A PMIPv6-based solution for Distributed Mobility Management
draft-bernardos-mext-dmm-pmip-01

Abstract

   The number of mobile users and their traffic demand is expected to be
   ever-increasing in future years, and this growth can represent a
   limitation for deploying current mobility management schemes that are
   intrinsically centralized, e.g., Mobile IPv6 and Proxy MIPv6. For
   this reason it has been waved a need for distributed and dynamic
   mobility management approaches, with the objective of reducing
   operators' burdens, evolving to a cheaper and more efficient
   architecture.

   This draft describes a solution to distribute the data forwarding
   plane on Proxy Mobile IPv6 domains, thus trying to overcome the
   suboptimal data path introduced when the LMA is traversed.

Status of this Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on January 12, 2012.

Table of Contents

1. Introduction

Current IP mobility solutions, standardized with the names of Mobile IPv6 [RFC3775], or Proxy Mobile IPv6 [RFC5213], just to cite the two most relevant examples, offer mobility support at the cost of handling operations at a cardinal point, the mobility anchor, and burdening it with data forwarding and control mechanisms for a great amount of users. As stated in [I-D.chan-distributed-mobility-ps], centralized mobility solutions are prone to several problems and limitations: longer (sub-optimal) routing paths, scalability problems, signaling overhead (and most likely a longer associated handover latency), more complex network deployment, higher vulnerability due to the existence of a potential single point of failure, and lack of granularity on the mobility management service (i.e., mobility is offered on a per-node basis, not being possible to define finer granularity policies, as for example per-application).

There are basically two main approaches being researched now: one aimed at making Mobile IPv6 work in a distributed way (a complete solution can be found in [I-D.bernardos-mext-dmm-cmip]), and another one doing the same exercise for Proxy Mobile IPv6. In this draft we describe a solution to achieve a DMM behavior for network-based mobility support (i.e., inspired by PMIPv6). This document is based on a research paper of the same authors, currently under submission, called "A Network-based Localized Mobility Solution for Distributed Mobility Management" [Net-basedDMM].

2. Terminology

The following terms used in this document are defined in the Proxy Mobile IPv6 specification [RFC5213]:

   Local Mobility Anchor (LMA)

   Mobile Access Gateway (MAG)

   Mobile Node (MN)

   Binding Cache Entry (BCE)

   Proxy Care-of Address (P-CoA)

   Proxy Binding Update (PBU)

   Proxy Binding Acknowledgement (PBA)

The following terms are defined and used in this document:

   MAAR (Mobility Anchor and Access Router). First hop routers where
      the mobile nodes attach to. They also play the role of mobility
      managers for the IPv6 prefixes they anchor.

   CMD (Central Mobility Database). Node that stores the BCEs for the
      MNs in the mobility domain.

   A-MAAR (Anchor MAAR). MAAR which was previously visited by the MN
      and is still involved in an active flow using an IPv6 prefix it
      has advertised to the MN (i.e., MAAR where that IPv6 prefix is
      anchored).

   S-MAAR (Serving MAAR). MAAR which the MN is currently attached to.

[3]. Description of the solution

   The purpose of Distributed Mobility Management approaches is to
   overcome the limitations of the traditional centralized mobility
   management by bringing the mobility anchor closer to the MN.
   Following this idea, in our proposal, the central anchor is moved to
   the edge of the network, being deployed in the default gateway of the
   mobile node. That is, the first elements that provide IP
   connectivity to a set of MNs are also the mobility managers for those
   MNs. In the following, we will call these Mobility Anchor and Access
   Routers (MAARs).

