MULTIMOB Working Group H. Asaeda
Internet-Draft Y. Uchida
Expires: April 28, 2011 Keio University
October 25, 2010

Tuning the Behavior of IGMP and MLD for Mobile Hosts and Routers
draft-asaeda-multimob-igmp-mld-optimization-04

Abstract

IGMP and MLD are the protocols used by hosts to report their IP multicast group memberships to neighboring multicast routers. This document describes the ways of IGMPv3 and MLDv2 protocol optimization for mobility, mainly for a query timer tuning.

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 April 28, 2011.

Table of Contents

1. Introduction

The Internet Group Management Protocol (IGMP) [2] for IPv4 and the Multicast Listener Discovery Protocol (MLD) [3] for IPv6 are the standard protocols for hosts to initiate joining or leaving multicast sessions. These protocols must be also supported by multicast routers or IGMP/MLD proxies [12] that maintain multicast membership information on their downstream interfaces. Conceptually, IGMP and MLD work on wireless networks. However, wireless access technologies operate on a shared medium or a point-to-point link with limited frequency and bandwidth. In many wireless regimes, it is desirable to minimize multicast-related signaling to preserve the limited resources of battery powered mobile devices and the constrained transmission capacities of the networks. A mobile host may cause initiation and termination of a multicast service in the new or the previous network upon its movement. Slow multicast service activation following a join may degrade reception quality. Slow service termination triggered by IGMP/MLD querying or by a rapid departure of the mobile host without leaving the group in the previous network may waste network resources.

To create the optimal multicast membership management condition, IGMP and MLD protocols could be tuned to "ease a mobile host's processing cost or battery power consumption by IGMP/MLD Query transmission timing coordination by routers" and "realize fast state convergence by successive monitoring whether downstream members exist or not".

This document describes the ways of tuning the IGMPv3 and MLDv2 protocol behavior for mobility, including a query and other timers tuning. The selective optimization that provides tangible benefits to the mobile hosts and routers is given by "keeping track of downstream hosts' membership status" and "varying IGMP/MLD Query types and values to tune the number of responses". A source filtering mechanism in a lightweight manner is also described for enabling Source-Specific Multicast. The proposed behavior interoperates with the IGMPv3 and MLDv2 protocols.

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 RFC 2119 [1].

3. Explicit Tracking of Membership Status

Mobile hosts use IGMP and MLD to request to join or leave multicast sessions. When the adjacent upstream routers receive the IGMP/MLD Report messages, they recognize the membership status on the link. To update the membership status, the routers send IGMP/MLD Query messages periodically as a soft-state approach does, and the member hosts reply IGMP/MLD Report messages upon reception. IGMP/MLD Query is therefore necessary to obtain the up-to-date membership information, but a large number of the reply messages sent from all member hosts may cause network congestion or consume network bandwidth.

The "explicit tracking function" [9] is the possible approach to reduce the transmitted number of IGMP/MLD messages and contribute to mobile communications. It enables the router to keep track of the membership status of the downstream IGMPv3 or MLDv2 member hosts.

Routers that enable the explicit tracking function can unicast IGMP/ MLD General Query messages. This is beneficial especially to mobile hosts that do not have enough battery power, since flooding IGMP/MLD messages on a wireless link makes all multicast members pay attention the messages and induces power consumption to the member hosts. This also allows the upstream router to proceed fast leaves, because the router can immediately converge and update the membership information, ideally. (The usage of unicast General Query is described in Section 4, and Section 6 show the potential timer value.)

In addition, the explicit tracking function reduces the chance of Group-Specific or Group-and-Source Specific Query transmission. Whenever a router that does not enable the explicit tracking function receives the State-Change Report, it sends the corresponding Group-Specific or Group-and-Source Specific Query messages to confirm whether the Report sender is the last member host or not. However, if a router enables the explicit tracking function, it does not need to ask Current-State Report message transmission to the receiver hosts since the router recognizes the (potential) last member hosts when it receives the State-Change Report.

This document therefore recommends to enable the explicit tracking function on adjacent upstream multicast routers. On the other hand, routers enabling the explicit tracking function may still need to maintain downstream membership status by sending IGMPv3/MLDv2 Query messages as IGMPv3/MLDv2 messages may be lost during transmission, while it gives the possibility to make several timers longer as described in Section 6. And the explicit tracking function requires additional processing capability and a possibly large memory for

routers to keep all membership status. Especially when a router needs to maintail a large number of receiver hosts, this resource requirement may be potentially-impacted. Operators may decide to disable this function when their routers do not have enough resources. See [9] for the detail.

4. IGMP/MLD Query Processing

IGMP and MLD are non-reliable protocols; to cover the possibility of a State-Change Report being missed by one or more multicast routers, a host retransmits the same State-Change Report [Robustness Variable] - 1 more times, at intervals chosen at random from the range (0, [Unsolicited Report Interval]) [2][3]. However, this manner does not guarantee that the State-Change Report is reached to the routers. The routers still need to refresh the downstream membership information by receiving Current-State Report periodically solicited by IGMP/MLD General Query, in order to be robust in front of host or link failures and packet loss. It supports the situation that mobile hosts turn off or move from the wireless network to other wireless network managed by the different router without any notification (e.g., leave request).

A multicast router periodically transmits IGMP/MLD General Query in the [Query Interval] sec., whose default value is 125 seconds [2][3]. In general, the all-hosts multicast address (224.0.0.1) or link-scope all-nodes multicast address (FF02::1) is used as the IP destination address of IGMP/MLD General Query. Unfortunately, flooding periodical message whose destination address is the all-hosts/ all-nodes multicast address consumes bettery power of mobile hosts. Only the active hosts that have been receiving multicast contents should respond the Query message.

IGMPv3 and MLDv2 specifications [2][3] describe that a host MUST accept and process any Query whose IP Destination Address field contains any of the addresses (unicast or multicast) assigned to the interface on which the Query arrives. According to the scenario, a router can unicast General Query to tracked member hosts in [Query Interval] (or [Unicast Query Interval] newly defined in [10]), if the router keeps track of membership information (Section 3). Unicasting IGMP/MLD General Query would be effective especially when a wireless link is heavily loaded.

If a multicast router attached to a wireless link enables an explicit tracking function and unicasts IGMP/MLD General Query for each member host, the General Query messages do not affect resources of non-member hosts. And since the router recognizes the (mostly) actual member hosts, whether IGMP/MLD General Query messages are transmitted by unicast or multicast, the router can configure longer [Query Interval] value and send IGMP/MLD Group-Specific and Group-and-Source Specific Queries when it recognizes the last member has left from the network. This will reduce the number of both IGMP/MLD General Query and Current-State Report messages.

Note that longer query interval may increase join latency and leave

latency when an unsolicited message with State-Change Record is not reached to the router.

There is another concern in unicast General Query. If a multicast router sends General Query "only" by unicast, it cannot discover potential member hosts whose join requests were lost. Since the hosts do not send the same join requests (i.e., unsolicited Report messages) again, they loose the chance to join the channels unless the upstream router asks the membership information by sending General Query by multicast. It will be solved by using both unicast and multicast General Queries and configuring the [Query Interval] timer value for multicast General Query and the [Unicast Query Interval] timer value for unicast General Query as proposed by [10]. However, using two different timers for General Queries may require the protocol extension this document does not focus on.

IGMP/MLD Group-Specific and Group-and-Source Specific Queries defined in [2][3] are sent to verify whether there are hosts that desire reception of the specified group or a set of sources or to rebuild the desired reception state for a particular group or a set of sources. These specific Queries build and refresh multicast membership state of hosts on an attached network. These specific Queries should be sent to each corresponding multicast address (not the all-hosts/all-nodes multicast address) as their IP destination addresses, because hosts that do not join the multicast session do not pay attention these specific Queries, and only active member hosts that have been receiving multicast contents with the specified address reply IGMP/MLD reports.

5. Interoperability

IGMPv3 [2] and MLDv2 [3] provide the ability for hosts to report source-specific subscriptions. With IGMPv3/MLDv2, a mobile host can specify a channel of interest, using multicast group and source addresses in its join request. Upon its reception, the upstream router that supports IGMPv3/MLDv2 establishes the shortest path tree toward the source without coordinating a shared tree. This function is called the source filtering function and required to support Source-Specific Multicast (SSM) [8].

Recently, the Lightweight-IGMPv3 (LW-IGMPv3) and Lightweight-MLDv2 (LW-MLDv2) [4] protocols have been proposed in the IETF. These protocols provide protocol simplicity for mobile hosts and routers, as they eliminate a complex state machine from the full versions of IGMPv3 and MLDv2, and promote the opportunity to implement SSM in mobile communications.

This document assumes multicast routers that deal with mobile hosts and mobile hosts MUST be IGMPv3/MLDv2 capable, regardless whether the protocols are the full or lightweight version. However, this document does not consider interoperability with older version protocols. The main reason not being interoperate with older IGMP/ MLD protocols is that the explicit tracking function does not work properly with older IGMP/MLD protocols.

6. Timers, Counters, and Their Default Values

The [Query Interval] is the interval between General Queries sent by the regular IGMPv3/MLDv2 querier, and the default value is 125 seconds [2][3]. By varying the [Query Interval], multicast routers can tune the number of IGMP messages on the network; larger values cause IGMP Queries to be sent less often.

This document proposes to inherit the default [Query Interval] value, 125 seconds, in case that a router enables the explicit tracking function and sends General Query to tracked member hosts by unicast. Nevertheless, the last-hop router may want to configure longer [Query Interval] value such as 150 seconds when operators expect the router will potentially maintains a number of member hosts such as more than 100 hosts on the wireless link.

This document also proposes 180 seconds for the [Query Interval] value in case that a router enables the explicit tracking function and sends General Query by multicast. This longer value contributes to minimizing traffic of Report messages and battery power consumption for mobile hosts, especially when operators expect the router will potentially maintains a large number of member hosts such as more than 500 hosts on the wireless link. Multicast General Query is necessary to update membership information if it is not correctly synchronized due to missing Reports.

The [Query Response Interval] is the Max Response Time (or Max Response Delay) used to calculate the Max Resp Code inserted into the periodic General Queries. Its default value is 10 seconds expressed by "100" for IGMPv3 [2] and "10000" for [3]. By varying the [Query Response Interval], multicast routers can tune the burstiness of IGMP/MLD messages on the network; larger values make the traffic less bursty as host responses are spread out over a larger interval, but will increase join latency when State-Change Report is missing.

This document proposes 5 seconds (expressed by "50" for IGMPv3 and "5000" for MLDv2) for the [Query Response Interval] value in case that a router enables the explicit tracking function and sends General Query by unicast. This shorter value introduces quick response, which is valuable for the use of unicast General Query. Note that the shorter [Query Response Interval] value may cause sudden increase in a traffic especially when a large number of member hosts attaches on the network; therefore it SHOULD NOT be 1 or smaller second.

In case that a router multicasts General Query, the default [Query Response Interval] value, 10 seconds, can be reasonablly used.

To cover the possibility of unsolicited reports being missed by multicast routers, unsolicited reports are retransmitted [Robustness Variable] - 1 more times, at intervals chosen at random from the defined range [2][3]. The QRV (Querier's Robustness Variable) field in IGMP/MLD Query contains the [Robustness Variable] value used by the querier. Routers adopt the QRV value from the most recently received Query as their own [Robustness Variable] value, whose range SHOULD be set between "1" to "7". The default [Robustness Variable] value defined in IGMPv3 [2] and MLDv2 [3] is "2". This document proposes "2" for the [Robustness Variable] value for mobility. And whether the explicit tracking function is enabled or not, the [Robustness Variable] value SHOULD NOT be bigger than "2".

The [Startup Query Interval] is the interval between General Queries sent by a Querier on startup. The default value is 1/4 of [Query Interval]; however, this document recommends the use of its shortended value such as 1 second since the shorter value would contribute to smooth handover for mobile hosts using e.g., PMIPv6 [13]. Note that the [Startup Query Interval] is a static value and cannot be changed by any external signal. Therefore operators who maintain routers and wireless links must properly configure this value.

7. Security Considerations

This document neither provides new functions or modifies the standard functions defined in [2][3][4]. Therefore there is no additional security consideration provided.

8. Acknowledgements

Marshall Eubanks, Gorry Fairhurst, Behcet Sarikaya, Thomas C. Schmidt, Stig Venaas, Jinwei Xia, and others provided many constructive and insightful comments.

9. References

9.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to indicate requirement levels", RFC 2119, March 1997.

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

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

[4] Liu, H., Cao, W., and H. Asaeda, "Lightweight IGMPv3 and MLDv2 Protocols", RFC 5790, February 2010.

[5] Deering, S., "Host Extensions for IP Multicasting", RFC 1112, August 1989.

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

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

[8] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[9] Asaeda, H. and Y. Uchida, "IGMP/MLD-Based Explicit Membership Tracking Function for Multicast Routers", draft-asaeda-mboned-explicit-tracking-01.txt (work in progress), October 2010.

9.2. Informative References

[10] Asaeda, H. and T. Schmidt, "IGMP and MLD Protocol Extensions for Mobility", draft-asaeda-multimob-igmp-mld-mobility-extensions-04.txt (work in progress), March 2010.

[11] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[12] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")",

RFC 4605, August 2006.

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

Authors' Addresses

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-0882
Japan

Email: asaeda@wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~asaeda/

Yogo Uchida
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-0882
Japan

Email: uchida@sfc.wide.ad.jp

MULTIMOB Working Group
INTERNET-DRAFT
Intended Status: Experimental
Expires: December 31, 2010

Luis M. Contreras
Carlos J. Bernardos
Ignacio Soto
Universidad Carlos III de Madrid
June 29, 2010

Rapid acquisition of the MN multicast subscription after handover
<draft-contreras-multimob-rams-00.txt>

Status of this Memo

Copyright and License Notice

Abstract

A new proposal is presented for speeding up the acquisition by the MAG of the MN's active multicast subscription information, in order to accelerate the multicast delivery to the MN during handover. To do that, an extension of the current PMIPv6 protocol is required. The solution described in this memo is not only applicable to the base multicast solution, but also it can be applied to other solutions envisioned as possible architectural evolutions of it. Furthermore, it is also independent of the role played by the MAG within the multicast metwork (either acting as MLD proxy or multicast router).

Table of Contents

1 Introduction

Recently, a base solution has been adopted for continuous multicast service delivery in PMIPv6 domains [4]. That solution specifies the basic functionality needed in the PMIPv6 entities to provide a multicast service, and supports the continuous delivery of multicast obtaining the on-going multicast subscription information directly from the MN after handover. Thus, once the MN attaches to a new MAG, the MN is interrogated by the MAG through an MLD General Query, which is sent just after any new link sets up, to get knowledge of any existing subscription, as specified in [2].

However, as highlighted by [5], the base solution must be improved to cover some performance requirements, especially those referred to the user perceived service quality, seriously affected by the disruption of multicast content forwarding to the MN during handovers.

The method used in the base solution to get knowledge of an existing multicast subscription relies on the intrinsics of the IGMP/MLD protocols. Both protocols send multicast membership interrogation messages when a new link is up. The answer to that request will report any existing multicast subscription by the MN.

Due to this behaviour, despite of being a straightforward method, the MAG can incur in a huge delay in receiving the corresponding MLD Report message caused by either the MLD query processing time or the radio transfer delays associated with this procedure.

The new approach proposed here consists of the extension of the current PMIPv6 signaling protocol defined in [1] by including a new multicast information option to update PMIPv6 entities during registration and de-registration processes, as well as new messages to trigger the transfer of such multicast information. No extension is considered for any of the multicast-related protocols (IGMP/MLD nor PIM protocols).

This proposal intends to provide a signaling method internal to the network to speed up the subscription information acquisition by the MAG, in order to accelerate the multicast delivery to the MN. By doing so, the knowledge by the MAG of the currently active multicast subscription becomes independent of the underlying radio technology dynamics and relaxes the requirement of a rapid response from the MN in processing MLD control messages. Issues like radio framing, radio access contention, channel reliability, IGMP/MLD timers optimisation for wireless environments, etc, are not relevant any more to determine multicast performance after handoff.

