An IPv6 Distributed Client Mobility Management approach using existing mechanisms
draft-bernardos-dmm-cmip-02

Abstract

   The use of centralized mobility management approaches -- such as
   Mobile IPv6 -- poses some difficulties to operators of current and
   future networks, due to the expected large number of mobile users and
   their exigent demands. All this has triggered the need for
   distributed mobility management alternatives, that alleviate
   operators' concerns allowing for cheaper and more efficient network
   deployments.

   This draft describes a possible way of achieving a distributed
   mobility behavior with Client Mobile IP, based on Mobile IPv6 and the
   use of Cryptographic Generated Addresses.

Requirements Language

   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].

Table of Contents

1. Introduction

    Most of the currently standardized IP mobility solutions, like Mobile
    IPv6 [RFC6275], or Proxy Mobile IPv6 [RFC5213] rely to a certain
    extent on a centralized mobility anchor entity. This centralized
    network node is in charge of both the control of the network entities
    involved in the mobility management (i.e., it is a central point for
    the control signalling), and the user data forwarding (i.e., it is
    also a central point for the user plane). This makes centralized

mobility solutions prone to several problems and limitations, as
identified in [RFC7333]: 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 that are being researched
now: one aimed at making Mobile IPv6 work in a distributed way, and
another one doing the same exercise for Proxy Mobile IPv6 (see the
document [RFC7429]). In this draft we describe a solution to achieve
a DMM behavior with a CMIP (MIPv6) solution. This document is based
on a research paper of the same authors, called "Flat Access and
Mobility Architecture: an IPv6 Distributed Client Mobility Management
solution" [GOB11].

2. Terminology

The following terms used in this document are defined in the Mobile
IPv6 specification [RFC6275]:

   Home Agent (HA)

   Home Link

   Home Address (HoA)

   Care-of Address (CoA)

   Binding Update (BU)

   Binding Acknowledgement (BA)

The following terms are defined and used in this document:

DAR (Distributed Anchor Router). First hop routers where the mobile
   nodes attach to. They also play the role of mobility managers for
   the IPv6 addresses they anchor.

HDAR (Home Distributed Anchor Router). DAR which plays the role of
   Home Agent for a particular IPv6 address (i.e., DAR where that
   IPv6 address is anchored).

3. Description of the solution

   Distributed Mobility Management approaches try to overcome the
   limitations of the traditional centralized mobility management, i.e.,
   Mobile IP, by bringing the mobility anchor closer to the MN.
   Following this idea, in our approach -- that we call Flat Access and
   Mobility Architecture (FAMA) -- the MIPv6 centralized home agent 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 access routers
   Distributed Anchor Routers (DARs).

   The diagram in Figure 1 depicts the operations of the proposed
   solution. When a mobile node attaches to a distributed anchor
   router, it gets an IPv6 address which is topologically anchored at
   the DAR (Pref1::addr1 - HoA1). In the scheme we assume the address
   configuration takes place through a Router Solicitation/Router
   Advertisement handshake. While attached to this DAR, the mobile can
   send and receive traffic using HoA1 without traversing any tunneling
   nor special packet handling.

