Distributed Mobility Anchoring
draft-seite-dmm-dma-05.txt

Abstract

   Most existing IP mobility solutions are derived from Mobile IP
   principles where a given mobility anchor maintains Mobile Nodes (MNs)
   binding up-to-date. Data traffic is then encapsulated between the
   mobility anchor and the MN or its Access Router. These approaches
   are usually implemented on a centralised architectures where both MN
   context and traffic encapsulation need to be processed at a central
   network entity, i.e. the mobility anchor. However, one of the trend
   in mobile network evolution is to "flatten" mobility architecture by
   confining mobility support in the access network, e.g. at the access
   routers level, keeping the rest of the network unaware of the
   mobility events and their support. This document discusses the
   deployment of a Proxy Mobile IP approach in such a flat architecture.
   The solution allows to dynamically distribute mobility functions
   among access routers for an optimal routing management. The goal is
   also to dynamically adapt the mobility support of the MN's needs by
   applying traffic redirection only to MNs' flows when an IP handover
   occurs.

Table of Contents

1. Terminology

Proxy Mobile IPv6 inherited terminology

   The following terms used in this document are to be interpreted as
   defined in the Proxy Mobile IPv6 specification [RFC5213]; Mobile
   Node (MN), home Network Prefix (HNP), Mobile Node Identifier (MN-
   Identifier), Proxy Binding Update (PBU), and Proxy Binding
   Acknowledgement (PBA).

Mobility capable Access Router (MAR)

   The Mobility capable Access Router is an access router which
   provides mobility management functions. It has both mobility
   anchoring and location update functional capabilities. A Mobility
   capable Access Router can act as a Home or as a Visited Mobility
   capable Access Router (respectively H-MAR and V-MAR). Any given
   MAR could act both as H-MAR and V-MAR for a given mobile node
   having different HNPs, either allocated by this MAR (H-MAR role)
   or another MAR on which the mobile node was previously attached
   (V-MAR role).

   * H-MAR: it allocates HNP for mobile nodes. Similarly to
      [RFC5213], the H-MAR is the topological anchor point for the
      mobile node's home network prefix(es) it has allocated. The
      H-MAR acts as a regular IPv6 router for HNPs it has allocated,
      and when a mobile node has moved away and attached to a V-MAR,
      the H-MAR is responsible for: tracking the mobile node location
      (i.e. the V-MAR where the mobile node is currently attached),
      and forwarding packets to the V-MAR where the mobile node is
      attached.
   * V-MAR: it manages the mobility-related signaling for a mobile
      node, using a HNP allocated by a MAR previously visited by the
      mobile node, that is attached to its access link.

2. Introduction

   Most existing IP mobility solutions are derived from Mobile IP
   [RFC3775] principles where a given mobility agent (e.g. the Home
   Agent (HA) in Mobile IP or the Local Mobility Agent (LMA) in Proxy
   Mobile IPv6 [RFC5213]) maintains Mobile Nodes (MNs) bindings. Data
   traffic is then encapsulated between the MN or its Access Router
   (e.g. the Mobile Access Gateway (MAG) in PMIPv6) and its mobility
   agent. In other words, these approaches rely on a centralised
   architecture where both MN mobility context and traffic encapsulation
   features need to be maintained at a central network entity, the
   mobility agent. Such centralised approach provides the ability to

route MN traffic whatever its localisation is, as well as to support handovers when it moves from access router to access router; however, when millions of MNs are communicating in a given cellular network, such a centralised network entity may cause bottlenecks and single point of failure issues, which requires costly network dimensioning and engineering to be fixed. In addition, tunnelling encapsulations impact the global network efficiency since they require the maintenance of MN's specific contexts in each tunnel end nodes and they incur delays in packet processing and transport functions. Besides, centralized mobility management might not take into account current network evolution where the trend is to cache and distribute content (e.g. CDN architecture) closer to the end-user. As a consequence, alternative mobility approaches are currently being discussed and a potential solution is the distribution of mobility anchors, as stated by requirement "REQ1" in [I-D.ietf-dmm-requirements].

Moreover, it is well established that a huge amount of mobile communications are set up while the MN remains attached to the same access router. For example, the user is being communicating at home, in his office, at a cafe, etc. and the mobility support is thus not required. Applying the aforementioned centralised principles leads then to maintain user's mobility contexts, whereas the MN remains motionless. So, to avoid such a waste of ressources, mobility management should come into play only when the mobile node changes the point of attachment (i.e. performs a handover) and when it needs the conservation of the current IP address. Actually, this is the requirement "REQ2" from [I-D.ietf-dmm-requirements].