The solution described in this memo is not only applicable to the

base solution defined in [4], but also it can be applied to other solutions envisioned as possible architectural evolutions of it, as those stated in [6] or [7]. Furthermore, it is also independent of the role played by the MAG within the multicast network (either acting as MLD proxy or multicast router).

1.1 Conventions and Terminology

This document uses the terminology referring to PMIPv6 components as defined in [1]. 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 finally registered in a handover event.

Reactive Handover
: A reactive handover is a handover event where the LMA receives the MN registration from the nMAG without having previously received the MN de-registration from the pMAG.

Predictive handover
: A predictive handover is a handover event where the LMA firstly receives the MN de-registration from the pMAG previously to receive the MN registration from the nMAG.

2 Overview

The LMA is a key element within the PMIPv6 infrastructure. It traces the MN reachability along the PMIPv6 domain, therefore the LMA is the best element to store and forward the multicast subscription information to the rest of entities within the PMIPv6, that is, to the MAGs, as the MN moves.

The LMA only requires to know the detailed subscription information (in terms of the IP addresses of both the multicast group subscribed, G, and the source delivering it, S) during the handover event. Apart from the handover event, it is not worthy to continuously inform the LMA about it. Such procedure would significantly increase the signaling load within the PMIPv6 domain without a clear benefit. The subscription information (S,G) is only critical during handover, 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.

To do that, it will be necessary to extend the PMIPv6 protocol in several ways. First of all, a new mobility option is needed to pack the IP addresses of the current multicast subscription. Furthermore, additional messages are required 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.

3 PMIPv6 extensions

This section outlines the extensions proposed to the PMIPv6 protocol specified in [1].

3.1 New "Active Multicast Subscription" mobility option

3.1.1 Option application rules

A new TLV-encoded mobility option, "Active Multicast Subscription" option is defined for use with the PBU and PBA messages exchanged between an LMA and a MAG to transfer the multicast subscription information. This option is used for exchanging the IP addresses of both the group subscribed by the MN, and the source delivering it as well. 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 taken place.

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

3.1.2 Option format

The format of this new option is as follows:

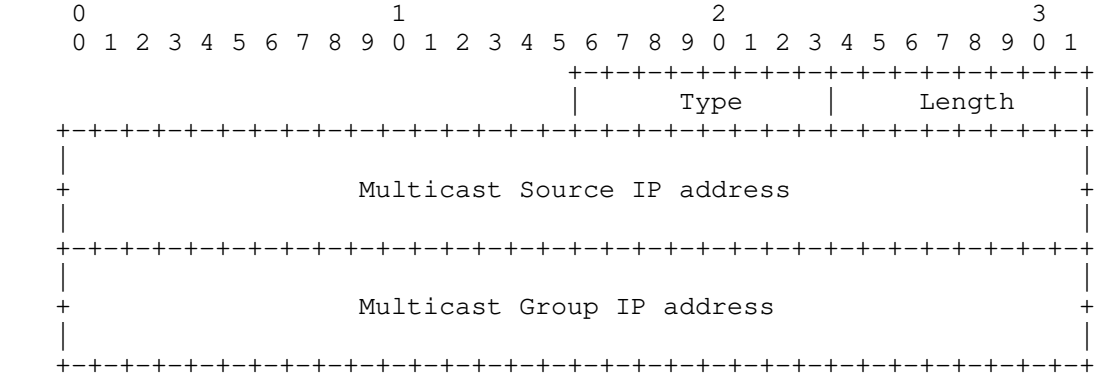

Type
To be defined

Length
8-bit unsigned integer indicating the length of the option in octects, excluding the type and length fields. This field must be set to the value 8 for IPv4, and 32 for IPv6.

Multicast Source IP address

Unicast IP address of the node which injects the multicast content in the network.

Multicast Group IP address
Multicast IP address identifying the content which the MN subscribes to.

3.2 New "multicast Signaling" flag on PBU/PBA message headers

3.2.1 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 subject of 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.

3.2.1.1 Registration process

These rules 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 predictive) 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, without any further meaning.

  * If S=1 in the corresponding PBU message, two sub-cases can happen

    o S=1 in the PBA message if the multicast subscription information is provided in this message for the MN. When S=1, if the MN maintains an active multicast session, the PBA

message will include the "Active Multicast Subscription" mobility option with the IP addresses of the subscribed group and the source providing it.

o S=0 in the PBA message if the multicast subscription information is not provided in this message for the MN. The PBA message will include the "Active Multicast Subscription" mobility option with the IP addresses of the group and the source set to 0. This case is useful to decouple unicast and multicast signaling for a MN being registered at nMAG. A way for obtaining later active multicast-subscription information is described later in this document.

3.2.1.2 De-registration process

These rules apply for the Binding De-registration process

o PBU message

* S=0, it indicates that the MN has no active multicast session.

* S=1, it indicates that the MN has an active multicast session, and the IP addresses of the subscribed group and the source providing it are transported in the "Active Multicast Subscription" mobility option.

o PBA message

The value of the flag in the PBA message should be a copy of the value received in the PBU message, without any further meaning.

3.2.2 New format of conventional PBU/PBA messages

3.2.2.1 Proxy Binding Update Message

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

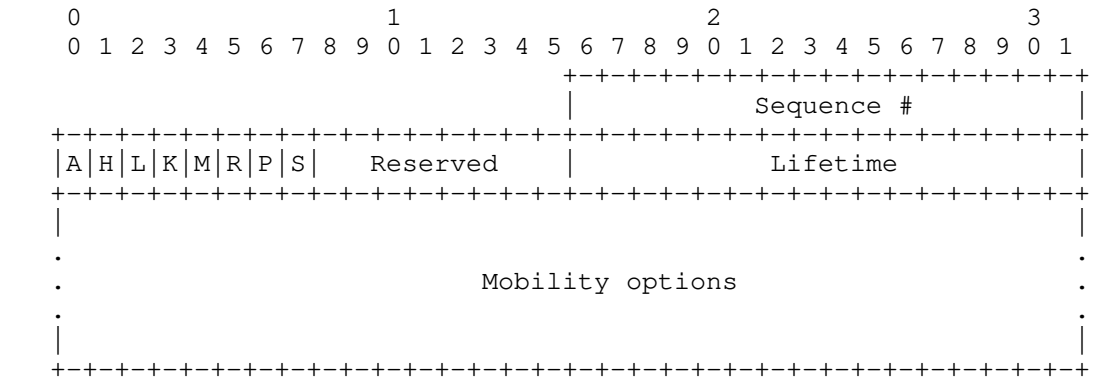

3.2.2.2 Proxy Binding Acknowledgement Message

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

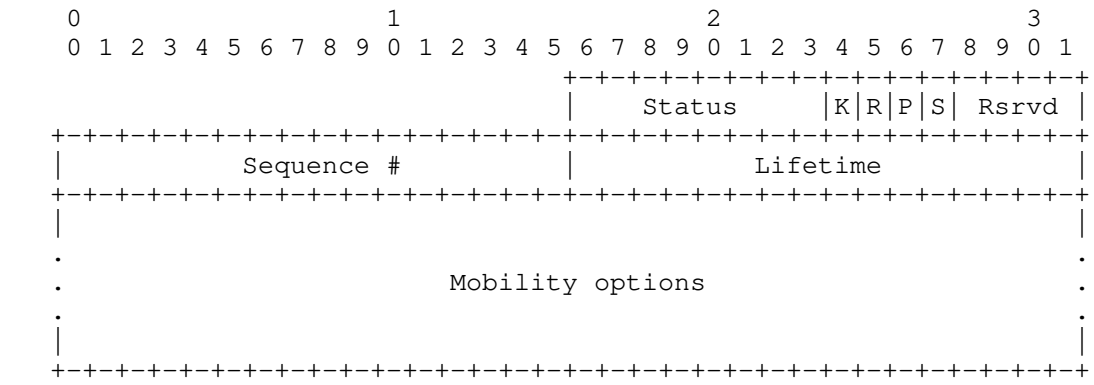

3.3 New optional "multicast Active" flag on LMA Binding Cache

3.3.1 Flag application rules

A new optional 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 interrogation 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 interrogate or not the pMAG (after receiving a PBU message from the nMAG) is as follows:

- 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 interrogate the pMAG, and registers the MN as attached to the nMAG as usual.

- 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 that multicast signaling is not supported by the nMAG for that MN, the LMA does not interrogate the pMAG, and registers the MN as attached to the nMAG as usual.

- 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 interrogate the pMAG, and registers the MN as attached to the nMAG as usual.

- 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 interrogates the pMAG to obtain the multicast subscription details in the form of (S,G) previously to complete the registration of the MN attached to the nMAG.

The flag A should be initialized to the value 0.

3.4 New messages for active multicast subscription interrogation

A new pair of messages is defined for interrogating 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 [3].

3.4.1 Subscription Query message

3.4.1.1 Message application rules

The Subscription Query message is sent by the LMA towards the pMAG to interrogate it about any existing multicast subscription of the MN which is being registered by the nMAG. This message is generated in case of the handover is of reactive type.

Additionally, this message is sent by the nMAG towards the LMA to interrogate it about the existing multicast subscription of the MN when the LMA acknowledges the PBU sent by the nMAG but the multicast information is not provided (in detail, when the PBU messages has set the flag S to 1, and the PBA message has set the flag S to 0).

3.4.1.2 Message format

The Subscription Query message has the following format.

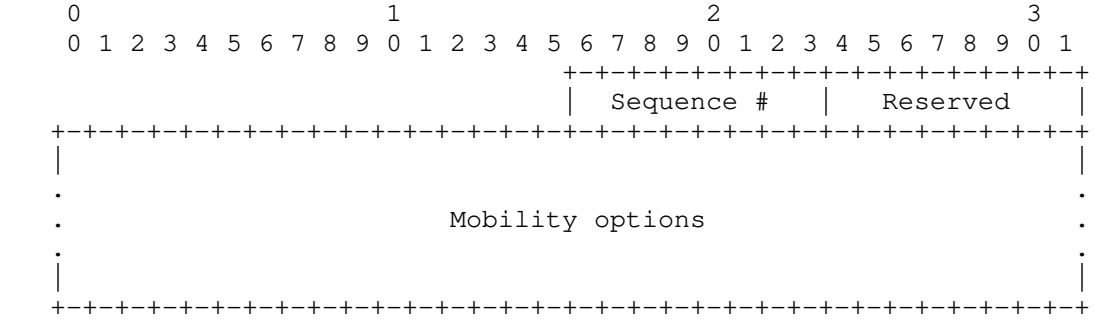

Sequence Number

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 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.

3.4.2 Subscription Response message

3.4.2.1 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.

3.4.2.2 Message format

The Subscription Response message has the following format.

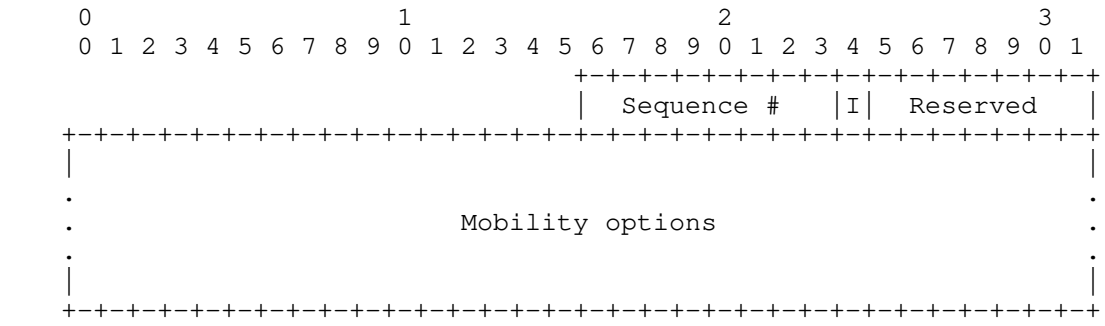

Sequence Number

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.

3.5 New 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 [3].

3.5.1 Multicast Activity Indication message

3.5.1.1 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 should be asked or not about multicast information corresponding to the MN being registered at the nMAG, in case of the handover is of reactive type.

3.5.1.2 Message format

The Multicast Activity Indication message has the following format.

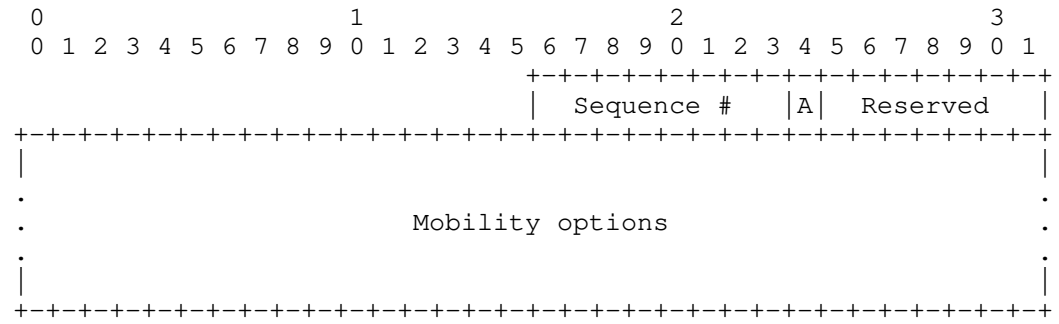

Sequence Number

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.

3.5.2 Multicast Activity Indication Acknowledge message

3.5.2.1 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.

3.5.2.2 Message format

The Multicast Activity Indication message has the following format.

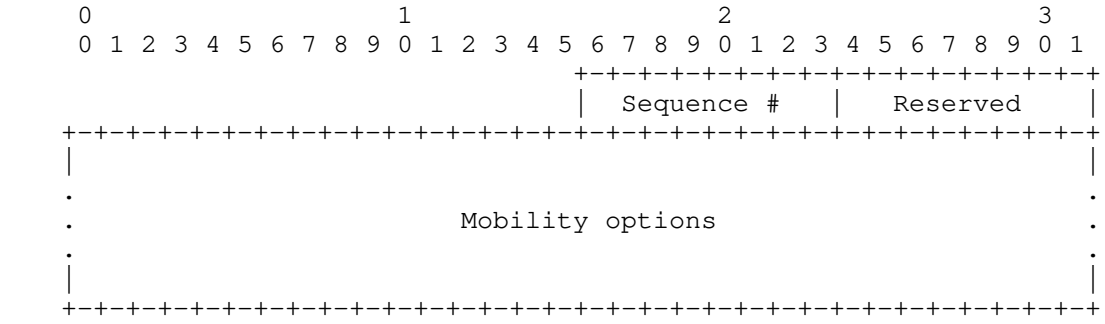

Sequence Number

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.

3.6 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 to decouple the unicast signaling from the multicast one.

This timer should be upper bounded by the constant defined in [3] INIT_BINDACK_TIMEOUT, which 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.

4 Signaling process description

As the MN moves from one access gateway (named previous-MAG, pMAG) to another (named new-MAG, nMAG), 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 should 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.

4.1 Handover of predictive type

4.1.1 Rationale

In the predictive case, the LMA firstly receives the MN de-registration from the pMAG previously to receive the MN registration from the nMAG.

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 IP addresses of the multicast subscription(s) active in 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 addresses of the group(s) subscribed by the MN, and the source(s) delivering it(them), 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.

4.1.2 Message flow description

The figure 1 summarizes this process.