   If the the mobile node moves to a different DAR, it gets a new IPv6
   address from the new access router (Pref2::addr2 - HoA2). In case
   the MN wants to keep the reachability of the IPv6 address(es) it
   obtained from the previous DAR (note that this decision is dynamic
   and it is out of scope of this document, it can be done on an
   application basis for example), the host has to involve its MIPv6
   stack, by sending a Binding Update to the DAR where the IPv6 address
   is anchored, using the address obtained from the current DAR as care-
   of address (in our example the MN binds HoA2 as CoA to DAR1).

```
     +-----+        +-----+        +-----+        +-----+        +-----+
     | MN  |        |DAR1 |        |DAR2 |        | CN1 |        | CN2 |
     +-----+        +-----+        +-----+        +-----+        +-----+
        |              |              |              |              |
        |-Rtr. Sol.->|               |              |              |
        |<-Rtr. Adv.-|               |              |              |
        |              |              |              |              |
  Pref1::addr1         |              |              |              |
  conf. (HoA1)         |              |              |              |
        |              |              |              |              |
        |<-----------+--IP flow 1------------->|              |
        |              |              |              |              |
  - o - o - o - o - o Handover to DAR2 o - o - o - o - o - o -
        |              |              |              |              |
        |--------Rtr. Sol.------->|               |              |
        |<-------Rtr. Adv.--------|               |              |
  Pref2::addr2         |              |              |              |
  conf. (HoA2)         |              |              |              |
        |----BU----->|               |              |              |
        | (CoA=HoA2) |               |              |              |
        |<---BA------|               |              |              |
        |(-tnl est.-)|               |              |              |
        |              |              |              |              |
        |(<---------)+--IP flow 1------------->|              |
        |              |              |              |              |
        |<-------------IP flow 2-+----------------------->|
        |              |              |              |              |
```
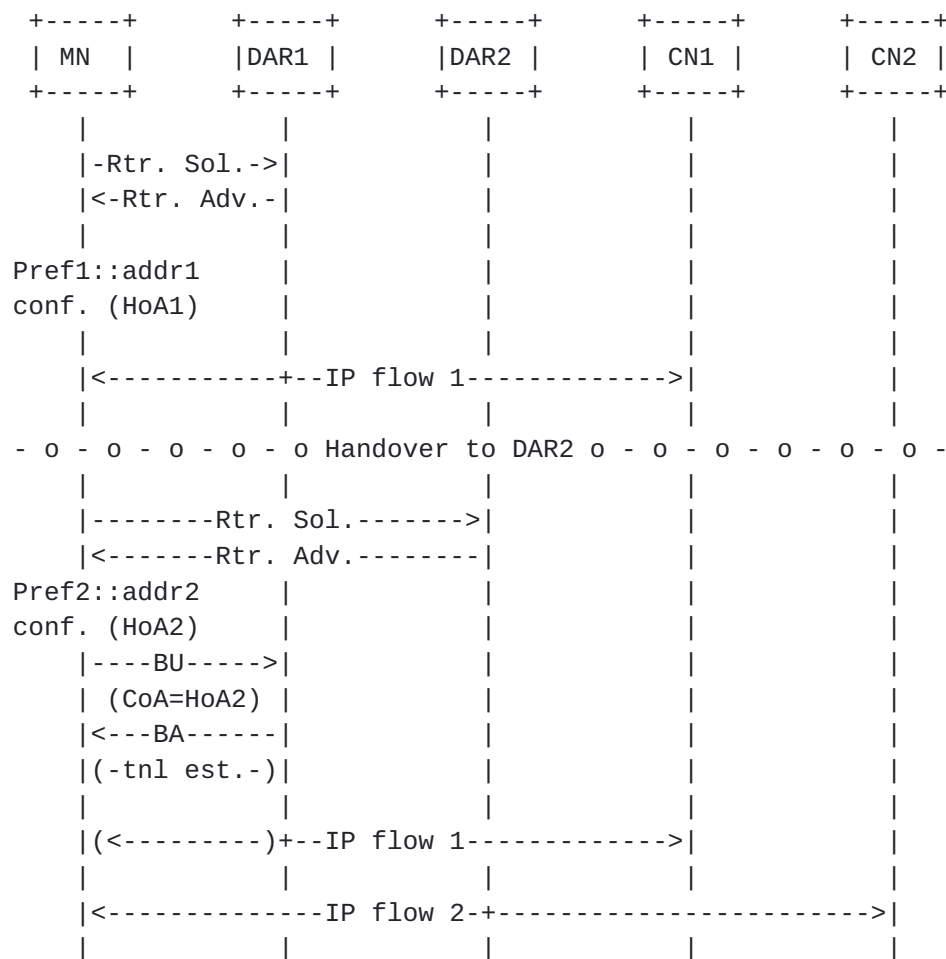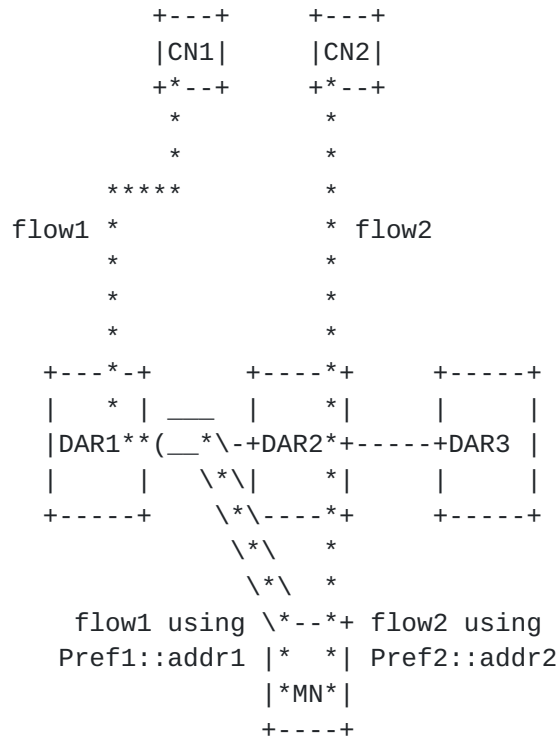
                Figure 1: Signalling after the first handover

