

MULTIMOB Working Group
Internet-Draft
Intended status: Experimental
Expires: August 29, 2013

LM. Contreras
Telefonica I+D
CJ. Bernardos
I. Soto
UC3M
February 25, 2013

PMIPv6 multicast handover optimization by the Subscription Information
Acquisition through the LMA (SIAL)
draft-ietf-multimob-handover-optimization-02

Abstract

This document specifies a multicast handover optimization mechanism for Proxy Mobile IPv6 to accelerate the delivery of multicast traffic to mobile nodes after handovers. The mechanism is based on speeding up the acquisition of mobile nodes' multicast context by the mobile access gateways. To do that, extensions to the current Proxy Mobile IPv6 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 being developed to avoid the tunnel convergence problem. Furthermore, they are also independent of the role played by the mobile access gateway within the multicast network (either acting as multicast listener discovery proxy or multicast router).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 29, 2013.

Copyright Notice

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

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

Table of Contents

- 1. Introduction 4
 - 1.1. Handover optimization requirements 5
- 2. Terminology 6
- 3. Overview 7
- 4. PMIPv6 extensions 9
 - 4.1. New mobility option 9
 - 4.1.1. Active Multicast Subscription mobility option 9
 - 4.1.1.1. Option application rules 9
 - 4.1.1.2. Option format 9
 - 4.1.1.3. Backward compatibility with MLDv1 10
 - 4.2. New flags 11
 - 4.2.1. Multicast Signaling flag on PBU/PBA message headers 11
 - 4.2.1.1. Flag application rules 11
 - 4.2.1.1.1. Registration process 11
 - 4.2.1.1.2. De-registration process 12
 - 4.2.1.2. New format of conventional PBU/PBA messages 13
 - 4.2.1.2.1. Proxy Binding Update message 13
 - 4.2.1.2.2. Proxy Binding Acknowledgement Message 13
 - 4.2.2. Multicast Active flag in the LMA Binding Cache and (optionally) on the MN's policy store 14
 - 4.2.2.1. Flag application rules 14
 - 4.3. New messages 15
 - 4.3.1. Messages for active multicast subscription query 15
 - 4.3.1.1. Subscription Query message 15
 - 4.3.1.1.1. Message application rules 15
 - 4.3.1.1.2. Message format 15
 - 4.3.1.2. Subscription Response message 16
 - 4.3.1.2.1. Message application rules 16
 - 4.3.1.2.2. Message format 16
 - 4.3.2. Messages for active multicast subscription indication 17
 - 4.3.2.1. Multicast Activity Indication message 18
 - 4.3.2.1.1. Message application rules 18
 - 4.3.2.1.2. Message format 18
 - 4.3.2.2. Multicast Activity Indication Acknowledge

- message 19
- 4.3.2.2.1. Message application rules 19
- 4.3.2.2.2. Message format 19
- 4.4. New PBA timer in the LMA 20
- 5. Signaling processes description 20
- 5.1. Multicast Activity signaling 20
- 5.1.1. Multicast Activity set to ON (A=1) 21
- 5.1.2. Multicast Activity set to OFF (A=0) 22
- 5.2. Handover signaling procedures 23
- 5.2.1. Handover of proactive type 23
- 5.2.1.1. Rationale 23
- 5.2.1.2. Message flow description 23
- 5.2.2. Handover of reactive type 25
- 5.2.2.1. Rationale 25
- 5.2.2.2. Message flow description 26
- 5.2.2.3. Further considerations for the reactive
handover signaling 29
- 5.2.3. LMA decision process 30
- 5.2.3.1. LMA processing of S flag on reception of PBU
messages 30
- 5.2.3.1.1. Proactive handover 30
- 5.2.3.1.2. Reactive handover 31
- 5.2.3.2. LMA set-up of S flag in PBA messages 32
- 5.2.4. Prevention of large delays of the binding
acknowledgement for unicast traffic 33
- 6. IPv4 support 36
- 6.1. Active Multicast Subscription for IPv4 36
- 6.2. Signaling procedures for IPv4 support 37
- 6.3. Binding Cache extensions for IPv4 support 38
- 7. Co-existence with PMIPv6 multicast architectural evolutions . 38
- 8. Security Considerations 38
- 9. IANA Considerations 40
- 10. Contributors 40
- 11. Acknowledgments 40
- 12. References 41
- 12.1. Normative References 41
- 12.2. Informative References 41
- Appendix A. Performance comparison with base solution 42
- A.1. Delay characterization of the base solution 43
- A.2. Delay characterization of the SIAL solution 43
- A.3. Performance comparison 44
- Appendix B. Change Log 46
- Authors' Addresses 47