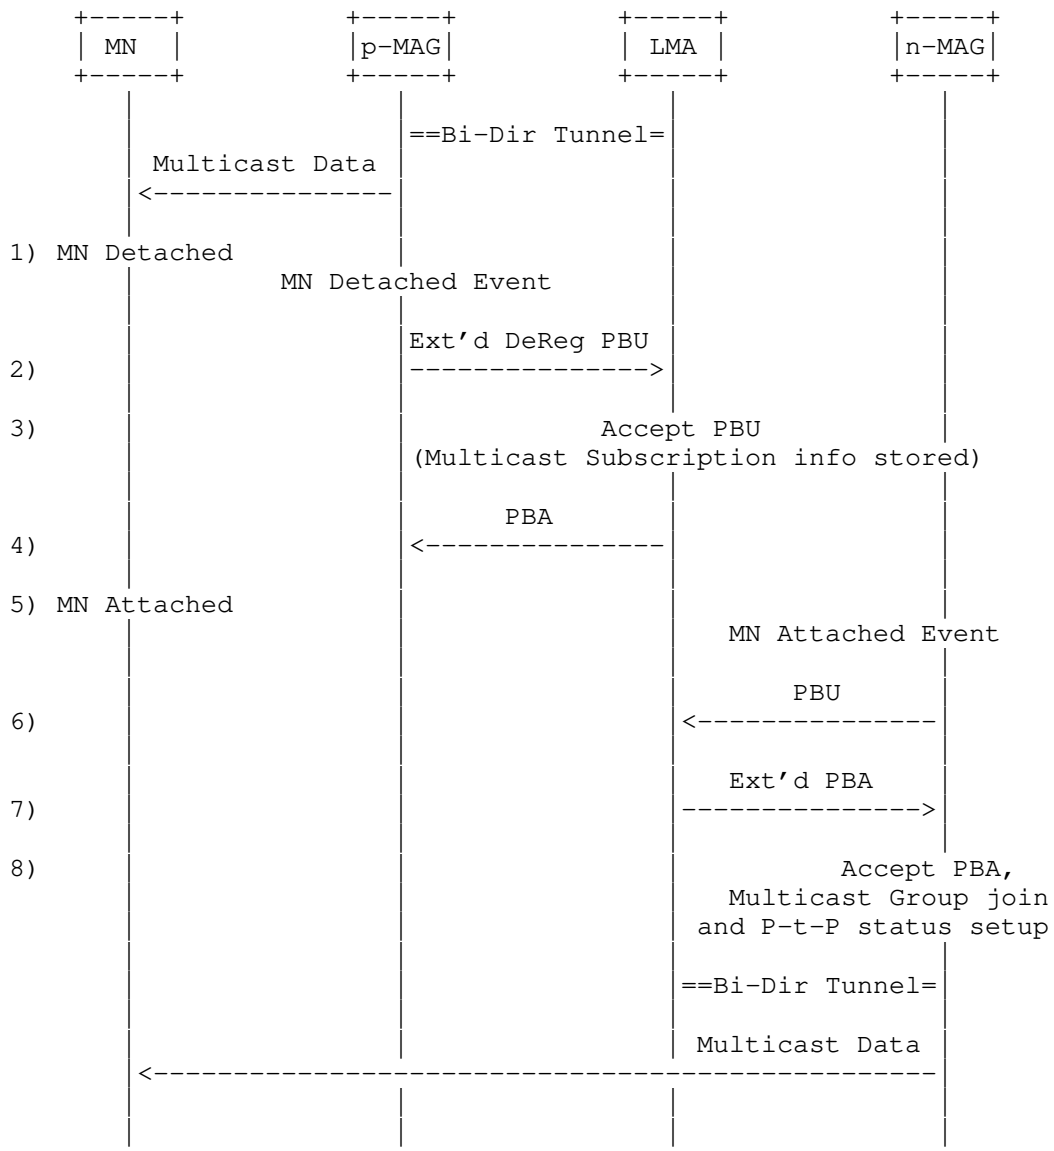

Figure 1. Predictive handover

The sequence of messages is the following:

1) A registered MN is receiving a multicast content which has been previously subscribed by sending an 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 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 IP addresses of the (S,G) pairs 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 when the handover has been 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 [1].

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 IP addresses of the (S,G) pairs of the active multicast subscriptions in the moment of handover.

8) The nMAG processes the PBA message, following all the standard procedures described in [1]. 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 present MN attached to it. From that instant, the multicast content is served to the MN.

4.2 Handover of reactive type

4.2.1 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.

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.

4.2.2 Message flow description

The set of figures 2a to 2d summarize this process.

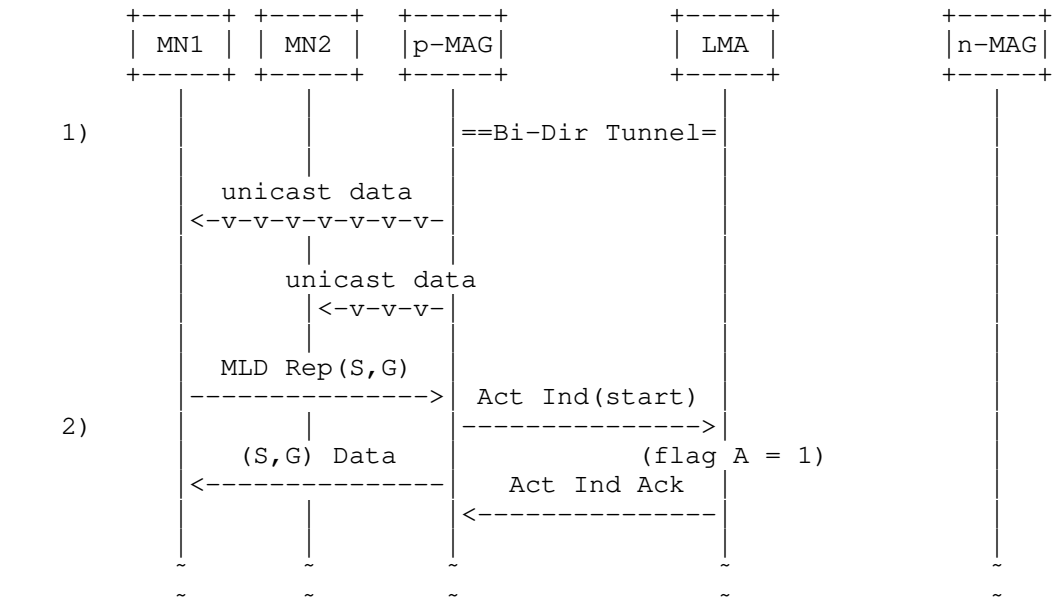

Figure 2a. Reactive handover (steps 1 to 2)

The sequence of messages is the following:

1) A pair of MNs, named MN1 and MN2, are attached to the pMAG. Both MNs are multicast-enabled nodes, and both MNs are only receiving unicast traffic as usual in PMIPv6 domains, with no multicast subscription yet. At some point of time, the MN1 request to the pMAG

to be subscribed to the content identified by the IP addresses (S,G), by sending an standard MLD report from the MN to the pMAG. The pMAG will keep the multicast status state of the point-to-point link with the MN. The multicast flow (S,G) is then forwarded by the pMAG to the MN1.

2) Due to this initial multicast subscription for the MN1, the pMAG 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 pMAG to acknowledge the previous indication.

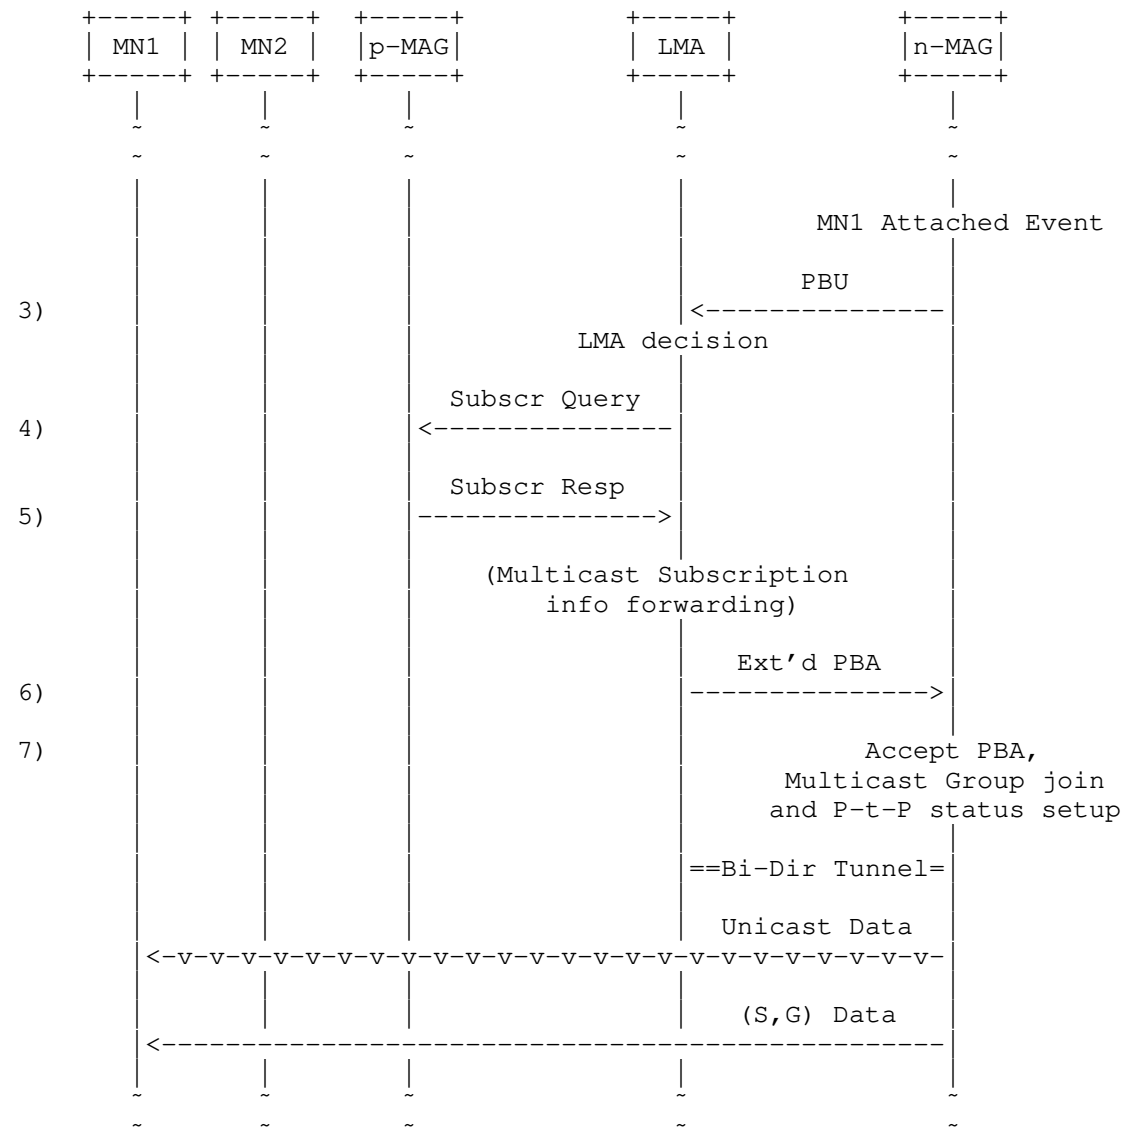

Figure 2b. Reactive handover (steps 3 to 7)

3) Some time later, the MN1 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.

4) 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 interrogates the pMAG about if there is any active multicast subscription for the MN1, by sending a Subscription Query message.

5) The pMAG answers the LMA with a Subscription Response message including the IP addresses of the existing subscriptions (the pair (S,G) in this case).

6) After processing the pMAG answer, the LMA acknowledges the PBU message, including the multicast subscription information within the new TLV-encoded option "Active Multicast Subscription". The nMAG then process the extended PBA message.

7) 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.

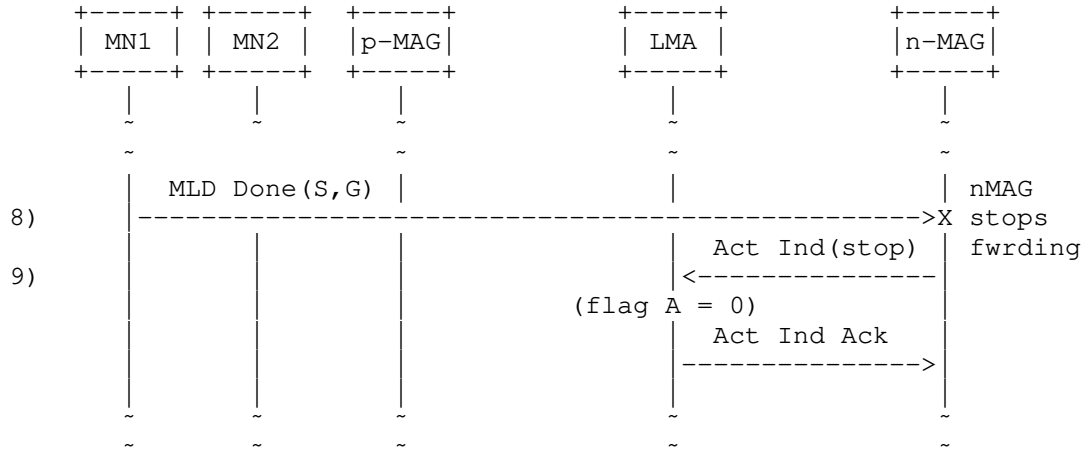

Figure 2c. Reactive handover (steps 8 to 9)

8) Some time later, the MN1 decides to totally stop all the active multicast subscriptions that it maintains. The MN1 will send an MLD Done message to nMAG to request the cease of the multicast traffic delivery. As consequence, the nMAG will stop all the multicast

traffic forwarding to the MN1.

9) After removing the active subscriptions for the MN1, the nMAG 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 nMAG to acknowledge the previous indication.

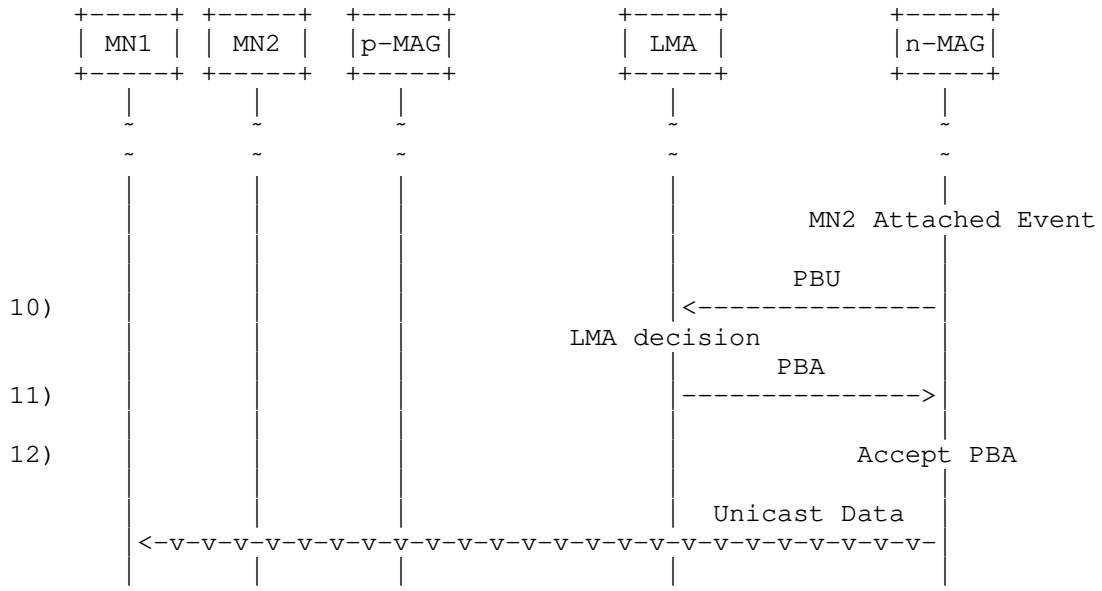

Figure 2d. Reactive handover(steps 10 to 12)

10) 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 neither 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.

11) 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 interrogate the pMAG, and acknowledges the PBU message. The nMAG then process the extended PBA message.

12) The nMAG is now ready to forward the unicast traffic to the MN2.

4.2.3 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 not loose 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.

4.2.4 Prevention of large delays of the binding acknowledgement for unicast traffic

Attending to the message sequences detailed above for reactive handovers, in case the LMA has to request the multicast subscription information to 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 to the LMA towards the MN, a mechanism should be implemented to decouple multicast from unicast traffic reception by the MN.

