Internet Engineering Task Force (IETF)
Request for Comments: 7161
Category: Experimental
ISSN: 2070-1721
March 2014

LM. Contreras
Telefonica I+D
CJ. Bernardos
I. Soto
UC3M

Proxy Mobile IPv6 (PMIPv6) Multicast Handover Optimization by the Subscription Information Acquisition through the LMA (SIAL)

Abstract

This document specifies an experimental multicast handover optimization mechanism for Proxy Mobile IPv6 (PMIPv6) to accelerate the delivery of multicast traffic to mobile nodes after handovers. The mechanism, called Subscription Information Acquisition through the LMA (SIAL), is based on speeding up the acquisition of mobile nodes' multicast context by the mobile access gateways. To do that, extensions to the current PMIPv6 protocol are proposed. These extensions are not only applicable to the base solution for multicast support in Proxy Mobile IPv6, but they can also be applied to other solutions developed to avoid the tunnel convergence problem. Furthermore, these extensions are also independent of the role played by the mobile access gateway within the multicast network (acting as either multicast listener discovery proxy or multicast router).

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 3
1.1. Handover Optimization Requirements 5
2. Terminology 6
3. Overview 7
4. Proxy Mobile IPv6 Extensions 8
4.1. Active Multicast Subscription Mobility Option 8
4.1.1. Option Application Rules 8
4.1.2. Option Format 9
4.1.3. Backward Compatibility with MLDv1 9
4.2. Multicast Signaling Flag on PBU/PBA Message Headers . . . 10
4.2.1. Flag Application Rules 10
4.2.1.1. Registration Process 11
4.2.1.2. De-registration Process 12
4.2.2. New Format of Conventional PBU/PBA Messages 12
4.2.2.1. Proxy Binding Update Message 12
4.2.2.2. Proxy Binding Acknowledgement Message 13
4.3. Messages for Active Multicast Subscription Query 13
4.3.1. Subscription Query Message 13
4.3.1.1. Message Application Rules 13
4.3.1.2. Message Format 14
4.3.2. Subscription Response Message 15
4.3.2.1. Message Application Rules 15
4.3.2.2. Message Format 15
4.4. New PBA Timer in the LMA 16
5. Handover Signaling Procedures 17
5.1. Handover of Proactive Type 17
5.1.1. Rationale 17
5.1.2. Message Flow Description 18
5.2. Handover of Reactive Type 20

5.2.1. Rationale 20
5.2.2. Message Flow Description 21
5.2.3. Further Considerations for the Reactive Handover Signaling 22
5.3. Prevention of Large Delays of the Binding Acknowledgement for Unicast Traffic 23
6. IPv4 Support 26
6.1. Active Multicast Subscription for IPv4 26
6.2. Signaling Procedures for IPv4 Support 27
6.3. Binding Cache Extensions for IPv4 Support 28
7. Coexistence with PMIPv6 Multicast Architectural Evolutions 28
8. Security Considerations 28
9. IANA Considerations 31
10. Contributors 31
11. Acknowledgments 31
12. References 32
12.1. Normative References 32
12.2. Informative References 32
Appendix A. Performance Comparison with Base Solution 34
A.1. Delay Characterization of the Base Solution 34
A.2. Delay Characterization of SIAL 35
A.3. Performance Comparison 35

1. Introduction

The base solution for providing continuous multicast service delivery in Proxy Mobile IPv6 (PMIPv6) domains is described in [RFC6224]. It specifies the basic functionality needed in the Proxy Mobile IPv6 [RFC5213] entities to provide a multicast service, so continuous delivery of multicast traffic is supported by obtaining, after each handover, the ongoing multicast subscription information directly from the Mobile Node (MN). When a mobile node attaches to a new Mobile Access Gateway (MAG), the mobile node is queried by the mobile access gateway through a Multicast Listener Discovery (MLD) General Query, which is sent just after any new link is set up, to learn of any existing subscription, as specified in [RFC2710] and [RFC3810].

However, the base solution needs to be improved to meet some performance requirements, especially those referring to the user-perceived service quality, which is seriously affected by the disruption of multicast content forwarding to the mobile node during handovers.

A mobile node with an active multicast subscription, moving from one point of attachment to another within a Proxy Mobile IPv6 domain, experiences a certain delay until it resumes receiving again the multicast content that it was receiving at the previous location.

Such delay causes a gap in the content reception. Two different actions can help mitigate such reception gap. One of them is to buffer at the previous mobile access gateway a copy of the multicast traffic destined to the mobile node and forward it to the new mobile access gateway, in order to deliver that traffic to the mobile node. The other possible (complementary) action is to reduce the time needed by the new mobile access gateway to learn of the active multicast subscription of the mobile node (i.e., the multicast context), so the new mobile access gateway can subscribe to the multicast group(s) on behalf of the mobile node as soon as possible.

While the first mechanism could potentially be accomplished by using some adaptation of [RFC5949] to multicast traffic (despite being only applicable in the case the underlying radio access technology supports Layer 2 (L2) triggers, thus requiring additional support on the mobile node), there is no generic standard solution for the accelerated acquisition of the ongoing multicast subscription of the mobile node.

The approach followed by the base solution [RFC6224] to learn of an existing multicast subscription relies on the behavior of the IGMP/MLD protocols. Both protocols send multicast membership query messages when a new link is up. The response to such a message reports any existing multicast subscriptions by the mobile node. While this is a straightforward approach, the mobile access gateway can incur in a non-negligible delay in receiving the corresponding MLD Report message. This delay is caused by the time needed for the detection of the attachment in the new link and the re-establishment of the data plane after the handover, the radio transfer delays associated with the signaling to the mobile node, and the MLD query response interval time required by this procedure (whose default value is 10 seconds as defined in [RFC2710] and [RFC3810], or between 5 and 10 seconds as considered in the best case wireless link scenario in [RFC6636]).

This document extends the Proxy Mobile IPv6 signaling protocol defined in the base protocol [RFC5213] by including a new multicast information option to update Proxy Mobile IPv6 entities during the registration and de-registration processes, and new messages to trigger the transfer of multicast information. No extension is required in any of the multicast-related protocols in use (IGMP/MLD or PIM protocols). Furthermore, this specification does not substitute the standard procedures defined in [RFC6224] (e.g., the mobile access gateway continues sending an MLD Query to the entering mobile node as soon as the point-to-point link is set up), but complements them for accelerating the acquisition of the multicast content by the mobile access gateway associated to the new point-of-attachment.

