Network Working Group	P. Seite	
Internet-Draft	P. Bertin	
Intended status: Informational	France Telecom - Orange	
Expires: November 20, 2010	May 19, 2010	

TOC

Dynamic Mobility Anchoring
draft-seite-netext-dma-00.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. 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.

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 November 20, 2010.

Copyright Notice

Table of Contents

---

1. Terminology

TOC

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] (Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” August 2008.); 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] (Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” August 2008.), 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

TOC

Most existing IP mobility solutions are derived from Mobile IP [RFC3775] (Johnson, D., Perkins, C., and J. Arkko, “Mobility Support in IPv6,” June 2004.) 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] (Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” August 2008.)) 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. These approaches lead to the implementation of centralised architectures where both MN context and traffic encapsulation need to be maintained at a central network entity, the mobility agent. Thus, when hundreds of thousands of MNs are communicating in a given cellular network, such a centralised network entity causes well-known bottlenecks and single point of failure issues, which requires costly network dimensioning and engineering to be fixed. In addition, tunnelling encapsulations impact the overall 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. 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.

It is however well established that a huge amount of mobile communications are set up while the user is not physically moving, i.e. its MN stays in the same radio cell. For example, the user is being communicating at home, in his office, at a café... Applying the aforementioned centralised principles leads then to aggregate user's contexts and traffic at a central node in the network for the sake of mobility support whereas the MN remains motionless. As this leads to the introduced scalability and performances issues, alternative approaches may consider a way to better adapt mobility support in the network to cope with MN's movements and its ongoing traffic flows' requirements. Typically, one of the trend in the evolution of mobile networks is to go on flat architecture with the distribution of network functions, including mobility functions [I-D.liu-distributed-mobility] (Liu, D. and Z. Cao, "Distributed mobility management Problem Statement," March 2010.). According to this principle, [I-D.chan-netext-distributed-lma] (Chan, H., Xia, F., Xiang, J., and H. Ahmed, "Distributed Local Mobility Anchors," March 2010.) proposes a deployment of Proxy Mobile IPv6 in a flat architecture by splitting the location management and routing functions of the LMA.

In this document, we propose a slightly different approach by dynamically distributing mobility handling among terminals and access routers. This document inherits from concepts introduced in [NTMS2008] (Bertin, P., "A Distributed Dynamic Mobility Management Scheme designed for Flat IP Architectures.," November 2008.). Our goal is twofold:

- 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;

- confining the mobility support at the access routers level, keeping the rest of the network unaware of mobility events and their support.

---

3. Use-case and requirements

TOC

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] (Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," September 2007.). 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. Flows mobility is then required only to support the necessary traffic indirection from the access router on which the flow has been initially set up to the access router the host is currently attached. Such IP based indirection can even be made independent from access technologies types, providing thus inherent inter-access mobility facilities.

Based on these considerations, IP flow mobility relies on the dynamic provision of flow based traffic indirection between access routers. Hence, any given IP flow can be considered as implicitly anchored on the current MN's access router when being set up. 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).

4. Solution Overview

TOC

4.1. Dynamic Mobility Anchoring

TOC

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 1 (Distributed Mobility Anchoring) where the Mobility capable Access Router (MAR) supports both traffic anchoring and MN's location management functionalities. The idea is that 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, the traffic will be routed through MAR1 without requiring any specific mobility operation. When MN1 moves away from MAR1 and attaches to MAR3, the traffic remains anchored to MAR1 and is tunneled between MAR1 and MAR3. MAR1 becomes the mobility anchor, but only for traffic initiated by MN1 when it was attached to MAR1.

Communications newly initiated, e.g. to CN2, while the mobile node is attached to MAR3 will be routed in a standard way via MAR3. So, MAR3 is both the mobility anchor, i.e. the H-MAR, for traffic newly initiated (i.e. when the mobile node is attached to MAR3) and the V-MAR for traffic initiated while being attached to MAR1. If the mobile node moves away from MAR3, while maintaining communications with both CN1 and CN2, two mobility anchors come into play: the data traffic will be anchored in MAR1 for communication with CN1 and in MAR3 for communication with CN2.

Summarizing the above mechanism, it is proposed to locate mobility anchoring for the same mobile node depending on where the flow is initially created. Accordingly, communications are expected to be initiated without requiring mobility anchoring and tunneling.

With this solution, 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 to forward HNP1; then a tunnel between MAR1 and MAR3 will be established.