The figures 3a and 3b show this mechanism:

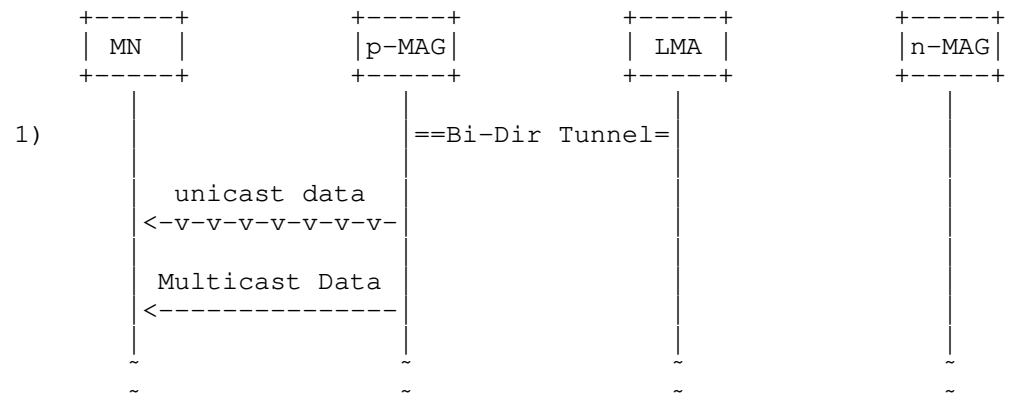

Figure 3a. Decoupling of unicast and multicast signaling (step 1)

The sequence of messages is the following:

1) An MN, named MN1, is attached to the pMAG. The MN is a multicast-enabled node, and it is receiving both unicast and multicast traffic simultaneously.

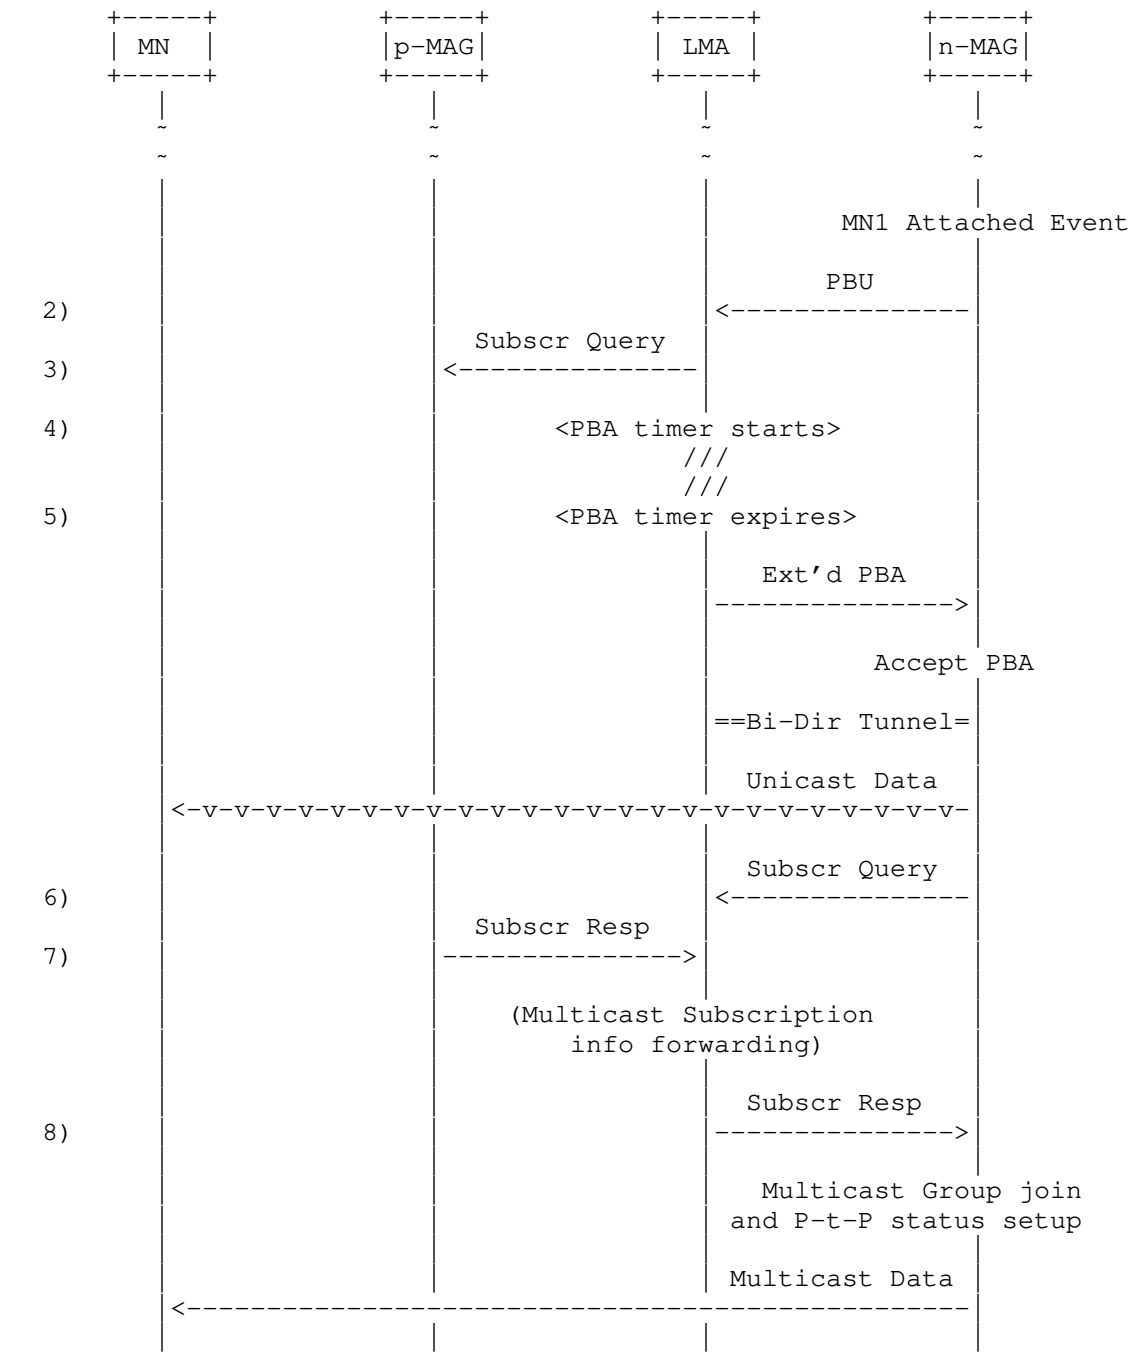

Figure 3b. Decoupling of unicast and multicast signaling (steps 2 to 8)

2) Some time later, the MN1 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 interrogate the pMAG about if there is any active multicast subscription for the MN1, 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 duration 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=0. 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 by another MN connection).

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 IP addresses of the existing subscriptions (the pair (S,G) in this case).

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 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 also be served to the MN.

5 Security Considerations

TBD.

6 IANA Considerations

This document defines the new following elements which values should be allocated:

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.

o Flags: the multicast Signaling (S), the multicast Information (I), and the multicast Active (A) flags.

7 References

7.1 Normative References

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

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

[3] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

7.2 Informative References

[4] TC. Schmidt, M. Waehlisch, and S. Krishnan, "A Minimal Deployment Option for Multicast Listeners in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-03.txt, (work in progress), June 2010.

[5] D. von Hugo, H. Asaeda, B. Sarikaya, and P. Seite, "Evaluation of further issues on Multicast Mobility: Potential future work for WG MultiMob", draft-von-hugo-multimob-future-work-02, (work in progress), June 2010.

[6] JC. Zuniga, G. Lu, and A. Rahman, "Support Multicast Services Using Proxy Mobile IPv6", draft-zuniga-multimob-smspmip-03, (work in progress), May 2010.

[7] S. Jeon and Y. Kim, "Mobile Multicasting Support in Proxy Mobile IPv6", draft-sijeon-multimob-mms-pmip6-02, (work in

progress), March 2010

8 Acknowledgments

The research of Carlos J. Bernardos leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 214994 (CARMEN project), and from the Ministry of Science and Innovation of Spain, under the QUARTET project (TIN2009-13992-C02-01).

Author's Addresses

Luis M. Contreras
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Email: luisc@it.uc3m.es

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Email: cjbc@it.uc3m.es

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

Email: isoto@it.uc3m.es

Internet Area
Internet Draft
Intended status: Informational
Expires: Jan 01, 2011

M.Hui
G.Chen
H.Deng
China Mobile

July 01, 2010

Fast Handover for Multicast in Proxy Mobile IPv6
draft-hui-multimob-fast-handover-02

Status of this Memo

This Internet-Draft will expire on Jan 01, 2011.

Copyright Notice

Abstract

This document specifies the predictive fast handover mechanism to solve the problem of handover latency and packet loss in Proxy Mobile IPv6 Multicast. Necessary extensions are specified for Handover Initiate (HI) and Handover Acknowledgement (HAck) messages to support multicast handover procedure.

Table of Contents

1. Introduction

Proxy Mobile IPv6 (PMIPv6) protocol provides local mobility management to a mobile node without requiring any modification of the mobile node. The Local Mobility Anchor (LMA) and Mobile Access Gateway (MAG) perform the mobility management signaling on behalf of the mobile node. Extensions for LMA and MAG are specified in [1] to support IP multicast in PMIPv6. Nevertheless, the basic performance including handover latency and packet loss is not considered different form that of PMIPv6.

Fast handover for Mobile IPv6 is specified in [2]. [3] extends the FMIPv6 and applies it to the PMIPv6 in order to decrease handover latency and packet loss as well as transfer of network-resident contexts. However, IP multicast is not considered in fast handover for PMIPv6.

We propose a fast handover mechanism to support multicast for PMIPv6. Necessary extensions are specified in HI and HAck message to transfer the multicast node’s context information and deliver the multicast data before the set up of tunnel between n-MAG and LMA.

2. Problem Statement

The existing solution for PMIPv6 multicast [1] specifies that, only after the bi-directional tunnel is built between n-MAG and LMA using extended PBU (PBU-M) message, the multicast packet can be continuously delivered to MN. It inevitably causes the latency and loss of packet during handover process.

The solution presents two ways to acquire the MN's profile, which includes MN' ID and multicast state information. One way is to use the Context Transfer Protocol (CXTP) [4] to transfer MN's profile from p-MAG to n-MAG. In the other way, if MN's profile is stored in a policy store [5], n-MAG obtains MN's multicast state by the same mechanism used to acquire MN' ID and profile during MN's attachment process [5].

In another PMIPv6 multicast solution [6], the author proposes normal handover and fast handover for proxy mobile multicast service. There is no any optimization in normal handover, the handover involves MN by running the MLDv2 [7] protocol with n-MAG to receive the related multicast packet. In the fast handover procedure, similar to the first method used in [1],the context transfer is used to provide multicast information. Although n-MAG can acquire the MN' multicast information before MN handovers to it, only after n-MAG joins the multicast group, it can receive the multicast data.

3. Terminology

This document refers to [1] [2] [3] for terminology. The following terms and abbreviations are additionally used in this document. The reference network is illustrated in Figure 1.

Previous Mobile Access Gateway (p-MAG):

The MAG that manages mobility related signaling for the MN before handover.

New Mobile Access Gateway (n-MAG):

The MAG that manages mobility related signaling for the MN after handover.

HO-Initiate:

A generic signaling that indicates the handover of the MN sent from the MN to the p-MAG. It is assumed that HO-Initiate can carry the information to identify the MN and to assist the p-MAG to resolve the n-MAG.

4. Protocol Operation

The architecture of fast handover for multicast in Proxy Mobile IPv6 is shown in Figure 1. A multicast tunnel is established to transfer the multicast data from p-MAG to n-MAG before the n-MAG joins the multicast group, so that whenever the MN handovers to the n-MAG, it can receive the multicast data from n-MAG.

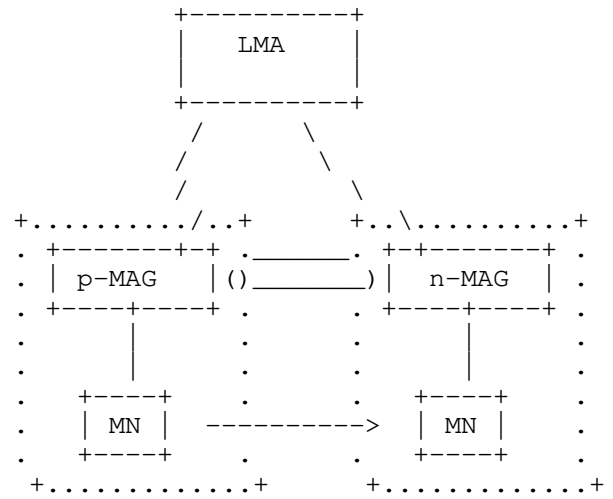

Figure 1: Reference network for fast handover

In order to decrease the handover latency and packet loss, this document specifies a bi-directional tunnel between the Previous MAG (p-MAG) and the New MAG (n-MAG). As the n-MAG needs the multicast node's context information to set up a bi-directional tunnel to continuous deliver multicast packet to mobile node, the HI and HAck messages are extended to support mobile multicast node's context transfer, in which parameters such as MN ID, MN Multicast State, are transferred from the p-MAG to the n-MAG. The sequence of events illustrating the fast handover for multicast is shown in Figure 2.