1. Introduction

The base solution describing how continuous multicast service delivery can be provided in Proxy Mobile IPv6 domains is described in [RFC6224]. It specifies the basic functionality needed in the Proxy Mobile IPv6 [RFC5213] entities to provide a multicast service, and supports the continuous delivery of multicast traffic by obtaining, after a handover, the on-going multicast subscription information directly from the mobile node. When a mobile node attaches to a new mobile access gateway, the mobile node is queried by the mobile access gateway through a multicast listener discovery General Query, which is sent just after any new link is set up, to get knowledge of any existing subscription, as specified in [RFC2710].

However, as highlighted by [I-D.von-hugo-multimob-future-work], the base solution needs to be improved to meet some performance requirements, especially those referred 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 to mitigate such reception gap. One of them is to buffer at the previous mobile access gateway the traffic with destination at 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 get knowledge 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 triggers, thus requiring additional support on the mobile node), there is no generic standard solution for the accelerated acquisition of the on-going multicast subscription of the mobile node.

The approach followed by the base solution [RFC6224] to get knowledge 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 subscription by the mobile node.

While this is a straightforward approach, it also causes that 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], 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 will continue sending an MLD Query to the entering MN 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 of 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 MLD control messages. Issues like radio framing, radio access contention, channel reliability, MN's capabilities (i.e., layer-2 triggering support), IGMP/MLD timers optimization for wireless environments, etc, will not impact on the observed multicast performance during handovers.

The solution described in this document can also be applied to the solutions described in [I-D.ietf-multimob-pmipv6-ropt]. Furthermore, it is also independent of the role played by the mobile access gateway within the multicast network (either acting as 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 this specification 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 in which the nMAG is informed about the on-going multicast subscription of the entering MN, mainly due to the fact that the mechanisms provided in the base solution rely 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 has to be applicable to any kind of MN, in such a way that any type of MN in a PMIPv6 domain being served with multicast traffic can benefit from the optimized solution.
- o The solution should be integrated as part of the PMIPv6 suite of protocols to ensure a smooth introduction of the new functionality in the network.
- o The solution does not have impact on existing multicast protocols.
- o The solution should optimize the handover performance 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 required in the network, aiming at an easier deployment.

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 [RFC2119].

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

Additionally, the following terms are defined.

pMAG: The previous MAG or pMAG is the MAG where the MN is initially registered in a handover event.

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

Reactive Handover: A reactive handover is a handover event in which the LMA receives the MN 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 MN is firstly de-registered on the LMA by the pMAG, and later on it is registered by the nMAG as consequence of changing the point of attachment.

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

3. Overview

The LMA is a key element within the PMIPv6 infrastructure, which traces the MN reachability along the PMIPv6 domain. Therefore the LMA is the best element to maintain the MNS' multicast subscription information updated and to forward it to the rest of PMIPv6 entities (i.e., to the MAGs) as needed when MNS move within the domain. The LMA has timely knowledge of the MNS' location, especially during handover events, and it is therefore able to quickly provide information to the new one point of attachment (querying the previous one if required). Figure 1 shows this idea.




Figure 1: High level solution description

The LMA only obtains the detailed subscription information or multicast context during a handover event. There is no need of 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, ready to be transferred to the new MAG where the MN has just attached. However it should be noted that some signaling is needed to differentiate what MNs are maintaining active subscriptions in order to restrict the optimization procedure to them in case of handover.

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.

Next sections provide the details of these Proxy Mobile IPv6 protocol extensions.