However such an 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. For example, if MN1 is anchored to both MAR1 and MAR3, these two mobility capable access routers would advertise respectively HNP1 and HNP3 for MN1.

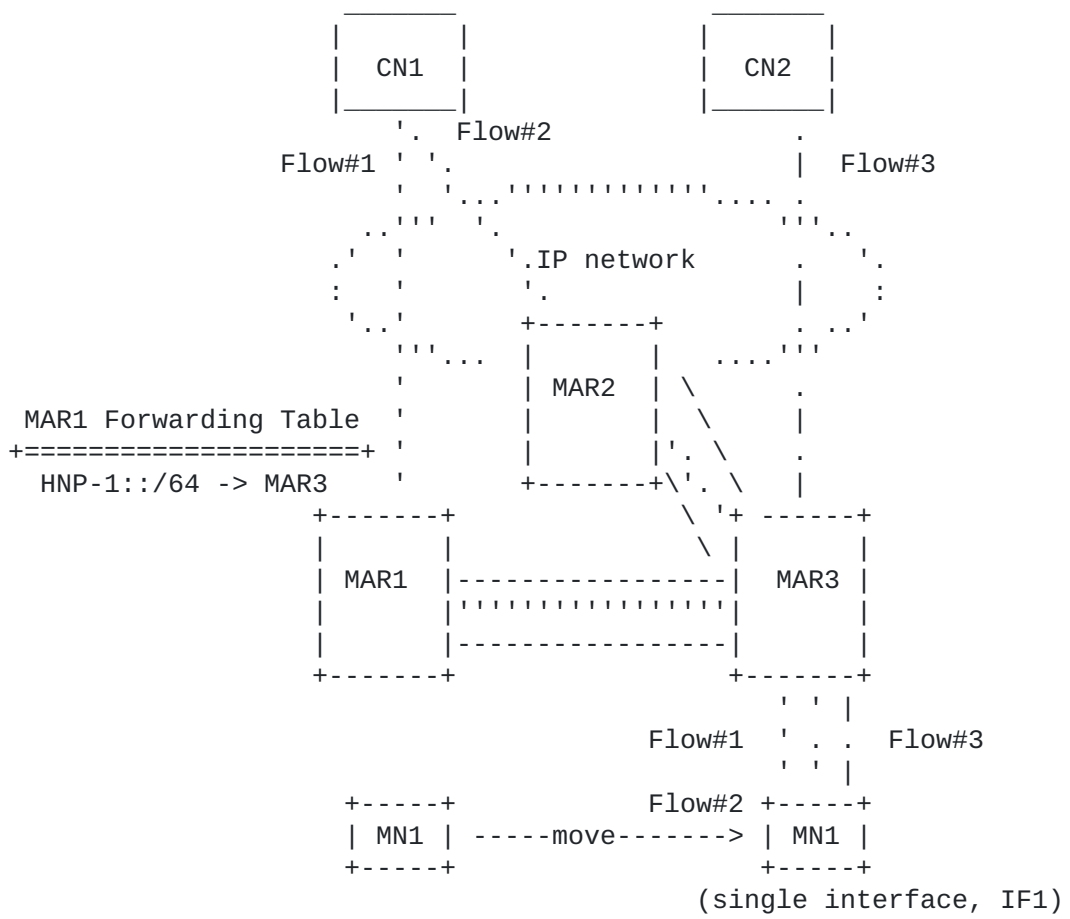

Figure 1: Distributed Mobility Anchoring

4.2. Protocol sequence for handover management

TOC

An example of handover management for a single interface mobile node is depicted on Figure 2 (Handover management with Distributed Mobility Anchoring). The mobile node, MN1, is assumed to move from MAR1 to MAR2. 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 handoffs to MAR2 which will thus act as V-MAR for HNP1: MAR2 retrieves the ongoing mobility sessions (e.g. from a policy store, as per [RFC5213] (Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” August 2008.)) for MN, then it proceeds to location update for HNP1 with MAR1 (H-MAR role), i.e., PBU/PBA exchange between MAR2 and MAR1.

4. MAR2 also allocates new HNPs for MN1; these HNPs are meant to be used by application flows initiated after the handoff.

5. MAR1, playing the H-MAR 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 can initiate and maintain 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.

