

Network-based Localized IP mobility Management: Proxy Mobile IPv6 and Current Trends in Standardization

Carlos J. Bernardos¹, Marco Gramaglia², Luis M. Contreras¹, Maria Calderon¹ and Ignacio Soto¹

¹ *Universidad Carlos III de Madrid
Leganes, Madrid, Spain*

{cjb, luisc, maria, isoto}@it.uc3m.es

² *Institute IMDEA Networks and
Universidad Carlos III de Madrid*

Leganes, Madrid, Spain

marco.gramaglia@imdea.org

Abstract

IP mobility support has been a hot topic over the last years, recently fostered by the role of IP in the evolution of the 3G mobile communication networks. Standardization bodies, namely IETF, IEEE and 3GPP are working on different aspects of the mobility aiming at improving the mobility experience perceived by users. Traditional IP mobility support mechanisms, Mobile IPv4 or Mobile IPv6, are based on the operation of the terminal to keep ongoing sessions despite the movement. The current trend is towards network-based solutions where mobility support is based on network operation. Proxy Mobile IPv6 is a promising specification that allows network operators to provide localized mobility support without relying on mobility functionality or configuration present in the mobile nodes, which greatly eases the deployment of the solution. This paper presents Proxy Mobile IPv6 and the different extensions that have been considered by the standardization bodies to enhance the basic protocol with interesting features needed to offer a richer mobility experience, namely, flow mobility, multicast and network mobility support.

1 Introduction

Mobility support in IP networks is a topic that has received considerable attention. This topic has gained more importance as IP networks are starting to be used to offer services, previously reserved for circuit switched ones, such as voice, and also because the role of IP in the evolution of the 3G mobile communication networks.

Traditional IP mobility support mechanisms, for example Mobile IPv4 [1] or Mobile IPv6 [2], are terminal based, meaning that terminals are aware of their mobility and have to do operations in order to be able to maintain their ongoing communication sessions. Nevertheless, an interest in network based solutions has appeared recently. Operators wish to support mobility without depending on functionality and configurations in the user terminal (Mobile Node in the usual IP mobility terminology). To face this requirement, the IETF¹ created the NetLMM (Network-based Localized Mobility Management) Working Group to develop a solution for providing network-based localized mobility support. The main objective was to develop a solution to provide, using functionality residing only in the network, mobility support to terminals moving and changing their point of attachment within a particular area (the Localized Mobility Domain, LMD). In this solution, terminals only perform the standard IP operations (e.g., Neighbor Discovery) without any particular functionality related to mobility at the IP layer. The base solution that has been developed by the NetLMM WG is the Proxy Mobile IPv6 protocol (PMIPv6) [3]. This protocol is based on the Mobile IPv6 protocol [2] but relocating the mobility related functionality of the terminal to network nodes.

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 1, number: 2/3, pp. 16-35

¹Internet Engineering Task Force: <http://www.ietf.org/>

In this paper we describe the PMIPv6 protocol (Section 2). This protocol provides the basic network-based localized mobility support for terminals, but there are also interesting and more advanced features that are required to allow operators offering a rich mobility experience using PMIPv6. We review in this paper three of the most interesting ones, looking at the motivation for the features, and the solutions that are being considered to provide them. The extensions analyzed in this paper are the ability for the terminal to receive flows through different network interfaces simultaneously attached: the so-called called flow mobility (Section 3); the support of multicast traffic and handover procedures in this situation (Section 4); and the support of mobile networks (Section 5). This will offer the reader a good understanding of the current status of the research and standardization work regarding network-based localized mobility support.

2 Network-based Localized Mobility Management: Proxy Mobile IPv6

When the IPv4 protocol [4] was designed, mobility was not considered a requirement, because user mobility was simply unthinkable in the Internet at that time. Without this requirement, a design decision was made to use the IP addresses for two different roles: locators and identifiers. More recently, IPv6 [5] took the same design decision of not separating the two roles. This is advantageous because if we know the name (identifier) of a node, we automatically have the locator for that node, which the routing system uses to reach it. Therefore, we do not need translation mechanisms nor a secure binding between identifiers and locators. Unfortunately, mobility requires separating the identifier role from the locator role: as an identifier, the address of a node should never change, but as a locator, it needs to change each time the node moves to a new location (IP subnetwork).

The IETF tackled this problem subsequently, designing the Mobile IP family of protocols that support mobility in IPv4 [6, 1] and in IPv6 [2]. These solutions provide the separation between the two roles by using two different addresses: the Home Address (HoA), used as identifier and the Care-of Address (CoA), used as locator. While the HoA is kept as the permanent identifier for a Mobile Node (MN) in the network, the CoA is updated every time the MN changes of subnetwork. An entity in the network, called the Home Agent (HA), is in charge of binding the various CoA with the unique HoA. In order to do this, the HA has to be placed where the HoA of that MN is topologically correct (i.e. the home network).

Although Mobile IP solutions enable the mobility of the MNs by allowing the change of point of attachment in an IP network while keeping a permanent identifier, they are not designed to provide efficient handovers. The handover performance is affected by the delay of the signaling exchanged between the MN and its HA. To reduce this delay, a family of solutions [7, 8] was proposed that used a local HA, closer to the MN and hence, providing mobility in a local domain.

All the mentioned solutions involve that the MN has to perform specific operations to support the IP mobility mechanisms. This has two important drawbacks: it increases the complexity of the MNs, and makes mobility support dependent on functionality and configurations (usually complex because of security related issues) that have to be present in the MN. Consequently, the IETF decided to develop a Network-based Localized Mobility Management (NetLMM) [9, 10] solution, an alternative approach for supporting mobility based on the philosophy of having all the required functionality in the network and providing the mobility support in the area covered by the deployment of the solution: the Localized Mobility Domain (LMD). In this solution, MNs do not require any special security configurations nor dedicated software, greatly simplifying the deployment. The protocol chosen by the IETF to offer Network-based localized mobility support is Proxy Mobile IPv6 (PMIPv6), which we will describe in detail shortly.

NetLMM can be applied in different scenarios [11], ranging from large campus WLANs where users move with their IP devices, to more complex ones such as automotive ones [12]. One of the main

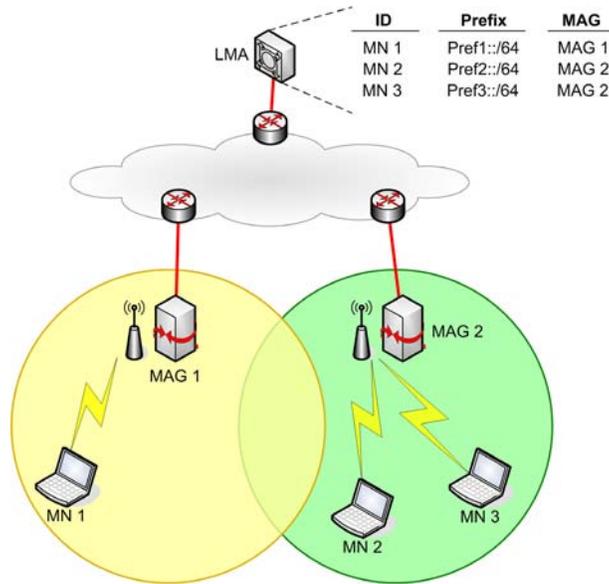


Figure 1: Network entities in Proxy Mobile IPv6.

application for a NetLMM protocol is its use in advanced beyond-3G Networks. Currently, Universal Mobile Telecommunications System (UMTS) and General Packet Radio Service (GPRS) networks use a proprietary network-based mobility mechanism based on the GPRS Tunneling Protocol (GTP) [10] to provide mobility support for user data traffic. However, a standardized protocol for any IP network will bring several advantages: reduced costs (either in terms of network management and equipment), easier extension to other technologies and less complex integration with other networks. For this reason, cellular networks operators have been a main driving force of the development of a NetLMM solution in the IETF.