4. PMIPv6 extensions

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

4.1. New mobility option

4.1.1. Active Multicast Subscription mobility option

4.1.1.1. Option application rules


A new TLV-encoded mobility option, "Active Multicast Subscription" option is defined for use with the PBU (Proxy Binding Update) and PBA (Proxy Binding Acknowledge) messages exchanged between an LMA and a MAG to transfer the multicast subscription information. This option is used for exchanging the multicast membership context. This information is carried by using directly 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 MN when the handover is taking place (i.e., one per multicast membership context).

This new option will be also used, with the same aim, by the new message Subscription Response described later in this document.

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

4.1.1.2. Option format

The format of this new option is as follows:



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

Type:

To be defined by IANA, for indication of an IPv6 Active Multicast Subscription 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 0x143.

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].

4.1.1.3. Backward compatibility with MLDv1

The following values are adopted when MLDv1 is used.


MLD type:

For MLDv1, the MLD Type value is 0x131.

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], in the following manner:



4.2. New flags

Two new flags are defined and used to handle the forwarding of multicast subscription information.

4.2.1. Multicast Signaling flag on PBU/PBA message headers

4.2.1.1. Flag application rules

A new flag S is added in both PBU and PBA message headers to advise about the MAG and the LMA capabilities of processing multicast-related signaling for the MN that caused the message.

This flag will govern the multicast-related signaling between the LMA and the MAG. As a general rule, the value of the flag in the PBA message SHOULD be a copy of the value received in the PBU message. Specific rules are described in next sub-sections.

4.2.1.1.1. 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.

- o PBU message
 - * S=0, it 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 which are not multicast enabled, for backward compatibility reasons.
 - * S=1, it 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 will take some actions, described later in this document.

- o PBA message
 - * If S=0 in the corresponding PBU message, the value of the flag in the PBA message SHOULD 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 sub-cases can happen
 - + 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 will include the "Active Multicast Subscription" mobility option. If the LMA is not able to provide such information during registration, the PBA message will 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).

4.2.1.1.2. 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, it 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, it indicates that the MN has an active multicast session, and the multicast context is 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).

4.2.1.2. New format of conventional PBU/PBA messages


4.2.1.2.1. Proxy Binding Update message

As result of the new defined flag, the PBU message results as follows:



4.2.1.2.2. Proxy Binding Acknowledgement Message

As result of the new defined flag, the PBA message results as follows:



4.2.2. Multicast Active flag in the LMA Binding Cache and (optionally) on the MN's policy store

4.2.2.1. Flag application rules

A new flag A is added in the LMA Binding Cache to retain the knowledge that the registered MN maintains or not an active multicast subscription. The basic use of this flag is to restrict the query of the pMAG only to the cases in which the MN certainly is maintaining an active subscription. The algorithm which is followed by the LMA to query or not the pMAG (after receiving a PBU message from the nMAG) is as follows:

- o Flag S=0 & flag A=0: this situation represents the case where the nMAG does not support multicast-related signaling for the MN being registered, and, additionally, the LMA is not aware of any active multicast subscription on-going. Then, the LMA does not query the pMAG, and registers the MN as attached to the nMAG as usual.
- o Flag S=0 & flag A=1: this situation represents the case where the nMAG does not support multicast-related signaling for the MN being registered, but the LMA is aware of one or more on-going MN's active multicast subscriptions. Due to the fact that multicast signaling is not supported by the nMAG for that MN, the LMA does not query the pMAG, and registers the MN as attached to the nMAG as usual.
- o Flag S=1 & flag A=0: this situation represents the case where the nMAG supports multicast-related signaling for the MN being registered, but the LMA is not aware of any active multicast subscription. Then, the LMA does not query the pMAG, and registers the MN as attached to the nMAG as usual.
- o Flag S=1 & flag A=1: this situation represents the case where the nMAG supports multicast-related signaling for the MN being registered, and, additionally, the LMA is aware of one or more on-going MN's active multicast subscriptions. Then, the LMA queries the pMAG to obtain the multicast context details previously to complete the registration of the MN attached to the nMAG.

The flag A SHOULD be initialized to the value 0.

Optionally, this flag can be also added to the MN's policy store, and dynamically updated by the LMA to signal that the MN has (or not) an active multicast subscription. By introducing this flag in the MN's policy profile, the nMAG can know in advance the existence of an active multicast session by the incoming MN.