The DMM working group has been chartered to address above issues by exploring the distribution of mobility management functions and, for the sake of pragmatism, it has been agreed to firstly focus on existing mobility protocols. The goal of this document is to address this concern and, thus, has no other ambition than to discuss the use of legacy IP mobility protocols in distributed anchoring architecture. Besides, it must be noted this document aims only to meet basic [I-D.ietf-dmm-requirements] requirements, namely:

o confining the mobility support at the access routers level,
   keeping the rest of the network unaware of mobility events and
   their support (REQ1);
o dynamically adapting mobility support to each of the MN's needs by
   applying traffic redirection only to MNs' flows that are already
   established when an IP handover occurs (REQ2).

3. Basics of Distributed Mobility Management

3.1. Fundamentals

   As stated in [I-D.ietf-dmm-requirements], mobility anchoring may be
   distributed to multiple locations in the access network. For
   example, mobility anchoring (MA) function could be co-located with
   the access router (AR) as shown on Figure 1. This architecture
   allows the traffic to be anchored closer to the mobile node and, for
   example, to provide optimal mobility support to distributed content
   (e.g. CDN based delivery architecture).

                +------+ +------+ +------+ +------+
                |AR/MA | |AR/MA | |AR/MA | |AR/MA |
                +------+ +------+ +------+ +------+
                                      |

                                   | MN |

               Figure 1: Distributed Mobility Management

   Mobility management may be partially distributed, i.e. only the data
   plane is distributed, or fully distributed, i.e. both the data plane
   and control plane are distributed [I-D.yokota-dmm-scenario]. If
   conceptual differences exist, these two approaches share common
   fundamentals and it is possible to describe the generic behavior of a
   DMM deployment. Note that the following focuses only on the two
   first requirements of [I-D.ietf-dmm-requirements] (i.e. distribution
   of mobile anchoring and dynamic mobility management)

   In a standard IPv6 network without specific mobility support, any
   host is able to set up communications flows using a global IPv6
   address acquired with the support of its current access router
   [RFC4862]. When the host moves from this access router to a new one,
   its ongoing IP sessions cannot be maintained without leveraging on IP
   mobility mechanisms. However, once attached to the new access
   router, the host can again acquire a routable global IPv6 address to
   be used for any new communication flow it sets up. Hence, a flow
   based mobility support may be restricted to provide traffic
   indirection to host's flows that are already ongoing during host's
   handovers between access routers. Any new flow being set up uses the
   new host's global address acquired on the new link available after
   the handover.

When a multiple-interface host moves between access routers of
different access technologies, such a simple approach can also be
applied, considering that each network interface provides dynamically
global IPv6 addresses acquired on current access routers.

Hence, any given IP flow can be considered as implicitly anchored on
the current MN's access router when being set up. Meaning that, if
the MN moves across more than one access router and initiates IP
communications while being attached to different access routers, the
MN might be served simultaneously by more than one mobility anchor.
While the MN is attached to its initial access router, the IP flow is
delivered as for any standard IPv6 node. The anchoring function at
the access router is thus needed only to manage traffic indirection
if the MN moves to a new access router and for subsequent movements
while the IP flow remains active), maintaining the flow communication
until it ends up.

Any flow's incoming packet toward the MN is routed in a standard way
to the access router anchoring the flow as the packet contains the
destination IP address issued from router prefix. Then, if the MN is
currently attached to the initial anchor access router, the incoming
packet is directly delivered over the access link. Otherwise, the
anchoring access router needs to redirect the packet to the current
(or one of the currents) MN's access router(s).

Any flow's outgoing packet from the MN is sent over either the
initial anchor access router link or another access router link it is
currently using. In the first case, the packet can be routed in a
standard way, i.e., without requiring networks mobility support
functions. In the second case, we consider its redirection to the
initial flows' anchor router, but it may be noticed that direct
routing by the current access router may be also allowed (yet this
may lead to more stringent security and policy considerations).