    In this way, the IPv6 address that the node wants to maintain in use
    (Pref1::addr1) plays the role of home address (HoA1), and the DAR
    from where that address was configured plays the role of Home Agent
    (for that particular address). In this scenario, old flows are
    anchored to the previous DAR (DAR1), which is in charge to
    encaspulate the packets and deliver them to the MN's CoA. The IP
    tunnel is bi-directional, so the MN does the same when sending
    packets with the old address (HoA1). Conversely, new IP flows are
    started using the address configured at the new DAR (HoA2). These
    flows are handled by the new DAR as a plain IPv6 router.

    Note that the FAMA approach basically enables a mobile node to
    simultaneously handle several IPv6 addresses -- each of them anchored
    at a different DAR -- ensuring their continuous reachability by using
    Mobile IPv6 in a distributed fashion (i.e., each access router is a
    potential home agent for the address it delegates, if required).
    Figure 2 illustrates the above case in which the MN is connected to

DAR2, but flow1 is anchored at DAR1, because it was started by the MN
using the IPv6 address Pref1::addr1, configured when the MN was
connected to DAR1. In the same example, the MN starts flow2 using
Pref2::addr2, assigned by DAR2.

```
                    +---+      +---+
                    |CN1|      |CN2|
                    +*--+      +*--+
                     *          *
                      *          *
                  *****          *
          flow1 *              * flow2
                *              *
                *              *
                *              *
           +---*-+       +----*+      +-----+
           |   * |  ___  |   *|      |     |
           |DAR1**(__*\-+DAR2*+-----+DAR3 |
           |     |  \*\|    *|      |     |
           +-----+   \*\----*+      +-----+
                      \*\    *
                       \*\   *
             flow1 using \*--*+ flow2 using
             Pref1::addr1 |*  *| Pref2::addr2
                          |*MN*|
                          +----+


       Operations sequence                  Packets flow
```