   However, MAARs leverage on the Central Mobility Database (CMD) to
   access and update information related to the MNs, stored as mobility
   sessions; hence, a centralized node maintains a global view on the
   status of the network. The CMD is queried whenever a MN is detected
   to join the mobility domain. It might be a fresh attachment or a
   handover, but as MAARs do not store any mobilty session, they contact
   the CMD to retrieve the data of interest and eventually take the
   appropriate action. The procedure adopted for the query and the
   messages exchange sequence might vary to optimize the update latency
   and/or the signaling overhead. Here is presented one method for the
   initial registration, and three different approaches to update the
   mobility sessions using PBUs and PBAs. Each approach assigns a
   different role to the CMD:

   o The CMD is a PBU/PBA relay

   o The CMD is only a MAAR locator

   o PBU/PBA is a PBU/PBA proxy

3.1. Initial registration

Upon the MN's attachment to a MAAR, say MAAR1, an IPv6 global prefix
belonging to the MAAR's prefix pool is reserved for it (Pref1). The
prefix is sent in a PBU with the MN's Identifier (MN-ID) to the CMD,
which, since the session is new, stores a Binding Cache Entry
containing as main fields the MN-ID, the MN's prefix and MAAR1's
address as Proxy-CoA. The CMD replies to MAAR1 with a PBA indicating
that the MN's registration is fresh and no past status is available.
MAAR1 sends a Router Advertisement (RA) in unicast to the MN
including the prefix reserved before, that can be used by the MN to
configure an IPv6 address (e.g., with stateless auto-configuration).
The address is routable at the MAAR, in the sense that it is on the
path of packets addressed to the MN. Moreover, the MAAR acts as
plain router for those packets, as no encapsulation nor special
handling takes place. Figure 1 illustrates this scenario.

```
   +-----+       +---+                    +--+
   |MAAR1|       |CMD|                    |CN|
   +-----+       +---+                    +*-+
      |            |                        *
     MN            |                        *      +---+
   attach.         |                    *****    _|CMD|_
  detection        |            flow1 *        / +-+-+ \
      |            |                  *       /    |     \
      |--- PBU -->|                   *      /     |      \
      |            |                   *    /      |       \
      |           BCE          +---*-+-'    +--+--+    `+-----+
      |         creation       |   * |      |    |    |     |
      |            |           |MAAR1+------+MAAR2+------+MAAR3|
      |<-- PBA ---|            |   * |      |    |    |     |
      |            |           +---*-+      +-----+    +-----+
      |            |                *
      |            |          Pref1 *
      |            |                +*-+
      |            |                |MN|
      |            |                +--+

   (Operations sequence)            (Packets flow)
```
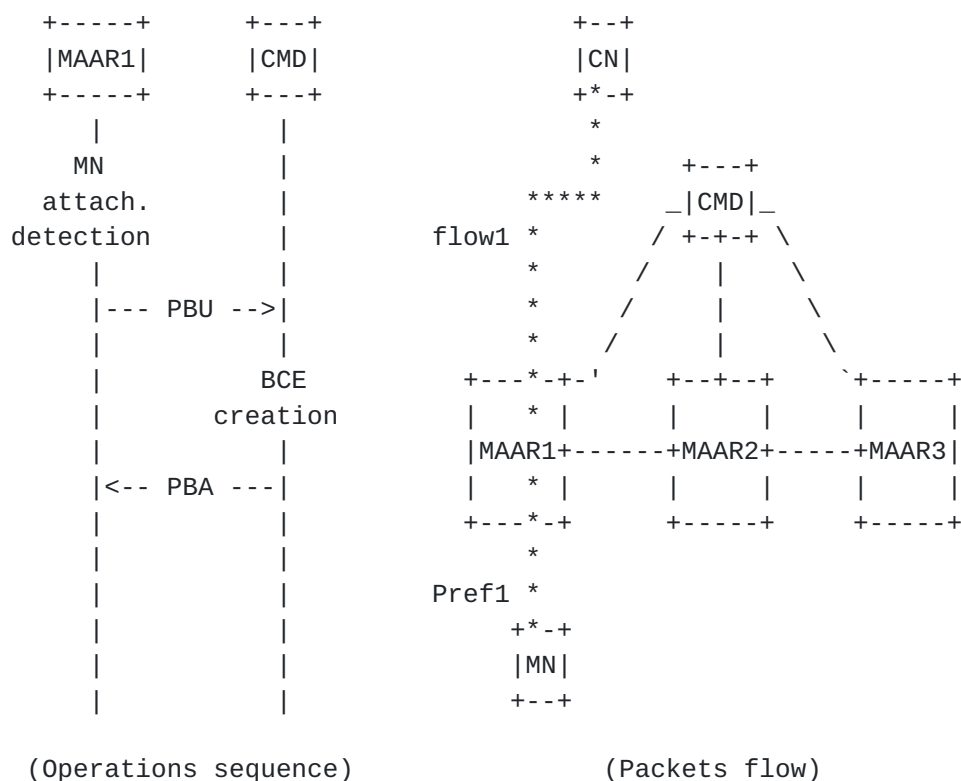
                 Figure 1: First attachment to the network