This document provides a signaling method internal to the network to speed up the subscription information acquisition by the mobile access gateway, in order to accelerate the multicast delivery to the mobile node after having completed a handover. By doing so, the knowledge by the mobile access gateway of the currently active multicast subscription becomes independent of the underlying radio technology dynamics and relaxes the requirement of a rapid response from the mobile node in processing IGMP/MLD control messages. Issues like radio framing, radio access contention, channel reliability, MN’s capabilities (i.e., L2 triggering support), IGMP/MLD timers optimization for wireless environments, etc., will not impact the observed multicast performance during handovers.

The mechanisms described in this document can also be applied to the solutions defined in [RFC7028]. Furthermore, it is also independent of the role played by the mobile access gateway within the multicast network (acting as either MLD proxy or multicast router).

1.1. Handover Optimization Requirements

A basic solution for providing support of multicast in a network-based mobility management environment has been specified in [RFC6224] without introducing changes on the original PMIPv6 specification [RFC5213]. The focus of the present document is on improving the efficiency of the base solution regarding handover performance.

One of the critical aspects of the base solution is the expected delay incurred by the mobile access gateway (where the mobile node is being attached to) to be informed about the ongoing multicast subscription of the entering MN, mainly due to the fact that the mechanisms provided in the base solution relay on the original MLD procedures, with long timing interactions not conceived for mobile environments. Then, the requirements to be covered by a handover optimization solution can be established in the following manner:

- o The solution MUST be applicable to any kind of MN (that is, not requiring any particular functionality such as, for example, L2 trigger capabilities), in such a way that any type of mobile node in a PMIPv6 domain being served with multicast traffic can benefit from the optimized solution.

- o The solution MUST NOT impact existing multicast protocols.

- o The solution MUST optimize the handover performance with respect to the performance achieved with the base solution for any kind of handover process (i.e., for proactive and reactive handovers).

- o The solution SHOULD minimize the number and extent of additional support (i.e., capabilities) required in the network, aiming at an easier deployment.

- o The solution MUST NOT impact deployments of legacy implementations of [RFC5213] and [RFC6224].

The present specification addresses all these requirements, as described in the following sections.

2. Terminology

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 [RFC2119].

This document uses the terminology referring to PMIPv6 components as defined in [RFC5213].

Additionally, the following terms are defined and used in this document.

pMAG: The previous MAG or pMAG is the mobile access gateway where the MN was initially registered before a handover event.

nMAG: The new MAG or nMAG is the mobile access gateway where the MN is registered at the end of the handover event.

Reactive Handover: A reactive handover is a handover event in which the Local Mobility Anchor (LMA) receives the mobile node registration from the nMAG without having previously received the MN de-registration from the pMAG.

Proactive Handover: A proactive handover is a handover event where the mobile node is firstly de-registered on the local mobility anchor by the pMAG, and later on it is registered by the nMAG as a consequence of changing the point of attachment.

Multicast Membership Context: In this document, multicast membership context makes reference to the information relative to the currently active multicast subscription of an MN in a handover event that is transferred between the PMIPv6 entities to support the handover optimization.

3. Overview

The local mobility anchor is a key element within the PMIPv6 infrastructure, which traces the mobile node reachability along the PMIPv6 domain. Therefore, the LMA is the best element to keep the MNs' multicast subscription information up-to-date and to forward it to the rest of PMIPv6 entities (i.e., to the mobile access gateways) as needed when MNs move within the domain. The LMA has timely knowledge of the MNs' locations, especially during handover events, and it is therefore able to quickly provide information to the new point of attachment (e.g., by querying the previous one). Figure 1 summarizes the main idea of the optimization.

```
                                    +------+
                                    | pMAG |   |
                                    +------+   |
                                   /           |
                                  /            |
                                 /             |
                                /              |
            -*-*-*-*-          /              (MN)
           (         )        /                |
          (           )   +-----+      +------+   |
         (  Internet   )--| LMA |------| nMAG |   v
          (           )   +-----+      +------+
           (         )
            -*-*-*-*-         Registration
                            <--------------

                           Registration Ack
                         & Multicast Context
                            -------------->
```

Figure 1: High-Level Description of the Solution

The local mobility anchor only obtains the detailed subscription information or multicast context during a handover event. There is no need for continuously informing the LMA about MNs' multicast state while the mobile nodes remain attached to the same mobile access gateway. Such a continuous updating procedure would significantly increase the signaling load within the PMIPv6 domain without a clear benefit. The multicast context is only critical during handovers: neither after nor before. Indicating the active subscription while the handover is ongoing guarantees that such information will be up to date and ready to be transferred to the new MAG where the mobile node has just attached. Therefore, this solution defines the Subscription Information Acquisition through the LMA (SIAL) as the procedure to inform the new MAG about the multicast subscriptions maintained by the entering MN.

To be able to transfer the multicast subscription information between PMIPv6 entities during a handover, this document extends the PMIPv6 protocol in several ways. First of all, a new mobility option is defined to carry the multicast context of the current subscription. Furthermore, additional messages are defined to manage the interchange of the multicast information among PMIPv6 entities. Finally, some flags are defined to govern the process.

4. Proxy Mobile IPv6 Extensions

This section outlines the extensions proposed to the PMIPv6 protocol specified in [RFC5213].

4.1. Active Multicast Subscription Mobility Option

4.1.1. Option Application Rules

A new TLV-encoded mobility option, Active Multicast Subscription option is defined for use with the Proxy Binding Update (PBU) and Proxy Binding Acknowledgement (PBA) messages exchanged between a local mobility anchor and a mobility access gateway to transfer the multicast subscription information. This option is used for exchanging the multicast membership context. This information is carried by directly using the format defined in the original MLD specifications. There can be multiple Active Multicast Subscription options present in the message, one for each active subscription maintained by the mobile node when the handover is taking place (i.e., one per multicast membership context).

This new option is also used for the same purposes by the new Subscription Response message defined later in this document.

MLDv2 [RFC3810] is the primary objective for the definition of the option format. MLDv1 [RFC2710] is also considered for backward compatibility.

4.1.2. Option Format

The format of this new option is as follows:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                    |      Type     |     Length    |    MLD Type   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    +                  Multicast Membership Context                 +
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The alignment requirement of this option is 8n+1.

Type:

57, which indicates the Active Multicast Subscription IPv6 option.