3.2. Considerations on Client based mobility management

Actually, there is no issue to implement a basic DMM (as described in
previous section) with vanilla Mobile IP protocol, e.g. [RFC3775],
as long as the MN can manage simultaneously different bindings to
different Home Agents (HA), i.e. manage simultaneously more than one
tunnel to the mobile anchors. Basically, nothing prevent to
implement the HA functionalities in the access routers, so that any
given IP flow can be considered as implicitly anchored on the current
host's access router when set up. The anchoring function at the
access router is acting only to manage traffic indirection while the
host moves to a new access router. When the MN moves to a new access
router, the MN implicitly considers the previous access router as the
HA for IP addresses allocated by this access router. Then, the MN

can perform the binding update to the previous access router for IP
session initiated on it. So, MN's current traffic remains attached
to the previous access router which is responsible for forwarding the
IP flows to the MN.

Figure 2 illustrates the use of Mobile IP in a distributed
architecture. For example, let's consider an IP flow, flow#1,
initiated by the mobile node, MN, when attached to AR1. Flow#1 is
routed in a standard way as long as the MN remains attached to AR1.
If the MN moves to AR2, the MN proceeds to the binding update to AR1,
which plays the role of HA, so that flow#1 remains anchored to AR1.
The home address is the IP address obtained from AR1 and the Care-of-
Address is the IP address obtained from AR2. If MN starts a new IP
communication, flow#2, while attached to AR2; flow#2 is routed in a
standard way as long as the MN remains attached to AR2. In this
situation, applications can use either the Home Address or the Care-
of-Address and the IP stack is supposed to make the source address
selection depending on the need for mobility support; in the example
of Figure 2, the Home Address shall be used as the source address for
flow#1 and the Care-of-Addresses for flow#2. Then, if the MN moves
to another access router, flow#1 and flow#2 will be respectively
anchored to AR1/HA and AR2/HA. Mobile IP resources (mobility context
and tunneling in both ARx/HA and MN) are released after IP
communcation stopped.

```
                    +---+             +---+
                    |CN1|             |CN2|
                    +---+             +--,+
                  _.- +----------.      \
                ,'      |          `---'-.
              ,-'       |flow#1        \ `-.
            ,'          |               '   `.
           (            |    IP Network   \
            `.          |                  '    ,'
             `-.        ;                   ,\'
               \_   ;                 _.----'   '
                 - +----------''          |
                   |                        '
              +---:---+             +-------+
              | AR1`--|------------|  AR2   |
              |  HA   |------------|HA      |
              +-------+             +-------+
                         flow#1       \\ \     flow#2
                         tunnelled      \\ '
              +-----+                    +--\--+
              | MN  | ----move------->   | MN  |
```

```
            +-----+                        +-----+
```

Figure 2: Distributed Client Based Mobility

3.3. Considerations on Network based mobility management

   It is also possible to go for DMM with Proxy Mobile IPv6 [RFC5213].
   For example, mobility functions, i.e. MAG and LMA, can be co-located
   with the access routers. The anchoring behavior might be similar to
   the client based solution; however there is an issue with the binding
   update management. In a network based solution, the MN is not
   supposed to participate to mobility signaling and the MAG is expected
   to know the mobility anchor serving the MN. This problem can be
   tricky in distributed mobility architecture because 1) the MN can be
   served by more than one LMA (see fundamentals in Section 3.1) and 2)
   the mobility anchor depends on point of attachment when the IP
   communication has been initiated. There are basically two ways to
   address the issue without modifying proxy mobile IP:

   1. Involve the MN in the mobility management process: during the
       attachment process to a new access router, the MN could
       communicate its ongoing mobility sessions (i.e. list of current
       HNP with associated mobility anchors) to the MAG. For example,
       this information could be provided in a dedicated router
       solicitation option.
   2. Rely on centralized part of the control plane: when the MN
       attaches to a new access router, the MAG function retrieves the
       mobility sessions, for that MN, from a centralized database.
       This database is expected to be updated each time a new prefix is
       allocated to the MN, and also when the prefix is released.

   Even if the first option does not introduce a new piece of protocol,
   it can be seen as a violation of the basic of the network based
   mobility approach where the MN must remain agnostic of the mobility
   support. So, in the following, we will go for the second option.

