

DMM Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 10, 2013

CJ. Bernardos
UC3M
JC. Zuniga
InterDigital
April 8, 2013

PMIPv6-based distributed anchoring
draft-bernardos-dmm-distributed-anchoring-02

Abstract

Distributed Mobility Management solutions allow for setting up networks so that traffic is distributed in an optimal way and does not rely on centralized deployed anchors to provide IP mobility support.

There are many different approaches to address Distributed Mobility Management, as for example extending network-based mobility protocols (like Proxy Mobile IPv6), or client-based mobility protocols (as Mobile IPv6), among others. This document follows the former approach, and proposes a solution based on Proxy Mobile IPv6 in which mobility sessions are anchored at the last IP hop router (called distributed gateway). The distributed gateway is an enhanced access router which is also able to operate as local mobility anchor or mobility access gateway, on a per prefix basis. The draft focuses on the required extensions to effectively support simultaneously anchoring several flows at different distributed gateways.

This draft introduces the concept of distributed logical interface (at the distributed gateway), which is a software construct that allows to easily hide the change of anchor from the mobile node. Additionally, the draft describes how to provide session continuity in inter-domain scenarios in which dynamic tunneling or signaling between distributed gateways from different operators is not allowed.

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].

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 October 10, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	4
3. Solution's overview	5
4. Simultaneous anchoring of multiple flows (single operator)	7
4.1. The Distributed Logical Interface (DLIF) concept	8
4.2. D-GW protocol operation	11
4.3. Message format	14
4.3.1. Proxy Binding Update	14
4.3.2. Proxy Binding Acknowledgment	14
4.3.3. Anchored Prefix Option	15
4.3.4. Local Prefix Option	16
4.3.5. DLIF Link-local Address Option	17
4.3.6. DLIF Link-layer Address Option	18
5. Simultaneous anchoring of multiple flows (multiple operators)	19
6. IANA Considerations	21
7. Security Considerations	21
8. References	21
8.1. Normative References	21
8.2. Informative References	22
Appendix A. Implementation experience	22
Authors' Addresses	23

1. Introduction

The Distributed Mobility Management (DMM) paradigm aims at minimizing the impact of currently standardized mobility management solutions, which are centralized (at least to a considerable extent).

Centralized mobility solutions, such as Mobile IPv6 or the different macro-level mobility management solutions of 3GPP EPS, base their operation on the existence of a central entity (e.g., HA, LMA, PGW or GGSN) that anchors the IP address used by the mobile node and is in charge of coordinating the mobility management (MM) (sometimes helped by a third entity like the MME or the HSS). This central anchor point is in charge of tracking the location of the mobile and redirecting its traffic towards its current topological location. While this way of addressing mobility management has been fully developed by the Mobile IP protocol family and its many extensions, there are also several limitations that have been identified [I-D.chan-distributed-mobility-ps]. Among them, we can just highlight sub-optimal routing, scalability problems (in the network and in the centralized anchor) and reliability [I-D.ietf-dmm-requirements].

Several DMM-based approaches are being proposed and explored now [I-D.ietf-dmm-best-practices-gap-analysis], [commag.dmm-standards]. One of them is based on extending network-based mobility protocols (such as Proxy Mobile IPv6 [RFC5213] or GTP) to operate in distributed fashion. This document proposes a solution that falls in this category, defining a new logical entity, called Distributed Gateway (D-GW) which basically encompasses the functionalities of plain IPv6 access router, MAG and LMA, on a per-IPv6 prefix basis. The main contribution of this draft is the definition of the mechanisms required to support the operation of such a network-based mobility solution when several flows are simultaneously anchored at different D-GWs, by introducing the concept of Distributed Logical Interface (DLIF). The document also defines the required PMIPv6 signaling extensions. Last, but not least, the solution is also extended to provide session continuity across different domains.

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 Acknowledgment (PBA)

The following terms are defined and/or used in this document:

D-GW (Distributed Gateway). First IP hop router used by the mobile node. It provides an IPv6 prefix (topologically anchored at the D-GW) to each attaching mobile node.