Length:

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

MLD type:

Field used to identify the IPv6 multicast membership protocol in use, and the corresponding format of the next Multicast Membership Context information field. This field maps the type codification used in the original MLD specifications for the Report message. For MLDv2, the MLD Type value is 143, as specified in [RFC3810].

Multicast Membership Context:

Multicast subscription information corresponding to a single subscribed multicast address. For MLDv2, the format of this field follows the Multicast Address Record format as defined in [RFC3810].

4.1.3. Backward Compatibility with MLDv1

The following values are adopted when MLDv1 is used.

MLD type:

For MLDv1, the MLD Type value is 131, as specified in [RFC2710].

Multicast Membership Context:

For MLDv1, the relevant information for multicast context is simply given, according to [RFC2710], by the multicast address of the subscribed content.

In consequence, the Multicast Membership Context is defined as a 4-octet reserved field and the Multicast Address of the subscribed content as in [RFC2710], as shown next.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Reserved                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
*                                                               *
|                                                               |
*                       Multicast Address                       *
|                                                               |
*                                                               *
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

4.2. Multicast Signaling Flag on PBU/PBA Message Headers

4.2.1. Flag Application Rules

A new flag S has been added in both the PBU and PBA message headers to advertise the mobile access gateway and the local mobility anchor capabilities of processing multicast-related signaling for the MN that caused the message.

This flag governs the multicast-related signaling between the LMA and the MAG. As a general rule, the value of the flag in the PBA message is a copy of the value received in the PBU message. Specific rules are described in next subsections.

4.2.1.1. Registration Process

During handover, the entities involved in this process are the nMAG and the LMA. These rules also apply for the initial binding registration process.

- PBU message

  * S=0 indicates that the MAG sending the PBU message does not accept multicast-related signaling for the MN being attached. This can be used to discriminate PMIPv6 nodes that are not multicast enabled, for backward compatibility reasons.

  * S=1 indicates that the MAG sending the PBU message accepts multicast-related signaling for the MN being attached. Depending on the type of handover (reactive or proactive) the LMA takes some actions, described later in this document.

- PBA message

  * If S=0 in the corresponding PBU message, the value of the flag in the PBA message MUST be a copy of the value received in the PBU message (thus S=0), without any further meaning.

  * If S=1 in the corresponding PBU message, two subcases are possible:

    + S=1 and Active Multicast Subscription mobility option in the PBA message. When the MN maintains an active multicast session, if the LMA is able to provide the multicast subscription information during registration, the PBA message MUST include the Active Multicast Subscription mobility option. If the LMA is not able to provide such information during registration, the PBA message MUST NOT include the Active Multicast Subscription mobility option. This case is useful to decouple unicast and multicast signaling for an MN being registered at nMAG. A way for obtaining later active multicast-subscription information is described later in this document.

    + S=0 in the PBA message if the MN does not maintain an active multicast subscription (note that for backward compatibility reasons, an LMA not supporting multicast related signaling would always send S=0).

4.2.1.2. De-registration Process

During handover, the entities involved in this process are the pMAG and the LMA. These rules apply for the binding de-registration process.

- o PBU message

  * S=0 indicates that the MN has no active multicast session (note that for backward compatibility reasons, a pMAG not supporting multicast related signaling would always send S=0).

  * S=1 indicates that the MN has an active multicast session, and the multicast context MUST be transported in the Active Multicast Subscription mobility option.

- o PBA message

  * The value of the flag in the PBA message SHOULD be 0, without any further meaning (note that for backward compatibility reasons, an LMA not supporting multicast related signaling would always send S=0).

4.2.2. New Format of Conventional PBU/PBA Messages

4.2.2.1. Proxy Binding Update Message

As result of the new defined flag, the PBU message format is updated as follows:

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                 |            Sequence #         |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |A|H|L|K|M|R|P|S|   Reserved    |            Lifetime           |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 .                                                               .
 .                        Mobility Options                       .
 .                                                               .
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

4.2.2.2. Proxy Binding Acknowledgement Message

As result of the new defined flag, the PBA message format is updated as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |    Status     |K|R|P|S| Rsrvd |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Sequence #          |           Lifetime            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                        Mobility Options                       .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

4.3. Messages for Active Multicast Subscription Query

A new pair of messages is defined for querying entities about the active multicast subscription of the MN when the handover is of reactive type.

These messages are sent using the Mobility Header as defined in [RFC6275].

4.3.1. Subscription Query Message

4.3.1.1. Message Application Rules

The Subscription Query message (value 22) is sent by the LMA towards the pMAG to query it about any existing multicast subscriptions of the MN that is being registered by the nMAG. This message is generated in case the handover is of reactive type.

Additionally, this message is sent by the nMAG towards the LMA to query it about the existing multicast subscriptions of the MN when the LMA acknowledges the PBU sent by the nMAG but the multicast context is not provided (namely, when the PBU message has set the flag S to 1, and the PBA message has set the flag S to 1 but the multicast context is missing).

4.3.1.2. Message Format

The Subscription Query message has the following format.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                               |  Sequence #   |   Reserved    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                        Mobility Options                       .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Sequence Number:

The Sequence Number field establishes the order of the messages sent in the Subscription Query / Subscription Response dialogue between the LMA and the MAG for a certain MN. The initial Sequence Number MUST be determined by the entity that creates the message (either LMA or MAG, depending on the scenario), which is responsible for managing this counter.

This Sequence Number comparison MUST be performed modulo 2**8, i.e., the number is a free-running counter represented modulo 256. A Sequence Number in a received Subscription Query message is considered less than or equal to the last received number if its value lies in the range of the last received number and the preceding 128 values, inclusive. For example, if the last received sequence number was 15, then messages with sequence numbers 0 through 15, as well as 143 through 255, would be considered less than or equal.

Reserved:

This field is unused for now. The value MUST be initialized to 0.

Mobility options:

This message carries one or more TLV-encoded mobility options. The valid mobility options for this message are the following:

* Mobile Node Identifier option [RFC4283] (mandatory).

* Home Network Prefix option [RFC5213] (optional).

There can be one or more instances of the Home Network Prefix option, but only one instance of the Mobile Node Identifier option.

4.3.2. Subscription Response Message

4.3.2.1. Message Application Rules

The Subscription Response message (value 23) is sent by the pMAG towards the LMA, or by the LMA towards the nMAG, to answer a previously received Subscription Query message, as described above.

4.3.2.2. Message Format

The Subscription Response message has the following format.