                 Figure 2: MN's flows forwarding in FAMA

The same operations take place if the MN moves to another DAR. The
MN obtains a new address (Pref3::addr3 - HoA3), which is indicated as
CoA in the BU messages sent by the MN to the previous DARs. This
distributed address anchoring is enabled on demand and on a per-
address granularity, which means that depending on the user needs, it
might be the case that all, some or none of the IPv6 addresses that a
mobile node configures while moving within a FAMA domain, are kept
reachable and used by the mobile. The scheme in Figure 3 depicts the
example where the MN updates all the previous DARs, mapping the
corresponding HoA with the new CoA

```
   +-----+        +-----+         +-----+         +-----+         +-----+
   | MN  |        |DAR1 |         |DAR2 |         |DAR3 |         | CNs |
   +-----+        +-----+         +-----+         +-----+         +-----+
      |              |               |               |               |
     (<-----------)--IP flow 1------------------------->|
      |              |               |               |               |
      |<-------------IP flow 2-+------------------------->|
      |              |               |               |               |
 - o - o - o - o - o Handover to DAR2 o - o - o - o - o - o -
      |              |               |               |               |
      |--------------Rtr. Sol.--------------->|               |
      |<-------------Rtr. Adv.---------------|               |
 Pref3::addr3        |               |               |               |
 conf. (HoA3)        |               |               |               |
      |----BU----->|               |               |               |
      | (CoA=HoA)  |               |               |               |
      |<---BA------|               |               |               |
      |(-tnl est.-)|               |               |               |
      |------------BU---------->|               |               |
      |          (CoA=HoA3)     |               |               |
      |<-----------BA-----------|               |               |
      |(---------tnl est.------)|               |               |
      |              |               |               |               |
      |(<---------)+-IP flow 1------------------------->|
      |              |               |               |               |
      |(<-----------IP flow 2-)+------------------------->|
      |              |               |               |               |
      |<-------------IP flow 3---------------+---------->|
      |              |               |               |               |
```
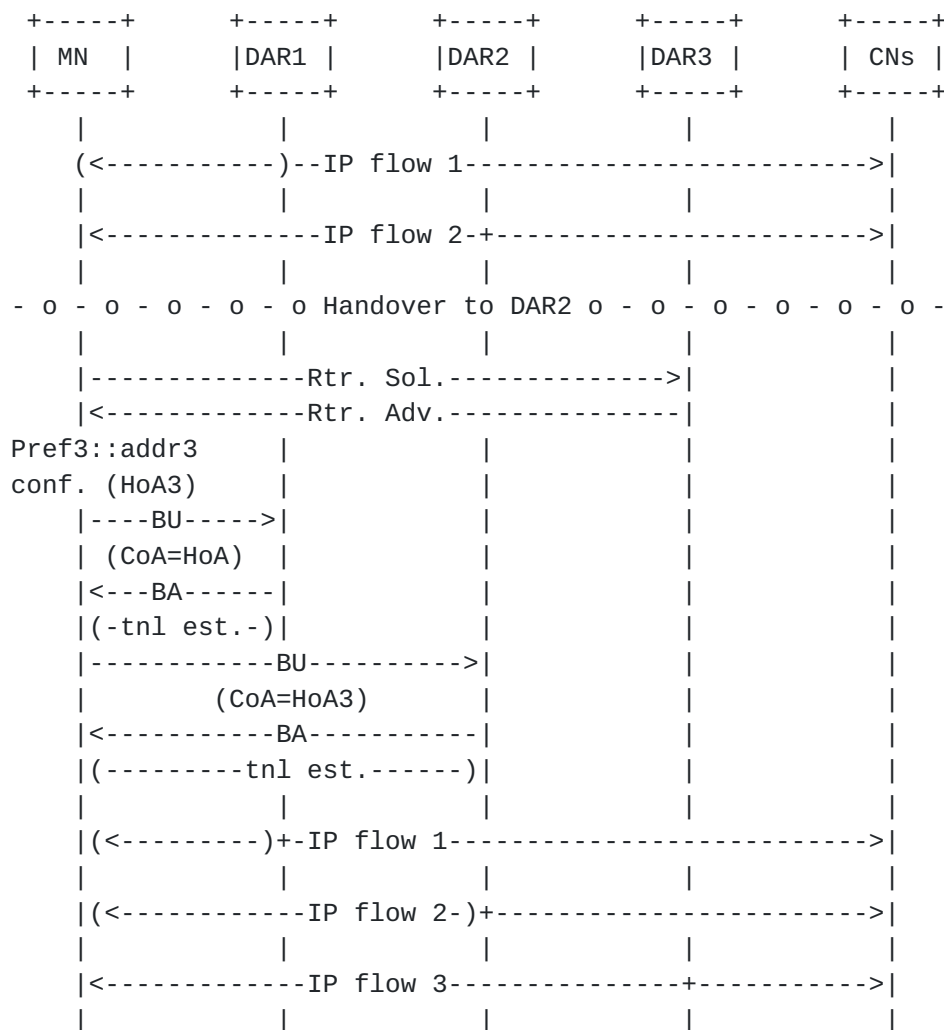
                 Figure 3: Signalling after a second handover

   In traditional Mobile IPv6, the communication between the MN and the
   HA is secured through IPsec [RFC4877]. Following a similar approach
   in FAMA is difficult due to the large number of security associations
   that would be required, since any gateway of the access network can
   play the role of home agent for any mobile node. In order to
   overcome this problem and provide authentication between the DAR and
   the MNs, we propose the use of Cryptographically Generated Addresses
   [RFC3972] (CGAs), as introduced in [I-D.laganier-mext-cga]. CGAs are
   a powerful mechanism allowing authentication of the packets and
   requires no public-key infrastructure, hence it is well-suited for
   this application.