Sequence Number will be determined by the entity which creates the message (either LMA or MAG, depending on the scenario), which will be responsible of managing this counter.

Reserved:

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

Mobility options:

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

- * Mobile Node Identifier option (mandatory)
- * Home Network Prefix option (optional)

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


4.3.1.2. Subscription Response message

4.3.1.2.1. Message application rules

The Subscription Response message 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.

4.3.1.2.2. Message format

The Subscription Response message has the following format.



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 if there is multicast subscription information available for the MN or not. 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 is 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 will carry one or more TLV-encoded mobility options. The valid mobility options for this message are the following:

- * Mobile Node Identifier option (mandatory)
- * Active Multicast Subscription option (mandatory) only when flag I=1, not present in any other case
- * Home Network Prefix option (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.

4.3.2. Messages for active multicast subscription indication

A new pair of messages is defined for setting up and down the optional A flag defined above.

These messages are sent using the Mobility Header as defined in

[RFC6275].


4.3.2.1. Multicast Activity Indication message

4.3.2.1.1. Message application rules

The Multicast Activity Indication message is sent by a MAG towards the LMA to set to 1 or 0 the A flag either to indicate the start or the complete cease of any multicast subscription by the MN. Through the use of this message, the LMA becomes aware that one or more multicast flows are being forwarded to a MN. This information is useful for the LMA during a handover to discriminate if the pMAG needs to be asked or not about multicast information corresponding to the MN being registered at the nMAG, in case that the handover is of reactive type.

4.3.2.1.2. Message format

The Multicast Activity Indication message has the following format.



Sequence Number:

The Sequence Number field establishes the order of the messages sent in the Activity Indication / Activity Indication Ack dialogue between the MAG and the LMA for a certain MN. The initial Sequence Number will be determined by the MAG, which will be responsible of managing this counter.

Activity indicator (A):

The Activity indicator flag A specifies if the MN multicast activity is on, that is, if the MN maintains one or more active multicast subscriptions at the MAG. The meaning is the following:

A=0: the multicast activity of the MN (identified by the Mobile Node Identifier option in this message) is OFF.

A=1: the multicast activity of the MN (identified by the Mobile Node Identifier option in this message) is ON.

Reserved:

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

Mobility options:

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

- * Mobile Node Identifier option (mandatory)
- * Home Network Prefix option (optional)

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


4.3.2.2. Multicast Activity Indication Acknowledge message

4.3.2.2.1. Message application rules

The Multicast Activity Indication Acknowledge message is sent by the LMA towards a MAG to confirm the reception of a previously sent Multicast Activity Indication message.

4.3.2.2.2. Message format

The Multicast Activity Indication message has the following format.



Sequence Number:

The value of the Sequence Number field in the Activity Indication Ack message MUST be a copy of the Sequence Number received in the

Activity Indication message.

Reserved:

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

Mobility options:

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

- * Mobile Node Identifier option (mandatory)

- * Home Network Prefix option (optional)

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

4.4. 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.

This timer SHOULD be upper bounded by the constant defined in [RFC6275] INIT_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 ensure that the nMAG does not enter the retry mode.

5. Signaling processes description

A number of new signaling processes are introduced with this solution. Next sections describe these new processes in detail.

5.1. Multicast Activity signaling

This solution makes use of the flag A to keep track of existing multicast activity in a certain MN. The idea behind this is to define a mechanism which helps the LMA to decide whether to query or not the pMAG about potential subscription information.

This signaling message is used to allow the LMA to distinguish among MNs with on-going multicast subscription, and MN without active multicast status. This differentiation further allows to apply the optimization procedure only to those MNs with active multicast subscription (no actions are taken for MN without active multicast subscription).

5.1.1. Multicast Activity set to ON (A=1)

Figure 2 summarizes this process.