2.1 Proxy Mobile IPv6 Operation

When designing a network-based mobility protocol, the main goal to achieve is that IP mobility operations should not involve the Mobile Node at all, so all the functionality has to be moved towards the network. Therefore, PMIPv6 [3] adds some extra functionality to existing network nodes to compensate for the absence of functionality in the MN. Specifically, it provides mobility support within a localized area by adding two logical entities to an IP network (see Figure 1):

- *Mobile Access Gateway (MAG)*. This entity takes care of the mobility related signaling on behalf of all the MNs attached to its links. Usually this role is done by the MN's access router (i.e. the first hop router for the MN in the LMD infrastructure). MAGs have to track MNs movements within the LMD and typically there will be several MAGs inside the same LMD.
- *Local Mobility Anchor (LMA)*. This entity is the anchor for the addresses used by the MNs in the LMD. It stores all the routing information needed to reach each of the MNs in the corresponding LMD by associating the MN with the MAG that the MN is using. A tunnel between the LMA and the MAG an MN is using allows the transfer of traffic from and to the MN.

As we already mentioned, PMIPv6 is totally transparent to the MN, that is able to move within the LMD without changing its IP address, being the IP layer of the MN unaware of the mobility. So, when an

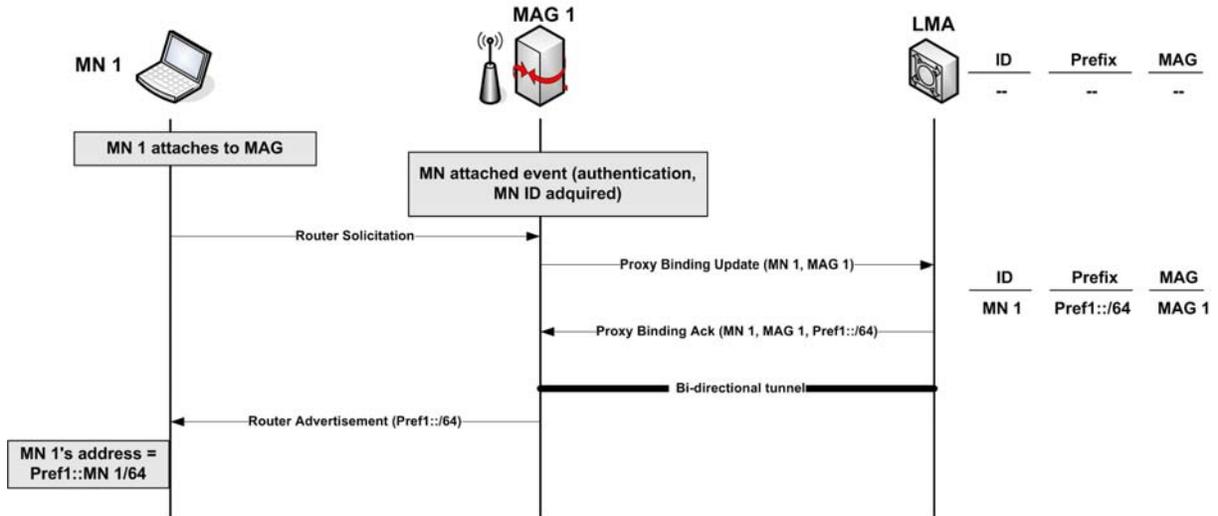


Figure 2: Signaling when an MN enters the PMIPv6 domain.

MN connects to an LMD, it just has to send a standard IPv6 Router solicitation², perhaps after performing some authentication operation to gain access to the network. The access router, that in PMIPv6 acts also as a MAG, will perform all the mobility related signaling on its behalf (see Figure2).

The MAG will perform all the necessary security checks in order to verify if the MN is authorized to use the network-based mobility management service. If the MN successfully passes this stage, the MAG has to use a stable and unique identifier for the MN, valid throughout the PMIPv6 domain. The MN's MAC address is a valid candidate for this purpose, and also some identifier related to the authentication operation. The standard does not specify what should be used as MN's identifier. The following step is to register the presence of the MN (attached to the MAG) to the LMA to let the MN use the network-mobility service. This registration is done by the MAG by sending a Proxy Binding Update (PBU, an extension of the Mobile IPv6 Binding Update message that is sent by the MN). The PBU includes the MN identifier and the Proxy CoA. This latter value is set by the MAG with the value of its egress interface because, as the MN is not involved in any mobility related signaling, the locator role carried out by the CoA in MIPv6 is assigned to the MAG in PMIPv6. The PBU message contains also additional information regarding the access link technology, the binding lifetime and a handover indicator that is used to differentiate handover from simultaneous attachment (multihoming). Upon sending the PBU, the MAG creates an entry inside its Binding Update List [2]. The Binding Update List is a data structure that is an extension from the one defined for MNs in Mobile IPv6. It is used by the MAG to keep track of each MN's bindings and stores information about the MN, the interface of the MAG to which the MN is connected, and the LMA serving it, among others.

When the LMA receives the PBU message it checks if the MN already has an active registration inside its Binding Cache (BC) using the MN's identifier. When an MN first enters the PMIPv6 domain, the LMA will not find any entry inside its BC, so it creates a new one. The PMIPv6 BC entry is an extended version of the MIPv6 one [2] that includes MN related information such as the MAG serving it. When the LMA accepts an MN inside a PMIPv6 domain, the LMA has to assign one or more network prefixes to it. The address(es) that the MN will eventually use inside the PMIPv6 domain will be configured from these prefixes (called Home Network Prefixes, HNPs). The MN identifier, the Proxy CoA, and the

²This is the default approach to enable the MAG detect the attachment of an MN. Nevertheless, other approaches, as layer-2 hints can also be used by the MAG.

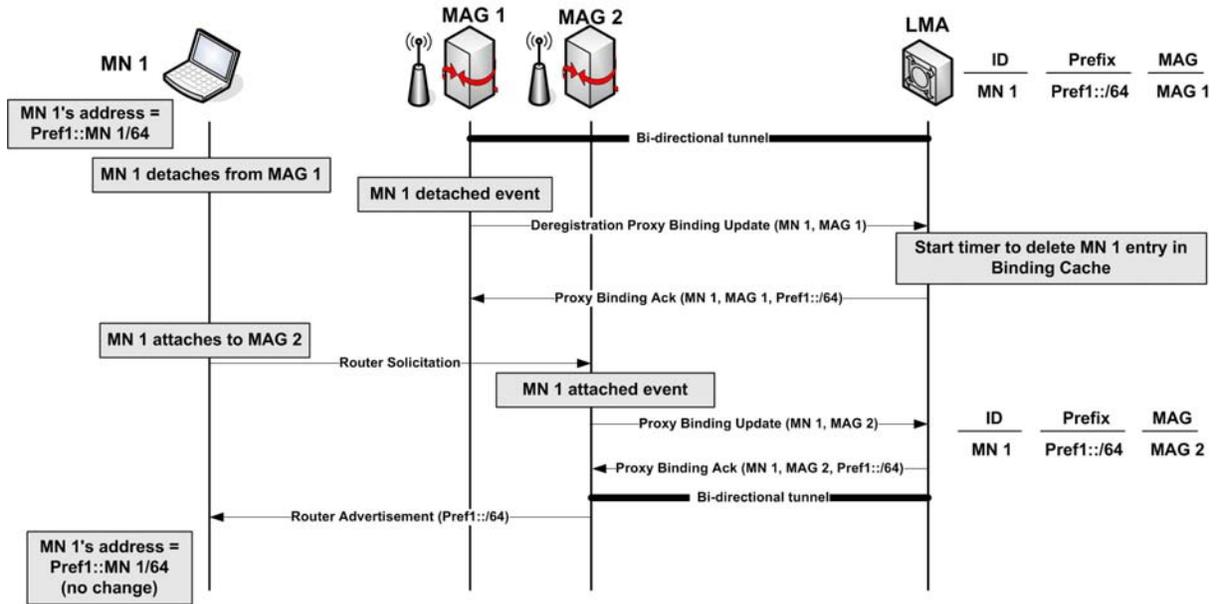


Figure 3: Signaling when an MN changes point of attachment.

HNPs allocated to the MN are used to create a new BC entry. The LMA also has to set up the right configuration to serve the MN traffic. So, in order to convey all the traffic destined to the MN, routes for the HNP are set accordingly to point to an IPv6-in-IPv6 tunnel from the LMA to the corresponding MAG. The LMA procedure is finished with a Proxy Binding Acknowledgment (PBA) that is sent from the LMA to the MAG, containing the HNP allocated to the MN.