4. Solution Overview for network based DMM

4.1. Distributed and Dynamic Mobility Anchoring

   The basic idea is to distribute mobility traffic management with
   dynamic user's traffic anchoring in access network nodes. The
   solution relies on a very simple flat architecture outlined in
   Figure 3 where the Mobility capable Access Router (MAR) supports both
   traffic anchoring and MN's location management functionalities. The

architecture relies on a centralized database storing ongoing
mobility sessions for the MNs (see Section 4.2 for details). This
database stores the HNPs currently allocated to the MN and their
respective anchoring point. This database could be an extension to
the policy store used in [RFC5213]; however, the detailed
specification of the interaction between MAGs and this database is
currenly out of the scope of this document.


```
                         +------------+
                         | session DB |
                        / +------------+
            +----------/--------|------\--+
            (       IP /Network |       \  )
            +--------/----------|------- \+
                /           |          \
           +-------+   +-------+    + ------+
           | MAR1  |___| MAR2  |____|  MAR3 |
           +-------+   +-------+    +-------+
                          |
                       +-----+
                       | MN1 |
                       +-----+
```
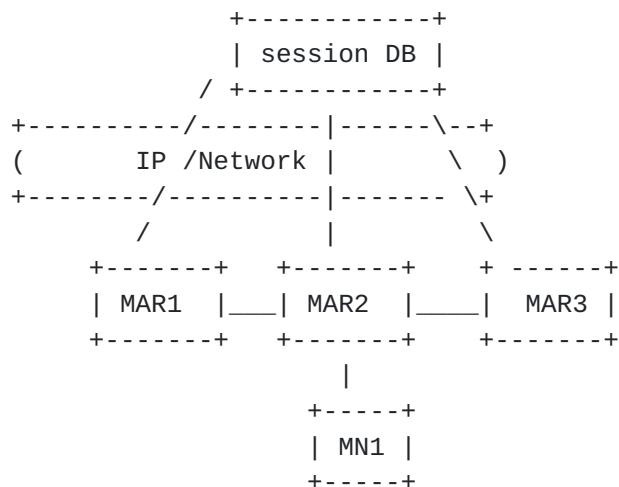

                Figure 3: Architecture for Distributed Mobility Anchoring

   Regular IPv6 routing applies when an IP communication is initiated.
   For instance, if the mobile node (e.g. MN1), being attached to MAR1,
   initiates a communication with CN1: flow#1; the traffic will be
   routed through MAR1 without requiring any specific mobility
   operation. When MN1 moves away from MAR1 and attaches to MAR2, the
   traffic remains anchored to MAR1 and is tunneled between MAR1 and
   MAR2. MAR1 becomes the mobility anchor, for IP sessions initiated by
   MN1 when it was attached to MAR1, and MAR2 plays the role of MAG for
   these sessions.

   Communications newly initiated, e.g. flow#2, while the mobile node is
   attached to MAR2 will be routed in a standard way via MAR2. But, if
   the mobile node moves away from MAR2 (e.g. attaches to MAR3), while
   maintaining both flow#1 and flow#2, two mobility anchors come into
   play: flow#1 and flow#2 will be respectively anchored in MAR1 and
   MAR2.

   Summarizing , it is proposed to dynamically locate mobility anchoring
   depending on where the flow is initially created. Accordingly,

communications are expected to be initiated without requiring
mobility anchoring and tunneling. Note that, even if a mobile node
is moving across several MARs, the tunnel endpoints are always on the
initial H-MAR and on the current V-MAR. In the case the mobile node
moves from MAR1 to MAR2 then to MAR3, a tunnel will be firstly
established between MAR1 and MAR2; then the tunnel will be moved
between MAR1 and MAR3.

However such architecture leads to new requirement on the HNP prefix
model. Actually, because the HNP is anchored to its mobility anchor
(i.e. H-MAR), a dynamic mobility anchoring requires that each MAR
must advertise different per-MN prefixes set.

```
                            _____                _____
                           |        |              |        |
                           |  CN1   |              |  CN2   |
                           |_____|              |_____|
                              '.  Flow#2               .
                      Flow#1 ' '.                      |  Flow#3
                            '  '...''''''''''''....  .
                          ..'''  '.                 '''..
                         .'  '        '.IP network    .   '.
                         :   '         '.             |    :
                          '..'        +-------+     . ..'
                           '''...  |       |   ....'''
                             '      | MAR2  | \      .
          MAR1 Forwarding Table '         |       | \     |
         +====================+ '         |       |'. \    .
           HNP-1::/64 -> MAR3    '       +-------+\'. \   |
                          +-------+              \ '+ ------+
                          |       |               \ |      |
                          | MAR1  |----------------|  MAR3 |
                          |       |'''''''''''''''|      |
                          |       |----------------|      |
                          +-------+              +-------+
                                                     ' ' |
                                          Flow#1  ' . .  Flow#3
                                                     ' ' |
                          +-----+          Flow#2 +-----+
                          | MN1 | -----move-------> | MN1 |
                          +-----+                +-----+
                                           (single interface, IF1)
```