Figure 2: Multicast Activity set to ON

The sequence of messages is the following:

1. An MN, named MN1, is attached to the MAG. The MN is a multicast-enabled node, and it is only receiving unicast traffic as usual in PMIPv6 domains, with no multicast subscription yet. At some point in time, the MN1 requests to the MAG to be subscribed to the content identified by the IP addresses (S,G), by sending a standard MLD report from the MN to the MAG. The MAG will keep the multicast status state of the point-to-point link with the MN. In case the MAG has not already subscribed to the multicast flow (S,G) it joins the content on behalf of MN. Multicast flow (S,G) is subsequently forwarded by the MAG to the MN1.
2. Due to this initial multicast subscription for the MN1, the MAG triggers the multicast Activity Indication message towards the LMA, to indicate that the MN1 multicast activity is ON. The LMA will set the flag A to 1. Afterwards, the LMA sends an Activity Indication Ack message to the MAG to acknowledge the previous

indication.

5.1.2. Multicast Activity set to OFF (A=0)

Figure 3 presents the corresponding flow.




Figure 3: Multicast Activity set to OFF

The message flow is as follows:

1. Some time later, the MN1 decides to totally stop all the active multicast subscriptions that it maintains. The MN1 will send an MLD Report message (containing an State Change Record for the last subscribed multicast group with a filter change record mode indicating INCLUDE mode and an empty source list) to the MAG to request the cease of the multicast traffic delivery. As a consequence, the MAG will stop all the multicast traffic forwarding to the MN1.
2. After removing the active subscriptions for the MN1, the MAG sends a multicast Activity Indication message to the LMA indicating that the MN1 multicast activity is OFF. The LMA will set the flag A to 0, its default value. Afterwards, the LMA sends an Activity Indication Ack message to the MAG to acknowledge the previous indication.

5.2. 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 and, 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.

5.2.1. Handover of proactive type

5.2.1.1. Rationale

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

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

The LMA will store that information in the corresponding binding cache. If, later on, the MN attaches to a nMAG, this information will be sent (using the same TLV option) to the nMAG as part of the PBA confirmation of the registration process (the PBU message sent by the nMAG SHOULD set the flag S to 1). On the other hand, if no further registration happens, the multicast information will be 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 if there is no other MN receiving it already at the nMAG. The multicast status can be also set in advance for the point-to-point link towards the MN.

Note that the 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 layer-2 capabilities.

5.2.1.2. Message flow description

Figure 4 summarizes this process.




Figure 4: Proactive handover

The message flow is as follows:

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

2. The MN perceives a better radio link and decides to initiate a handover process over a radio access controlled by a new MAG, nMAG. As a consequence, pMAG determines a detach event corresponding to this MN, and updates the attachment status of this MN to the LMA by sending an extended Proxy Binding Update message, including a new TLV-encoded option, named "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 on-going multicast subscription(s) when the detachment is initiated. This information will be kept until a new registration of the MN is completed by another MAG, or till the Binding Cache expiration, according to [RFC5213].
4. The LMA acknowledges to the pMAG the previous PBU message.
5. As a result of the handover process, the MN attaches to another MAG, called nMAG.
6. The nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the LMA.
7. After the analysis of the PBU message, the LMA sends an extended PBA including the new "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 will set up the multicast status of the point-to-point link between the nMAG and the MN, and will join 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.

5.2.2. Handover of reactive type

5.2.2.1. Rationale

In the reactive case, the LMA receives the MN 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

MN, the nMAG will start a conventional registration process, by sending a normal PBU message (with flag S set to 1) towards the LMA.

After receiving the PBU message from the nMAG, the LMA will take the decision of interrogating or not the pMAG regarding any existing multicast subscription for that MN. This decision is taken following a procedure that is described later in section Section 5.2.3.

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 MN receiving it already at the nMAG. The multicast status can be also set in advance for the point-to-point link towards the MN.

5.2.2.2. Message flow description

Figure 5 and Figure 6 summarize this process.




Figure 5: Reactive handover (steps 1 to 5)

Consider as starting point the situation where a couple of MNs, named MN1 and MN2, are attached to the pMAG, both MNs being multicast-enabled nodes, but only MN1 maintains an active multicast subscription at this moment. As consequence, the value of the A flag in the LMA is set to 1 for MN1, and set to 0 for MN2.