After receiving the PBA, the MAG can reply to the initial Router Solicitation sent by the MN. The MAG will send to the MN a Router Advertisement (RA) containing the Home Network Prefix(es). At the same time, the MAG completes the configuration to serve the MN traffic by setting up the appropriate forwarding rules for the downlink/uplink traffic to/from the MN. PMIPv6 assumes the use of point-to-point access links. If PMIPv6 is used in a broadcast access link (e.g., a WLAN) it has to emulate a point to point link, because although RAs are intrinsically multicast, in PMIPv6 they convey configuration information that is specific for each MN. This issue can be solved by sending the RA to the MN’s IPv6 link-local address only (instead of to the IPv6 all-nodes multicast address) or using the MN’s MAC address as layer-2 destination address.

The procedure is slightly different when dealing with an MN that changes its point of attachment (see Figure 3). Although an explicit deregistration message does not exist, by employing link layer techniques or catching an IPv6 Neighbor Unreachability Detection event, the MAG can recognize that an MN left its access link. In that case, the MAG will send a Deregistration Proxy Binding update to the LMA (a PBU with lifetime set to 0). When receiving this kind of message, the LMA goes into a “wait mode” where it removes the routes associated to the MN and drops all the traffic going to the MN, but still does not erase the BC entry. This waiting time allows the MN to move from one access link to another and attach to the new MAG. Hence, the new MAG will take care of the MN and will register it in the LMA. As the LMA was still keeping the information from the previous registration, it has just to update the Proxy CoA (the new MAG egress interface address) and continue with the standard procedure, but using the HNP already assigned to the MN in the LMD. In this way the MN will not change the address it is using in the LMD and will be, at the IP layer, unaware of the movement. The new MAG has to use, when communicating to the MN, the same IPv6 link-local and link layer address of the previous MAG to avoid

the MN detecting a change of access router. A way to achieve this behavior consists in configuring the same fixed IPv6 link-local and link layer address to all the MAGs inside the PMIPv6 domain. Another option is to generate these two addresses at the LMA for each BC entry and send them to the MAG. Any of the proposed techniques are also helpful to guarantee address uniqueness in the access links of the PMIPv6 domain, but there still can be a Duplicated Address issue regarding the IPv6 link-local address used by an MN when first enters the PMIPv6 domain. In these cases PMIPv6 relies on the standard DAD procedure, with the constraint of configuring the PMIPv6 registration procedure in such a way that it is likely to be completed before the default waiting time of a DAD procedure.

Similarly to MIPv6, PMIPv6 signaling is very sensitive to security threats. PMIPv6 requires that the signaling between MAGs and the LMA is protected by IPsec. Security associations between the LMA and the MAGs are needed but this is an affordable problem, because they belong to the same administrative domain, i.e. operator.

Using PMIPv6, the MN can move across the LMD and change its access link, while keeping the same IP address, and be still reachable from any Correspondent Node on the Internet. The LMA forwards traffic destined to the Home Network Prefixes to the correct MAG using the configured tunnel. The MAG decapsulates the packets and forwards them to the MN transparently. In the opposite direction, traffic coming from the MN is encapsulated by the MAG and decapsulated by the LMA that will route it towards the final destination. If the destination is inside the LMD the MAG can take a local routing decision and forward it directly.

The main advantages of PMIPv6 compared with the standard MIPv6 solution are:

- *Reduced handover latency*, as the LMA is a local network entity and sending signaling to it produces less delay than sending signaling to a remote HA.
- *Reduced overhead*, as the tunnel ends at the MAG instead of at the MN as in MIPv6. Packets inside the access link are sent without any extra header. This is relevant because access links are likely wireless and bandwidth resources are scarcer over the air interface than in the wired backhaul network.
- *Reduced complexity and configuration requirements in the MNs*, terminals do not require any mobility related configuration nor new specific software, greatly simplifying the deployment.
- *Improved location privacy*, keeping the MN's IP address fixed over the PMIPv6 domain considerably reduces the chance that an attacker can deduce the exact location of the MN.

3 Flow Mobility

This section is devoted to flow mobility in PMIPv6 domains. We first introduce what flow mobility is and why it is needed. Then we explain why current Proxy Mobile IPv6 protocol does not fully support flow mobility and introduce ongoing standardization efforts to tackle this problem. Finally, we take a look at other approaches – not based on PMIPv6 – to support flow mobility.

3.1 Problem Statement

The success of the wireless mobile world is a fact that nobody can argue against. The number of wireless mobile subscribers accessing data services does not stop increasing. This is motivated by a variety of different reasons: 3G access is widely available (coverage reaches almost 100% of dense populated

areas in developed countries) and affordable by users (most mobile handsets are 3G capable³, USB modems are quite cheap and operators offer flat rates to their customers). Besides, the number and popularity of applications designed for smart-phones that make use of Internet connectivity is getting higher every day, contributing to the amplification of the penetration of these devices (e.g., iPhone, Android, Blackberry and Windows Mobile phones), which results in bigger demands for 3G connectivity everywhere. Due to the huge connectivity needs from users, 3G operators are challenged to enhance their network deployments to be able to cope with the users' traffic load.

Driven by this continuous growth on the users' demand for connectivity and the high costs of 3G deployment (mainly caused because the radio spectrum is limited), the use of disparate heterogeneous access technologies – what is commonly referred to as 4G [13] – is considered as a mechanism to expand network capacity. This extension is not only achieved in terms of effective coverage (i.e. one particular access technology might not be offered in certain locations, while others could be deployed as an alternative way of accessing the network) but also in terms of simultaneously available bandwidth (i.e. the effective data rate that could be achieved by using two or more access technologies at the same time). User devices equipped with multiple radios (also known as multi-mode terminals) would be potentially capable of improving the connectivity experience they provide by simultaneously using more than one single access technology. Mobile operators see today an opportunity of reducing the average cost per offered Megabyte (and therefore an increase of the revenue) by introducing an intelligent resource management mechanism that allows to offload traffic from the 3G network into other access candidate networks (mainly WLAN due to its high penetration) when available. This optimizes the operator's network use, while keeping the users' Quality of Experience (QoE).

Fully exploiting heterogeneity in the network access – e.g., enabling 3G offload – has proved to be difficult. Most of existing solutions in use nowadays enable the use of different technologies (e.g., 3G and WLAN) by adopting one of the following approaches (or a combination of them): *a*) manual user-based switching, or *b*) application-based switching. In the former case, users decide to switch on a network interface based on their preferences (e.g., cost, required bandwidth for the applications being used, WLAN availability, etc.), while in the latter, applications decide to turn on and off interfaces based on predefined preferences and network availability. Both approaches involve a change on the IP address seen by the applications, and therefore rely on them surviving that change (or re-establishing the session). Operators are not satisfied with any of these approaches, as they leave the mobility control on the final users and/or the application developers. Additionally, the QoE obtained by users in this case may not be good enough, as it depends on the application behavior or requires the session to be restarted.

A much richer solution can be achieved by enabling true flow mobility. Flow mobility refers to the movement of selected flows from one access technology to another, minimizing the impact on the users' QoE.

3.2 Flow Mobility for PMIPv6

A first step required in order to support flow mobility is the capacity to use several physical network interfaces. Proxy Mobile IPv6 allows an MN to connect to the same PMIPv6 domain through different interfaces, though in a very limited way. There are three possible scenarios [14]:

- *Unique set of prefixes per interface.* This is the default mode of operation in PMIPv6. Each attached interface is assigned a different set of prefixes, and the LMA maintains a mobility session (i.e. a binding cache entry) per MN's interface. PMIPv6 only allows to transfer all the prefixes assigned to a given interface to another one attached to the same PMIPv6 domain, and does not