Anchoring D-GW. A previously visited D-GW anchoring an IPv6 prefix which is still used by a mobile node.

Serving D-GW. The D-GW the MN is currently attached to.

DLIF (Distributed Logical Interface). It is a logical interface at the IP stack of the D-GW. For each active prefix used by the mobile node, the serving D-GW has a DLIF configured (associated to the anchoring D-GW). In this way, a serving D-GW exposes itself towards each MN as multiple routers, one per active anchoring D-GW.

HSS (Home Subscriber Server). In a 3GPP architecture, it is the master user database that contains the subscription-related information (subscriber profiles), performs authentication and authorization of the user, and can provide information about the subscriber's location and IP information.

3. Solution's overview

A new logical network entity, called Distributed Gateway (D-GW) is introduced at the edge of the network, close to the MN. It implements the functionality of a plain IPv6 access router (AR), a mobile access gateway (MAG) and a local mobility anchor (LMA), on a per-MN and per-IPv6-prefix, as described later.

The solution basically extends Proxy Mobile IPv6 [RFC5213] to behave in a distributed fashion, similarly as what has been proposed in [I-D.seite-dmm-dma] and [I-D.bernardos-dmm-pmip]. This is achieved by the D-GW logically behaving as a distributed mobility anchor, which comprises the following:

- o When a mobile node attaches to a D-GW (initial attachment or handover), the D-GW provides an IPv6 prefix to the MN, acting as a regular IPv6 router (with the only difference that the delegated prefix is only assigned to one single MN, not being shared with any other node). The D-GW that the mobile node is currently attached to is called "serving D-GW".
- o When a mobile node performs a handover, it attaches to a new D-GW and configures a new IPv6 address out of the prefix provided and anchored by the new serving D-GW. As before, the serving D-GW behaves as a plain IPv6 router for that particular MN and the delegated (locally anchored) prefix. If the MN has active traffic using addresses anchored by other D-GWs (which are called "anchoring D-GWs") or it just needs to keep the reachability of these addresses, the current serving D-GW also acts as MAG, by sending the required proxy binding update (PBU) to the corresponding anchoring D-GWs. The anchoring D-GWs therefore behave as LMA for this particular MN and the IPv6 prefixes they are anchoring, replying with a PBA.
- o Once the PBU/PBA signaling is completed, a bidirectional tunnel is established between the serving D-GW and the anchoring D-GW (one per D-GW anchoring an active prefix used by the MN). These tunnels are used to provide IP address continuity to prefixes that are not anchored at the serving D-GW.
- o The means for a serving D-GW to obtain the information about the prefixes that a locally attached mobile node wants to keep reachable, and the associated anchoring D-GWs are out of the scope of this draft. Among the possible mechanisms that can be used to let the D-GW know about the prefixes that should be kept reachable, we can cite for instance layer-2 triggers/signaling. Regarding the mapping of IPv6 prefixes to anchoring D-GWs, there might be either fully distributed mechanisms in place, or the information can be maintained in a centralized repository (e.g., in the HSS, using a centralized LMA [I-D.bernardos-dmm-pmip], etc.).

The basic operation of the solution is shown with an example in Figure 1. MN1 attaches to D-GW1 (thus becoming its serving D-GW) and configures an IPv6 address (prefA::MN1) out of a prefix locally anchored at D-GW1 (prefA::/64). At this point, MN1 can communicate with any correspondent node of the Internet, being the traffic anchored at D-GW1. If later on MN1 moves to D-GW2, a new IPv6 address (PrefB::MN1) is configured by the mobile node, this time out of prefB::/64, which is anchored at D-GW2 (which becomes the new serving D-GW). D-GW2 also exchanges the required PBU/PBA signaling to ensure that data traffic using prefA::MN1 still reaches the mobile

node, by setting up a bidirectional tunnel between D-GW1 (anchoring D-GW) and D-GW2 (serving D-GW).




Figure 1: Basic operation of the solution

The next sections of this draft focus on the detailed operation of the D-GWs when a mobile node has multiple flows anchored at different distributed gateways.