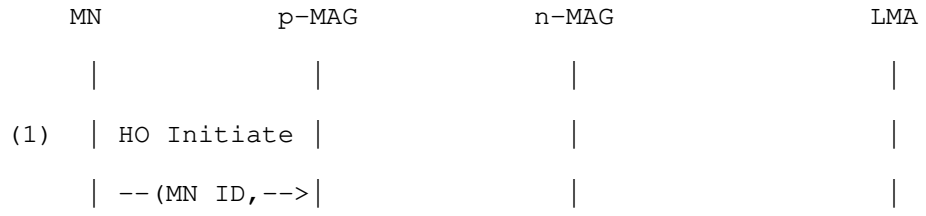

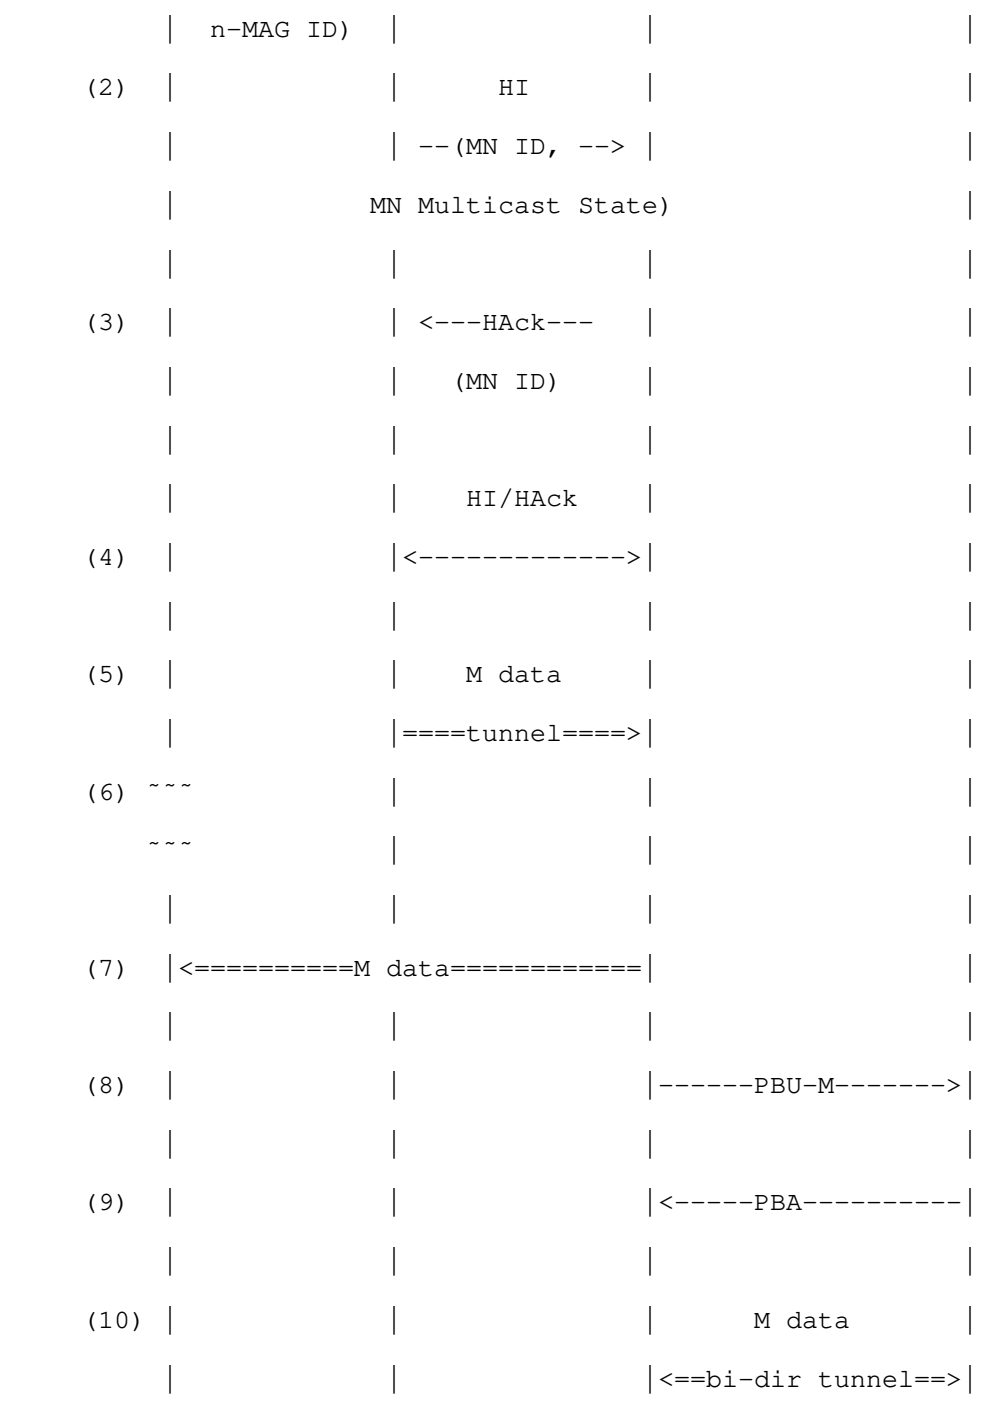
n-MAG ID)
(2)
HI
--(MN ID, -->
MN Multicast State)
(3)
<---HAck---
(MN ID)
HI/HAck
(4)
(5)
M data
====tunnel====>
(6)
(7)
<==========M data============
(8)
------PBU-M------->
(9)
<-----PBA----------
(10)
M data
<==bi-dir tunnel==>

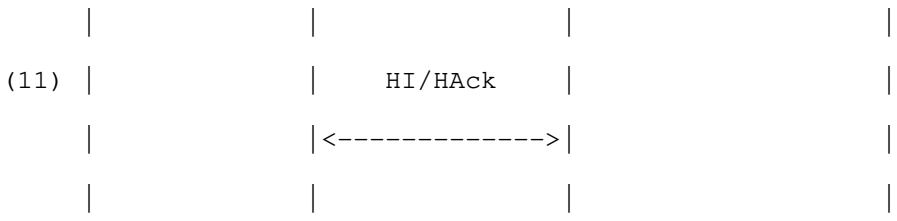

Figure 2: Fast handover for PMIPv6 multicast

The detailed descriptions are as follows:

(1) The MN detects that a handover is imminent and reports the MN ID and n-MAG ID.

(2) The p-MAG sends the HI to the n-MAG. The HI message includes MN ID and MN Multicast State.

(3) The n-MAG sends the HAck back to the p-MAG.

(4) The n-MAG requests the p-MAG to forward multicast packets by setting F flags in the HI message.

(5) A tunnel is established between the p-MAG and n-MAG and multicast packets destined for the MN are forwarded from the p-MAG to the n-MAG over this tunnel.

(6) The MN undergoes handover to n-MAG.

(7) The n-MAG starts to forward multicast packets destined for the MN.

(8) The n-MAG sends the Proxy Binding Update with multicast extension (PBU-M)(proposed in [1]) to the LMA.

(9) The LMA sends back the Proxy Binding Acknowledgment (PBA) to the n-MAG.

(10) A bi-directional tunnel is set up for forwarding corresponding multicast data.

(11)Multicast packet forwarding is completed between p-MAG and n-MAG.

5. Message Format

This document defines new Mobility Header messages for the extended HI and HAck and new mobility options for delivering context information.

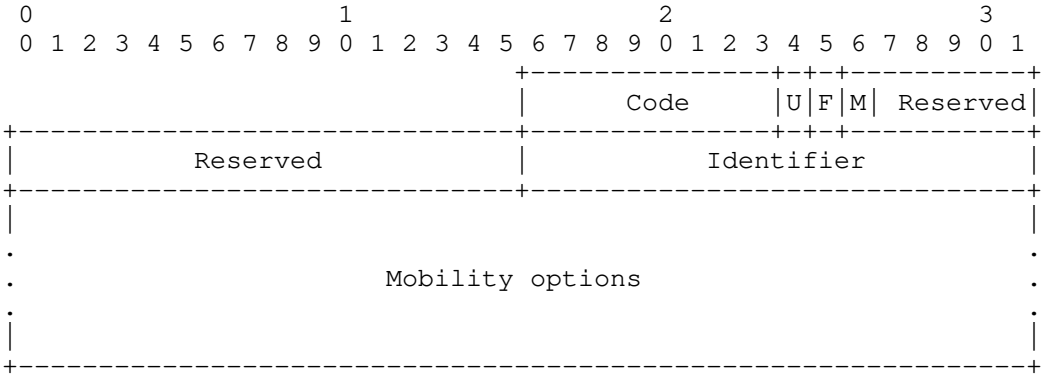

Figure 3: HI Mobility Header message with multicast extension

A new flag (M) is included in the HI Mobility Header message with multicast extension. The rest of the message format remains the same as defined in [3].

When (M) flag is specified in HI Mobility Header message, the mobility options field needs to be extended to include the multicast addresses.

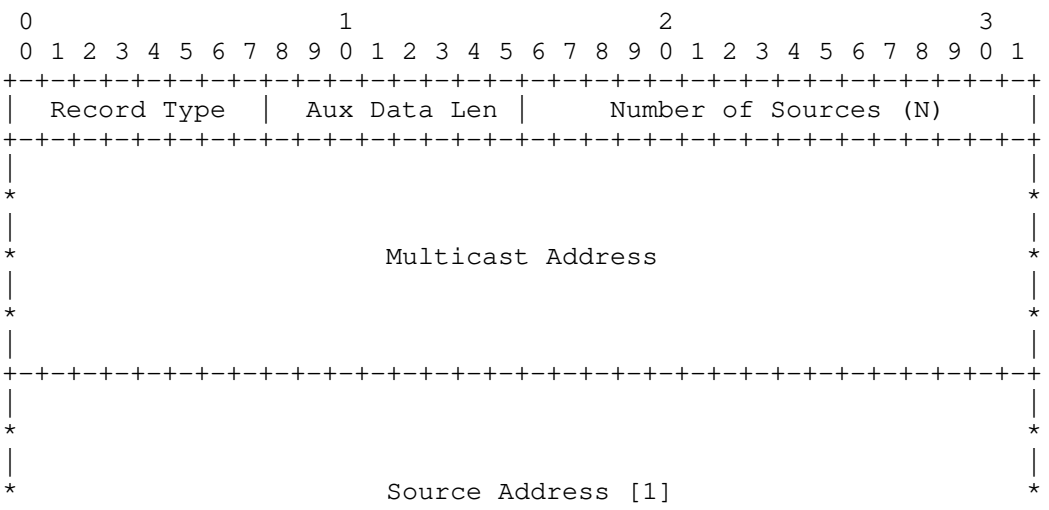

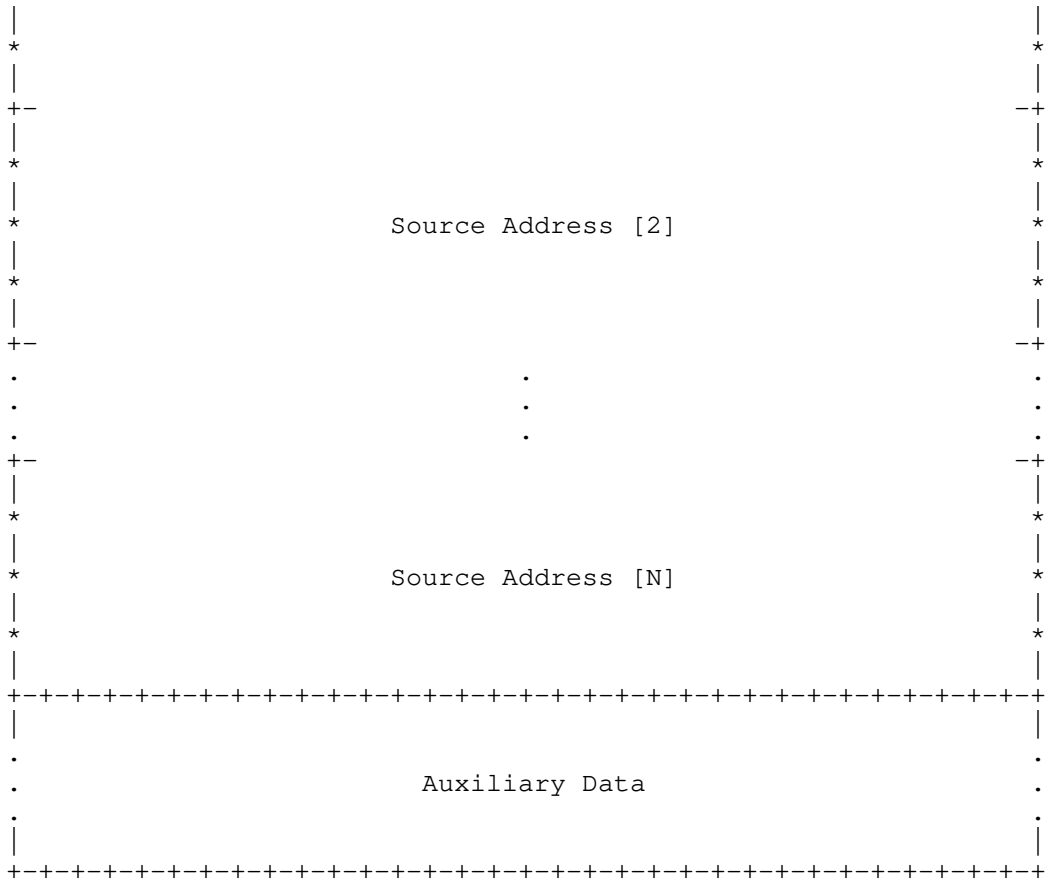
Source Address [2]
Source Address [N]
Auxiliary Data

6. Security Considerations

TBD.

7. IANA Considerations

This document does not require any IANA action.

8. References

8.1. Normative References

[1] H. Asaeda, P. Seite, J. Xia, "PMIPv6 Extensions for Multicast", draft-asaeda-multimob-pmip6-extension-01.txt (work in progress), July 2009.

[2] Koodli, R., Ed., "Mobile IPv6 Fast Handovers", RFC 5268, June 2008.

[3] H. Yokota, K. Chowdhury, R. Koodli, B. Patil, F. Xia, "Fast Handovers for Proxy Mobile IPv6", draft-irtf-mipshop-pfmipv6-00.txt (work in progress), October 2008.

[4] Loughney, Ed., J., Nakhjiri, M., Perkins, C., and R. Koodli,"Context Transfer Protocol (CXTP)", RFC 4067, July 2005.

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

[6] Y K. ZHAO, P. Seite, "The Solution for Pmipv6 Multicast Service", draft-zhao-multimob-pmipv6-solution-02.txt (work in progress), November 2008.

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

8.2. Informative References

Author's Addresses

Min Hui
China Mobile
53A,Xibianmennei Ave.,
Xuanwu District,
Beijing 100053
China
Email: huimin.cmcc@gmail.com

Gang Chen
China Mobile
53A,Xibianmennei Ave.,
Xuanwu District,
Beijing 100053
China
Email: phdgang@gmail.com

Hui Deng
China Mobile
53A,Xibianmennei Ave.,
Xuanwu District,
Beijing 100053
China
Email: denghui02@gmail.com

MULTIMOB Group
Internet-Draft
Intended status: BCP
Expires: April 28, 2011

T C. Schmidt
HAW Hamburg
M. Waehlisch
link-lab & FU Berlin
S. Krishnan
Ericsson
October 25, 2010

Base Deployment for Multicast Listener Support in PMIPv6 Domains
draft-ietf-multimob-pmipv6-base-solution-06

Abstract

This document describes deployment options for activating multicast listener functions in Proxy Mobile IPv6 domains without modifying mobility and multicast protocol standards. Similar to Home Agents in Mobile IPv6, Local Mobility Anchors of Proxy Mobile IPv6 serve as multicast subscription anchor points, while Mobile Access Gateways provide MLD proxy functions. In this scenario, Mobile Nodes remain agnostic of multicast mobility operations. A support for mobile multicast senders is outside the scope of this document.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

Proxy Mobile IPv6 (PMIPv6) [RFC5213] extends Mobile IPv6 (MIPv6) [RFC3775] by network-based management functions that enable IP mobility for a host without requiring its participation in any mobility-related signaling. Additional network entities called the Local Mobility Anchor (LMA), and Mobile Access Gateways (MAGs), are responsible for managing IP mobility on behalf of the mobile node (MN).

With these entities in place, the mobile node experiences an exceptional access topology towards the static Internet in the sense that the MAG introduces a routing hop also in situations, were the LMA architecturally acts as the next hop (or designated) router for the MN. In the particular case of multicast communication, group membership management as signaled by the Multicast Listener Discovery protocol (MLD) [RFC3810], [RFC2710] requires dedicated treatment at the network side .

Multicast routing functions need to be placed carefully within the PMIPv6 domain to augment unicast transmission with group communication services. [RFC5213] does not explicitly address multicast communication. Bi-directional home tunneling, the minimal multicast support arranged by MIPv6, cannot be directly transferred to network-based management scenarios, since a mobility-unaware node will not initiate such a tunnel after movement. Consequently, even a minimal multicast listener support in PMIPv6 domains requires an explicit deployment of additional functions.

This document describes options for deploying multicast listener functions in Proxy Mobile IPv6 domains without modifying mobility and multicast protocol standards. Similar to Home Agents in Mobile IPv6, PMIPv6 Local Mobility Anchors serve as multicast subscription anchor points, while Mobile Access Gateways provide MLD proxy functions. Mobile Nodes in this scenario remain agnostic of multicast mobility operations. This document does not address specific optimizations and efficiency improvements of multicast routing for network-based mobility discussed in [RFC5757], as such solutions would require changes to the base PMIPv6 protocol [RFC5213]. A support for mobile multicast senders is outside the scope of this document, as well.

2. Terminology

This document uses the terminology as defined for the mobility protocols [RFC3775], [RFC5213] and [RFC5844], as well as the multicast edge related protocols [RFC3376], [RFC3810] and [RFC4605].

3. Overview

The reference scenario for multicast deployment in Proxy Mobile IPv6 domains is illustrated in Figure 1.