                    Figure 4: Distributed Mobility Anchoring

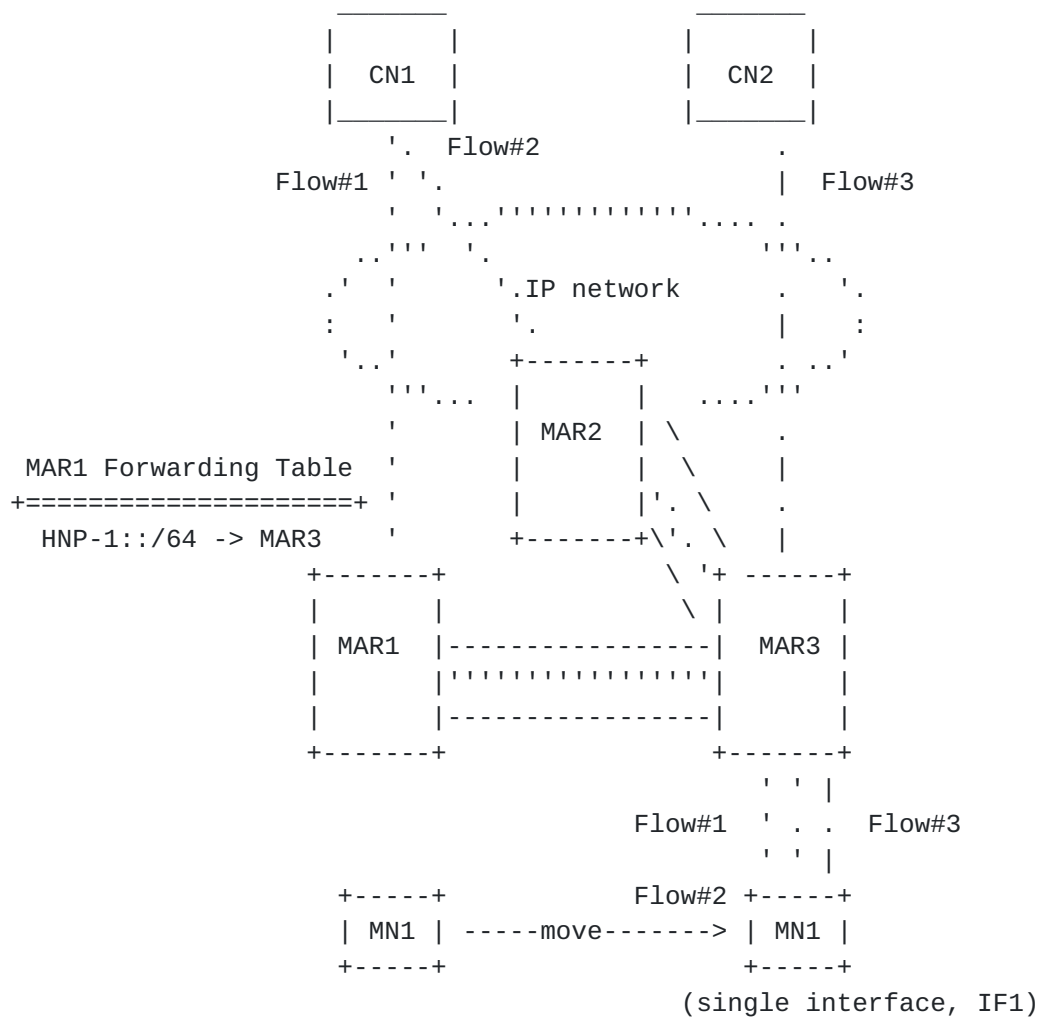## 4.2. Protocol sequence for handover management