4. Simultaneous anchoring of multiple flows (single operator)

In this section we describe the mechanisms required in the network to enable simultaneous anchoring of several flows at different D-GWs within the same operator.

4.1. The Distributed Logical Interface (DLIF) concept

One of the main challenges of a network-based DMM solution is how to allow a mobile node to simultaneously send/receive traffic which is anchored at different D-GWs, and how to influence on the preference of the mobile selecting the source IPv6 address for a new communication, without requiring special support on the mobile node stack. This document defines the Distributed Logical Interface (DLIF), which is a software construct that allows to easily hide the change of anchor from the mobile node.




Figure 2: DLIF: exposing multiple routers (one per active anchoring D-GW)

The basic idea of the DLIF concept is the following. Each serving D-GW exposes itself towards a given MN as multiple routers, one per active anchoring D-GW associated to the MN. Let's consider the example shown in Figure 2, MN1 initially attaches to D-GW1, configuring an IPv6 address (prefA::MN1) from a prefix locally anchored at D-GW1 (prefA::/64). At this stage, D-GW1 plays both the role of anchoring and serving D-GW, and also it behaves as a plain IPv6 access router. D-GW1 creates a distributed logical interface to communicate (point-to-point link) with MN1, exposing itself as a (logical) router with a specific MAC (e.g., 00:11:22:33:01:01) and IPv6 addresses (e.g., prefA::DGW1/64 and fe80:211:22ff:fe33:101/64)

Figure 3 shows the logical interface concept in more detail. The figure shows two D-GWs and three MNs. D-GW1 is currently serving MN2 and MN3, while D-GW2 is serving MN1. MN1, MN2 and MN3 have two active anchoring D-GWs: D-GW1 and D-GW2. Note that a serving D-GW always plays the role of anchoring D-GW for the attached (served) MNs. Each D-GW has one single physical wireless interface.

As introduced before, each MN always "sees" multiple logical routers -- one per active anchoring D-GW -- independently of to which serving D-GW the MN is currently attached. From the point of view of the MN, these D-GWs are portrayed as different routers, although the MN is physically attached to one single interface. The way this is achieved is by the serving D-GW configuring different logical interfaces. If we focus on MN1, it is currently attached to D-GW2 (i.e., D-GW2 is its serving D-GW) and, therefore, it has configured an IPv6 address from D-GW2's pool (e.g., prefB::/64). D-GW2 has set-up a logical interface (mnlDgw2) on top of its wireless physical interface (phy if D-GW2) which is used to serve MN1. This interface has a logical MAC address (LMAC6), different from the hardware MAC address (HMAC2) of the physical interface of D-GW2. Over the mnlDgw2 interface, D-GW2 advertises its locally anchored prefix prefB::/64. Before attaching to D-GW2, MN1 visited D-GW1, configuring also an address locally anchored at this D-GW, which is still being used by the MN1 in active communications. MN1 keeps "seeing" an interface connecting to D-GW1, as if it were directly connected to the two D-GWs. This is achieved by the serving D-GW (D-GW1) configuring an additional distributed logical interface: mnlDgw1, which behaves exactly as the logical interface configured by the actual D-GW1 when MN1 was attached to it. This means that both the MAC and IPv6 addresses configured on this logical interface remain the same regardless of the physical D-GW which is serving the MN. The information required by a serving D-GW to properly configure this logical interfaces can be obtained in different ways: as part of the information conveyed in the PBA, from an external database (e.g., the HSS) or by other means. As shown in the figure, each D-GW may have several logical interfaces associated to each attached MN, having always at least one (since a serving D-GW is also an anchoring D-GW for the attached MN).

In order to enforce the use of the prefix locally anchored at the serving D-GW, the router advertisements sent over those logical interfaces playing the role of anchoring D-GWs (different from the serving one) include a zero prefix lifetime. The goal is to deprecate the prefixes delegated by these D-GWs (which will be no longer serving the MN). Note that on-going communications keep on using those addresses, even if they are deprecated, so this only affects to new sessions.