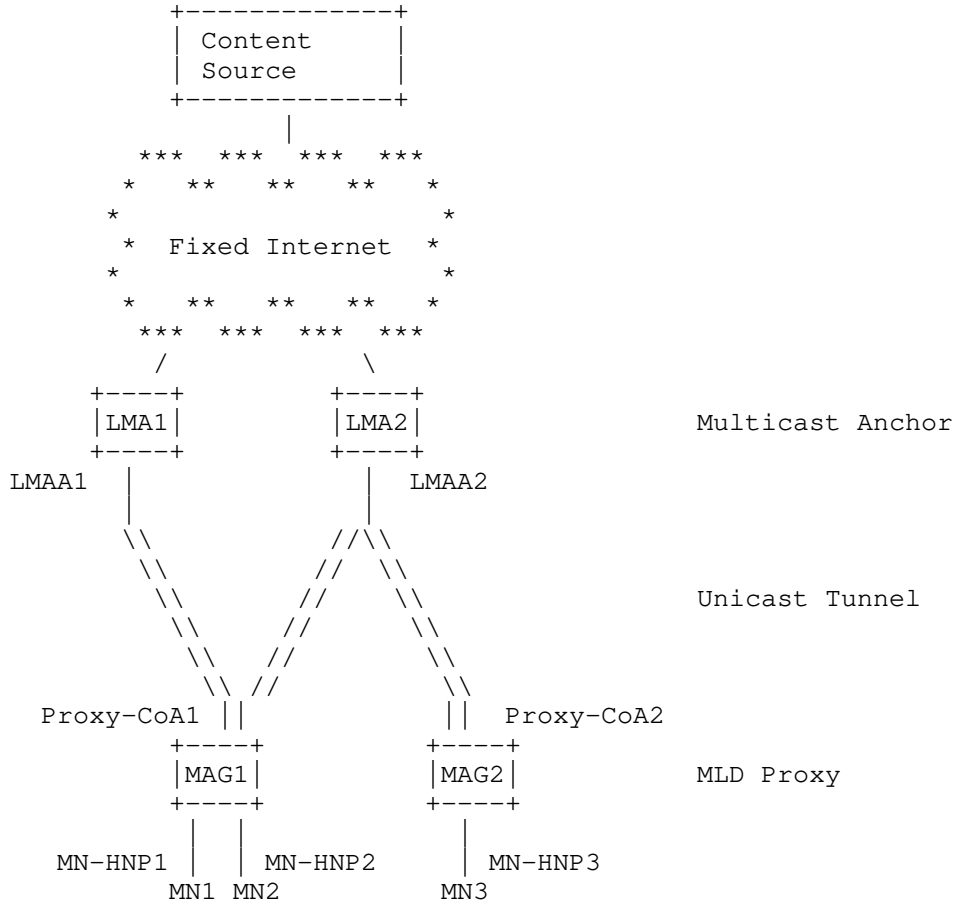

Figure 1: Reference Network for Multicast Deployment in PMIPv6

An MN in a PMIPv6 domain will decide on multicast group membership management completely independent of its current mobility conditions. It will submit MLD Report and Done messages, based on application triggers, using its link-local source address and multicast destination addresses according to [RFC3810], or [RFC2710]. These link-local signaling messages will arrive at the currently active MAG via one of its downstream local (wireless) links. A multicast unaware MAG would simply discard these MLD messages.

To facilitate multicast in a PMIPv6 domain, an MLD proxy function

[RFC4605] needs to be deployed on the MAG that selects the tunnel interface corresponding to the MN's LMA for its upstream interface (cf., section 6 of [RFC5213]). Thereby, each MAG-to-LMA tunnel interface defines an MLD proxy domain at the MAG, and it contains all downstream links to MNs that share this specific LMA. According to standard proxy operations, MLD Report messages will be aggregated and then forwarded up the tunnel interface to its corresponding LMA.

Serving as the designated multicast router or an additional MLD proxy, the LMA will transpose any MLD message from a MAG into the multicast routing infrastructure. Correspondingly, the LMA will create appropriate multicast forwarding states at its tunnel interface. Traffic of the subscribed groups will arrive at the LMA, and the LMA will forward this traffic according to its group/source states. In addition, the LMA will act as an MLD querier, seeing its downstream tunnel interfaces as multicast enabled links.

At the MAG, MLD queries and multicast data will arrive on the (tunnel) interface that is assigned to a group of access links as identified by its Binding Update List (cf., section 6.1 of [RFC5213]). As specified for MLD proxies, the MAG will forward multicast traffic and initiate related signaling down the appropriate access links to the MNs. Hence all multicast-related signaling and the data traffic will transparently flow from the LMA to the MN on an LMA-specific tree, which is shared among the multicast sources.

In case of a handover, the MN (unaware of IP mobility) will not send unsolicited MLD reports. Instead, the MAG is required to maintain group memberships in the following way. On observing a new MN on a downstream access link, the MAG sends a General MLD Query. Based on its outcome and the multicast group states previously maintained at the MAG, a corresponding Report will be sent to the LMA aggregating group membership states according to the proxy function. Additional Reports can be omitted when the previously established multicast forwarding states at the new MAG already cover the subscriptions of the MN.

In summary, the following steps are executed on handover:

1. The MAG-MN link comes up and the MAG discovers the new MN.

2. Unicast address configuration and PMIPv6 binding are performed after the MAG determines the corresponding LMA.

3. Following IPv6 address configuration, the MAG SHOULD send an (early) MLD General Query to the new downstream link as part of its standard multicast-enabled router operations.

4. The MAG SHOULD determine whether the MN is admissible to multicast services, and stop here otherwise.

5. The MAG adds the new downstream link to the MLD proxy instance with up-link to the corresponding LMA.

6. The corresponding Proxy instance triggers an MLD General Query on the new downstream link.

7. The MN Membership Reports arrive at the MAG, either in response to the early Query or to that of the Proxy instance.

8. The Proxy processes the MLD Report, updates states and reports upstream if necessary.

After Re-Binding, the LMA is not required to issue a General MLD Query on the tunnel link to refresh forwarding states. Multicast state updates SHOULD be triggered by the MAG, which aggregates subscriptions of all its MNs (see the call flow in Figure 2).

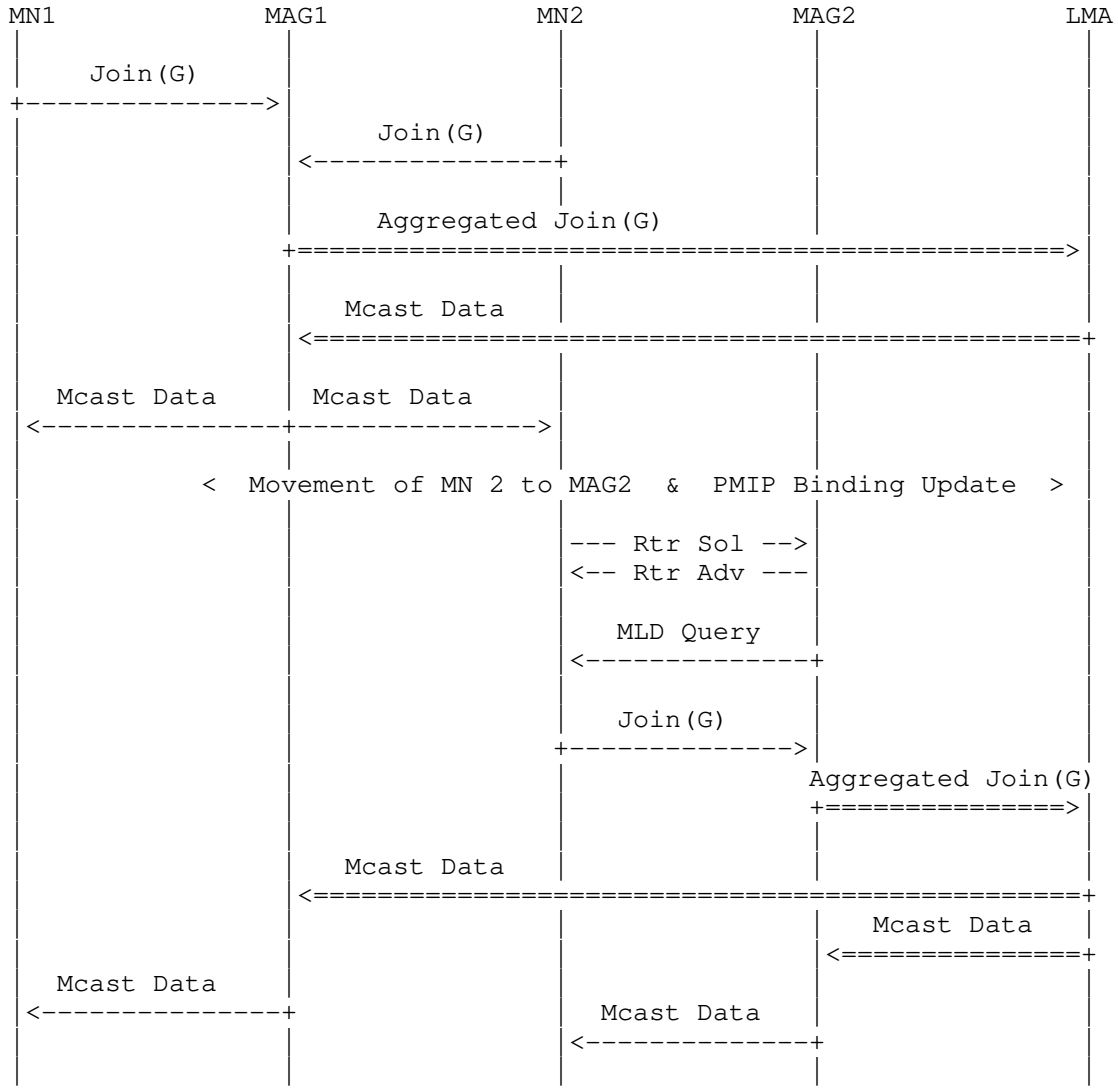

Figure 2: Call Flow of Multicast-enabled PMIP with "MLD Membership Report" abbreviated by "Join"

These multicast deployment considerations likewise apply for mobile nodes that operate with their IPv4 stack enabled in a PMIPv6 domain. PMIPv6 can provide IPv4 home address mobility support [RFC5844]. Such mobile nodes will use IGMP [RFC2236],[RFC3376] signaling for multicast, which is handled by an IGMP proxy function at the MAG in an analogous way.

Following these deployment steps, multicast management transparently inter-operates with PMIPv6. It is worth noting that MNs - while being attached to the same MAG, but associated with different LMAs - can subscribe to the same multicast group. Thereby data could be distributed redundantly in the network and duplicate traffic could arrive at a MAG. Additionally in a point-to-point wireless link model, a MAG might be forced to transmit the same data over one wireless domain to different MNs. However, multicast traffic arriving at one interface of the MN will always remain unique, i.e., the mobile multicast distribution system will never cause duplicate packets arriving at an MN (see Appendix C for further considerations).

4. Deployment Details

Multicast activation in a PMIPv6 domain requires to deploy general multicast functions at PMIPv6 routers and to define their interaction with the PMIPv6 protocol in the following way:

4.1. Operations of the Mobile Node

A Mobile Node willing to manage multicast traffic will join, maintain and leave groups as if located in the fixed Internet. No specific mobility actions nor implementations are required at the MN.

4.2. Operations of the Mobile Access Gateway

A Mobile Access Gateway is required to assist in MLD signaling and data forwarding between the MNs which it serves, and the corresponding LMAs associated to each MN. It therefore needs to implement an instance of the MLD proxy function [RFC4605] for each upstream tunnel interface that has been established with an LMA. The MAG decides on the mapping of downstream links to a proxy instance (and hence an upstream link to an LMA) based on the regular Binding Update List as maintained by PMIPv6 standard operations (cf., section 6.1 of [RFC5213]). As links connecting MNs and MAGs change under mobility, MLD proxies at MAGs MUST be able to dynamically add and remove downstream interfaces in its configuration.

On the reception of MLD reports from an MN, the MAG MUST identify the corresponding proxy instance from the incoming interface and perform regular MLD proxy operations: it will insert/update/remove multicast forwarding state on the incoming interface, and will merge state updates into the MLD proxy membership database. It will then send an aggregated Report via the upstream tunnel to the LMA when the membership database (cf., section 4.1 of [RFC4605]) changes. Conversely, on the reception of MLD Queries, the MAG proxy instance

will answer the Queries on behalf of all active downstream receivers maintained in its membership database. Queries sent by the LMA do not force the MAG to trigger corresponding messages immediately towards MNs. Multicast traffic arriving at the MAG on an upstream interface will be forwarded according to the group/source-specific forwarding states as acquired for each downstream interface within the MLD proxy instance. At this stage, it is important to note that IGMP/MLD proxy implementations capable of multiple instances are expected to closely follow the specifications of section 4.2 in [RFC4605], i.e., treat proxy instances in isolation of each other while forwarding. In providing isolated proxy instances, the MAG will uniquely serve its downstream links with exactly the data that belong to whatever group is subscribed on the particular interface.

After a handover, the MAG will continue to manage upstream tunnels and downstream interfaces as specified in the PMIPv6 specification. It MUST dynamically associate new access links to proxy instances that include the upstream connection to the corresponding LMA. The MAG detects the arrival of a new MN by receiving a router solicitation message and by an upcoming link. To learn about multicast groups subscribed by a newly attaching MN, the MAG SHOULD send a General Query to the MN's link. Querying an upcoming interface is a standard operation of MLD queriers (see Appendix A) and is performed immediately after address configuration. In addition, an MLD query SHOULD be initiated by the proxy instance, as soon as a new interface has been configured for downstream. In case, the access link between MN and MAG goes down, interface-specific multicast states change. Both cases may alter the composition of the membership database and this will trigger corresponding Reports towards the LMA. Note that the actual observable state depends on the access link model in use.

An MN may be unable to answer MAG multicast membership queries due to handover procedures, or its report may arrive before the MAG has configured its link as proxy downstream interface. Such occurrences are equivalent to a General Query loss. To prevent erroneous query timeouts at the MAG, MLD parameters SHOULD be carefully adjusted to the mobility regime. In particular, MLD timers and the Robustness Variable (see section 9 of [RFC3810]) MUST be chosen to be compliant with the time scale of handover operations and proxy configurations in the PMIPv6 domain.

In proceeding this way, the MAG is able to aggregate multicast subscriptions for each of its MLD proxy instances. However, this deployment approach does not prevent multiple identical streams arriving from different LMA upstream interfaces. Furthermore, a multipoint channel forwarding into the wireless domain is prevented by the point-to-point link model in use.

4.3. Operations of the Local Mobility Anchor

For any MN, the Local Mobility Anchor acts as the persistent Home Agent and at the same time as the default multicast querier for the corresponding MAG. It implements the function of the designated multicast router or a further MLD proxy. According to MLD reports received from a MAG (on behalf of the MNs), it establishes/maintains/removes group/source-specific multicast forwarding states at its corresponding downstream tunnel interfaces. At the same time, it procures for aggregated multicast membership maintenance at its upstream interface. Based on the multicast-transparent operations of the MAGs, the LMA treats its tunnel interfaces as multicast enabled downstream links, serving zero to many listening nodes. Multicast traffic arriving at the LMA is transparently forwarded according to its multicast forwarding information base.

After a handover, the LMA will receive Binding De-Registrations and Binding Lifetime Extensions that will cause a re-mapping of home network prefix(es) to a new Proxy-CoA in its Binding Cache (see section 5.3 of [RFC5213]). The multicast forwarding states require updating, as well, if the MN within an MLD proxy domain is the only receiver of a multicast group. Two different cases need to be considered:

1. The mobile node is the only receiver of a group behind the interface at which a De-Registration was received: The membership database of the MAG changes, which will trigger a Report/Done sent via the MAG-to-LMA interface to remove this group. The LMA thus terminates multicast forwarding.

2. The mobile node is the only receiver of a group behind the interface at which a Lifetime Extension was received: The membership database of the MAG changes, which will trigger a Report sent via the MAG-to-LMA interface to add this group. The LMA thus starts multicast distribution.

In proceeding this way, each LMA will provide transparent multicast support for the group of MNs it serves. It will perform traffic aggregation at the MN-group level and will assure that multicast data streams are uniquely forwarded per individual LMA-to-MAG tunnel.

4.4. IPv4 Support

An MN in a PMIPv6 domain may use an IPv4 address transparently for communication as specified in [RFC5844]. For this purpose, LMAs can register IPv4-Proxy-CoAs in its Binding Caches and MAGs can provide IPv4 support in access networks. Correspondingly, multicast membership management will be performed by the MN using IGMP. For

multicast support on the network side, an IGMP proxy function needs to be deployed at MAGs in exactly the same way as for IPv6. [RFC4605] defines IGMP proxy behaviour in full agreement with IPv6/ MLD. Thus IPv4 support can be transparently provided following the obvious deployment analogy.