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                    |  Sequence #   |I|  Reserved   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    .                                                               .
    .                        Mobility Options                       .
    .                                                               .
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Sequence Number:

The value of the Sequence Number field in the Subscriber Response message MUST be a copy of the Sequence Number received in the Subscription Query message.

Multicast Information (I):

The multicast Information flag I specifies whether or not there is multicast subscription information available for the MN. The meaning is the following:

I=0: there is no multicast subscription information available for the MN identified by the Mobile Node Identifier option in this message.

I=1: there is multicast subscription information available for the MN identified by the Mobile Node Identifier option in this message. The multicast subscription information MUST be carried on one or more instances of the Active Multicast Subscription option in this message (one instance for each active subscription).

Reserved:

This field is unused for now. The value MUST be initialized to 0.

Mobility options:

This message carries one or more TLV-encoded mobility options. The valid mobility options for this message are the following:

* Mobile Node Identifier option [RFC4283] (mandatory).

* Active Multicast Subscription option (mandatory) only when flag I=1; it MUST NOT be present in any other case.

* Home Network Prefix option [RFC5213] (optional).

There can be one or more instances of the Home Network Prefix option (in all cases) and the Active Multicast Subscription option (only when I=1), but only one instance of the Mobile Node Identifier option.

4.4. New PBA Timer in the LMA

A new timer named "PBA timer" is used in the LMA to define the maximum waiting time before the PBA message is sent to the nMAG in case the multicast subscription information relative to the MN is not yet available. The aim of this timer is to prevent potential large delays in the forwarding of unicast traffic towards the MN being registered at the nMAG. This timer allows decoupling the unicast signaling from the multicast one in the SIAL solution.

This timer SHOULD be upper bounded by the constant defined in [RFC6275] INITIAL_BINDACK_TIMEOUT, whose default value is 1 s. This constant sets the time when the nMAG will retry the MN registration by sending again the PBU message. The "PBA timer" has to be set to a value that ensures that the nMAG does not enter the retry mode. Operational experience is needed on how to set up the PBA timer, and therefore it is RECOMMENDED to set the "PBA timer" to zero, except for experimental purposes.

5. Handover Signaling Procedures

As the MN moves from one access gateway to another, the mobility-related signaling due to the handover event is carried out independently by the pMAG and the nMAG. That signaling process is not synchronized; thus, two scenarios need to be considered depending on the order in which the LMA receives notification of the MN registration and de-registration in the nMAG and the pMAG, respectively.

5.1. Handover of Proactive Type

5.1.1. Rationale

In the proactive case, the MN is firstly de-registered by the pMAG, and later on it is registered by the nMAG as a consequence of changing the point of attachment.

Only for those MNs that maintain an active multicast subscription, the pMAG includes the Active Multicast Subscription mobility option carrying the multicast context of the MN at that moment as part of the PBU message (with flag S set to 1).

The local mobility anchor stores that information in the corresponding binding cache. If later on the MN attaches to an nMAG, this information is sent (using the same TLV option) to the nMAG as part of the PBA confirmation of the registration process (if the PBU message sent by the nMAG has the flag S set to 1). On the other hand, if no further registration happens, the multicast information is removed together with the rest of binding database for that MN.

After receiving the multicast context, the nMAG can subscribe to the multicast flow(s) on behalf of the MN in case there is no other MN already receiving it at the nMAG. The multicast status can also be set in advance for the point-to-point link towards the MN.

Note that the SIAL solution described here does not prevent benefiting from extended support in the mobile node / network that facilitates the proactive mode operation of the solution, e.g., based on L2 capabilities.

5.1.2. Message Flow Description

Figure 2 summarizes this process.

```
          +-----+          +----+           +-----+          +----+
          | MN  |          |pMAG|           | LMA |          |nMAG|
          +-----+          +----+           +-----+          +----+
             |                |                |                |
             |                |==Bi-Dir Tunnel=|                |
             | Multicast Data |                |                |
             |<---------------|                |                |
             |                |                |                |
      1) MN Detached          |                |                |
             |         MN Detached Event       |                |
             |                |                |                |
             |                |Ext'd DeReg PBU |                |
      2)     |                |--------------->|                |
             |                |                |                |
      3)     |                |            Accept PBU           |
             |                |(Multicast Subscription info stored)
             |                |                |                |
             |                |      PBA       |                |
      4)     |                |<---------------|                |
             |                |                |                |
      5) MN Attached          |                |                |
             |                |                |   MN Attached Event
             |                |                |                |
             |                |                |       PBU      |
      6)     |                |                |<---------------|
             |                |                |                |
             |                |                |   Ext'd PBA    |
      7)     |                |                |--------------->|
             |                |                |                |
      8)     |                |                |          Accept PBA,
             |                |                |   Multicast Group join
             |                |                | and P-t-P status setup
             |                |                |                |
             |                |                |==Bi-Dir Tunnel=|
             |                |                |                |
             |                |                | Multicast Data |
             |<-------------------------------------------------|
             |                |                |                |
             |                |                |                |
```

Figure 2: Proactive Handover

The message flow is as follows:

1. A registered MN is receiving a multicast content that has been previously subscribed to by sending a standard MLD report from the mobile node to the currently serving mobile access gateway, pMAG. The pMAG keeps the multicast state of the point-to-point link with the MN.

2. The MN initiates a handover process (e.g., because of better radio conditions) over a radio access controlled by a new MAG. As a consequence, pMAG determines a detachment event corresponding to this mobile node, and updates the attachment status of this MN to the local mobility anchor by sending an extended Proxy Binding Update message, including the Active Multicast Subscription, which contains the multicast context of the active multicast subscriptions in the moment of handover.

3. The LMA processes the PBU message. Additionally, the LMA stores in the binding cache the information regarding the ongoing multicast subscription(s) when the detachment is initiated. This information is kept until a new registration of the MN is completed by another MAG, or until the binding cache expiration, according to [RFC5213].

4. The local mobility anchor acknowledges to the pMAG the previous PBU message.

5. As a result of the handover process, the mobile node attaches to another mobility access gateway, called nMAG.

6. The nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the local mobility anchor.

7. After the analysis of the PBU message, the LMA sends an extended PBA including the Active Multicast Subscription option, which contains the multicast context of the active subscriptions in the moment of handover.

8. The nMAG processes the PBA message following all the standard procedures described in [RFC5213]. Additionally, with the new information relative to multicast subscription, the nMAG sets up the multicast status of the point-to-point link between the nMAG and the MN, and joins the content identified by (S,G) on behalf of the MN in case the nMAG is not receiving already such content due to a previous subscription ordered by another MN attached to it. From that instant, the multicast content is served to the MN.