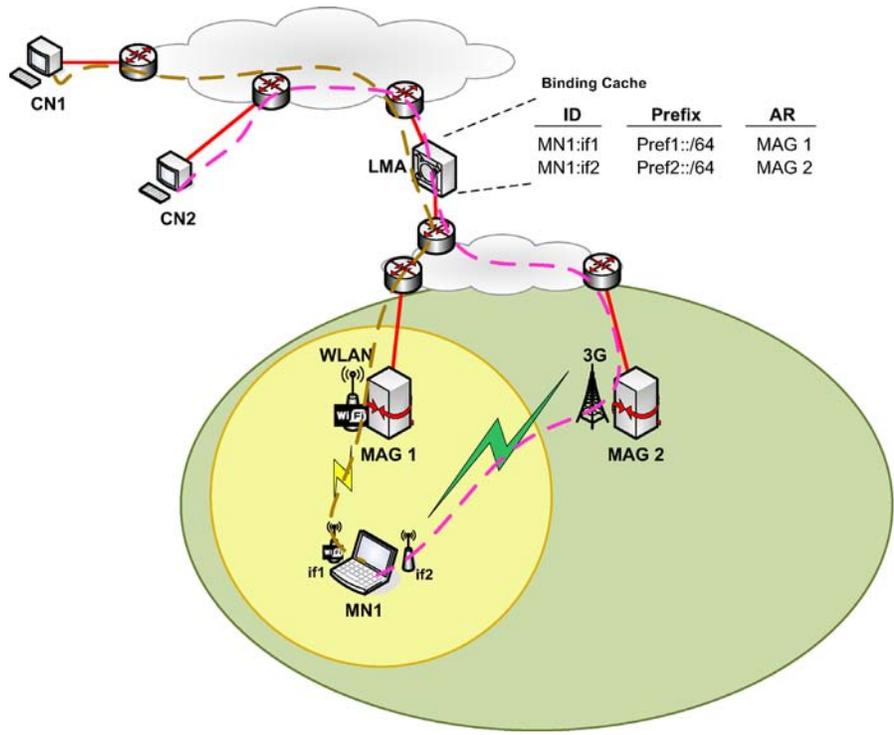
³As just one example, the number of HSPA devices has growth a 48% since October 2009, as reported by the Global mobile Suppliers Association (GSA).

fully specify how a MAG can figure out if a new mobile node wants to get a new set of prefixes assigned (i.e. having simultaneous access via multiple interfaces) or if the mobile node is performing a handover (i.e. the MN wants to transfer the prefixes bound to a previous interface to the new one).

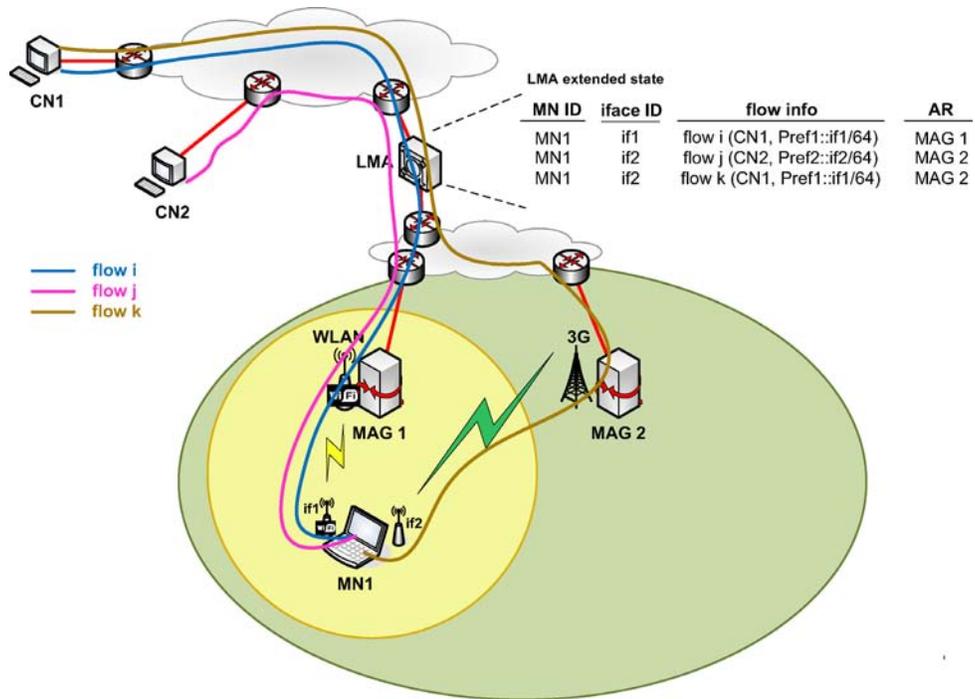
- *Same prefix but different global addresses per interface.* In this case the same prefix is assigned to multiple interfaces, though a different address is configured on each interface. This mode is not completely supported by PMIPv6. It either requires two different mobility sessions (as in the previous scenario) or only one but two separate host route entries. In any case, this scenario creates a multi-link subnet as the same prefix is advertised over different point-to-point links. This kind of scenario presents some issues as documented in [15].
- *Shared address across multiple interfaces.* In this scenario, the MN is assigned the same IP address across multiple interfaces. This enables applications on the terminal to see and use only one address, and therefore the MN could be able to benefit from transparent mobility of flows between interfaces. This scenario is not supported by current PMIPv6, it requires one mobility session per terminal and some kind of flow filters/routes at the LMA to be able to forward packets via the appropriate MAG. Besides, ensuring that multiple IP interfaces of the same device configure the same IP address is not easy to achieve (e.g., IPv6 specifications assume that unique IPv6 addresses are configured per interface, as guaranteed by running Duplicate Address Detection, DAD) nor to operate (not all Operating Systems support assigning the same IP address to multiple interfaces, and the multi-link subnet issue also appears here). One approach to mitigate this is to make use of link layer implementations that can hide the actually used physical interfaces from the IP stack [16]. For instance, the *logical interface* solution at the IP layer may enable packet transmission and reception over different physical media [17, 18].

PMIPv6 as defined in [3] cannot provide flow mobility in any of the previously described scenarios. We next identify and describe what functionality is missing from PMIPv6 to support flow mobility, by making use of an example. Figure 4 shows a potential use case of interest involving a multi-mode terminal attached to a PMIPv6 domain. The MN is attached to MAG1 through its WLAN interface (*if1*), and to MAG2 through its 3G interface (*if2*). With current PMIPv6 specification (*plain* PMIPv6, see Figure 4(a)), each interface is assigned a different prefix by the LMA (to allow simultaneous access) and two different mobility sessions (i.e. two separate binding cache entries) are maintained at the LMA. PBU/PBA signalling is used to keep alive the bindings at the LMA or to completely transfer the whole set of assigned prefixes from one interface to another. In order to support flow mobility, the state at the LMA needs to be extended (*extended* PMIPv6, see Figure 4(b)), so the LMA is able to group mobility bindings referring to the same MN. Additionally, flow state should be introduced at the LMA, so it can forward packets differently (i.e. through different MAGs) on a per-flow basis. The MAG behavior needs also to be modified, since the MAG should be aware of all the MNs' IP addresses that are reachable through the point-to-point link it has set up with the MN. In order to transfer this information, the PMIPv6 signalling between the MAG and the LMA may need to be extended as well.

The mobile node behavior needs also to be considered. In the plain PMIPv6 scenario, the IPv6 addresses assigned to *if1* (*addr1*) and *if2* (*addr2*) are different ($\text{Pref1}::\text{if1}/64$ and $\text{Pref2}::\text{if2}/64$, respectively). Packets addressed to *addr1* will always arrive via *if1* (and the same for packets addressed to *addr2*, arriving via *if2*). In a flow mobility-enabled scenario, *addr1* and *addr2* may belong to different prefixes, belong to the same one, or even be the same IP address. Moreover, packets addressed to *addr1* may arrive at *if2* (and the other way around), and should be processed by the MN normally.



(a) Plain PMIPv6 (as defined in RFC 5213)



(b) Extended (flow mobility enabled) PMIPv6

Figure 4: Flow mobility in PMIPv6: what is missing?

3.3 Standardization Work

Flow mobility is currently a hot topic in the standards community. The Internet Engineering Task Force (IETF), Institute of Electrical and Electronics Engineers (IEEE) and the 3rd Generation Partnership Project (3GPP) are all working in different aspects towards the standardization of flow mobility support. Among them, the IETF is particularly focusing – within the NETEXT WG – on extensions for Proxy Mobile IPv6.