For a dual-stack IPv4/IPv6 access network, the MAG proxy instances SHOULD choose multicast signaling according to address configurations on the link, but MAY submit IGMP and MLD queries in parallel, if needed. It should further be noted that the infrastructure cannot identify two data streams as identical when distributed via an IPv4 and IPv6 multicast group. Thus duplicate data may be forwarded on a heterogeneous network layer.

A particular note is worth giving the scenario of [RFC5845] in which overlapping private address spaces of different operators can be hosted in a PMIP domain by using GRE encapsulation with key identification. This scenario implies that unicast communication in the MAG-LMA tunnel can be individually identified per MN by the GRE keys. This scenario still does not impose any special treatment of multicast communication for the following reasons.

MLD/IGMP signaling between MNs and the MAG is on point-to-point links (identical to unicast). Aggregated MLD/IGMP signaling between the MAG proxy instance and the LMA remains link-local between the routers and independent of any individual MN. So the MAG-proxy and the LMA SHOULD not use GRE key identifiers, but plain GRE encapsulation to exchange MLD queries and reports. Similarly, multicast traffic sent from an LMA to MAGs proceeds as router-to-router forwarding according to the multicast forwarding information base (MFIB) of the LMA and independent of MN's unicast addresses, while the MAG proxy instance distributes multicast data down the point-to-point links (interfaces) according to its own MFIB, independent of MN's IP addresses.

It remains an open issue how communication proceeds in a multi-operator scenario, i.e., from which network the LMA pulls multicast traffic. This could be any mobility Operator itself, or a third party. However, this backbone routing in general is out of scope of the document, and most likely a matter of contracts.

4.5. Multihoming Support

An MN can connect to a PMIPv6 domain through multiple interfaces and experience transparent unicast handovers at all interfaces (cf., section 5.4 of [RFC5213]). In such simultaneous access scenario, it can autonomously assign multicast channel subscriptions to individual interfaces (see [RFC5757] for additional details). While doing so, multicast mobility operations described in this document will

transparently preserve the association of channels to interfaces in the following way.

Multicast listener states are kept per interface in the MLD state table. An MN will answer to an MLD General Query received on a specific (re-attaching) interface's state table. Thereafter, multicast forwarding is resumed for channels identical to those under subscription prior to handover. Consequently, an MN in a PMIPv6 domain MAY use multiple interfaces to facilitate load balancing or redundancy, but cannot follow a 'make-before-break' approach to service continuation on handovers.

4.6. Multicast Availability throughout the Access Network

There may be deployment scenarios, where multicast services are available throughout the access network independent of the PMIPv6 infrastructure. Direct multicast access at MAGs may be supported through native multicast routing within a flat access network that includes a multicast router, via dedicated (tunnel or VPN) links between MAGs and designated multicast routers, or by deploying AMT [I-D.ietf-mboned-auto-multicast].

Multicast deployment can be simplified in these scenarios. A single proxy instance at MAGs with up-link into the multicast cloud, for instance, could serve group communication purposes. MAGs could operate as general multicast routers or AMT gateways, as well.

Common to these solutions is that mobility management is covered by the dynamics of multicast routing, as initially foreseen in the Remote Subscription approach sketched in [RFC3775]. Care must be taken to avoid avalanche problems or service disruptions due to tardy multicast routing operations, and to adapt to different link-layer technologies [RFC5757]. The different possible approaches should be carefully investigated beyond the initial sketch in Appendix C. Such work is beyond the scope of this document.

4.7. A Note on Explicit Tracking

An IGMPv3/MLDv2 Querier may operate in combination with explicit tracking as described in Appendix 2 of [RFC3376], or Appendix 2 of[RFC3810]. This mechanism allows routers to monitor each multicast receiver individually. Even though this procedure is not standardized yet, it is widely implemented by vendors as it supports faster leave latencies and reduced signaling.

Enabling explicit tracking on downstream interfaces of the LMA and MAG would track a single MAG and MN respectively per interface. It may be used to preserve bandwidth on the MAG-MN link.

5. Message Source and Destination Address

This section describes source and destination addresses of MLD messages and encapsulating outer headers when deployed in the PMIPv6 domain. This overview is for clarification purposes, only, and does not define a behavior different from referenced standards in any way.

The interface identifier A-B denotes an interface on node A, which is connected to node B. This includes tunnel interfaces. Destination addresses for MLD/IGMP messages SHALL be as specified in Section 8. of [RFC2710] for MLDv1, and Section 5.1.15. and Section 5.2.14. of [RFC3810] for MLDv2.

5.1. Query

Interface	Source Address	Destination Address	Header
LMA-MAG	LMAA	Proxy-CoA	outer
	LMA-link-local	[RFC2710], [RFC3810]	inner
MAG-MN	MAG-link-local	[RFC2710], [RFC3810]	--

5.2. Report/Done

Interface	Source Address	Destination Address	Header
MN-MAG	MN-link-local	[RFC2710], [RFC3810]	--
MAG-LMA	Proxy-CoA	LMAA	outer
	MAG-link-local	[RFC2710], [RFC3810]	inner

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

This draft does not introduce additional messages or novel protocol operations. Consequently, no new threats are introduced by this

document in addition to those identified as security concerns of [RFC3810], [RFC4605], [RFC5213], and [RFC5844].

However, particular attention should be paid to implications of combining multicast and mobility management at network entities. As this specification allows mobile nodes to initiate the creation of multicast forwarding states at MAGs and LMAs while changing attachments, threats of resource exhaustion at PMIP routers and access networks arrive from rapid state changes, as well as from high volume data streams routed into access networks of limited capacities. In addition to proper authorization checks of MNs, rate controls at replicators MAY be required to protect the agents and the downstream networks. In particular, MLD proxy implementations at MAGs SHOULD carefully procure for automatic multicast state extinction on the departure of MNs, as mobile multicast listeners in the PMIPv6 domain will not actively terminate group membership prior to departure.

8. Acknowledgements

This memo follows initial requirements work presented in draft-deng-multimob-pmip6-requirement, and is the outcome of extensive previous discussions and a follow-up of several initial drafts on the subject. The authors would like to thank (in alphabetical order) Jari Arkko, Luis M. Contreras, Greg Daley, Gorry Fairhurst, Dirk von Hugo, Seil Jeon, Jouni Korhonen, Guang Lu, Sebastian Meiling, Liu Hui, Akbar Rahman, Imed Romdhani, Behcet Sarikaya, Pierrick Seite, Stig Venaas, and Juan Carlos Zuniga for advice, help and reviews of the document. Funding by the German Federal Ministry of Education and Research within the G-LAB Initiative is gratefully acknowledged.

9. References

9.1. Normative References

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

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

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

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[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.

9.2. Informative References

[I-D.ietf-mboned-auto-multicast] Thaler, D., Talwar, M., Aggarwal, A., Vicisano, L., and T. Pusateri, "Automatic IP Multicast Without Explicit Tunnels (AMT)", draft-ietf-mboned-auto-multicast-10 (work in progress), March 2010.

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

[RFC5757] Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey", RFC 5757, February 2010.

[RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", RFC 5845, June 2010.

Appendix A. Initial MLD Queries on Upcoming Links

According to [RFC3810] and [RFC2710] when an IGMP/MLD-enabled multicast router starts operating on a subnet, by default it considers itself as Querier and sends several General Queries. Such initial query should be sent by the router immediately, but could be delayed by a (tunable) Startup Query Interval (see Sections 7.6.2. and 9.6. of [RFC3810]).

Experimental tests on Linux and Cisco systems have revealed immediate IGMP Queries following a link trigger event (within a fraction of 1

ms), while MLD Queries immediately followed the autoconfiguration of IPv6 link-local addresses at the corresponding interface.

Appendix B. State of IGMP/MLD Proxy Implementations

The deployment scenario defined in this document requires certain proxy functionalities at the MAGs that implementations of [RFC4605] need to contribute. In particular, a simultaneous support of IGMP and MLD is needed, as well as a configurable list of downstream interfaces that may be altered during runtime, and the deployment of multiple proxy instances at a single router that can operate independently on separated interfaces.

A brief experimental trial undertaken in February 2010 revealed the following divergent status of selected IGMP/MLD proxy implementations.

Cisco Edge Router Software-based commodity edge routers (test device from the 26xx-Series) implement IGMPv2/v3 proxy functions only in combination with PIM-SM. There is no support of MLD Proxy. Interfaces are dynamically configurable at runtime via the CLI, but multiple proxy instances are not supported.

Linux igmpproxy IGMPv2 Proxy implementation that permits a static configuration of downstream interfaces (simple bug fix required). Multiple instances are prevented by a lock (corresponding code re-used from a previous DVMRP implementation). IPv6/MLD is unsupported. Project page: http://sourceforge.net/projects/igmpproxy/.

Linux gproxy IGMPv3 Proxy implementation that permits configuration of the upstream interface, only. Downstream interfaces are collected at startup without dynamic extension of this list. No support of multiple instances or MLD. Project page: http://potiron.loria.fr/projects/madynes/internals/perso/lahmadi/igmpv3proxy/.

Linux ecmh MLDv1/2 Proxy implementation without IGMP support that inspects IPv4 tunnels and detects encapsulated MLD messages. Allows for dynamic addition of interfaces at runtime and multiple instances. However, downstream interfaces cannot be configured. Project page: http://sourceforge.net/projects/ecmh/

Appendix C. Comparative Evaluation of Different Approaches

In this section, we briefly evaluate two orthogonal PMIP concepts for multicast traffic organization at LMAs: In scenario A, multicast is provided by combined unicast/multicast LMAs as described in this document. Scenario B directs traffic via a dedicated, central multicast router ("LMA-M") that tunnels packets to MAGs independent of unicast hand-offs.

Both approaches do not establish native multicast distribution between the LMA and MAG, but use tunneling mechanisms. In scenario A, a MAG is connected to different multicast-enabled LMAs, and can receive the same multicast stream via multiple paths depending on the group subscriptions of MNs and their associated LMAs. This problem, a.k.a. tunnel convergence problem, may lead to redundant traffic at the MAGs. Scenario B in contrast configures MAGs to establish a tunnel to a single, dedicated multicast LMA for all attached MNs and relocates overhead costs to the multicast anchor. This eliminates redundant traffic, but may result in an avalanche problem at the LMA.

We quantify the costs of both approaches based on two metrics: The amount of redundant traffic at MAGs and the number of simultaneous streams at LMAs. Realistic values depend on the topology and the group subscription model. To explore scalability in a large PMIP domain of 1,000,000 MNs, we consider the following two extremal multicast settings.

1. All MNs participate in distinct multicast groups.

2. All MNs join the same multicast groups.

A typical PMIP deployment approximately allows for 5,000 MNs attached to one MAG, while 50 MAGs can be served by one LMA. Hence 1,000,000 MNs require approx. 200 MAGs backed by 4 LMAs for unicast transmission. In scenario A, these LMAs also forward multicast streams, while in scenario B one additional dedicated LMA (LMA-M) serves multicast. In the following, we calculate the metrics described above. In addition, we display the number of packet streams that cross the interconnecting (wired) network within a PMIPv6 domain.

Setting 1:

PMIP multicast scheme	# of redund. streams at MAG	# of simul. streams at LMA/LMA-M	# of total streams in the network
Combined Unicast/ Multicast LMA	0	250,000	1,000,000
Dedicated Multicast LMA	0	1,000,000	1,000,000

1,000,000 MNs are subscribed to distinct multicast groups

Setting 2:

PMIP multicast scheme	# of redund. streams at MAG	# of simul. streams at LMA/LMA-M	# of total streams in the network
Combined Unicast/ Multicast LMA	3	200	800
Dedicated Multicast LMA	0	200	200

1,000,000 MNs are subscribed to the same multicast group

These considerations of extremal settings show that packet duplication and replication effects apply in changing intensities for different use cases of multicast data services. However, tunnel convergence, i.e., duplicate data arriving at a MAG, does cause much smaller problems in scalability than the stream replication at LMAs (avalanche problem). For scenario A, it should be also noted that the high stream replication requirements at LMAs in setting 1 can be attenuated by deploying additional LMAs in a PMIP domain, while scenario B does not allow for distributing the LMA-M, as no handover management is available at LMA-M.

Appendix D. Change Log

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-05.

1. Clarification and section-based reference to destination addresses in MLD in response to WG feedback.

2. Removed reference to individual draft-zuniga-multimob-smspmip in Appendix C and added explanations in response to WG feedback.

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-04.

1. Clarifications and editorial improvements in response to WG feedback.

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-03.

1. Clarifications and editorial improvements in response to WG feedback.

2. Added pointers and explanations to Explicit Tracking and GRE tunneling in the IPv4 scenario (RFC 5845).

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-02.

1. Clarifications and editorial improvements in response to WG feedback.

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-01.

1. Editorial improvements in response to WG feedback.

The following changes have been made from version draft-ietf-multimob-pmipv6-base-solution-00.

1. Added section on multihoming.

2. Updated security section.

3. Several editorial improvements and minor extensions.

The following changes have been made from the previous individual version draft-schmidt-multimob-pmipv6-mcast-deployment-04.

1. Updated references.

2. Corrected typos.

3. Adjusted title & document name.

The following changes have been made from draft-schmidt-multimob-pmipv6-mcast-deployment-03.

1. Detailed outline of multicast reconfiguration steps on handovers added in protocol overview (section 3).

2. Clarified the details of proxy operations at the MAG along with the expected features of IGMP/MLD Proxy implementations (section 4.2).

3. Clarified querying in dual-stack scenarios (section 4.4).

4. Subsection added on the special case, where multicast is available throughout the access network (section 4.5).

5. Appendix on IGMP/MLD behaviour added with test reports on current Proxy implementations.

The following changes have been made from draft-schmidt-multimob-pmipv6-mcast-deployment-02.

1. Many editorial improvements, in particular as response to draft reviews.

2. Section on IPv4 support added.

3. Added clarifications on initial IGMP/MLD Queries and supplementary information in appendix.

4. Appendix added an comparative performance evaluation regarding mixed/dedicated deployment of multicast at LMAs.

Authors' Addresses

Thomas C. Schmidt
HAW Hamburg
Berliner Tor 7
Hamburg 20099
Germany

Email: schmidt@informatik.haw-hamburg.de
URI: http://inet.cpt.haw-hamburg.de/members/schmidt

Matthias Waehlisch
link-lab & FU Berlin
Hoenower Str. 35
Berlin 10318
Germany

Email: mw@link-lab.net

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Email: suresh.krishnan@ericsson.com

MULTIMOB Group
Internet-Draft
Intended status: Standards Track
Expires: May 12, 2011

T C. Schmidt
HAW Hamburg
M. Waehlisch
link-lab & FU Berlin
R. Koodli
Cisco Systems
G. Fairhurst
University of Aberdeen
November 08, 2010

Multicast Listener Extensions for MIPv6 and PMIPv6 Fast Handovers
draft-schmidt-multimob-fmipv6-pfmipv6-multicast-03

Abstract

Fast handover protocols for MIPv6 and PMIPv6 define mobility management procedures that support unicast communication at reduced handover latencies. Fast handover base operations do not affect multicast communication, and hence do not accelerate handover management for native multicast listeners. Many multicast applications like IPTV or conferencing, though, are comprised of delay-sensitive real-time traffic and will benefit from fast handover execution. This document specifies extension of the Mobile IPv6 Fast Handovers (FMIPv6) and the Fast Handovers for Proxy Mobile IPv6 (PFMIPv6) protocols to include multicast traffic management in fast handover operations.