   Following the ideas presented above, every time an MN attaches to a
   DAR, it configures a CGA from a prefix anchored at the DAR (e.g., by
   using stateless address auto-configuration mechanisms). This address

can then be used by the MN to establish a communication with a remote
Correspondent Node (CN) while attached to that particular DAR. If
the mobile then moves to a new DAR (nDAR), the following two cases
are possible: i) there is no need for the address that was configured
at the previous DAR (pDAR) to survive the movement: in this case
there is no further action required; ii) the mobile wants to keep the
reachability of the address configured at pDAR: in this case Mobile
IPv6 is triggered, and the MN sends a Binding Update (BU) message to
the pDAR, using the address configured at the previous DAR as home
address, and the address configured at the new DAR as care-of
address. This BU includes the CGA parameters and signature
[I-D.laganier-mext-cga], which are used by the receiving DAR to
identify the MN as the legitimate owner of the address. Although the
use of CGAs does not impose a heavy burden in terms of performance,
depending on the number of MNs handled at the DAR, the processing of
the CGAs can be problematic. To reduce the complexity of the
proposed protocol, we suggest an alternative mechanism to
authenticate any subsequent signaling packets exchanged between the
MN and the DAR (in case the mobile performs a new attachment to a
different DAR). This alternative method relies on the use of a
Permanent Home Keygen Token (PHKT), which will be used to generate
the Authorization option that the MN has to include in all next
Binding Update messages. This token is forwarded to the MN in the
Binding Acknowledgment message, sent on reply to the BU. The
procedure is depicted in Figure 4. Once the signaling procedure is
completed, a bi-directional tunnel is established between the mobile
node and the DAR where the IPv6 address is anchored (the "home" DAR
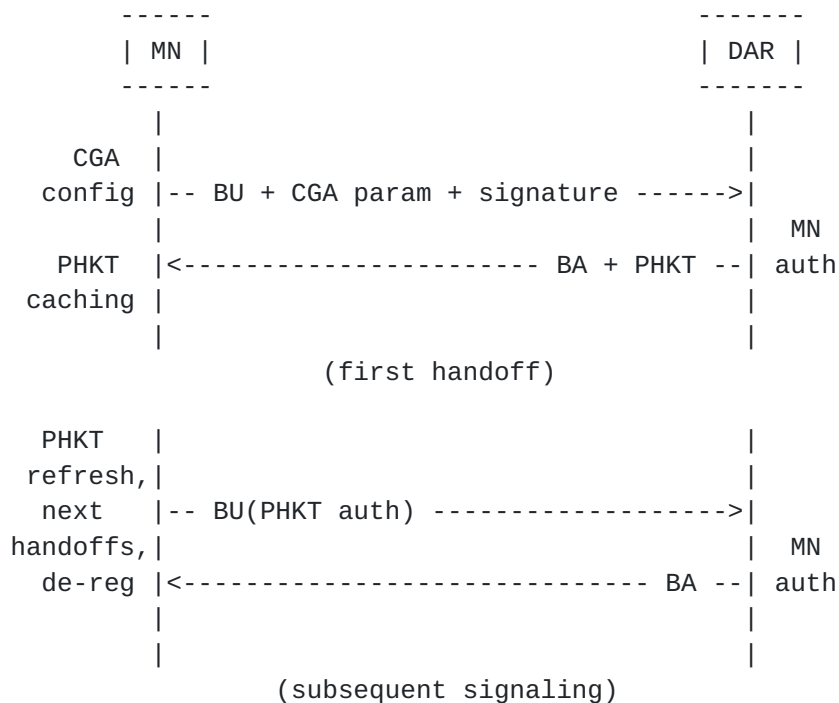-- HDAR -- for that particular address), so the mobile can continue
using the IPv6 address.

```
              ------                           -------
             | MN |                           | DAR |
              ------                           -------
                |                                |
        CGA     |                                |
      config  |-- BU + CGA param + signature ------>|
                |                                | MN
      PHKT    |<--------------------- BA + PHKT --| auth
     caching  |                                |
                |                                |
                      (first handoff)

      PHKT    |                                |
     refresh,|                                |
      next    |-- BU(PHKT auth) ------------------>|
    handoffs,|                                | MN
      de-reg |<---------------------------- BA --| auth
                |                                |
                |                                |
                    (subsequent signaling)
```

Figure 4: Signaling between the MN and the DAR

In case the MN performs any subsequent movements and it requires to
maintain the reachability of an address for which it has already sent
a BU, the following BU messages can be secured using the PHKT
exchanged before, reducing the computational load at the receiving
DAR.