   Handover management for a single interface mobile node is depicted on
   Figure 5 where the mobile node, MN1, is assumed to move from MAR1 to
   MAR2.


```
          MN1              MAR2             MAR1     CN1      CN2
           |                |                |        |        |
           |                |                |        |        |
        L2 Attach           |                |        |        |
           |                |                |        |        |
    (1)  |---------------RS---------->|        |        |
           |                |                | MAR1 allocates  |
           |                |                | and advertises HNP1
           |                |                | MAR1 updates MN's
           |                |                | mobility session
           |<--------------RA-----------| up to the database
           |                |                |        |        |
        comm. to CN1 using HNP1           |        |        |
    (2)  |<---------------data-flow#1--------->|        |
           |                |                |        |        |
        handover            |                |        |        |
        to MA2              |                |        |        |
    (3)  |-----RS----->| MAR2 allocates|        |        |
           |                | HNP2 for new communications     |
           |                | and retrieve the anchoring point
           |                | from the centralized database   |
           |                |----pBU------->|        |        |
           |                |                |        |        |
           |                |<---pBA--------|        |        |
    (4)  |<---RA-------|                |        |        |
           |                |                |        |        |
        handover            |                |        |        |
        completed           |                |        |        |
           |                |                |        |        |
    (5)  |<---flow#1 --|<===tunnel===>|------->|        |
           |                |                |        |        |
        comm. to CN2 using HNP2           |        |        |
           |                |                |        |        |
    (6)  |<----------------data-flow#2----------------->|
           |                |                |        |        |
           |                |                |        |        |
```
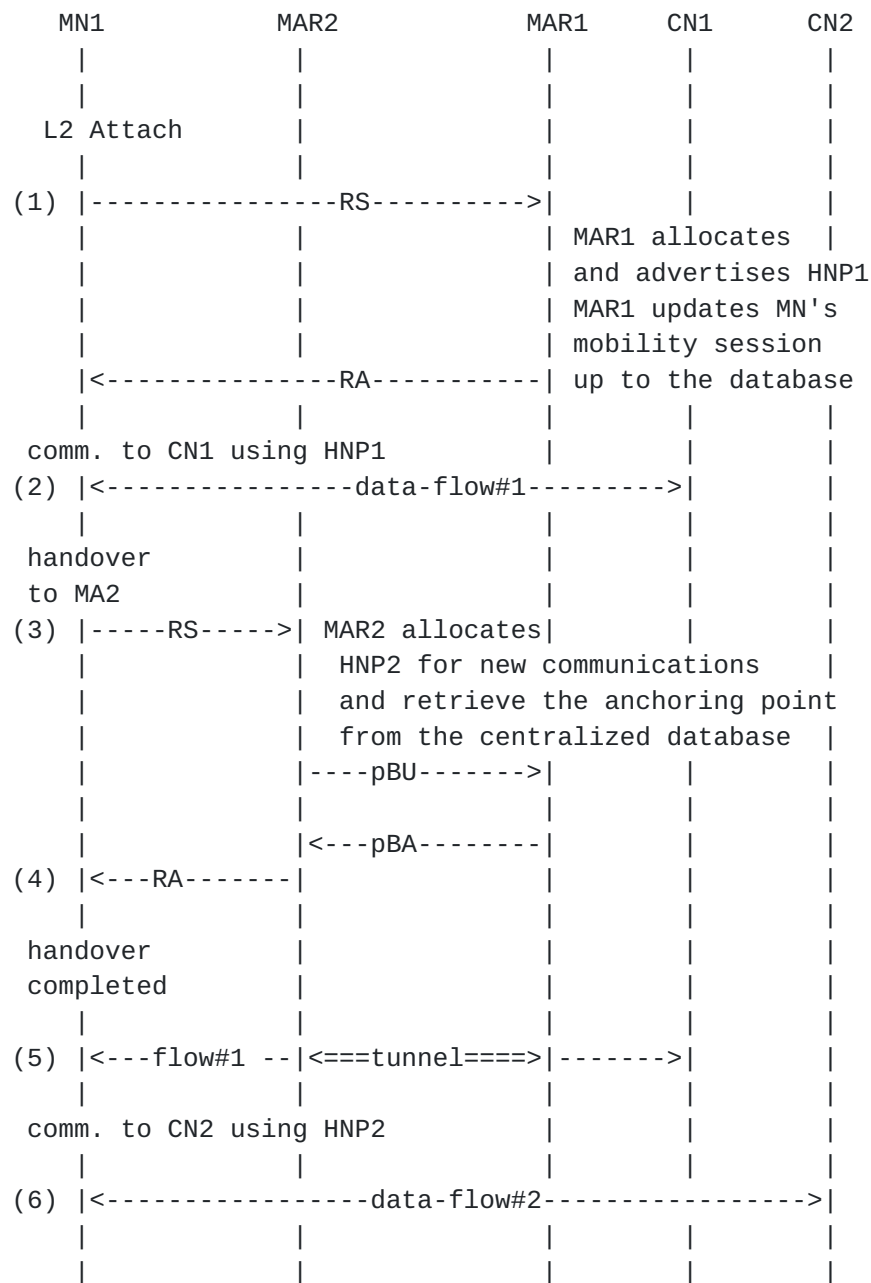
Figure 5: Handover management with Distributed Mobility Anchoring