3.2. The CMD as PBU/PBA relay

When the MN moves from its current access, it associates to MAAR2
(now the S-MAAR), which delegates another IPv6 prefix (Pref2) and
sends it to the CMD for registration. The CMD has already an entry
for the MN, binding the MN-ID to its former location; thus, it

forwards the PBU to the MAAR indicated as Proxy CoA, in this case
MAAR1 (now the A-MAAR), and it updates the P-CoA field with the
S-MAAR's address. Upon PBU reception, MAAR1 replies to the CMD with
a PBA to ensure that the new location has successfully changed,
containing the prefix anchored at MAAR1. The CMD updates the BCE
adding the P-MAAR address in the list of old P-CoAs and forwards the
PBA to the new S-MAAR, containing the previous Proxy-CoA and the
prefix anchored to it, so that a tunnel can be established between
the two MAARs and new routes are set appropriately to recover the
flow(s).

Now packets destined to Pref1 are first received by MAAR1,
encapsulated into the tunnel and forwarded to MAAR2, which finally
delivers them to their destination. In uplink, when the MN transmits
packets using Pref1 for the source address, they are sent to MAAR2,
as it is MN's new default gateway, then tunneled to MAAR1 which
routes them towards the next hop to destination. Conversely, packets
carrying Pref2 are routed by MAAR2 without any special packet
handling both for uplink and downlink. The procedure is depicted in
Figure 2.

```
+-----+       +---+      +-----+            +--+            +--+
|MAAR1|       |CMD|      |MAAR2|            |CN|            |CN|
+-----+       +---+      +-----+            +*-+            +*-+
   |            |           |                *               *
   |            |          MN                *     +---+      *
   |            |        attach.       *****  _|CMD|_    *
   |            |          det.   flow1 *      / +-+-+ \   *flow2
   |            |<-- PBU ---|            *      /    |   \  *
   |          BCE           |            *     /     | *******
   |         check+         |            *    /      | *    \
   |         update         |       +---*-+-'    +--+-*+    `+-----+
   |<-- PBU ---|            |       |   * |      |    *|     |     |
  route        |            |       |MAAR1|_____|MAAR2+-----+MAAR3|
  update       |            |       |   **(_____)** *|     |     |
   |--- PBA -->|            |       +-----+      +-*--*+     +-----+
   |          BCE           |                       *   *
   |         update         |              Pref1 *   *Pref2
   |          |--- PBA -->|                    +*--*+
   |          |         route        ---move-->|*MN*|
   |          |        update                  +----+

        (Operations sequence)            (Packets flow)
```