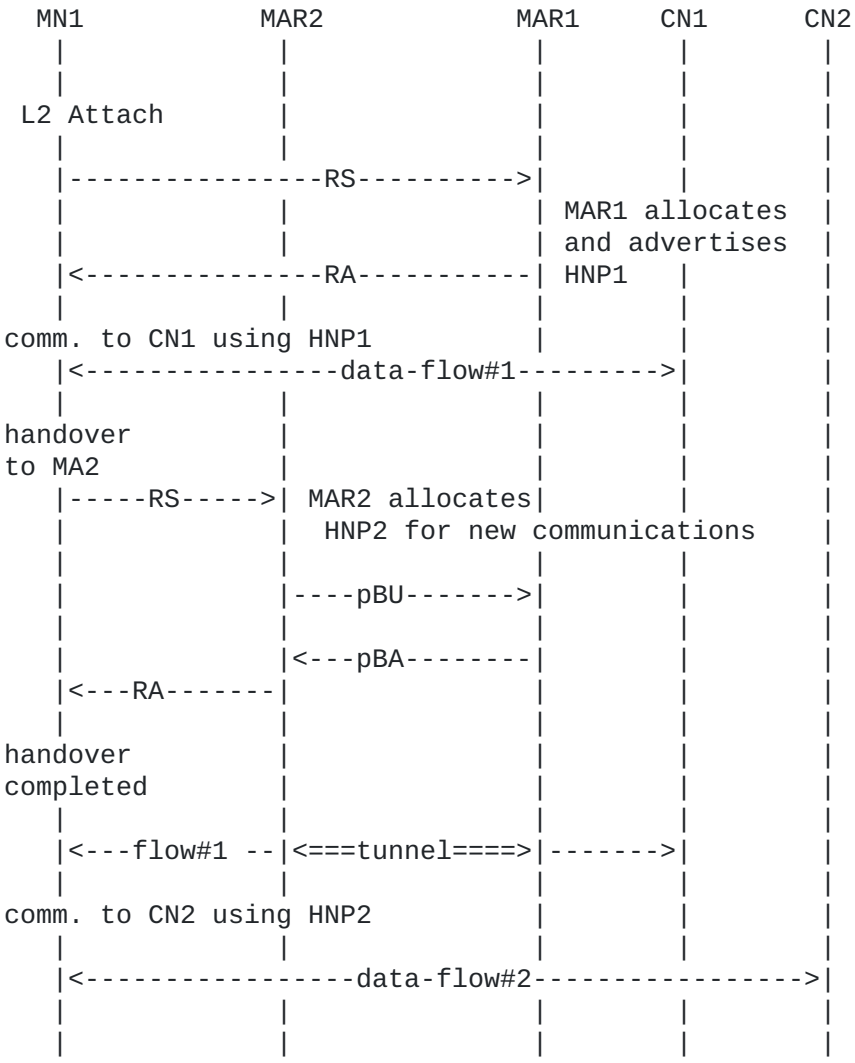

Figure 2: Handover management with Distributed Mobility Anchoring

4.3. Difference with Proxy Mobile IPv6

TOC

A V-MAR is required to advertise new per-MN HNP set for new IP communications to be initiated, while Proxy Mobile IPv6 advertises the same HNPs when roaming from MAG to MAG. So, while Proxy Mobile IPv6 is based on the per-MN prefix model, this proposal leverages on a per-MN and per-MAR prefix model. It is not required to statically allocate different set of HNPs per MAR. Actually, at a given time, only active MARs for an MN (i.e. access routers on which the mobile

node is currently attached to) need to share the per-MN HNPs set. So, for the sake of scalability, per-MN HNPs should be dynamically shared out among MN's active MARs.

A mobile node may be served simultaneously by more than one mobility anchor at the same time. Each MAR anchors the IP traffic initiated when the mobile node was attached to it.

4.4. IP flow mobility support

TOC

The distribution of mobility functions can also apply in the context of multiple-interfaces terminals and IP flow mobility. 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.

In the following, it is assumed that flow mobility consists in transferring a subset of prefixes from one access to another (i.e. a given prefix is associated to a given IP flow). This scenario is described in [I-D.jeyatharan-netext-multihoming-ps] (Jeyatharan, M. and C. Ng, “Multihoming Problem Statement in NetLMM,” March 2010.) and implemented in [I-D.yokota-netlmm-pmipv6-mn-itho-support] (Yokota, H., Gundavelli, S., Trung, T., Hong, Y., and K. Leung, “Virtual Interface Support for IP Hosts,” March 2010.). However, providing specific extensions to mobility signalling (extensions to be defined), the solution could also matches the scenario where a same prefix is shared across multiple interfaces (scenario described in [I-D.jeyatharan-netext-multihoming-ps] (Jeyatharan, M. and C. Ng, “Multihoming Problem Statement in NetLMM,” March 2010.)). In this case, a prefix is still anchored to one MAR but redirected IP flows are routed by the H-MAR using flow filtering mechanism.