The distributed logical interface concept also enables the following use case. Suppose that access to a local IP network is provided by a given D-GW (e.g., D-GW1 in the example shown in Figure 2) and that the resources available at that network cannot be reached from outside the local network (e.g., cannot be accessed by an MN attached to D-GW2). This is similar to the LIPA scenario currently being considered by 3GPP. The goal is to allow an MN to be able to roam while still being able to have connectivity to this local IP network. The solution adopted to support this case makes use of RFC 4191 [RFC4191] more specific routes when the MN moves to a D-GW different from the one providing access to the local IP network (D-GW1 in the example). These routes are advertised through the distributed logical interface representing the D-GW providing access to the local network (D-GW1 in this example). In this way, if MN1 moves from D-GW1 to D-GW2, any active session that MN1 may have with a node of the local network connected to D-GW1 will survive, being the traffic forwarded via the tunnel between D-GW1 and D-GW2. Also, any potential future connection attempt towards the local network will be supported, even though MN1 is no longer attached to D-GW1.

4.2. D-GW protocol operation

This section describes the D-GW operation in more detail.

Figure 4 shows an example of the D-GW operation:

1. MN1 attaches to D-GW1. This event is detected by D-GW1 (based on layer 2 signaling/triggers or the reception of a Router Solicitation sent by MN1).
2. An IPv6 prefix from the pool of locally anchored prefixes is selected by D-GW1 to be delegated to MN1 (prefA::/64). D-GW1 sets up a distributed logical interface aimed at interfacing with MN1, called mnldgw1. D-GW1 starts sending router advertisements to MN1, including the delegated prefix.
3. D-GW1 learns if it is an attachment due to a handover (how this is done is out-of-scope of this draft). In this case it is an initial attachment, so nothing else is required.
4. The DLIF mnldgw1 is used by D-GW1 to advertise the locally anchored prefix (prefA::/64) to MN1. Using this prefix, MN1 configures an IPv6 address (prefA::MN1/64) that can be used to start new sessions (which will be anchored at D-GW1). Traffic using the address prefA::MN1 is received at the interface mnldgw1 and directly forwarded by D-GW1 towards its destination. Traffic between MN1 and the local network reachable via D-GW1 (localnet) is handled normally by D-GW1 (as MN1 is locally attached).

5. MN1 performs a handover to D-GW2. This event is detected by D-GW2.
6. An IPv6 prefix from the pool of locally anchored prefixes is selected by D-GW2 to be delegated to MN1 (prefB::/64). D-GW2 sets up a distributed logical interface aimed at interfacing with MN1, called mnlDgw2. D-GW2 starts sending router advertisements to MN1, including the delegated prefix. Traffic using the address prefB::MN1 is received at the interface mnlDgw2 and directly forwarded by D-GW2 towards its destination.
7. D-GW2 learns that this is a handover of MN1, and that it previously visited D-GW1. D-GW2 sends a PBU to D-GW1, which replies with a PBA. This PBA MAY include information about the prefix(es) anchored at D-GW1, the parameters needed by D-GW2 to set-up the DLIF mnlDgw1, and the prefixes of local networks reachable via D-GW (if any). Alternatively, this information MAY be obtained using a different approach (such as storing it in the HSS or some other external repository). A bi-directional tunnel between D-GW1 and D-GW2 is set-up, as well as the required routing entries.
8. D-GW2 sets up the DLIF mnlDgw1, aimed at "logically" resembling D-GW1, so MN1 does not detect any change at layer-3. D-GW2 starts sending router advertisements to MN1 through mnlDgw2, which include the prefix anchored at D-GW1 (prefA::/64) with zero lifetime to deprecate the prefix (or alternatively it MAY include a low Default Router Preference [RFC4191] if communication to this D-GW is still needed in the future). In this way, prefA::MN1 is not preferred for new communications. The RAs MAY also include a Route Information Option (RIO) [RFC4191] with the prefix of localnet, which is the network that is only locally reachable via D-GW1 (e.g., as in the LIPA scenarios considered by the 3GPP), so MN1 picks D-GW1 (the "logical" version of it portrayed by D-GW2) when sending traffic to that network, including the delegated prefix. Traffic using the address prefA::MN1 is received at the interface mnlDgw1 and forwarded via the tunnel with D-GW1, which then forwards it towards its destination. Traffic between MN1 and the network locally reachable via D-GW1 (localnet) is also handled via mnlDgw1 and sent through the tunnel.