5.2. Handover of Reactive Type

5.2.1. Rationale

In the reactive case, the LMA receives the mobile node registration from the nMAG without having previously received the MN de-registration from the pMAG.

As the nMAG is not aware of any active multicast subscription of the mobile node, the nMAG starts a conventional registration process, by sending a normal PBU message (with flag S set to 1) towards the local mobility anchor.

In the reactive handover case, after MN registration at the nMAG, the local mobility anchor SHOULD generically query the pMAG to retrieve the multicast context of the ongoing multicast subscription of the mobile node. However, the LMA may know in advance if the pMAG supports multicast signaling based on the value of the flag S received during the MN registration in pMAG. Specifically, in case the pMAG does not support multicast signaling (e.g., the S flag value received from pMAG at the time of registering the mobile node was 0), the LMA MAY decide not to query pMAG even in the case of receiving an nMAG indication of supporting multicast signaling.

Once the multicast subscription information is retrieved from the pMAG, the LMA encapsulates it in the PBA message by using the TLV option Active Multicast Subscription and forwards the PBA message to the nMAG. Then, the nMAG can subscribe the multicast flow on behalf of the MN, if there is no other mobile node receiving it already at the nMAG. The multicast status can be also set in advance for the point-to-point link towards the mobile node.

5.2.2. Message Flow Description

Figure 3 summarizes this process.

```
       +-----+          +----+           +-----+          +----+
       | MN  |          |pMAG|           | LMA |          |nMAG|
       +-----+          +----+           +-----+          +----+
          |                |                |                |
          |                |                |         MN Attached Event
          |                |                |                |
          |                |                |       PBU      |
   1)     |                |                |<---------------|
          |                |                |                |
          |                |  Subscr Query  |                |
   2)     |                |<---------------|                |
          |                |                |                |
          |                |  Subscr Resp   |                |
   3)     |                |--------------->|                |
          |                |                |                |
          |                |    (Multicast Subscription      |
          |                |        info forwarding)         |
          |                |                |                |
          |                |                |   Ext’d PBA    |
   4)     |                |                |--------------->|
          |                |                |                |
   5)     |                |                |           Accept PBA,
          |                |                |      Multicast Group join
          |                |                |     and P-t-P status setup
          |                |                |                |
          |                |                |==Bi-Dir Tunnel=|
          |                |                |                |
          |                |                |   (S,G) Data   |
          |<-----------------------------------------------|
          |                |                |                |
          |                |                |                |
```

Figure 3: Reactive Handover

We next take as starting point the situation where an MN is attached to the pMAG, being multicast enabled and maintaining an active multicast subscription at this moment.

The sequence of messages for the handover of the mobile node is the following (as depicted in Figure 3):

1. At a certain time, the MN initiates a handover process (e.g., because of better radio conditions) over a radio access controlled by a new MAG. Then, the nMAG triggers a registration

process by sending a PBU message (with flag S set to 1) to the local mobility anchor. As it is a reactive case, the pMAG is not aware of the detachment process.

2. Prior to acknowledging the received PBU message, the LMA queries the pMAG about if there is any active multicast subscription for the MN, by sending a Subscription Query message.

3. The pMAG answers the LMA with a Subscription Response message including the multicast context of the existing subscriptions.

4. After processing the pMAG answer, the LMA acknowledges (with flag S set to 1) the PBU message, including the multicast subscription information within the Active Multicast Subscription option. The nMAG then processes the extended PBA message.

5. The nMAG processes the PBA message, and it proceeds to set up the multicast status of the point-to-point link between the nMAG and the mobile node, and to join the content identified by (S,G) on behalf of the MN in case the nMAG is not receiving already such content. The bidirectional tunnel is also set up between the nMAG and the local mobility anchor if it has not been established before by another MN connection. At this moment, the multicast content can be served to the MN. The unicast traffic for the mobile node can be forwarded as well.

5.2.3. Further Considerations for the Reactive Handover Signaling

A handover event is managed independently by the pMAG and nMAG. It is not a synchronized process. In a reactive handover, the LMA receives a registration PBU from nMAG before a de-registration PBU is received from pMAG.

In the message flows detailed above, it could be the case that the LMA receives a de-registration PBU from pMAG just after sending the Subscription Query message, but before receiving the Subscription Response message. That de-registration PBU message from pMAG carries the multicast subscription information required to assist the MN in the handover, so such valuable information SHOULD be kept by the LMA. Furthermore, it is possible that once the Subscription Query message arrives to pMAG, the pMAG could have already removed the multicast related information for the MN.

In order to avoid losing the multicast subscription information sent in the de-registration PBU message, the local mobility anchor SHOULD store it, and SHOULD include it in the PBA message towards the nMAG in case the Subscription Response message from the pMAG does not contain multicast subscription information for the mobile node.

5.3. Prevention of Large Delays of the Binding Acknowledgement for Unicast Traffic

According to the message sequences described for the reactive handover case, in case the LMA has to request the multicast subscription information from the pMAG, the binding request sent by the nMAG is maintained on-hold until the local mobility anchor receives, processes and includes the multicast subscription information into the extended PBA message. As a consequence, the unicast traffic may then suffer an extra delay motivated by the multicast-related signaling. During that time, the unicast traffic with destination the MN being registered by the nMAG MAY be buffered by the local mobility anchor.

In order to avoid any potential large delay in the forwarding of unicast traffic arriving at the LMA towards the MN, a mechanism SHOULD be implemented to decouple multicast from unicast traffic reception by the MN. Figure 4 shows this mechanism.

```
       +-----+          +----+           +-----+          +----+
       | MN  |          |pMAG|           | LMA |          |nMAG|
       +-----+          +----+           +-----+          +----+
   1)     |                |==Bi-Dir Tunnel=|                |
          |  unicast data  |                |                |
          |<-v-v-v-v-v-v-v-|                |                |
          |                |                |                |
          | Multicast Data |                |                |
          |<---------------|                |                |
          |                |                |        MN Attached Event
          |                |                |       PBU      |
   2)     |                |                |<---------------|
          |                |  Subscr Query  |                |
   3)     |                |<---------------|                |
          |                |                |                |
   4)     |                |       <PBA timer starts>        |
          |                |               ///               |
          |                |               ///               |
   5)     |                |       <PBA timer expires>       |
          |                |                |                |
          |                |                |   Ext’d PBA    |
          |                |                |--------------->|
          |                |                |                |
          |                |                |          Accept PBA
          |                |                |                |
          |                |                |==Bi-Dir Tunnel=|
          |                |                |                |
          |                |                |  Unicast Data  |
          |<-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-v-|
          |                |                |                |
          |                |                |  Subscr Query  |
   6)     |                |                |<---------------|
          |                |  Subscr Resp   |                |
   7)     |                |--------------->|                |
          |                |                |                |
          |                |    (Multicast Subscription      |
          |                |        info forwarding)         |
          |                |                |                |
          |                |                |  Subscr Resp   |
   8)     |                |                |--------------->|
          |                |                |                |
          |                |                |   Multicast Group join
          |                |                | and P-t-P status setup
          |                | Multicast Data |                |
          |<-------------------------------------------------|
          |                |                |                |
```