Lets consider a simple example to illustrate the dynamic per-flow mobility anchoring. Figure 3 (Distributed IF flow Mobility Anchoring) 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 the the 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.

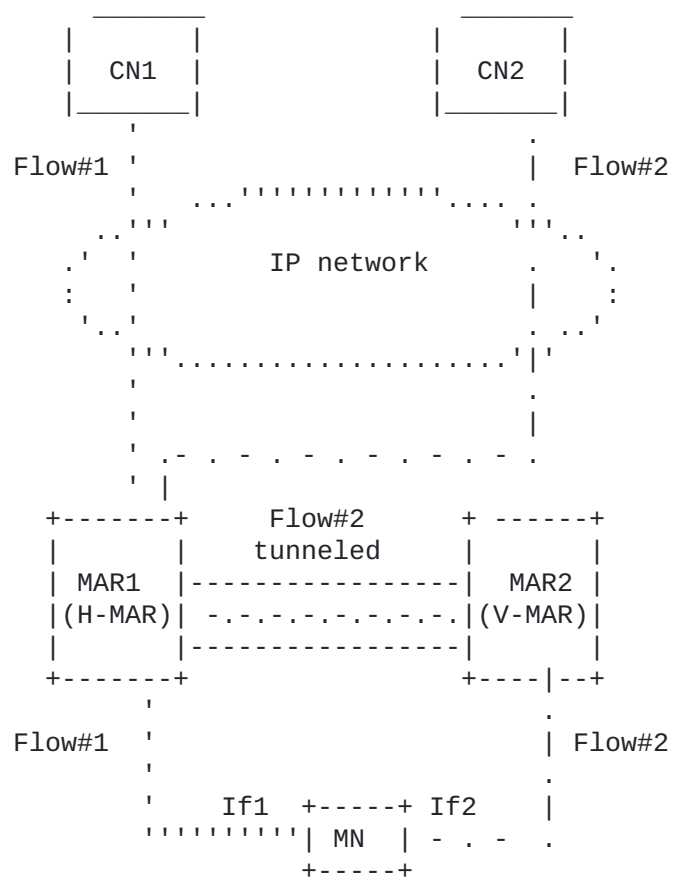

Figure 3: Distributed IF flow Mobility Anchoring

In case of the handover of an IP flow, initially adressed to one interface, the mobile node must be able to process that traffic also on the target interface. In order to meet that requirement, the mobile node could support the weak host model, as per [RFC1122] (Braden, R., “Requirements for Internet Hosts - Communication Layers,” October 1989.), [I-D.bernardos-mif-pmip] (Bernardos, C., Melia, T., Seite, P., and J. Korhonen, “Multihoming extensions for Proxy Mobile IPv6,” March 2010.). By supporting the weak host model, the mobile node can accept traffic, addressed to one IP address, on any of its interfaces.

Another solution for the host to support the handover from one interface to another, is to hide the inter-access handover to layers above IP. The mobile node can support this scenario by using a virtual IP interface. The applicability of that approach is discussed on [I-D.bernardos-netext-ll-statement] (Bernardos, C., Zuniga, J.,

and T. Melia, “Applicability Statement on Link Layer implementation/ Logical Interface over Multiple Physical Interfaces,” March 2010.) and [I-D.yokota-netlmm-pmipv6-mn-itho-support] (Yokota, H., Gundavelli, S., Trung, T., Hong, Y., and K. Leung, “Virtual Interface Support for IP Hosts,” March 2010.) describes a solution.