Figure 4: D-GW protocol operation



Distributed Gateway Flag (D)

The Distributed Gateway Flag is set to indicate that the sender of the message supports operating as a distributed gateway.


Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The encoding and format of defined options are described in Section 6.2 of [RFC6275]. The distributed gateway MUST ignore and skip any options that it does not understand.

4.3.3. Anchored Prefix Option

A new Anchored Prefix option is defined for use with the Proxy Binding Update and Proxy Binding Acknowledgment messages exchanged between distributed gateways. This option is used for exchanging the mobile node's prefix anchored at the anchoring D-GW. There can be multiple Anchored Prefix options present in the message.

The Anchored Prefix Option has an alignment requirement of 8n+4. Its format is as follows:



Type

To be assigned by IANA.

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the type and length fields. This field MUST be set to 18.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

Prefix Length

8-bit unsigned integer indicating the prefix length of the IPv6 prefix contained in the option.

Anchored Prefix


A sixteen-byte field containing the mobile node's IPv6 Anchored Prefix.

4.3.4. Local Prefix Option

A new Local Prefix option is defined for use with the Proxy Binding Update and Proxy Binding Acknowledgment messages exchanged between distributed gateways. This option is used for exchanging a prefix of a local network that is only reachable via the anchoring D-GW. There can be multiple Local Prefix options present in the message.

The Local Prefix Option has an alignment requirement of 8n+4. Its

format is as follows:



Type

To be assigned by IANA.

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the type and length fields. This field MUST be set to 18.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

Prefix Length


8-bit unsigned integer indicating the prefix length of the IPv6 prefix contained in the option.

Local Prefix

A sixteen-byte field containing the IPv6 Local Prefix.

4.3.5. DLIF Link-local Address Option

A new DLIF Link-local Address option is defined for use with the Proxy Binding Update and Proxy Binding Acknowledgment messages exchanged between distributed gateways. This option is used for exchanging the link-local address of the DLIF to be configured on the serving D-GW so it resembles the DLIF configured on the anchoring D-GW.



Type

To be assigned by IANA.

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the type and length fields.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

DLIF Link-layer Address

A variable length field containing the link-layer address of the logical interface to be configured on the serving distributed gateway.

The content and format of this field (including byte and bit ordering) is as specified in Section 4.6 of [RFC4861] for carrying link-layer addresses. On certain access links, where the link-layer address is not used or cannot be determined, this option cannot be used.

5. Simultaneous anchoring of multiple flows (multiple operators)

An MN may roam between D-GWs that do not belong to the same operator, and therefore might end up having multiple simultaneous flows, anchored at different operators. Since dynamically setting up tunnels between different operators (i.e., between D-GWs belonging to different operators) is usually not supported, a solution should be devised to ensure session continuity in this scenario, even if it is at the cost of a sub-optimal routing.

establishes one tunnel with the centralized LMA to send/receive traffic using `prefA::/64`. From the point of view of D-GW2, the operation is just as if the LMA was the D-GW anchoring this prefix. Analogously, the LMA establishes one tunnel with D-GW1 (from the point of view of D-GW1, the LMA is the current serving D-GW of MN1). Regarding the signaling, it is similar to the intra-operator scenario, though in this case the PBU/PBA sequence is performed twice, once between D-GW2 and the LMA, and another one between the LMA and D-GW1 (i.e., because two different tunnels are created).

6. IANA Considerations

This document defines new mobility options that require IANA actions.