The sequence of messages for the handover of MN1 and MN2 is the following (as depicted in Figure 5):

1. At certain time, the MN1 perceives a better radio link and decides to attach at a new MAG, nMAG, in a handover process (as it is a reactive case, the pMAG is not aware of the detachment process). Then, the nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the LMA.
2. Prior to acknowledge the received PBU message, the LMA checks the status of the A flag for this MN. Due that the flag A=1, the LMA queries the pMAG about if there is any active multicast subscription for the MN1, 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 new TLV-encoded option "Active Multicast Subscription". 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 MN1, and to join the content identified by (S,G) on behalf of the MN1 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 by another MN connection). At this moment, the multicast content can be served to the MN1. The unicast traffic for the MN1 can be forwarded as well.




Figure 6: Reactive handover (steps 6 to 8)

6. In parallel, the MN2 perceives a better radio link and decides to attach also to the nMAG in a reactive handover process as well (the pMAG is not aware of this detachment process either). Then, the nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the LMA.
7. Prior to acknowledge the received PBU message, the LMA checks the status of the A flag for this MN. Due that the flag A=0, the LMA does not query the pMAG, and acknowledges the PBU message (with flag S set to 0). The nMAG then processes PBA message.
8. The nMAG is now ready to forward the unicast traffic to the MN2.

5.2.2.3. 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 will receive a registration PBU from nMAG before a de-registration PBU from pMAG, if any.

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 will carry 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 LMA SHOULD store it, and 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 MN.

5.2.3. LMA decision process

A key point of the solution proposed in this document resides on the LMA decision of interrogating the pMAG about a potential active subscription of the MN entering the nMAG. Several variables take place, and it is required to define a mechanism for assisting the LMA in its decision process.

Basically two flags will be used. One flag, the named "multicast Signaling" or S flag, is used to signal the multicast capabilities of the MAGs and the transport of the multicast subscription information within the PBU/PBA messages. The other one, the named "multicast Activity" or A flag, is used to register on the LMA whether the MN is maintaining an active multicast subscription or not.

The following sections summarize the use of these flags on the LMA decision process.

5.2.3.1. LMA processing of S flag on reception of PBU messages

5.2.3.1.1. Proactive handover

In the event of proactive handover, the pMAG has previously informed the LMA about any potential subscription information currently active in the MN. The actions to be carried out by the LMA once it receives the PBU message from the nMAG are summarized in the table below.

multicast signaling flag S in PBU	multicast activity flag A in LMA	Meaning	LMA action
S=0	A=0	- Multicast not supported by nMAG - No active subscription by MN	- MN registration as in RFC 5213 (S=0 in PBA)
	A=1	- Multicast not supported by nMAG - Active subscription by MN	- LMA stores multicast subscription info - MN registration as in RFC 5213 (S=0 in PBA)
S=1	A=0	- Multicast supported by nMAG - No active subscription by MN	- MN registration as in RFC 5213 (S=0 in PBA)
	A=1	- Multicast supported by nMAG - Active subscription by MN	- LMA stores multicast subscription info - MN registration conveys multicast subscription info (S=1 in PBA)

Figure 7: Flag processing on LMA for proactive handover

5.2.3.1.2. Reactive handover

In the event of reactive handover, the LMA is not aware about any potential subscription information currently active in the MN. The actions to be carried out by the LMA once it receives the PBU message from the nMAG are summarized in the table below.

multicast signaling flag S in PBU	multicast activity flag A in LMA	Meaning	LMA action
S=0	A=0	<ul style="list-style-type: none"> - Multicast not supported by nMAG - No active subscription by MN 	<ul style="list-style-type: none"> - MN registration as in RFC 5213 (S=0 in PBA)
	A=1	<ul style="list-style-type: none"> - Multicast not supported by nMAG - Active subscription by MN 	<ul style="list-style-type: none"> - LMA does not query pMAG - MN registration as in RFC 5213 (S=0 in PBA)
S=1	A=0	<ul style="list-style-type: none"> - Multicast supported by nMAG - No active subscription by MN 	<ul style="list-style-type: none"> - LMA does not query pMAG - MN registration as in RFC 5213 (S=0 in PBA)
	A=1	<ul style="list-style-type: none"> - Multicast supported by nMAG - Active subscription by MN 	<ul style="list-style-type: none"> - LMA queries pMAG to obtain multicast subscription - MN registration conveys multicast subscription info (S=1 in PBA)

Figure 8: Flag processing on LMA for reactive handover

5.2.3.2. LMA set-up of S flag in PBA messages

Once the LMA decision process is finished, the LMA builds the PBA message to complete the registration process triggered by the nMAG. The value of the S flag in the PBA message will be set according to the data specified in the table below.

