This results in some of previously identified issues (e.g. non-optimality in the path that both the subscription and multicast traffic traverse). It can be stated that the MLD Proxy nature doesn’t comply with the aforementioned requirement, leading to the subscription of any multicast flow using the same multicast mobility data path.

This problem is depicted in Figure 8, where both multicast flow 1 and flow 2 reach MAR2 from MAR1, being flow 2’s optimal routing path affected by the mobility status of the MN, and in particular by the order in which the multicast flows were subscribed. While this issue is not exclusively related to mobile multicast sources, it is better depicted and its’ impact in the routing is more obvious when considering one.

```
                             +----------------+
                             | Multicast Tree |
                           * +----------------+#
                         *                      #
                       *                          #
                     *                              #
                   *                                  #
           +----------+  (-->)                      +----------+
           |   MAR1   |---------            -------|   MAR2   |
           |          |#*#*#*#*#............#*#*#*#|          |
           | (MLD-P)  |---------            -------|  (MLD-P) |
           +----------+  Tun.                       +----------+
               *                                      *    #
               *                          downstream) #    # (upstream)
               *                                      *    #
               *                                      #    #
             +---+                move              +---+  +---+
             | L |            ---------->           | L |  | S |
             +---+                                  +---+  +---+

             * : Multicast flow 1
             # : Multicast flow 2 (subscribed after some time in MAR2)
```

Figure 8 Non-optimal routing due to single MLD Proxy uplink

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

TBD

6. References

6.1. Normative References

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

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

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

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

[RFC4605] B. Fenner, H. He, B. Haberman, and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) Based Multicast Forwarding ("IGMP/MLD Proxying")", IETF RFC 4605, August 2006.

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

6.2. Informative References

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

[DMMREQ] H. Chan, "Requirements of distributed mobility management", draft-ietf-dmm-requirements-01 (work in progress), July 2012.

[SENDER] T C. Schmidt et al, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains", draft-ietf-multimob-pmipv6-source-00 (work expired), January 2012.

[PM-HOME] S. Jeon, N. Kang, and Y. Kim, "Mobility Management based on Proxy Mobile IPv6 for Multicasting Services in Home Networks," IEEE Transactions on Consumer Electronics (TCE), vol. 55, no. 3, pp. 1227-1232, August 2009.

Authors' Addresses

Sergio Figueiredo
Universidade de Aveiro
3810-193 Aveiro, Portugal

E-mail: sfigueiredo@av.it.pt

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

E-mail: seiljeon@av.it.pt

Rui L. Aguiar
Universidade de Aveiro
3810-193 Aveiro, Portugal

E-mail: ruilaa@ua.pt