5. Implementation feedback

TOC

The solution proposed in this document has been implemented and tested on a Linux based testbed and for a single interface terminal. When several IPv6 addresses are available, Linux (at least the distribution we use) leverages on [RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.) default rules to select the source address. The problem is that, on a single interface host and when several global addresses are available, any of the [RFC3484] (Draves, R., “Default Address Selection for Internet Protocol version 6 (IPv6),” February 2003.) source address selection rules applies. So, in this case, Linux selects the more recent address registered among the list of potential source address. In our context, it leads to the following situation:

> A mobile node (MN1) attaches to a mobility capable access router (MAR1) advertising the prefix HNP1; so MN1 generates the IP address IP1. If MN1 attaches to a new mobility capable access router (MAR2) advertising the prefix HNP2, MN1 generates a new IP address IP2. At this stage, MN1 has two IP addresses: IP1 and IP2. If the mobile node comes back to MAR1, the more recent IP address, IP2, will be used to start new application. This behaviour brings issue with regards to the expected prefix management (described in Section 4.1 (Dynamic Mobility Anchoring)); actually applications are meant to use prefixes advertised on the current access link to start new data flow. In this example, MN1 must use IP1, and not IP2, to start new applications when coming back to MAR1.

In order to address the above issue, we have modified Linux source address selection algorithm. The modification overtake Linux mechanism and consists in always selecting the source address corresponding to the prefix advertised on the current access.

6. Security Considerations

TOC

TBD.

7. IANA Considerations

TOC

This document has no actions for IANA.

8. Acknowledgements

TOC

The authors would like to acknowledge Philippe Quenard and Carole Bonan who have implemented the solution decribed here. The authors would also like to express their gratitude to Lucian Suciu, Servane Bonjour and Karine Guillouard for their suggestions and reviews of this document.

Last but not least, the authors would like to acknowledge Dapeng Liu, Anthony Chan and Julien Laganier for having shared thoughts on the concept of distributed mobility.

9. References

TOC

9.1. Normative References

TOC

[RFC4862]	Thomson, S., Narten, T., and T. Jinmei, “IPv6 Stateless Address Autoconfiguration,” RFC 4862, September 2007 (TXT).

9.2. Informative References

TOC

[I-D.bernardos-mif-pmip]	Bernardos, C., Melia, T., Seite, P., and J. Korhonen, “Multihoming extensions for Proxy Mobile IPv6,” draft-bernardos-mif-pmip-02 (work in progress), March 2010 (TXT).
[I-D.bernardos-netext-ll-statement]	Bernardos, C., Zuniga, J., and T. Melia, “Applicability Statement on Link Layer implementation/Logical Interface over Multiple Physical Interfaces,” draft-bernardos-netext-ll-statement-01 (work in progress), March 2010 (TXT).
[I-D.chan-netext-distributed-lma]	Chan, H., Xia, F., Xiang, J., and H. Ahmed, “Distributed Local Mobility Anchors,” draft-chan-netext-distributed-lma-03 (work in progress), March 2010 (TXT, PDF).
	Jeyatharan, M. and C. Ng, “Multihoming Problem Statement in NetLMM,” draft-jeyatharan-netext-

[I-D.jeyatharan-netext-multihoming-ps]	multihoming-ps-02 (work in progress), March 2010 (TXT).
[I-D.liu-distributed-mobility]	Liu, D. and Z. Cao, "Distributed mobility management Problem Statement," draft-liu-distributed-mobility-01 (work in progress), March 2010 (TXT).
[I-D.yokota-netlmm-pmipv6-mn-itho-support]	Yokota, H., Gundavelli, S., Trung, T., Hong, Y., and K. Leung, "Virtual Interface Support for IP Hosts," draft-yokota-netlmm-pmipv6-mn-itho-support-03 (work in progress), March 2010 (TXT).
[NTMS2008]	Bertin, P., "A Distributed Dynamic Mobility Management Scheme designed for Flat IP Architectures.," NTMS'2008 , November 2008.
[RFC1122]	Braden, R., "Requirements for Internet Hosts - Communication Layers," STD 3, RFC 1122, October 1989 (TXT).
[RFC3484]	Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)," RFC 3484, February 2003 (TXT).
[RFC3775]	Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," RFC 3775, June 2004 (TXT).
[RFC5213]	Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6," RFC 5213, August 2008 (TXT).

Authors' Addresses

TOC

	Pierrick Seite
	France Telecom - Orange
	4, rue du Clos Courtel, BP 91226
	Cesson-Sevigne 35512
	France
Email:	pierrick.seite@orange-ftgroup.com
	Philippe Bertin
	France Telecom - Orange
	4, rue du Clos Courtel, BP 91226
	Cesson-Sevigne 35512
	France
Email:	philippe.bertin@orange-ftgroup.com