Figure 2: Scenario after a handover, CMD as relay

For next MN's movements the process is repeated except for the number
of P-MAARs involved, that rises accordingly to the number of prefixes

that the MN wishes to maintain. Indeed, once the CMD receives the
first PBU from the new S-MAAR, it forwards copies of the PBU to all
the A-MAARs indicated in the BCE as current P-CoA (i.e., the MAAR
prior to handover) and old P-CoAs. They reply with a PBA to the CMD,
which aggregates them into a single one to notify the S-MAAR, that
finally can establish the tunnels with the A-MAARs. It should be
noted that this design separates the mobility management at the
prefix granularity, and it can be tuned in order to erase old
mobility sessions when not required, while the MN is reachable
through the latest prefix acquired. Moreover, the latency associated
to the mobility upadate is bound to the PBA sent by the furthest
A-MAAR, that takes the longest time to reach the CMD. The drawback
can be mitigated introducing a timeout at the CMD, by which, after
its expiration, all the PBAs so far collected are transmitted, and
the remaining are sent later upon their arrival.

3.3. The CMD as MAAR locator

The latency experienced in the approach shown before can be mitigated
if the A-MAARs are allowed to signal directly their information to
the new S-MAAR. This procedure reflect what was described in
Section 3.2 up to the moment the A-MAAR receives the PBU. At that
point an A-MAAR is aware of the new MN's location (i.e., S-MAAR) and,
besides sending a PBA to the CMD, it also sends a PBA to the S-MAAR
including the prefix it is anchoring. The CMD is relieved from
forwarding the PBA to the S-MAAR, as it receives a copy directly from
the A-MAAR with the necessary information to build the tunnel and set
the appropriate routes. In Figure 3 is illustrated the new messages
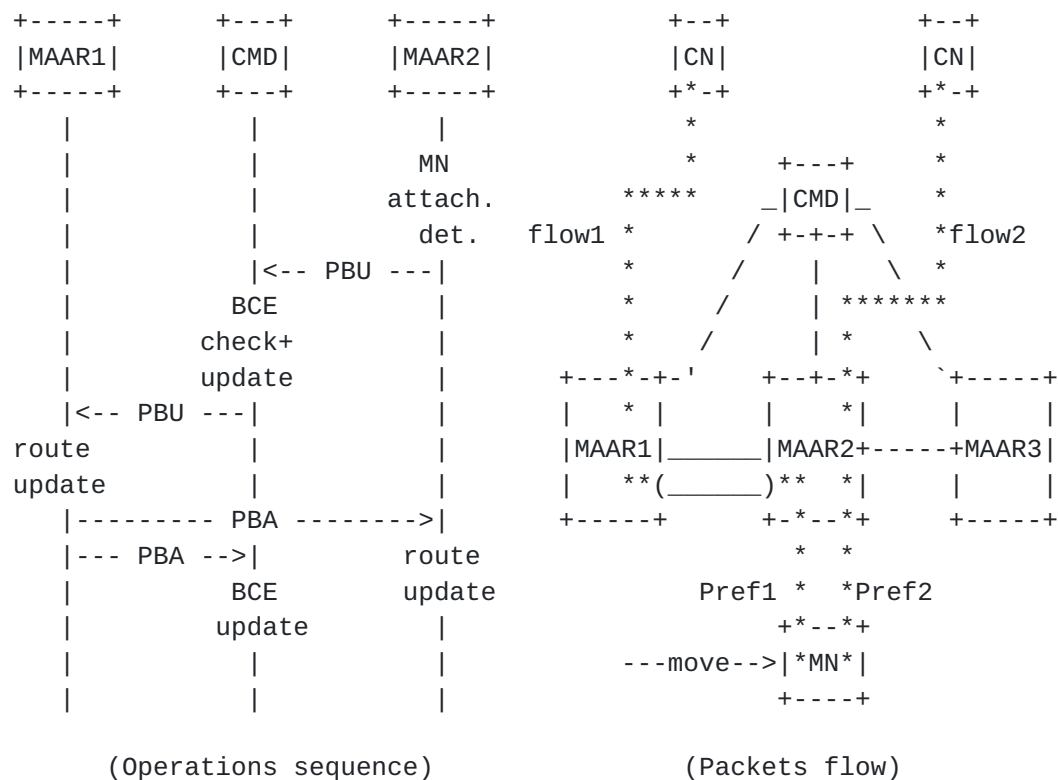sequence, while the data forwarding is unaltered.

```
    +-----+        +---+        +-----+            +--+              +--+
    |MAAR1|        |CMD|        |MAAR2|            |CN|              |CN|
    +-----+        +---+        +-----+            +*-+              +*-+
       |             |             |                *                 *
       |             |             MN               *       +---+     *
       |             |           attach.         *****     _|CMD|_    *
       |             |             det.   flow1 *        / +-+-+ \   *flow2
       |             |<-- PBU ---|               *       /   |   \  *
       |            BCE            |              *      /    | ******
       |           check+         |              *     /     | *    \
       |           update         |         +---*-+-'     +--+-*+    `+-----+
       |<-- PBU ---|              |         |   * |       |    *|     |     |
     route         |              |         |MAAR1|_____|MAAR2+-----+MAAR3|
     update        |              |         |  **(_____)**  *|     |     |
       |--------- PBA -------->|            +-----+       +-*--*+     +-----+
       |--- PBA -->|           route                       *  *
       |          BCE          update                  Pref1 *  *Pref2
       |         update          |                          +*--*+
       |           |             |               ---move-->|*MN*|
       |           |             |                          +----+

          (Operations sequence)              (Packets flow)
```