Before sketching how IETF is addressing the problem of flow mobility for PMIPv6, it is worth noticing that there are two different types of IPv6 mobile nodes that can be supported by a flow mobility solution (shown in Figure 5):

1. *Terminals with a single interface visible from the IP stack.* Certain link-layer implementations can hide the use of multiple physical interfaces from the IP stack [16]. The *logical interface* [17], [18] at the IP layer is the most complete approach, as it allows both sequential and simultaneous use of different physical media. From the perspective of the IP stack and the applications, a logical interface is just another interface. A host does not see any difference between a logical and a physical interface. All interfaces are represented as software objects to which IP address configuration is bound (see Figure 5(a)).
2. *Terminals with multiple IP interfaces.* In case the mobile terminal does not implement the logical interface concept (or an alternative link-layer approach that hides the use of multiple media to the IP layer), it is still possible to enable full flow mobility if the terminal follows the *weak host* model [19, 20]. This model (see Figure 5(b)) does not limit the traffic reception at a host to only those IP packets whose destination address matches the IP address assigned to the interface receiving the packets, but allows the host to receive and process packets whose IP destination address corresponds to that of any of the local interfaces of the host. We have performed some tests with different operating systems, and the results show that both Linux (tested with Linux-2.6.26) and Mac OS X (tested with Leopard version) implement the weak host model for both IPv4 and IPv6 traffic. We have not performed tests with Windows, but some results have been reported in [21]. Windows XP and Windows Server 2003 use the weak host model for all IPv4 interfaces and the strong host model for all IPv6 interfaces, not being possible to modify that behavior. The Next Generation TCP/IP stack in Windows Vista and Windows Server 2008 supports the strong host model for both IPv4 and IPv6 by default on all interfaces. In this case, the stack can be configured to use the weak host model.

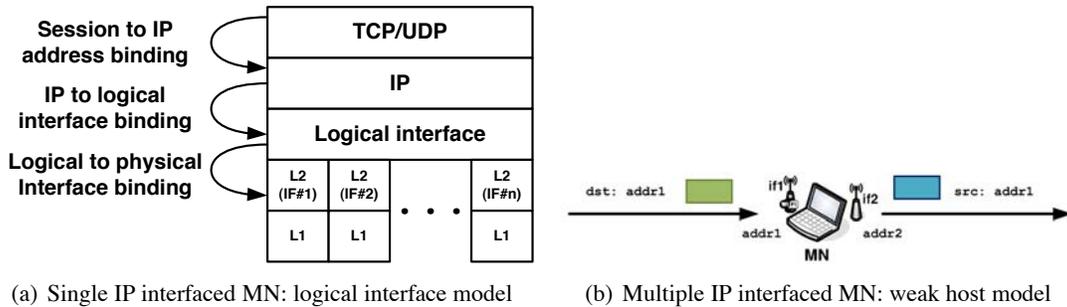


Figure 5: Types of IPv6 Mobile Node.

The NETEXT WG at the IETF is currently chartered to work on solutions that assume that the mobile node supports the logical interface model. This model is not the only one that could potentially be

followed to support flow mobility, though it is shown in [16] to be the one that meets all the requirements to enable flow mobility while hiding the existence of multiple physical network interfaces to the IP stack. There is still no solution standardized by the IETF, but the WG is currently working on protocol extensions that – controlled by the LMA – allow the routing of different flows through different physical interfaces of the MN. In order to do so, new signalling messages are being defined to convey flow mobility operations from the LMA to the MAGs, as well as extensions to the binding cache data structures to support the logical grouping of different physical interfaces of the same mobile node [22].

Current IETF work does not deal with how the network decides to move a particular flow, triggers could be potentially received from non-PMIPv6 network entities, or even from the MN itself. As an example of mobile-based triggers, the LMA could receive input (e.g., by means of a layer 2.5 signalling [23]) from the MN detecting changes in the mobile wireless environment (e.g., weak radio signal, new network detected, etc.). Upon receiving these triggers, the LMA can initiate the flow mobility procedures. For instance, when the mobile node only supports single-radio operation (i.e. one radio transmitting at a time), only sequential (i.e. not simultaneous) attachment to different MAGs over different media is possible. In this case, link layer signalling can be used to perform the inter-access technology handover and communicate to the LMA the desired target access technology, MN-ID, Flow-ID and prefix. The IEEE is currently evaluating if this type of extensions to the Media Independent Handover (MIH) Services [24] could be developed, to aid the MN take handover decisions and inform the network about events with a flow mobility granularity.

3.4 Alternative Approaches to Support Flow Mobility

The concept of flow mobility has been extensively analyzed for client-based mobility protocols, and there already exist standardized solutions, such as the use of flow bindings extensions for Mobile IPv6 [25] and Dual Stack Mobile IP (DSMIP) [26]. The use of this kind of client-based solution has been proposed as a mechanism to enable mobile operators to offload data from their 3G networks [27], and there even exist approaches based on the IP Multimedia Subsystem (IMS) framework [28]. We argue that client-based solutions have several disadvantages, since they require to modify the users' devices to include an IP mobility stack, which also has to be provisioned with proper configuration and security credentials (in addition to those required to access the operator's network). This additional requirements might limit the usability of a solution due to the difficulties involved in its deployment.

4 Multicast Support

IP multicast allows to support efficiently group communication services (one-to-many or many-to-many) over existing IP networks. Applications like video-conference, distant learning, interactive gaming or TV distribution may take advantage of this extension of the IP protocol which basically facilitates the delivery of a single copy of a data stream to multiple listeners or receivers interested in receiving the same content simultaneously. Multicast enabled routers in the network are in charge of replicating the data packets to reach the receivers, usually spread around the whole network, creating then a multicast distribution tree over the network. Thus, IP multicast helps to save resources in terms of both data source processing (allowing the source to deliver just one copy of data independently of the number of listeners), and network bandwidth (the network duplicates traffic only when the receivers are reachable via different links).

The increasing generalization in the use of multicast has triggered the need for supporting a multicast receptor which is mobile or attached to a network that is moving [29]. Group communication is out of the scope of the PMIPv6 standard specification. This fact produces inefficiencies when distributing contents

to multiple receivers in these scenarios (individual copies per MN of a common data stream between the LMA and the MAG) as well as when supporting MN mobility (dynamic adaptation of multicast distribution tree in case of handover).

Several solutions have been already proposed supporting multicast for Mobile IP networks [30]. However, they are not directly applicable to PMIPv6 because of the specificities of network-localized management environments, where the MN is not aware of network layer changes. A new approach is then needed to provide multicast in PMIPv6 domains.

With such aim, the MULTIMOB working group has been chartered at the IETF to specify a solution for multicast listener mobility compatible with current PMIPv6 and multicast protocol standards, that is, without any additional protocol extensions. The focus of the working group at its initial stage has been limited to both the study of the adaptation of the multicast group management protocols originally defined for wired networks, and the definition of a primary base solution to deploy multicast in the current Proxy Mobile IP architecture, as the multicast receiver MN moves.

An additional constraint is observed regarding the network from where the MN obtains the multicast stream it subscribes to. Basically, two kind of solutions can be differentiated: the remote subscription, where the MN gets the multicast data stream from the LMA, and the local subscription, where the MN directly obtains the multicast stream from the visited network. According to the current MULTIMOB WG charter terms, only the remote subscription case is considered by now.

4.1 Multicast Basics

Conceptually, IP multicast represents a receiver centric architecture. A number of receivers, located anywhere in the network, subscribe to a content in the form of a multicast session group. The content is distributed by means of a particular data stream creating a multicast flow. A single copy of such data stream is carried on every link in the network along the multicast path dynamically created to reach the interested receivers. The data stream is replicated on the routers where the multicast path topologically diverges.

The multicast source injecting the data stream to the network does not maintain any subscription list of interested receivers. The source simply sends the data stream to an arbitrary group of hosts represented by an IP multicast address. The receivers indicate their interest on receiving certain content by explicitly joining the multicast group. The Multicast Listener Discovery (MLD) protocol for IPv6 (with its versions MLDv1 [31] and MLDv2 [32]) defines the control messages for managing the group membership process. Multicast protocols distinguish between multicast receiver (or host part) and multicast router (or network part) functionalities. Basically, the host part is devoted to the group subscription management, while the router part is focused on building and maintaining the multicast tree.

A multicast-enabled router in the receiver's subnetwork will capture the control messages for joining or leaving the multicast group. The router will periodically check if there are receivers interested on a previously subscribed group in the subnet by sending out queries asking for some intended receiver for that group. The router will stop the data forwarding of a certain group when receives an explicit indication of the last receiver interested in such group, or as result of not receiving any confirmation of interest from the periodical queries.

In some cases, the router can act as a proxy [33] for the group membership indications of the receivers connected to it, instead of the multicast router role described above. This typically occurs in aggregation networks, where the first-hop router concentrates the traffic of a huge number of receivers. The proxy will perform the router portion of the group membership protocol on each downstream interface, while it will play the host role on the upstream interface towards the next multicast-enabled backbone router. The proxy is in charge of summarizing the subscription demand of the receivers, acting as a unique host towards the upstream multicast router.