Note that on every attachment of a node to a DAR, the terminal also
obtains a new IPv6 address which is topologically anchored at that
DAR, and that this address can be used for new communications
(avoiding in this way the tunneling required when using an address
anchored at a different DAR). A mobile can keep multiple IPv6
addresses active and reachable at a given time, and that requires to
send -- every time the MN moves -- a BU message to all the previous
DARs that are anchoring the IP flows that the MN wish to maintain.

4. IANA Considerations

TBD.

5. Security Considerations

Although the approach documented in this document is attractive for
the reduced signaling overhead caused by the mobility support, it can
be misused in some particular scenarios by malicious nodes that wish
to export an incorrect CoA in the BU message, since it does provide

proof of the MN's reachability at the visited network. Indeed, the
CGA approach assures that the BU message has been sent by the
legitimate HoA's owner but it does not make sure that same MN to be
reachable at the CoA indicated. This requires further analysis.

A possible approach to provide a more secure solution is the
following: a Return Routability procedure similar to the one defined
in MIPv6 Route Optimization can be used to mitigate the
aforementioned security issue. The Return Routability procedure
starts after the handoff. Instead of sending the BU message, the MN
sends a Care-of Test Init message (CoTI). This message is replied by
the DAR with a Care-Of Test message containing a CoA Keygen Token.
The MN can now send a BU using both Home and CoA Keygen tokens to
proof its reachability at both the HoA and the CoA. The message and
the knowledge of both tokens is a proof that the MN is the legitimate
node who has sent the BU and also is reachable at the CoA indicated.
As all security improvements, the one proposed incurs in a
performance penalty, in this case an increase in the handover delay.
Specifically this enhanced security approach requires four messages
to be exchanged between the MN and the DAR instead of the two
messages of the original solution. In terms of handover delay, it
increases it by a factor of two, as the new solution requires to two
Round Trip Times (RTTs) to conclude, instead of one.

6. References

6.1. Normative References

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
              Requirement Levels", BCP 14, RFC 2119, March 1997.

   [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)",
              RFC 3972, March 2005.

   [RFC4877] Devarapalli, V. and F. Dupont, "Mobile IPv6 Operation with
              IKEv2 and the Revised IPsec Architecture", RFC 4877, April
              2007.

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

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

6.2. Informative References

[GOB11]   Giust, F., de la Oliva, A., and CJ. Bernardos, "Flat
            Access and Mobility Architecture: an IPv6 Distributed
            Client Mobility Management solution", 3rd IEEE
            International Workshop on Mobility Management in the
            Networks of the Future World (MobiWorld 2011), colocated
            with IEEE INFOCOM 2011 , 2011.

[I-D.laganier-mext-cga]
            Laganier, J., "Authorizing Mobile IPv6 Binding Update with
            Cryptographically Generated Addresses", draft-laganier-
            mext-cga-01 (work in progress), October 2010.

[RFC7333]   Chan, H., Liu, D., Seite, P., Yokota, H., and J. Korhonen,
            "Requirements for Distributed Mobility Management", RFC
            7333, August 2014.

[RFC7429]   Liu, D., Zuniga, JC., Seite, P., Chan, H., and CJ.
            Bernardos, "Distributed Mobility Management: Current
            Practices and Gap Analysis", RFC 7429, January 2015.

Appendix A. Comparison with Requirement document

In this section we descrbe how our solution addresses the DMM
requirements listed in [RFC7333].

A.1. Distributed mobility management

"IP mobility, network access solutions, and forwarding solutions
provided by DMM MUST enable traffic to avoid traversing a single
mobility anchor far from the optimal route."