                Figure 3: Scenario after a handover, CMD as locator

3.4. The CMD as PBU/PBA proxy

   A further enhancement of previous solutions can be achieved when the
   CMD sends the PBA to the new S-MAAR before notifying the A-MAARs of
   the location change. Indeed, when the CMD receives the PBU for the
   new registration, it is already in possess of all the information
   that the new S-MAAR requires to set up the tunnel and the routes.
   Thus the PBA is sent to the S-MAAR immediately after a PBU is
   received. In parallel, a PBU is sent by the CMD to the A-MAARs to
   notify them about the new MN's location, so they receive the
   information to establish the tunnel and routes on their side. When
   A-MAARs complete the update, they send a PBA to the CMD to indicate
   that the operation is concluded and the information are updated in
   all network nodes. This scheme is depicted in Figure 4, where,
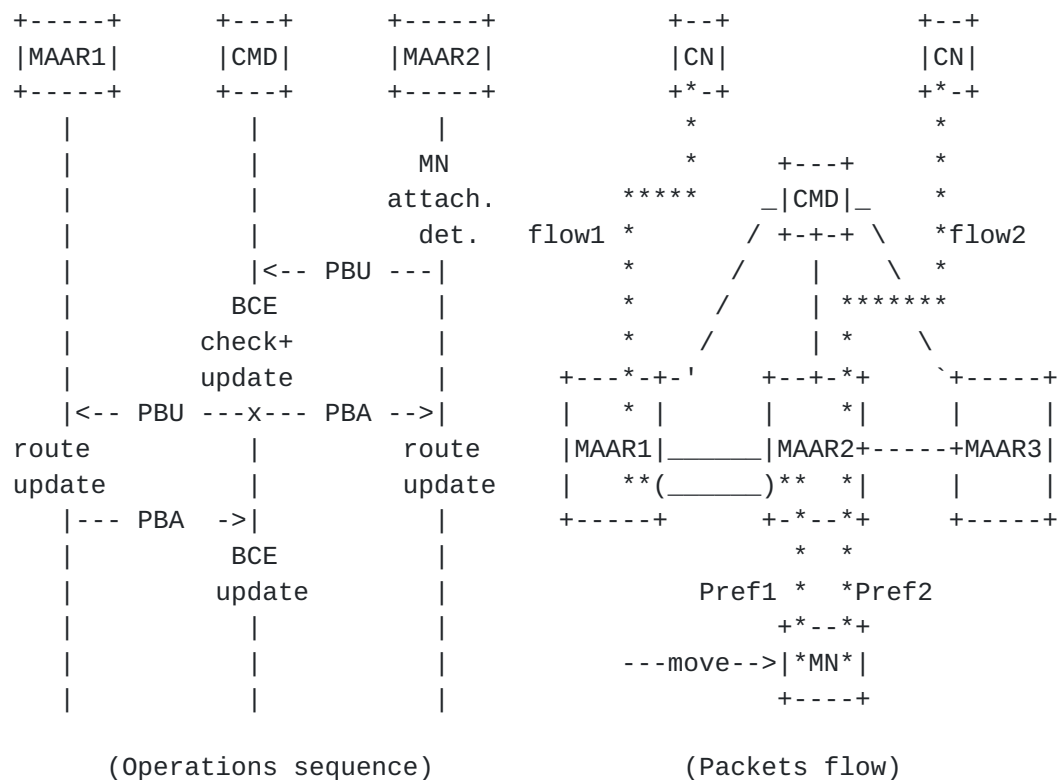   again, the data forwarding is kept untouched.

```
   +-----+       +---+      +-----+           +--+              +--+
   |MAAR1|       |CMD|      |MAAR2|           |CN|              |CN|
   +-----+       +---+      +-----+           +*-+              +*-+
      |            |           |               *                 *
      |            |          MN               *       +---+      *
      |            |         attach.    *****   _|CMD|_     *
      |            |          det.   flow1 *      / +-+-+ \   *flow2
      |          |<-- PBU ---|           *      /    |    \  *
      |           BCE         |           *     /     | ******
      |          check+       |           *    /      | *    \
      |          update       |      +---*-+-'    +--+-*+   `+-----+
   |<-- PBU ---x--- PBA -->|      |   * |     |    *|    |     |
  route          |          route |MAAR1|_____|MAAR2+-----+MAAR3|
  update         |          update |   **(_____)**  *|    |     |
   |--- PBA  ->|            |      +-----+      +-*--*+    +-----+
      |          BCE         |                     *   *
      |          update      |                  Pref1 *  *Pref2
      |           |          |                     +*--*+
      |           |          |           ---move-->|*MN*|
      |           |          |                     +----+

       (Operations sequence)              (Packets flow)
```