Following are the main steps of the handover management process:

1.   The mobile node, MN1, attaches to MAR1 which is responsible for
     allocating the MN-HNP, e.g. HNP1 for MN1.
2.   Hence, the mobile node can initiate and maintain data transport
     sessions (with CN1 in the picture), using IP addresses derived
     from HNP1, in a standard way while it remains attached to MAR1,
     i.e. mobility functions do not come into play.
3.   The MN attaches to MAR2 which will thus acts as V-MAR for HNP1.
     Fistly, MAR2 retrieves the ongoing MN's mobility sessions from
     the centralized sessions database; here only one mobility session
     is ongoing: (MN::HNP1,MAR1). Then MAR2 proceeds to location
     update for HNP1 with MAR1, which plays the LMA role, i.e., PBU/
     PBA exchange between MAR2 and MAR1. MAR2 also allocates new
     prefix (HNP2) for MN1; this prefix is meant to be used by
     application flows initiated after the handover.
4.   In response to MN's router solicitation, MAR2 is expected to
     advertise both HNP1 and HNP2 to the MN, for respectively, the IP
     communications initiated when the MN was attached to MAR1 and the
     IP communications which will be initiated while attached to MAR2.
     An IP address derived from HNP1 must not be used for new IP
     communications; so, prefix HNP1 is announced as deprecated. The
     MN could also make the prefix selection relying on prefix
     properties [I-D.korhonen-dmm-prefix-properties] if supported.
5.   MAR1, playing the LMA role for HNP1, encapsulates MN1's traffic
     and tunnels it to the V-MAR, i.e. MAR2, where packets are
     decapsulated and delivered to the MN.
6.   The mobile node initiates and maintains new data transport
     sessions, e.g. with CN2, using IP addresses derived from HNP2.
     This traffic is routed in a standard way while the mobile node
     remains attached to MAR2.

4.3. Multiple Interfaces support

The distribution of mobility functions can also apply in the context
of multiple-interfaces terminals. In such a case, any given IP flow
can be considered as implicitly anchored on the current host's access
router when set up. Until the host does not move from the initial
access router (H-MAR), the IP flow is delivered as for any standard
IPv6 node. The anchoring function at the H-MAR is thus managing
traffic indirection only if one, or several, IP flow(s) are moved to
another interface, and for subsequent movements while the initial
anchored flows remain active. This anchoring is performed on a per-
flow basis and each H-MAR needs to track all possible V-MARs for a
given host on the move. The H-MAR must also manage different tunnels
for a given mobile node providing that the node is multihomed and it

simultaneously processes different IP flows on its interfaces.