In our solution, a DAR is responsible to handle the mobility for
those IP flows started when the MN is attached to it. As long as the
MN remains connected to the DAR's access links, the IP packets of
such flows can benefit from the optimal path. When the MN moves to
another DAR, the path becomes non-optimal for ongoing flows, as they
are anchored to the previous DAR, but newly started IP sessions are
forwarded by the new DAR through the optimal path.

A.2. Bypassable network-layer mobility support for each application session

"DMM solutions MUST enable network-layer mobility, but it MUST be
possible for any individual active application session (flow) to not
use it. Mobility support is needed, for example, when a mobile host
moves and an application cannot cope with a change in the IP address.

Mobility support is also needed when a mobile router changes its IP
address as it moves together with a host and, in the presence of
ingress filtering, an application in the host is interrupted.
However, mobility support at the network layer is not always needed;
a mobile node can often be stationary, and mobility support can also
be provided at other layers. It is then not always necessary to
maintain a stable IP address or prefix for an active application
session."

Our DMM solution operates at the IP layer, hence upper layers are
totally transparent to the mobility operations. In particular,
ongoing IP sessions are not disrupted after a change of access
network. The routability of the old address is ensured by the IP
tunnel with the old DAR. New IP sessions are started with the new
address. From the application's perspective, those processes which
sockets are bound to a unique IP address do not suffer any impact.
For the other applications, the sockets bound to the old address are
preserved, whereas next sockets use the new address.

A.3. IPv6 deployment

"DMM solutions SHOULD target IPv6 as the primary deployment
environment and SHOULD NOT be tailored specifically to support IPv4,
particularly in situations where private IPv4 addresses and/or NATs
are used."

The DMM solution we propose targets IPv6 only.

A.4. Existing mobility protocols

"A DMM solution MUST first consider reusing and extending IETF
standard protocols before specifying new protocols."

This DMM solution is derived from the operations and messages
specified in [RFC6275], [RFC3972], and [I-D.laganier-mext-cga].

A.5. Coexistence with deployed networks/hosts and operability across
     different networks

"A DMM solution may require loose, tight, or no integration into
existing mobility protocols and host IP stacks. Regardless of the
integration level, DMM implementations MUST be able to coexist with
existing network deployments, end hosts, and routers that may or may
not implement existing mobility protocols. Furthermore, a DMM
solution SHOULD work across different networks, possibly operated as
separate administrative domains, when the needed mobility management
signaling, forwarding, and network access are allowed by the trust
relationship between them"

The proposed solution can provide a fallback mechanism employing
legacy Mobile IPv6, for instance forcing the MN to use only one DAR.
Moreover, this solution applies when the MN is connected to an
administrative domain not supporting trust relationships. Indeed,
all the IP sessions can remain anchored to the DARs of the "home"
domain. Our solution can be deployed across different domains with
trust agreements.

A.6. Operation and management considerations

"A DMM solution needs to consider configuring a device, monitoring
the current operational state of a device, and responding to events
that impact the device, possibly by modifying the configuration and
storing the data in a format that can be analyzed later.

The proposed solution can re-use existing mechanisms defined for the
operation and management of Mobile IPv6.

A.7. Security considerations

"A DMM solution MUST support any security protocols and mechanisms
needed to secure the network and to make continuous security
improvements. In addition, with security taken into consideration
early in the design, a DMM solution MUST NOT introduce new security
risks or amplify existing security risks that cannot be mitigated by
existing security protocols and mechanisms."

The proposed solution uses a CGA-based security system to enable
authentication and authorization of mobile hosts.

A.8. Multicast

"DMM SHOULD enable multicast solutions to be developed to avoid
network inefficiency in multicast traffic delivery."

This solution does not include multicast traffic in its scope.
Nevertheless, it allows combining multicast support solutions, such
as local subscription at each DAR, which would result in a flexible
distribution scenario.

Appendix B. Open DMM platform

The client-based DMM solution described in this document will be soon
available at the Open Distributed Mobility Management (ODMM) project
(http://www.odmm.net). The ODMM platform is intended to foster DMM
development and deployment, by serving as a framework to host open
source implementations.

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

   Phone: +34 91624 8859
   Email: fgiust@it.uc3m.es
   URI: http://www.it.uc3m.es/fgiust/