S flag received in PBU message	S flag sent in PBA message	Meaning
S=0 (multicast not supported by nMAG)	S=0	No further meaning
	S=1	N/A
S=1 (multicast is supported by nMAG)	S=0	No active subscription on MN
	S=1	- Mcast context available: Multicast subscription info is conveyed in the PBA message - Mcast context not available: It has to be requested by using the Subscription Query message.

Figure 9: S flag configuration in PBA messages

5.2.4. 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 till the LMA receives, processes and includes the multicast subscription information into the extended PBA message. As 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 MUST be buffered or discarded by the LMA.

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 10 shows this mechanism:




Figure 10: 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 perceives a better radio link and decides to attach at a new MAG, nMAG, in a handover process (as a reactive case, the pMAG is not aware of the detachment process). Then, the nMAG triggers a registration process by sending a PBU message (with flag S set to 1) to the LMA.
3. Prior to acknowledge the received PBU message, the LMA decides to query the pMAG about if there is any active multicast subscription for the MN, by sending a Subscription Query message. The LMA decision is based on the checking of flag A when the reactive handover manages the multicast activity indication.
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 will allow the MN 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 MN.
7. The pMAG answers the Subscription Query message originally sent by the LMA, 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 new TLV-encoded option "Active Multicast Subscription". 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 MN, 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 MN.


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 MAG to request multicast contents to the mobility anchor on behalf of the mobile nodes attached to it.

In order for this specification to support that, the following additions are needed.

6.1. Active Multicast Subscription for IPv4

The Active Multicast Subscription option defined in Section 4.1.1, which transports the multicast membership context of the MN during handover, should be compatible with IGMP-based formats. Specifically, the option format is defined for IPv4-based MNs as follows:



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:

To be defined by IANA, for indication of an IPv4 Active Multicast Subscription 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 MN, the format of the Multicast 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 on-going multicast subscription for a MN during the handover process, the following extensions have to be considered.

- 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 is 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 is 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) is (are) accepted independently of the type indication.

6.3. Binding Cache extensions for IPv4 support

Additionally, since the 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. Co-existence with PMIPv6 multicast architectural evolutions

Along this document, it has been considered that the LMA entity is in charge of delivering both unicast and multicast traffic to a certain MN through the bi-directional tunnels connecting to the MAG where the MN is attached, as specified in the base solution defined in [RFC6224]. However, the solution described in this memo is not only applicable to the base solution, but it can also be applied to the solutions described in [I-D.ietf-multimob-pmipv6-ropt] to solve the tunnel convergence problem.

The Multicast Tree Mobility Anchor (MTMA) solution in [I-D.ietf-multimob-pmipv6-ropt] makes use of a separate entity to serve multicast traffic through distinct tunnels connected to the MAGs. The tunnels for multicast traffic could not be set up in advance if they are dynamical in nature.

When the "multicast activity" flag is also present in the MN's policy store, the nMAG knows in advance the multicast activity of the incoming MN. Consequently, the nMAG can trigger the multicast tunnel set up in parallel to the registration process, including the acquisition of the active multicast subscription details (the multicast context), saving time on serving the multicast flow to the incoming MN. The concrete procedure for multicast tunnel establishment is out of the scope of this document.

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, the Subscription Response Message, the Multicast Activity Indication and the Multicast Activity Indication

Acknowledge, 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, Subscription Response Message, Multicast Activity Indication and Multicast Activity Indication Acknowledge). 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

9. IANA Considerations

This document defines the new following elements which values to be allocated by IANA:

- o Mobility Header types: the Subscription Query and Subscription Response, and the Multicast Activity Indication and Multicast Activity Indication Acknowledge mobility header types.
- o Mobility options: the Active Multicast Subscription mobility option for both IPv4 and IPv6 modes of operation.
- o Flags: the multicast Signaling (S), the multicast Information (I), and the multicast Active (A) flags.