On the network side, the routers make use of different protocols to dynamically build and maintain the multicast distribution tree from the multicast source to the final set of receivers. Protocol Independent Multicast (PIM) is the most commonly deployed protocol in commercial networks. PIM uses the unicast routing information (independently of the conventional routing protocol used to obtain it, i.e., OSPF, BGP, IS-IS, etc) to build a loop-free multicast distribution tree. The distribution trees can be further classified as source trees, where the multicast source is the root of the tree (forming in this case a shortest path tree between source and receiver), and shared trees, where the root of the tree is some predetermined router in the network (known as rendezvous-point) which collects the incoming traffic from different sources (if there are more than one) and distributes it to the interested receivers. Source trees fit the one-to-many group communication model, or Source-Specific Multicast (SSM), whereas the shared trees can support both the one-to-many and the many-to-many group communication model, also known as Any-Source Multicast (ASM), at the cost of a higher complexity. PIM-SSM [34] and PIM-SM [35] are the corresponding protocol specifications.

4.2 Base Solution Rationale for Multicast Listener Mobility in PMIPv6 Domains

The base solution [36], as in the simple unicast mobility case, intends to provide a solution to manage multicast traffic delivery to MNs that do not participate in the IP mobility signaling process. Since the MN is, at the IP layer, not aware of the mobility, the network should implement some mechanisms to ensure the multicast reception by the MN while it moves.

When forwarding packets to the MN, the case for multicast traffic is slightly different to the unicast one. The IP destination address of multicast packets does not identify the receiver, so the LMA would need additional information to be able to forward the multicast flow towards an MN requiring some content. To this end, new functionality is required in PMIPv6 entities to support multicast. The goal of the base solution is to introduce such functionality without modifying or extending current mobility and multicast standards.

The MN will express its interest in joining or leaving a multicast group by sending MLD control messages to the MAG, which acts as first hop on the point-to-point link. The MAG will maintain the individual multicast status of the interface for that link and will handle the multicast traffic towards the MN accordingly to the MLD messages received (a non multicast-enabled MAG would simply ignore the group membership messages). In the base solution the MAG incorporates MLD proxy functionality. As a proxy, the MAG is responsible of summarizing the group subscription requests of the MNs connected to it, alleviating the processing load of control messages by the first multicast router, located deeper in the network.

As direct consequence of the remote subscription model, the multicast traffic reaches the MNs passing through the corresponding LMA (note that there might be multiple LMAs deployed in the same LMD). A distinct MLD proxy instance is then defined per LMA connected to the MAG, in such a way that every MAG-LMA tunnel is part of a separate MLD proxy domain. For every proxy instance in the MAG, the tunnel interface pointing to the LMA becomes the proxy upstream interface, whereas the point-to-point links towards the MNs associated with a specific LMA, as defined in the Binding Update List, are the corresponding downstream interfaces of each instance.

The LMA at the end of the tunnel can play the role of being either the first multicast router for the associated MNs or a further MLD proxy. This is transparent for the MNs and has no impact on the multicast traffic reaching them, which will flow from the LMA to the MAG using the existing tunnel between both. Independently of the role played by the LMA, this entity, being the addresses anchor point agent for every MN, will be in charge of interacting with the multicast infrastructure out of the PMIPv6 domain.

According to this architecture, the MAG capabilities for summarizing control messages upstream

apply per set of MNs associated with a certain LMA, as the different proxy instances of the same MAG work isolated one from each other. The LMA will maintain the multicast status state of every tunnel interface connecting to a MAG, and will be in charge of sending the periodical queries on the tunnel to verify that there are at least one interested receiver per forwarded group. Such multicast status reflects the summarized view offered by the MAG on behalf of the MNs attached to it and bound to the LMA. A multicast data stream will be delivered over the tunnel interface or removed from it according to the aggregated behavior of the group of MNs attached to the MAG.

The LMA, the MAG and the tunnel linking both entities are all part of the multicast delivery tree built to distribute the multicast traffic. This branch will be common to every multicast tree providing content subscribed by an MN in a MAG and associated to a particular LMA. It makes possible to send just a single copy of a concrete data stream per group of MNs demanding the same content.

As the MN moves and connects to different points of attachment, the network should be able to guarantee the delivery of the multicast traffic following the MN movement, since the MN will not participate on any signaling process to keep the multicast subscription at the new location because it is not aware of the mobility. The handover process involves the set up of a new point-to-point link at the new MAG where the MN attaches, and the release of the old one at the previous MAG where the MN comes from. In such a mobile environment, the MLD proxy instances in the MAGs have to be able of adding and removing downstream interfaces dynamically. A proxy instance will add a new downstream interface when a new MN connects to the MAG, while it will remove the downstream interface when an MN leaves the MAG.

In a handover event, the MN enters an access network under responsibility of a distinct multicast-enabled MAG. Following the expected standard behavior of MLD for a new link set up, once the MAG has completed the configuration of the interface on the new point-to-point link against the MN just attached (that is, after receiving the PBA message from the LMA for this MN), the MAG kernel immediately sends an MLD General Query. At this moment the point-to-point link has not been yet configured as part of the downstream set of interfaces of any MLD proxy instance. The decision of what proxy instance corresponds to the new attached MN will be based on the LMA serving the MN, as there is a instance per LMA. Once the point-to-point link is configured as a new downstream interface of the corresponding proxy instance, the proxy instance sends an additional MLD Query for getting knowledge of any active multicast subscription by the MN.

When an MLD Query is received by the MN, it will provide information about the active memberships it has, if any, in the form of an MLD Report message. The MAG, upon receiving the MLD Report message, will pass it to the assigned proxy instance for further processing, considering that the assignment to a proxy instance has been already configured for the point-to-point link (if not, the MLD Report message coming from the MN is discarded). Finally, the proxy instance will set up the multicast status of the downstream interface adding group membership information to the multicast forwarding information base if one or more active subscriptions are maintained by the MN, or will not take any further action in other case.

If the content required by the newly connected MN is currently being received by another MN managed by the same MLD proxy instance, the content is directly delivered to the receiver, without triggering a MLD Report upstream, towards the LMA. However, if the content subscribed by the entering MN is new within the framework of the MLD proxy instance, a change in the group membership database is produced, and the MAG sends an MLD Report upstream to the LMA requiring such flow. The LMA will also update the multicast status of its downstream interfaces and will proceed to inject the data stream on the tunnel pointing to the new MAG.

Similarly, in the MAG where the MN was previously attached, the point-to-point interface is removed from the assigned proxy instance when the MAG verifies that the MN is no longer active. If that MN was the unique MN in a proxy instance demanding a certain data stream, the MAG will send an MLD Done

message to stop the forwarding of the content to the MAG from the LMA because no more receivers are interested in it. The LMA will update its forwarding states and will stop the delivery of the flow onto the tunnel to the previous MAG.

In this way, the multicast distribution tree will be re-built by the PMIPv6 network entities accompanying the MN movement.

4.3 Future Steps

The base solution described above represents a first step towards the integration of multicast distribution within PMIPv6 domains. With the aim of preserving current protocol specifications, the solution suffers some inefficiencies [37]. Issues such as the fact that MNs associated to different LMAs but subscribed to the same multicast content will cause that multiple copies of the same data stream arrive to the MAG, or the huge delay the MN can incur in receiving the multicast flow after handovers caused by either the MLD Query and subsequent Report messages processing times by the MN, or the radio transfer delays (i.e., access to radio resources among competing MNs, radio frame structure timing, retransmissions caused by radio channel unreliability, etc), makes it necessary to revisit the solution proposal and provide it with the needed mechanisms to improve the expected performance. Other aspects such as local multicast subscription for direct multicast routing, dedicated LMA for multicast traffic, or membership protocol tuning for wireless environments can contribute to a better and more efficient architecture.

Additionally, topics such as multicast source mobility or multicast flow management on multi-homed nodes are also open points that need to be addressed.