          Figure 4: Decoupling of Unicast and Multicast Signaling

The sequence of messages is the following:

1. An MN is attached to the pMAG. The MN is a multicast-enabled node, and it is receiving both unicast and multicast traffic simultaneously.

2. Some time later, The MN initiates a handover process (e.g., because of better radio conditions) over a radio access controlled by a new mobile access gateway. Then, the nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the local mobility anchor. As it is a reactive case, the pMAG is not aware of the detachment process.

3. Prior to acknowledging the received PBU message, the LMA decides to query the pMAG about if there is any active multicast subscription for the mobile node, by sending a Subscription Query message.

4. Immediately after sending the Subscription Query message, the LMA starts the timer "PBA timer", which determines the maximum waiting time before the PBA is sent to avoid any potential large delay in the forwarding of unicast traffic towards the MN.

5. In case the "PBA timer" expires, the LMA acknowledges the PBU message, by sending the PBA message with flag S=1, without the multicast context information. The nMAG then processes the extended PBA message. Such acknowledgement allows the mobile node to receive the unicast traffic from that time on. The bidirectional tunnel is also set up between the nMAG and the LMA if it has not been established before.

6. In parallel, the nMAG sends a Subscription Query message to the LMA requesting the multicast-subscription details yet unknown for the mobile node.

7. The pMAG answers the Subscription Query message originally sent by the local mobility anchor, including the multicast context.

8. After processing the pMAG answer, the LMA sends a Subscription Response message to the nMAG, including the multicast subscription information within the Active Multicast Subscription option. The nMAG processes the PBA message, and it proceeds to set up the multicast status of the point-to-point link between the nMAG and the mobile node, and to join the content identified by (S,G) on behalf of the MN in case the nMAG is not receiving already such content. The bidirectional tunnel is also set up

between the nMAG and the LMA if it has not been established before. At this moment, the multicast content can also be served to the mobile node.

The "PBA timer" in the LMA determines if the signaling flow follows Figure 3 or Figure 4 in a reactive handover. A value of 0 for the "PBA timer" guarantees that the unicast traffic does not suffer any delay (according to the Figure 4 signaling flow), because the PBA is sent immediately after the LMA receives the PBU from the nMAG. A small non-zero "PBA timer" value MAY be used to reduce the signaling load in the LMA and MAGs (as shown in the signaling flow of Figure 3 if the Subscription Response message from the pMAG is received at the LMA before the "PBA timer" expires), but this has to be carefully balanced against added delay to the unicast traffic.

6. IPv4 Support

IPv4-based mobile nodes (being either IPv4/IPv6 dual-stack or IPv4-only enabled) can be supported in a PMIPv6 domain according to [RFC5844]. When referring to multicast membership protocols and procedures, this means that IGMP functionality has to be also supported between the PMIPv6 entities, as documented in [RFC6224], to allow the mobile access gateway requesting multicast contents to the mobility anchor on behalf of the mobile nodes attached to it.

6.1. Active Multicast Subscription for IPv4

The Active Multicast Subscription option defined in Section 4.1, which transports the multicast membership context of the mobile node during handover, should be compatible with IGMP-based formats. Specifically, the option format is defined for IPv4-based MNs as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
               |      Type     |    Length     |  IGMP Type    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                  Multicast Membership Context                 +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

IGMPv3 is the primary objective for the definition of the option format. IGMPv1 and IGMPv2 are also considered for backward compatibility. The alignment requirement of this option is 4n+1.

Type:

56, which indicates the Active Multicast Subscription IPv4 option.

Length:

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

IGMP type:

Field used to identify the IPv4 multicast membership protocol in use, and the corresponding format of the next Multicast Membership Context information field. This field maps the type codification used in the original IGMP specifications for the Report message.

0x12: Use of IGMPv1 multicast membership protocol.

0x16: Use of IGMPv2 multicast membership protocol.

0x22: Use of IGMPv3 multicast membership protocol.

Multicast Membership Context:

Multicast subscription information corresponding to a single subscribed multicast address. Depending on the IGMP version being used by the mobile node, the format of the Multicast Membership Context could follow the following formats:

* For IGMPv1, the Group Address format as defined in [RFC1112].

* For IGMPv2, the Group Address format as defined in [RFC2236].

* For IGMPv3, the Group Record format as defined in [RFC3376].

6.2. Signaling Procedures for IPv4 Support

Generic signaling procedures for the support of IPv4 in PMIPv6 domains have been already specified in [RFC5844]. In order to prevent errors while signaling the ongoing multicast subscription for a mobile node during the handover process, the following extensions have to be considered in SIAL.

o If the registration/de-registration process in a handover is for an IPv6-only MN, and the type of the received Active Multicast Subscription option indicates IPv4, then the multicast membership context received MUST be silently discarded.

- o If the registration/de-registration process in a handover is for an IPv4-only MN, and the type of the received Active Multicast Subscription option indicates IPv6, then the multicast membership context received MUST be silently discarded.

- o If the registration/de-registration process in a handover is for a dual stack MN, the received Active Multicast Subscription option (or options) MUST be accepted independently of the type indication.

6.3. Binding Cache Extensions for IPv4 Support

Additionally, since the multicast membership information is temporally stored in the mobility anchor under some circumstances (e.g., proactive handover), the binding cache entry for an IPv4-based multicast-enabled MN should be extended for storing the IGMP-based context formats mentioned above, including the IGMP version indicator.