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, 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 review 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.
- [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

- [I-D.ietf-multimob-pmipv6-ropt]
Zuniga, J., Contreras, L., Bernardos, C., Jeon, S., and Y. Kim, "Multicast Mobility Routing Optimizations for Proxy Mobile IPv6", draft-ietf-multimob-pmipv6-ropt-03 (work in progress), February 2013.

- [I-D.von-hugo-multimob-future-work]
Hugo, D., Asaeda, H., Sarikaya, B., and P. Seite,
"Evaluation of further issues on Multicast Mobility:
Potential future work for WG MultiMob",
draft-von-hugo-multimob-future-work-02 (work in progress),
June 2010.
- [Papagiannaki, et al.]
Papagiannaki, K., Moon, S., Fralieggh, 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.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery
Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [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., Waehlich, 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.
- [Verizon] "LTE: The Future of Mobile Broadband Technology", Verizon
White Paper ([http://opennetwork.verizonwireless.com/pdfs/
VZW_LTE_White_Paper_12-10.pdf](http://opennetwork.verizonwireless.com/pdfs/VZW_LTE_White_Paper_12-10.pdf)) , 2010.
- [Y.1541] "Network performance objectives for IP-based services",
ITU-T Recommendation Y.1541 , December 2011.

Appendix A. Performance comparison with base solution

This appendix 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 will focus 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, then 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 getting knowledge 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 will set 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:

- o D_{bh} , which is the unidirectional (one way) delay encountered in the transmission path between the nMAG and the wireless point of attachment
- o 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.
- o D_{mld} , which is the delay incurred by the MN to answer the MLD Query.

The total observed delay can be then formulated as:

$$D_{base} = 2 \times (D_{bh} + D_{radio}) + D_{mld}$$

A.2. Delay characterization of the SIAL solution

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 will store the multicast subscription information, which will be 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 will query 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:

- o 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:

$$D_{sial} = 2 \times D_{net}$$

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 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 will wait 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], or 5 s in the best case described in [RFC6636]. Then, in average, it can be expected a potential delay of 5 or 2,5 s, respectively.

As we have seen, D_{base} is, on average, greater than 2,5 sec 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 it is 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 sec, D_{base} would be of around $2 \times (5 \text{ ms} + 15 \text{ ms}) = 40 \text{ ms}$ 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 5 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, et al.] and 3 ms in the worst case [Y.1541]. With this in mind, a total delay D_{sial} of around $2 \times 5 \times 3 \text{ ms} = 30 \text{ ms}$ can be expected in the worst case.

Then, as conclusion, in a typical deployment, it can be stated that 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 deviation from the delay achievable with the SIAL proposal.

Appendix B. Change Log

The following changes has been made from -00 version.

1. Multicast Address Record format defined in [RFC3810] has been adopted for transferring multicast subscription information in the Active Multicast Subscription mobility option.

The following changes has been made from -01 version.

1. A new appendix has been created to include a performance comparison between this proposal and the base solution.
2. Comments from Akbar Rahman review has been addressed.

The following changes has been made from -02 version.

1. Minor editorial corrections.

The following changes has been made from draft-ietf-multimob-fast-handover-03 version.

1. The name of the draft file has been changed to draft-ietf-multimob-handover-optimization as requested by the chairs.

The following changes has been made from draft-ietf-multimob-handover-optimization-00 version.

1. New Section 1.1 has been introduced to list a number of basic requirements to be covered for an optimization solution to the handover process in [RFC6224].
2. New Section 6 has been introduced to define the support of IPv4-based MNs in this specification.
3. Section 4.1.1 has been reworked for facilitating the alignment with the IPv4 format defined in new Section 6.
4. Minor editorial corrections.

The following changes has been made from draft-ietf-multimob-handover-optimization-01 version.

1. Clarification statements added in several sections according to the comments from B. Sarikaya and S. Venaas.

2. Minor editorial corrections.

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