Lets consider a simple example to illustrate the dynamic per-flow
mobility anchoring. Figure 6 depicts the IP flow mobility management
for a mobile node with two interfaces. The IP data flows, Flow#1 and
Flow#2, have been initiated on if1. Thus, Flow#1 and Flow#2, using
respecively prefixes HNP1 and HNP2, are anchored to MAR1. Referring
to the picture, Flow#1 has not been moved; so Flow#1 is delivered in
a standard IPv6 way. Flow#2 has been transferred from If1 to If2, so
Flow#2 packets, corresponding to HNP2, are tunneled from MAR1 to
MAR2. In other words, MAR1 and MAR2 are respectively the H-MAR
anchor and the V-MAR for flow#2.

```
                   _____               _____
                  |        |             |       |
                  |  CN1   |             |  CN2  |
                  |_____|             |_____|
                     '                       .
              Flow#1 '                       |  Flow#2
                     '    ...''''''''''''.... .
                   ..'''                     ''...
                  .'   '        IP network      .    '.
                  :    '                     |      :
                   '..'                      .  ..'
                      '''...................'|'
                       '                        .
                       '                        |
                       ' .- . - . - . - . - . - .
                      ' |
                  +-------+      Flow#2       + ------+
                  |       |     tunneled      |       |
                  | MAR1  |----------------|  MAR2 |
                  |(H-MAR)| -.-.-.-.-.-.-.|(V-MAR)|
                  |       |----------------|       |
                  +-------+                +----|--+
                      '                         .
                 Flow#1 '                       | Flow#2
                      '                         .
                      '   If1  +-----+ If2    |
                      ''''''''''| MN  | - . -   .
                                +-----+
```
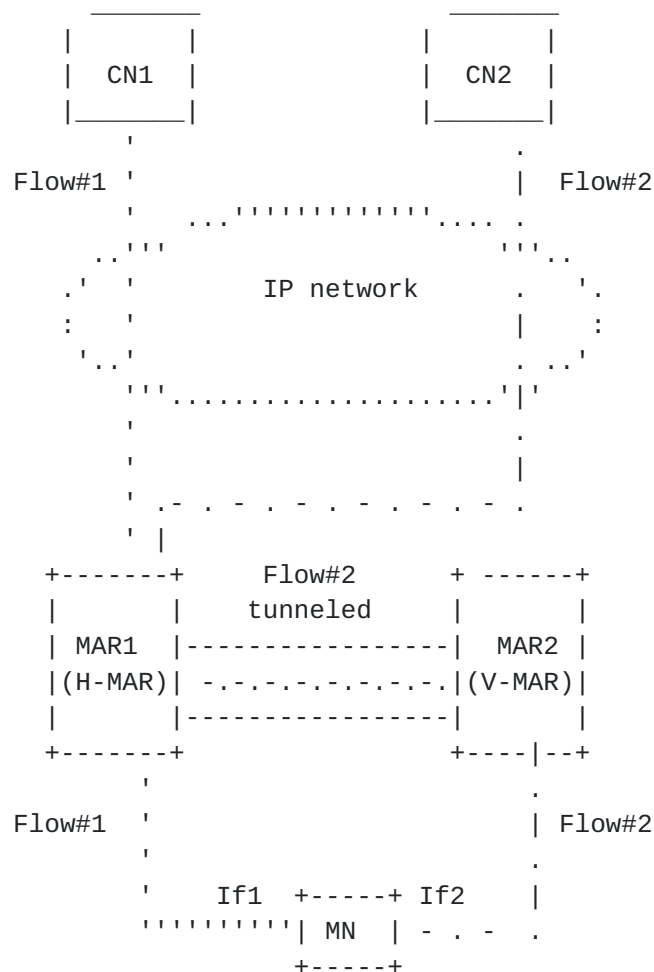
                Figure 6: Distributed IF flow Mobility Anchoring

In case of the handover of an IP flow between interfaces, the mobile

node must rely on the logical interface support, as per
[I-D.ietf-netext-logical-interface-support].

5. Security Considerations

   TBD.

6. IANA Considerations

   This document has no actions for IANA.

7. Acknowledgements

   The authors would also like to express their gratitude to Hidetoshi
   Yokota, Telemaco Melia, Dapeng Liu, Anthony Chan, Julien Laganier,
   Lucian Suciu and many others for having shared thoughts on the
   concept of distributed mobility.

   This document inherits from concepts introduced in [NTMS2008], co-
   signed by Philippe Bertin, Servane Bonjour, Jean-Marie Bonnin, Karine
   Guillouard.

8. References

8.1. Normative References

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

   [RFC4862]   Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
               Address Autoconfiguration", RFC 4862, September 2007.

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

8.2. Informative References

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

   [I-D.ietf-netext-logical-interface-support]
               Melia, T. and S. Gundavelli, "Logical Interface Support

                   for multi-mode IP Hosts",
                   draft-ietf-netext-logical-interface-support-05 (work in
                   progress), April 2012.

     [I-D.korhonen-dmm-prefix-properties]
                   Korhonen, J., Patil, B., Gundavelli, S., Seite, P., and D.
                   Liu, "IPv6 Prefix Mobility Management Properties",
                   draft-korhonen-dmm-prefix-properties-02 (work in
                   progress), July 2012.

     [I-D.yokota-dmm-scenario]
                   Yokota, H., Seite, P., Demaria, E., and Z. Cao, "Use case
                   scenarios for Distributed Mobility Management",
                   draft-yokota-dmm-scenario-00 (work in progress),
                   October 2010.

     [NTMS2008]
                   Bertin, P., "A Distributed Dynamic Mobility Management
                   Scheme designed for Flat IP Architectures.", NTMS'2008 ,
                   November 2008.

Authors' Addresses

     Pierrick Seite
     France Telecom - Orange
     4, rue du Clos Courtel, BP 91226
     Cesson-Sevigne 35512
     France

     Email: pierrick.seite@orange.com

     Philippe Bertin
     France Telecom - Orange
     4, rue du Clos Courtel, BP 91226
     Cesson-Sevigne 35512
     France

     Email: philippe.bertin@orange.com