5 Network Mobility and PMIPv6

The NEMO Basic Support protocol (NEMO B.S.) [38] extends Mobile IPv6 to enable Network Mobility. In this protocol, a mobile network (known also as Network that Moves - NEMO) is defined as a network whose attachment point to the Internet varies with time. The router within the NEMO that connects to the Internet is called the Mobile Router (MR) [39]. It is assumed that the NEMO has a Home Network where it resides when it is not moving. Since the NEMO is part of the Home Network, the Mobile Network has configured addresses belonging to one or more address blocks anchored at the Home Network: the Mobile Network Prefixes (MNPs). These addresses remain assigned to the NEMO when it is away from home. Of course, these addresses only have topological meaning when the NEMO is at home. When the NEMO is away from home, packets addressed to the Mobile Network Nodes (MNNs) will still be routed to the Home Network. Additionally, when the NEMO is away from home, i.e. it is in a visited network, the MR acquires an address from the visited network, called the Care-of Address (CoA), where the routing architecture can deliver packets without additional mechanisms.

The goal of the network mobility support mechanisms [40] is to preserve established communications between the MNNs and external Correspondent Nodes (CNs) despite movement. Packets of such communications will be addressed to the MNNs addresses, which belong to the MNP, so additional mechanisms to forward packets between the Home Network and the NEMO are needed. The basic solution for network mobility support [38] essentially creates a bi-directional tunnel between the Home Agent and the Care-of Address of the MR.

When a NEMO is connected to a PMIPv6 domain, two different scenarios are possible as determined by the role played by the PMIPv6 domain: NEMO's Home Network or NEMO's visited Network.

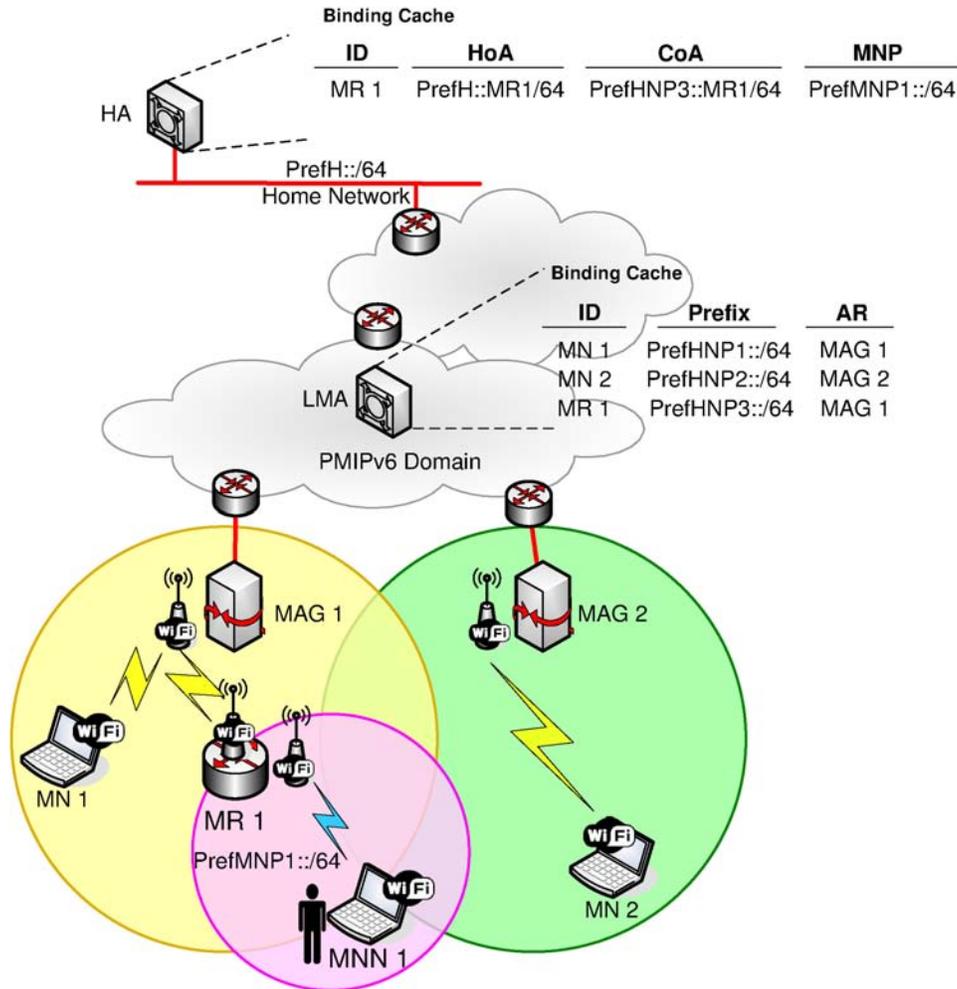


Figure 6: NEMO visiting a PMIPv6 domain.

5.1 PMIPv6 Domain as Visited Network of the Mobile Network

In this scenario (see Figure 6), when the mobile network, i.e. the MR, enters the PMIPv6 domain, it gets its Home Network Prefix(es) (HNPs) from the MAG, builds an address (to be used as its CoA) and it registers this address in its HA, binding the MNPs managed by the MR (that is, the IPv6 prefixes used in the mobile network) to its current CoA. Since this IPv6 address is provided by the PMIPv6 domain, it does not change while the MR moves within the domain, and therefore these movements are transparent to the MR’s mobility software. The mobile network is able to roam within the LMD but also to move outside the domain, thanks to the NEMO B.S. operation. With the current standards if a node in the NEMO wanted to detach from the MR and attach directly to the PMIPv6 domain, i.e. to a MAG, it would need MIPv6 to manage its mobility by itself.

5.2 PMIPv6 Domain as Home Network of the Mobile Network

Having the PMIPv6 domain as Home Network of the NEMO requires some extensions to the PMIPv6 operation. First, the MR needs to get not only a HNP to configure its home address but also the Mobile Network Prefix(es) to be used within the mobile network (i.e. by the MNNs connected to the MR).

Second, the MAG and the LMA should also have forwarding state for the MNPs, given that the LMA has not only to intercept and redirect the packets for the MR, but also intercept and redirect the packets for the MNNs belonging to the MNP(s). There are several works in the IETF extending the operation of PMIPv6 to add this functionality [41, 42].

It is worth noticing that with current standards, nodes attached to the mobile network without MIPv6 cannot detach from the MR and attach to an access network without changing their IP address and therefore breaking their ongoing session. This functionality would be useful in some interesting scenarios as, for example, the following: users move around a large area (e.g., an airport, an exhibition site, a fair-ground or even a metropolitan area covered by different public transportation systems), where attachment points to the Internet might be available both in fixed locations (such as coffee shops, airport terminals or train stations) or in mobile platforms, such as vehicles (e.g., buses that move between pavilions at a fair or a train that moves from one terminal to another at an airport). Users demand the ability to keep their ongoing communications while changing their point of attachment to the network as they move around (e.g., when a user leaves a coffee shop and gets on a bus).

A possible extension to PMIPv6 is presented in [12] where a new functional entity named moving MAG (mMAG) is proposed to be part of the PMIPv6 domain. An mMAG moves within the domain changing the MAG it is connected to, and extends the PMIPv6 domain by providing IPv6 prefixes belonging to this domain to mobile nodes and by forwarding their packets through the LMA.

6 Conclusion

Telecommunication operators are pushing for the standardization of network-based localized IP mobility support solutions. These solutions allow them to offer mobility support in their IP networks without depending on functionality or mobility related configuration in user terminals. This paper has described the main efforts along these lines. PMIPv6 is the base protocol standardized by the IETF to provide network-based localized mobility support in IP networks. Both the research community and different standardization bodies – mainly the IETF – are working on additional features to enhance the possibilities of using PMIPv6, including flow mobility, multicast support and network mobility support. This paper has described the motivation for each feature and the solutions that are being considered to provide them.

PMIPv6 and its potential extensions are being positioned as the solution of choice to provide mobility support in IP networks. The inherent benefits of using PMIPv6 and, the interest and value of the additional proposed features allow us to envision a significant impact of these technologies in future operator networks.

6.1 Acknowledgments

The research leading to the results presented in this paper has received funding from the Spanish MICINN through the I-MOVING project (TEC2010-18907) and from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement 258053 (MEDIEVAL project).