7. Security Considerations

The protocol extensions defined in this document share the same security concerns of Proxy Mobile IPv6 [RFC5213]. It is recommended that the signaling messages, Proxy Binding Update and Proxy Binding Acknowledgment, exchanged between the distributed gateways, or between a distributed gateway and a centralized local mobility anchor, are protected using IPsec using the established security association between them. This essentially eliminates the threats related to the impersonation of a distributed gateway or the local mobility anchor.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 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.

8.2. Informative References

[I-D.bernardos-dmm-pmip]

Bernardos, C., Oliva, A., Giust, F., Melia, T., and R. Costa, "A PMIPv6-based solution for Distributed Mobility Management", draft-bernardos-dmm-pmip-01 (work in progress), March 2012.

[I-D.chan-distributed-mobility-ps]

Chan, A., "Problem statement for distributed and dynamic mobility management", draft-chan-distributed-mobility-ps-05 (work in progress), October 2011.

[I-D.ietf-dmm-best-practices-gap-analysis]

Liu, D., Zuniga, J., Seite, P., Chan, A., and C. Bernardos, "Distributed Mobility Management: Current practices and gap analysis", draft-ietf-dmm-best-practices-gap-analysis-00 (work in progress), February 2013.

[I-D.ietf-dmm-requirements]

Chan, A., Liu, D., Seite, P., Yokota, H., and J. Korhonen, "Requirements for Distributed Mobility Management", draft-ietf-dmm-requirements-03 (work in progress), December 2012.

[I-D.seite-dmm-dma]

Seite, P., Bertin, P., and J. Lee, "Distributed Mobility Anchoring", draft-seite-dmm-dma-06 (work in progress), January 2013.

[commag.dmm-standards]

Zuniga, JC., Bernardos, CJ., de la Oliva, A., Melia, T., Costa, R., and A. Reznik, "Distributed mobility management: a standards landscape", IEEE Communications Magazine Vol. 51, No. 3, March 2013.

Appendix A. Implementation experience

The DLIF concept can be easily implemented using features that are usually available on several OSs. Among the possible mechanisms that can be used to do it, the Linux macvlan support allows the creation of different logical interfaces over the same physical one. Each logical interface appears as a regular interface to the Linux OS (which can be configured normally), and it supports configuring the MAC address exposed by the logical interface. The destination MAC

address is used by the OS to decide which logical interface (configured on top of a physical interface) is in charge of processing an incoming L2 frame.

The EU FP7 MEDIEVAL project implemented a prototype of the DLIF concept using the Linux macvlan support, the radvd daemon, the Linux Advanced Routing and Traffic Control features and the standard iproute2 collection of utilities:

- o The macvlan support enables iproute2 tools to be able to create, destroy and configure DLIFs on demand over a single physical interface. One of the important features that needs to be configured is the logical MAC address exposed by the DLIF, as well as the IPv6 addresses, as they should remain the same regardless of the serving D-GW where the DLIF is configured.
- o Since the distributed logical interfaces created using the macvlan support appear as regular network interfaces, they can be used normally in the radvd configuration file. Then, by dynamically modifying the radvd configuration file and reloading it, we can control the router advertisements sent to each MN (e.g., advertizing new IPv6 prefixes, deprecating prefixes anchored at other serving D-GWs, announcing RFC 4191 specific routes or changing router preferences).
- o Each time a DLIF is created, it is also needed to properly configure source-based IPv6 routes, as well as tunnels (in case of handover). This is supported by the Linux Advanced Routing and Traffic Control features.
- o Last, but not least, current Linux kernels support the configuration of RFC 4191 specific routes (by processing Route Information Options contained in RAs). The kernel support can be easily enabled by using the `net.conf.ipv6.*.accept_ra_rt_info_max_plen` kernel configuration parameter.
- o A prototype implementation of the solution described in this document, together with the solution specified in [I-D.bernardos-dmm-pmip] was shown during IETF 83rd, in Paris in March 2012.

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/>

Juan Carlos Zuniga
InterDigital Communications, LLC
1000 Sherbrooke Street West, 10th floor
Montreal, Quebec H3A 3G4
Canada

Email: JuanCarlos.Zuniga@InterDigital.com
URI: <http://www.InterDigital.com/>