7. Coexistence with PMIPv6 Multicast Architectural Evolutions

Throughout this document, the base solution for multicast support in Proxy Mobile IPv6, described in [RFC6224], has been implicitly considered, i.e., both unicast and multicast traffic addressing a mobile node is delivered via the standard PMIPv6 bidirectional tunnel between LMA and MAG. While here all multicast traffic is assumed to be delivered via the local mobility anchor, the SIAL approach described in this document can be also applied to other solutions in which the multicast content is served from other entities in the PMIPv6 domain, as described in [RFC7028] to solve the tunnel convergence problem.

In this case, the transfer of the multicast context would also pass through the local mobility anchor, as described here. However, the nMAG subscribes to the multicast content through the node in charge of distributing multicast according to the adopted solution for multicast distribution in the PMIPv6 domain.

8. Security Considerations

This proposal does not pose any additional security threats to those already identified in [RFC5213]. All the security considerations in [RFC5213] are directly applicable to this protocol. The signaling messages, Proxy Binding Update, and Proxy Binding Acknowledgement (extended with the new options defined in this document), the Subscription Query Message, and the Subscription Response Message exchanged between the mobile access gateway and the local mobility anchor, MUST be protected using end-to-end security association(s) offering integrity and data origin authentication.

The mobile access gateway and the local mobility anchor MUST implement the IPsec security mechanism mandated by Proxy Mobile IPv6 [RFC5213] to secure the signaling described in this document. In the following, we describe the Security Policy Database (SPD) and Security Association Database (SAD) entries necessary to protect the new signaling introduced by this specification (Subscription Query Message and Subscription Response Message). We use the same format used by [RFC4877]. The SPD and SAD entries are only example configurations. A particular mobile access gateway implementation and a local mobility anchor home agent implementation could configure different SPD and SAD entries as long as they provide the required security of the signaling messages.

For the examples described in this document, a mobile access gateway with address "mag_address_1", and a local mobility anchor with address "lma_address_1" are assumed.

```
mobile access gateway SPD-S:
- IF local_address = mag_address_1 &
remote_address = lma_address_1 &
proto = MH & (remote_mh_type = Subscription Query |
local_mh_type = Subscription Response |
remote_mh_type = Multicast Activity Indication Ack.|
local_mh_type = Multicast Activity Indication)
Then use SA1 (OUT) and SA2 (IN)

mobile access gateway SAD:
- SA1(OUT, spi_a, lma_address_1, ESP, TRANSPORT):
local_address = mag_address_1 &
remote_address = lma_address_1 &
proto = MH
- SA2(IN, spi_b, mag_address_1, ESP, TRANSPORT):
local_address = lma_address_1 &
remote_address = mag_address_1 &
proto = MH

local mobility anchor SPD-S:
- IF local_address = lma_address_1 &
remote_address =mag_address_1 &
proto = MH & (remote_mh_type = Subscription Response |
local_mh_type = Subscription Query |
remote_mh_type = Multicast Activity Indication |
local_mh_type = Multicast Activity Indication Ack.)
Then use SA2 (OUT) and SA1 (IN)

local mobility anchor SAD:
- SA2(OUT, spi_b, mag_address_1, ESP, TRANSPORT):
local_address = lma_address_1 &
remote_address = mag_address_1 &
proto = MH
- SA1(IN, spi_a, lma_address_1, ESP, TRANSPORT):
local_address = mag_address_1 &
remote_address = lma_address_1 &
proto = MH
```

While in the base solution the LMA has learned of the subscribed multicast groups per MAG, in this specification the LMA is aware (during a handover process) of the multicast groups to which an MN visiting the PMIP domain is subscribed.

9. IANA Considerations

This document establishes new assignments to the IANA mobility parameters registry.

- o Mobility Header types: the Subscription Query (22) and Subscription Response (23) mobility header types. The Type value for these Headers has been assigned from the "Mobility Header Types - for the MH Type field in the Mobility Header" registry defined in <http://www.iana.org/assignments/mobility-parameters>.

- o Mobility options: the Active Multicast Subscription mobility option for both IPv4 (56) and IPv6 (57) modes of operation. The Type value for these Mobility options has been assigned from the "Mobility Options" registry defined in <http://www.iana.org/assignments/mobility-parameters>.

- o Flags: this document reserves a new multicast Signaling flag (S). This flag has been reserved as value 0x0020 in the "Binding Update Flags" registry and value 0x04 in the "Binding Acknowledgment Flags" registry. These registries appear on <http://www.iana.org/assignments/mobility-parameters>.

10. Contributors

Dirk Von Hugo (Telekom Innovation Laboratories, Dirk.von-Hugo@telekom.de) extensively contributed to this document.

11. Acknowledgments

The authors would like to thank (in alphabetical order) Hitoshi Asaeda, Sergio Figueiredo, Georgios Karagiannis, Marco Liebsch, Behcet Sarikaya, Thomas C. Schmidt, and Stig Venaas for their valuable comments and discussions on the MULTIMOB mailing list. The authors are also grateful with Hitoshi Asaeda, Akbar Rahman, Behcet Sarikaya, and Stig Venaas for their reviews of this document.

The research of Carlos J. Bernardos leading to these results has received funding from the European Community’s Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project), being also partially supported by the Ministry of Science and Innovation (MICINN) of Spain under the QUARTET project (TIN2009-13992-C02-01).

The research of Ignacio Soto has also received funding from the Spanish MICINN through the I-MOVING project (TEC2010-18907).

12. References

12.1. Normative References

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.

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

[RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC4283] Patel, A., Leung, K., Khalil, M., Akhtar, H., and K. Chowdhury, "Mobile Node Identifier Option for Mobile IPv6 (MIPv6)", RFC 4283, November 2005.

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

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

[RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.

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

12.2. Informative References

[Papagiannaki] Papagiannaki, K., Moon, S., Fraliegh, C., Thiran, P., and C. Diot, "Measurement and Analysis of Single-Hop Delay on an IP Backbone Network", IEEE Journal on Selected Areas in Communications, vol. 21, no. 6, August 2003.

[RFC5949] Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", RFC 5949, September 2010.

[RFC6224] Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

[RFC6636] Asaeda, H., Liu, H., and Q. Wu, "Tuning the Behavior of the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) for Routers in Mobile and Wireless Networks", RFC 6636, May 2012.

[RFC7028] Zuniga, JC., Contreras, LM., Bernardos, CJ., Jeon, S., and Y. Kim, "Multicast Mobility Routing Optimizations for Proxy Mobile IPv6", RFC 7028, September 2013.

[Verizon] Verizon, "LTE: The Future of Mobile Broadband Technology", Verizon White Paper, 2010, <http://opennetwork.verizonwireless.com/pdfs/VZW_LTE_White_Paper_12-10.pdf>.