References

- [1] C. Perkins. IP Mobility Support for IPv4. RFC 3344, August 2002.
- [2] D. Johnson, C. Perkins, and J. Arkko. Mobility Support in IPv6. RFC 3775, June 2004.
- [3] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil. Proxy Mobile IPv6. RFC 5213, August 2008.
- [4] J. Postel. Internet Protocol. RFC 791, September 1981.
- [5] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460, December 1998.
- [6] W. Stallings. Mobile IP. *The Internet Protocol Journal*, 4(2):2–14, June 2001.
- [7] E. Fogelstroem, A. Jonsson, and C. Perkins. Mobile IPv4 Regional Registration. RFC 4857, June 2007.
- [8] H. Soliman, C. Castelluccia, K. ElMalki, and L. Bellier. Hierarchical Mobile IPv6 (HMIPv6) Mobility Management. RFC 5380, October 2008.
- [9] J. Kempf. Goals for Network-Based Localized Mobility Management (NETLMM). RFC 4831, April 2007.
- [10] 3GPP TS 29.060. GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface, 2009.
- [11] J. Kempf. Problem Statement for Network-Based Localized Mobility Management (NETLMM). RFC 4830, April 2007.
- [12] Ignacio Soto, Carlos J. Bernardos, Maria Calderon, Albert Banchs, and Arturo Azcorra. NEMO-enabled Localized Mobility Support for Internet Access in Automotive Scenarios. *IEEE Communications Magazine*, 47(5):152–159, May 2009.
- [13] Suk Yu Hui and Kai Hau Yeung. Challenges in the Migration to 4G Mobile Systems. *IEEE Communications Magazine*, 41(12):54–59, December 2003.
- [14] V. Devarapalli, N. Kant, H. Lim, and C. Vogt. Multiple Interface Support with Proxy Mobile IPv6. draft-devarapalli-netext-multi-interface-support-00.txt (work-in-progress), March 2009.
- [15] D. Thaler. Multi-Link Subnet Issues. RFC 4903, June 2007.
- [16] T. Melia, S. Gundavelli, H. Yokota, and C. Bernardos. Logical Interface Support for multi-mode IP Hosts. draft-melia-netext-logical-interface-support-01.txt (work-in-progress), July 2010.
- [17] R. Wakikawa, S. Kiriyaama, and S. Gundavelli. The applicability of virtual interface for inter-technology handoffs in Proxy Mobile IPv6. *Wireless Communications and Mobile Computing*, 2009.
- [18] H. Yokota, S. Gundavelli, T. Trung, Y. Hong, and K. Leung. Virtual Interface Support for IP Hosts. draft-yokota-netlmm-pmipv6-mn-itho-support-03.txt (work-in-progress), March 2010.
- [19] R. Braden. Requirements for Internet Hosts - Communication Layers. RFC 1122, October 1989.
- [20] D. Thaler. Evolution of the IP Model. draft-thaler-ip-model-evolution-01.txt (work-in-progress), July 2008.
- [21] M. Wasserman. Current Practices for Multiple Interface Hosts. draft-ietf-mif-current-practices-02.txt (work-in-progress), June 2010.
- [22] C. Bernardos (ed) et al. Proxy Mobile IPv6 Extensions to Support Flow Mobility. draft-bernardos-netext-pmipv6-flowmob-00.txt (work-in-progress), July 2010.
- [23] T. Melia, G. Bajko, S. Das, N. Golmie, and JC. Zuniga. IEEE 802.21 Mobility Services Framework Design (MSFD). RFC 5677, December 2009.
- [24] IEEE. IEEE Standard for Local and Metropolitan Area Networks - Part 21: Media Independent Handover Services. IEEE LAN/ MAN Std 802.21-2008, January 2009.
- [25] H. Soliman, G. Tsirtsis, N. Montavont, G. Giaretta, and K. Kuladinithi. Flow Bindings in Mobile IPv6 and NEMO Basic Support. draft-ietf-mext-flow-binding-06.txt (work-in-progress), March 2010.
- [26] Ed. H. Soliman. Mobile IPv6 Support for Dual Stack Hosts and Routers. RFC 5555, June 2009.
- [27] Qualcomm. 3G/Wi-Fi Seamless Offload (whitepaper), March 2010.
- [28] Naoki Imai, Manabu Isomura, and Akira Idoue. Coordination path control method to reduce traffic load on NGN access gateway. In *13th International Conference on Intelligence in Next Generation Networks, 2009 (ICIN 2009)*, October 2009.
- [29] T. Schmidt, M. Waehlich, and G. Fairhurst. Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey. RFC 5757, February 2010.

- [30] Imed Romdhani, Mounir Kellil, Hong-Yon Lach, Abdelmadjid Bouabdallah, and Hatem Bettahar. IP Mobile Multicast: Challenges and Solutions. *IEEE Communications Surveys & Tutorials*, 6(1), 2004.
- [31] S. Deering, W. Fenner, and B. Haberman. Multicast Listener Discovery (MLD) for IPv6. RFC 2710, October 1999.
- [32] R. Vida and L. Costa. Multicast Listener Discovery Version 2 (MLDv2) for IPv6. RFC 3810, June 2004.
- [33] B. Fenner, H. He, B. Haberman, and H. Sandick. Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding (“IGMP/MLD Proxying”). RFC 4605, August 2006.
- [34] H. Holbrook and B. Cain. Source-Specific Multicast for IP. RFC 4607, August 2006.
- [35] B. Fenner, M. Handley, H. Holbrook, and I. Kouvelas. Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised). RFC 4601, August 2006.
- [36] T.C. Schmidt, M. Waehlich, and S. Krishnan. A Minimal Deployment Option for Multicast Listeners in PMIPv6 Domains. draft-ietf-multimob-pmipv6-base-solution-05.txt (work-in-progress), July 2010.
- [37] D. von Hugo, H. Asaeda, B. Sarikaya, and P. Seite. Evaluation of further issues on Multicast Mobility: Potential future work for WG MultiMob. draft-von-hugo-multimob-future-work-02.txt (work-in-progress), June 2010.
- [38] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. Network Mobility (NEMO) Basic Support Protocol. RFC 3963, January 2005.
- [39] T. Ernst and H-Y. Lach. Network Mobility Support Terminology. RFC 4885, July 2007.
- [40] T. Ernst. Network Mobility Support Goals and Requirements. RFC 4886, July 2007.
- [41] Hong-Ke Zhang, Zhi-Wei Yan, Hua-Chun Zhou, Jian-Feng Guan, and Si-Dong Zhang. Consideration of Network Mobility in PMIPv6. IETF, draft-zhang-netlmm-nemo-01.txt (work-in-progress), March 2010.
- [42] J-C Lee and D. Kaspar. NEMO in PMIPv6 domain. draft-lee-netlmm-ps-01.txt (work-in-progress), August 2007.



Carlos J. Bernardos received a Telecommunication Engineering degree in 2003, and a PhD in Telematics in 2006, both from the University Carlos III of Madrid (UC3M), where he worked as a research and teaching assistant from 2003 to 2008 and, since then, has worked as an Associate Professor. His current work focuses on vehicular networks and IP-based mobile communication protocols. He has published over 30 scientific papers in prestigious international journals and conferences, and he is also an active contributor to the Internet Engineering Task Force (IETF).



Marco Gramaglia received a BsC and a MsC in Computer Science from the Polytechnic University of Turin (Politecnico di Torino) in 2006 and 2008 respectively, and a MsC in Telematics from the University Carlos III of Madrid (UC3M) in 2009. Currently he is a Research Assistant at Institute IMDEA Networks and a PhD candidate in Telematics at University Carlos III of Madrid.



Luis M. Contreras is a Ph.D. student at the Universidad Carlos III. He holds a Telecom Engineer degree from the Polytechnic University of Madrid (1997). He firstly joined Alcatel Spain taking several positions (RD, standardization, systems design and customer engineering) in both wireless and wired network fields. Since 2006 he is with Orange Spain being part of the Network Strategy Planning group, and serves as Peering manager for Orange Spain.



Maria Calderon is an associate professor at the Telematics Engineering Department of University Carlos III of Madrid. She received a computer science engineering degree in 1991 and a Ph.D. degree in computer science in 1996, both from the Technical University of Madrid. She has published over 40 papers in the fields of advanced communications, reliable multicast protocols, programmable networks and IPv6 mobility. Her current work focuses on vehicular networks and IP-based mobile communication protocols.



Ignacio Soto received a telecommunication engineering degree in 1993, and a Ph.D. in telecommunications in 2000, both from the University of Vigo, Spain. He was a research and teaching assistant in telematics engineering at the University of Valladolid from 1993 to 1999. In 1999 he joined University Carlos III of Madrid, where he has been an associate professor since 2001. His research activities focus on mobility support in packet networks and heterogeneous wireless access networks.