              Figure 4: Scenario after a handover, CMD as proxy

4. IANA Considerations

   TBD.

5. Security Considerations

   TBD.

6. Acknowledgments

   The authors would like to thank Marco Liebsch for his comments and
   discussion on this document.

7. References

7.1. Normative References

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

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

7.2. Informative References

[I-D.bernardos-mext-dmm-cmip]
            Bernardos, C. and F. Giust, "A IPv6 Distributed Client
            Mobility Management approach using existing mechanisms",
            draft-bernardos-mext-dmm-cmip-00 (work in progress),
            March 2011.

[I-D.chan-distributed-mobility-ps]
            Chan, A., "Problem statement for distributed and dynamic
            mobility management",
            draft-chan-distributed-mobility-ps-03 (work in progress),
            July 2011.

[Net-basedDMM]
            Giust, F., de la Oliva, A., Bernardos, CJ., and RP.
            Ferreira Da Costa, "FA Network-based Localized Mobility
            Solution for Distributed Mobility Management", under
            submission , 2011.

Authors' Addresses

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

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

Antonio de la Oliva
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 8803
Email: aoliva@it.uc3m.es
URI: http://www.it.uc3m.es/aoliva/

Fabio Giust
Institute IMDEA Networks and Universidad Carlos III de Madrid
Av. del Mar Mediterraneo, 22
Leganes, Madrid 28918
Spain

Phone: +34 91481 6979
Email: fabio.giust@imdea.org

Telemaco Melia
Alcatel-Lucent Bell Labs
Route de Villejust
Nozay, Ile de France 91620
France

Email: telemaco.melia@alcatel-lucent.com