[Y.1541] ITU-T, "Network performance objectives for IP-based services", ITU-T Recommendation Y.1541, December 2011.

Appendix A. Performance Comparison with Base Solution

This informative annex briefly analyzes and compares the performance improvement provided by the fast handover extensions specified in this document with the base multicast solution defined in [RFC6224]. The main aim is to determine the potential delay reduction in the acquisition of the multicast subscription information by the nMAG during the MN handover. To do that, the analysis focuses on the delay additional to the unicast handover due to the multicast operation in both cases.

Different delay components have to be taken into account for this comparison. Since the interaction between the actors during the handover process (MN, pMAG, nMAG, LMA) is different for each of the solutions, different sources of delay can be expected for each of them.

A.1. Delay Characterization of the Base Solution

The base solution relies on the standard MLD procedures to obtain the multicast subscription information directly from the MN. Once the nMAG completes the configuration of point-to-point link to the attaching MN (the configuration of this link as downstream interface of an MLD proxy instance can run in parallel), it immediately sends an MLD General Query towards the MN for learning of any active multicast subscription by the MN. When the MN receives the MLD Query, the MN provides information about the active memberships it maintains in the form of an MLD Report message. After successful transmission of this information via the wireless point of attachment to nMAG, the corresponding MLD proxy instance at the nMAG sets up the multicast status of the downstream interface. According to this process, the delay is originated on the MAG-MN communication.

The delay components to be considered for the base solution are the following:

- D_bh, which is the unidirectional (one-way) delay encountered in the transmission path between the nMAG and the wireless point of attachment.

- D_radio, which is the unidirectional delay due to the transfer of MLD control messages over the radio channel (user plane) between the wireless point of attachment and the MN, for the MLD Query and Report messages.

- D_mld, which is the delay incurred by the MN to answer the MLD Query.

The total observed delay can be then formulated as:

D_base = 2 x (D_bh + D_radio) + D_mld

A.2. Delay Characterization of SIAL

As described in this document, it is possible to distinguish two scenarios depending on the order in which the LMA receives the notifications of the MN registration and de-registration in the nMAG and the pMAG, respectively.

In the proactive case, the MN is firstly de-registered by the pMAG, and later on it is registered by the nMAG. As specified in this document, the LMA stores the multicast subscription information, which is provided to the nMAG during the MN registration process. Since the registration process necessarily happens before the MLD Query and Report process described in the base solution, the proactive case is inherently faster than the base solution. In fact, since the multicast subscription information is acquired properly during the registration process, the delay incurred is null.

In the reactive case, the LMA receives the MN registration from the nMAG without having previously received the MN de-registration from the pMAG. In case the MN maintains an active subscription, the LMA queries the pMAG to retrieve the multicast subscription information, which is forwarded to the nMAG. According to this process, the delay is originated on the MAG-LMA communication.

The delay components to be considered for the base solution are the following:

- D_net, which is the unidirectional delay found in the network path between the LMA and the MAG.

The total observed delay can be then formulated as:

D_sial = 2 x D_net

A.3. Performance Comparison

The performance of the base solution is highly dependent on the radio technology used by the MN to attach to the PMIPv6 domain. Different radio technologies have distinct properties in terms of radio framing, radio access contention or collision avoidance, channel reliability, etc.

New radio access technologies, such as the one specified in new Long Term Evolution (LTE) standards intend to reduce the latency in order to provide high-speed communications. Even though, typical one-way latencies in the LTE radio access will stay around 15 ms [Verizon].

The backhaul delay characterization becomes problematic. In a real network, there are several solutions for the backhaul connection in terms of network topology (ring, star, point-to-point, etc.) and technology (optical fiber, microwave transmission, xDSL-based accesses, etc.), all of them having distinct properties in terms of performance, reliability, and delay. These solutions commonly coexist in a real mobile network, in such a way that an MN changing the point of attachment can pass smoothly from one solution to another. A value of D_bh = 5 ms can be established as the typical value for the backhaul latency in modern networks.

Finally, the MLD induced delay is intrinsic to the MLD protocol specification. A host receiving an MLD Query message waits a random time in the range (0, Maximum Response Delay) to send the MLD Report message. The default value of the Maximum Response Delay (configurable through the Query Response Interval in MLD) is 10 s in [RFC2710] and [RFC3810]. In [RFC6636] the effect of tuning the value of the Query Response Interval is analyzed and 5 s is the smallest value recommended (best case). Then, on average, a potential delay of 5 s or 2.5 s, default and best case respectively, can be expected.

As we have seen, D_base is, on average, greater than 2.5 s with the best case of the values of Query Response Interval in MLD that are recommended in [RFC6636]. That means that the handover delay of the base solution is on the order of seconds, while in the solution presented in this specification it is on the order of milliseconds (as shown below). To improve the performance of the base solution, we could further reduce the value of Query Response Interval, but the implications of doing so would need to be carefully analyzed. Even if we assume that Query Response Interval is 0 s, D_base would be around 2 x (5 ms + 15 ms) = 40 ms for last-generation systems. Note that this calculation does not take into account the necessary time to re-establish the data plane after the handover to make possible the MLD Query reception. The expected delay will get much worse for older generation systems (e.g., 3G-based radio systems can suffer radio delays in the order of hundreds of ms).

For the SIAL case, the delay in the MAG-LMA communication will be derived from the network diameter (i.e., the number of hops found between the MAG and the LMA in the PMIPv6 domain). This is largely influenced by the internal network planning. An administrative domain can typically have in the order of five hops from access to the interconnection gateway providing connectivity to other networks.

Even if the LMA plays a central role topologically in the PMIPv6 domain, such number of hops seems reasonable in a common nation-wide network. Each hop in the path between MAG and LMA will add a certain delay, which can be estimated to be around 1 ms in the best case [Papagiannaki] and 3 ms in the worst case [Y.1541]. With this in mind, a total delay D_sial of around 2 x 5 x 3 ms = 30 ms can be expected in the worst case.

Then, in conclusion, in a typical deployment, it can be stated that the SIAL proposal, even for the worst-case consideration, will perform better than the best-case situation for the base solution, which consists of the last-generation radio technology, LTE. For any other radio technology, the base solution will show even larger deviations from the delay achievable with the SIAL solution.

Authors' Addresses

Luis M. Contreras
Telefonica I+D
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
Madrid 28050
Spain

EMail: lmcm@tid.es

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/

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

EMail: isoto@it.uc3m.es