Table of Contents

1. Introduction

Mobile IPv6 [RFC3775] defines a network layer mobility protocol involving mobile nodes participation, while Proxy Mobile IPv6 [RFC5213] provides a mechanism without requiring mobility protocol operations at a Mobile Node (MN). Both protocols introduce traffic disruptions on handovers that may be intolerable in many application scenarios. Mobile IPv6 Fast Handovers (FMIPv6) [RFC5568], and Fast Handovers for Proxy Mobile IPv6 (PFMIPv6) [RFC5949] improve these handover delays for unicast communication to the order of the maximum delay needed for link switching and signaling between Access Routers (ARs) or Mobile Access Gateways (MAGs) [FMIPv6-Analysis].

No dedicated treatment of seamless multicast data reception has been proposed by any of the above protocols. MIPv6 only roughly defines multicast for Mobile Nodes using a remote subscription approach or a home subscription through bi-directional tunneling via the Home Agent (HA). Multicast forwarding services have not been specified at all in [RFC5213], but are subject to current specification [I-D.ietf-multimob-pmipv6-base-solution]. It is assumed throughout this document that mechanisms and protocol operations are in place to transport multicast traffic to ARs. These operations are referred to as 'JOIN/LEAVE' of an AR, while the explicit techniques to manage multicast transmission are beyond the scope of this document.

Mobile multicast protocols need to serve applications such as IPTV with high-volume content streams to be distributed to potentially large numbers of receivers, and therefore should preserve the multicast nature of packet distribution and approximate optimal routing [RFC5757]. It is undesirable to rely on home tunneling for optimizing multicast. Unencapsulated, native multicast transmission requires establishing forwarding state, which will not be transferred between access routers by the unicast fast handover protocols. Thus multicast traffic will not experience expedited handover performance, but an MN - or its corresponding MAG in PMIPv6 - can perform remote subscriptions in each visited network.

This document specifies extensions of FMIPv6 and PFMIPv6 for including multicast traffic management in fast handover operations. The solution common to both underlying protocols defines the per-group transfer of multicast contexts between ARs or MAGs. The protocol defines corresponding message extensions necessary for carrying group context information independent of the particular handover protocol. ARs or MAGs are then enabled to treat multicast traffic according to fast unicast handovers and with similar performance. No protocol changes are introduced that prevent a multicast unaware node from performing fast handovers with multicast aware ARs or MAGs.

This specification is applicable when a mobile node has joined and maintains one or several multicast group subscriptions prior to undergoing a fast handover. It does not introduce any requirements on the multicast routing protocols in use, nor are the ARs or MAGs assumed to be multicast routers. It assumes network conditions, though, that allow native multicast reception in both, the previous and new access network. Methods to bridge regions without native multicast connectivity are beyond the scope of this document.

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 RFC 2119 [RFC2119]. The use of the term, "silently ignore" is not defined in RFC 2119. However, the term is used in this document and can be similarly construed.

This document uses the terminology of [RFC5568], [RFC5949], [RFC3775], and [RFC5213]. In addition, the following terms are introduced:

3. Protocol Overview

The reference scenario for multicast fast handover is illustrated in Figure 1.

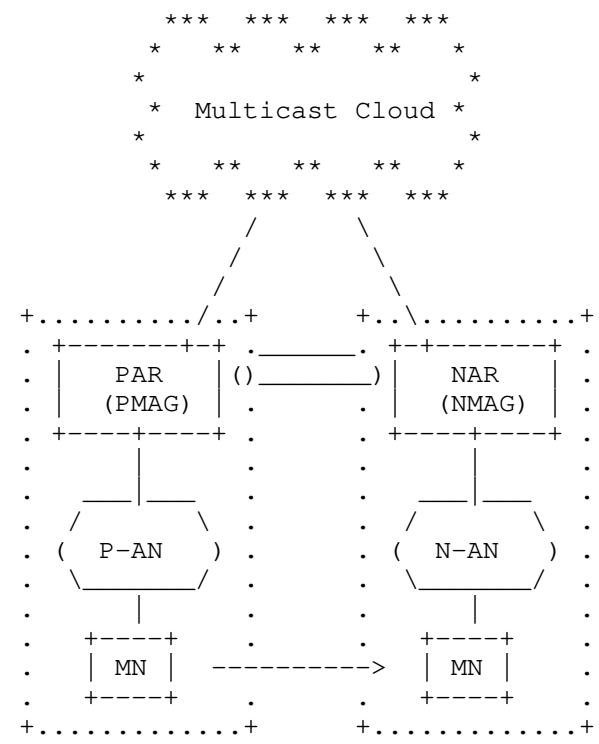

Figure 1: Reference Network for Fast Handover

3.1. Multicast Context Transfer between Access Routers

In a fast handover scenario (cf. Figure 1), ARs/MAGs establish a mutual binding and provide the capability to exchange context information concerning the MN. This context transfer will be triggered by detecting MN's forthcoming move to a new AR and assist the MN to immediately resume communication on the new subnet link using its previous IP address. In contrast to unicast, multicast stream reception does not primarily depend on address and binding cache management, but requires distribution trees to adapt so that traffic follows the movement of the MN. This process may be significantly slower than fast handover management [RFC5757]. Multicast listeners at handover may take the twofold advantage of including the multicast groups under subscription in context transfer. First, the NAR can proactively join the desired groups as soon as it gains knowledge of them. Second, multicast streams may be included in traffic forwarding via the tunnel established from PAR to NAR.

There are two modes of operation in FMIPv6 and in PFMIPv6. The

predictive mode allows for AR-binding and context transfer prior to an MN handover, while in the reactive mode, these steps are executed after detection that the MN has re-attached to NAR. Details of the signaling schemes differ between FMIPv6 and PFMIPv6 and are outlined in Section 3.2 and Section 3.3.

In a predictive fast handover, the access router (i.e., PAR (PMAG) in Figure 1) learns about the impending movement of the MN and simultaneously about the multicast group context as specified in Section 3.2 and Section 3.3. Thereafter, PAR will initiate an AR-binding and context transfer by transmitting a HI message to NAR (NMAG). HI is extended by multicast group states carried in mobility header options as defined in Section 5.3. On reception of the HI message, NAR returns a multicast acknowledgement in its HACK answer that indicates its ability to support each requested group (see Section 5.4). NAR (NMAG) expresses its willingness to receive multicast traffic from forwarding by PAR using standard MLD signaling. There are several reasons to waive forwarding, e.g., the group may already be under native subscription or capacity constraints may hinder decapsulation of additional streams at the NAR. For the groups requested, PAR will add the tunnel interface to its multicast forwarding database, so that multicast streams can be forwarded in parallel to unicast traffic. NAR, taking the role of an MLD proxy [RFC4605] with upstream router PAR, will submit an MLD report on this upstream tunnel interface to request the desired groups, but will terminate multicast forwarding [RFC3810] from PAR, as soon as group traffic natively arrives. In addition, NAR immediately joins all groups that are not already under subscription using its native multicast upstream interface and loopback as downstream. It starts to downstream multicast forwarding after the MN has arrived.

In a reactive fast handover, PAR will learn about the movement of the MN, after the latter has re-associated with the new access network. Also from the new link, it will be informed about the multicast context of the MN. As group membership information are present at the new access network prior to context transfer, MLD join signaling can proceed in parallel to HI/HACK exchange. Following the context transfer, multicast data can be forwarded to the new access network using the PAR-NAR tunnel of the fast handover protocol. Depending on the specific network topology though, multicast traffic for some groups may natively arrive before it is forwarded from PAR.

In both modes of operation, it is the responsibility of the PAR (PMAG) to properly react on the departure of the MN in the context of local group management. Depending on the multicast state management, link type and MLD parameters deployed (cf., [RFC5757]), it is requested to take appropriate actions to adjust multicast service to

requirements of the remaining nodes.

In this way, the MN will be able to participate in multicast group communication with a handover performance comparable to that for unicast, while network resources consumption is minimized.

3.2. Protocol Operations Specific to FMIPv6

ARs that provide multicast support in FMIPv6 will advertise this general service by setting an indicator bit (M-bit) in its PrRtAdv message as defined in Section 5.1. Additional details about the multicast service support, e.g., flavors and groups, will be exchanged within HI/HACK dialogs later at handovers.

An MN operating FMIPv6 will actively initiate the handover management by submitting a fast binding update (FBU). The MN, which is aware of the multicast groups it wishes to maintain, will attach mobility options containing its group states (see Section 5.3) to the FBU, and thereby inform ARs about its multicast context. ARs will use these multicast context options for inter-AR context transfer.

In predictive mode, FBU is issued on the previous link and received by PAR as displayed in Figure 2. PAR will extract the multicast context options and append them to its HI message. From the HACK message, PAR will redistribute the multicast acknowledgement by adding the corresponding mobility options to its FBACK message. From receiving FBACK, the MN will learn about a per group multicast support in the new access network. If some groups or a multicast flavour are not supported, it MAY decide on taking actions to compensate the missing service. Note that the proactive multicast context transfer may proceed successfully, even if the MN misses the FBACK message on the previous link.

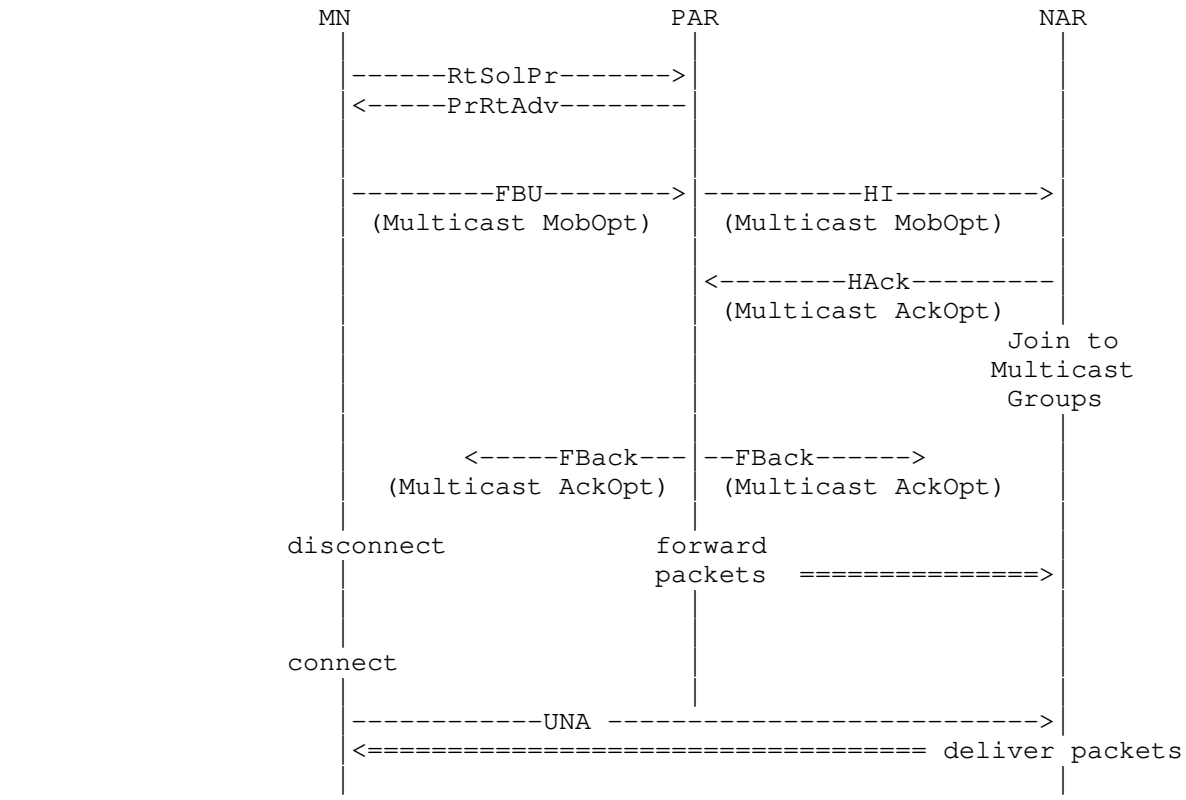

Figure 2: Predictive Multicast Handover for FMIPv6

The call flow for reactive mode is visualized in Figure 3. After attaching to the new access link and performing an unsolicited neighbor advertisement (UNA), the MN issues an FBU which NAR forwards to PAR without processing. At this time, the MN is able to re-join all desired multicast groups without relying on AR assistance. Nevertheless, multicast context options are exchanged in the HI/HACK dialog to facilitate intermediate forwarding of requested streams. Note that group traffic may already arrive from a MN's subscription at the time NAR receives the HI message. Such streams may be transparently excluded from forwarding by setting an appropriate multicast acknowledge option. In any case, NAR MUST ensure that not more than one stream of the same group is forwarded to the MN.

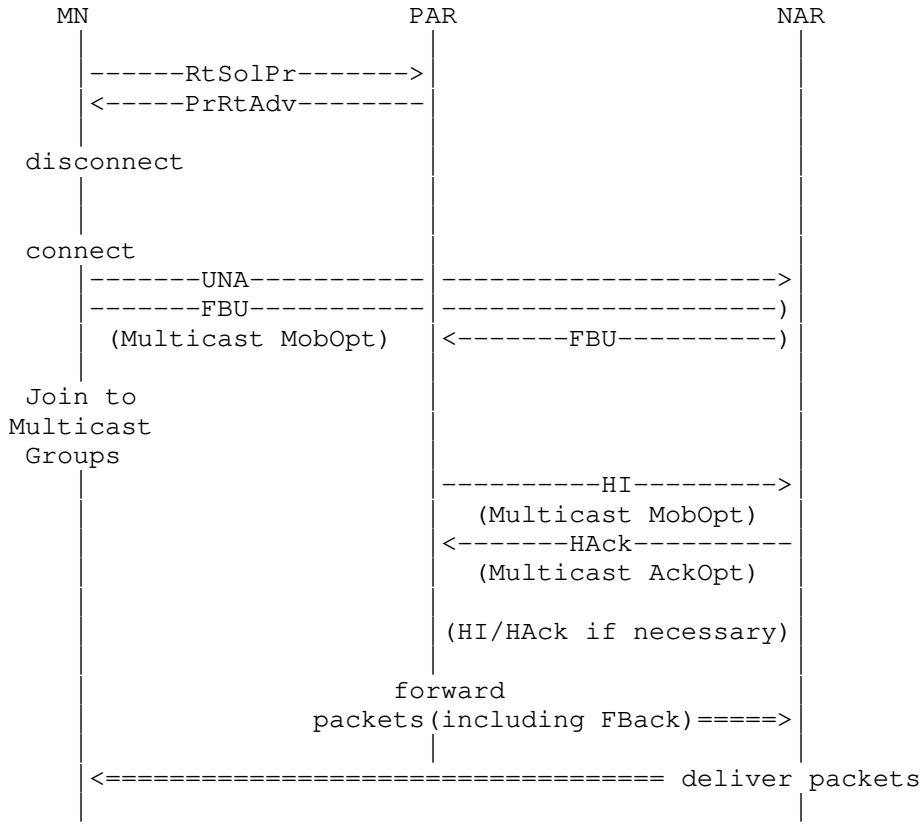

Figure 3: Reactive Multicast Handover for FMIPv6

3.3. Protocol Operations Specific to PFMIPv6

In a proxy mobile IPv6 environment, the MN remains agnostic of network layer changes, and fast handover procedures are operated by the access routers or MAGs. The handover initiation, or the re-association respectively are managed by the access networks. Consequently, access routers need to be aware of multicast membership state at the mobile node. There are two ways to obtain record of MN's multicast membership. First, MAGs MAY perform an explicit tracking (cf., [RFC4605], [I-D.ietf-multimob-pmipv6-base-solution]) or extract membership status from forwarding states at node-specific point-to-point links. Second, routers can perform general queries at handovers. Both methods are equally applicable. However, a router that does not operate explicit tracking MUST query its downstream links subsequent to handovers. In either case, the PAR will become knowledgeable about multicast group subscriptions of the MN.

In predictive mode, the PMAG (PAR) will learn about the upcoming movement of the mobile node. Without explicit tracking, it will immediately submit a general MLD query and learn about the multicast groups under subscription. As displayed in Figure 4, it will initiate binding and context transfer with the NMAG (NAR) by issuing a HI message that is augmented by multicast contexts in the mobility options defined in Section 5.3. NAR will extract multicast context information and act as described in Section 3.1.

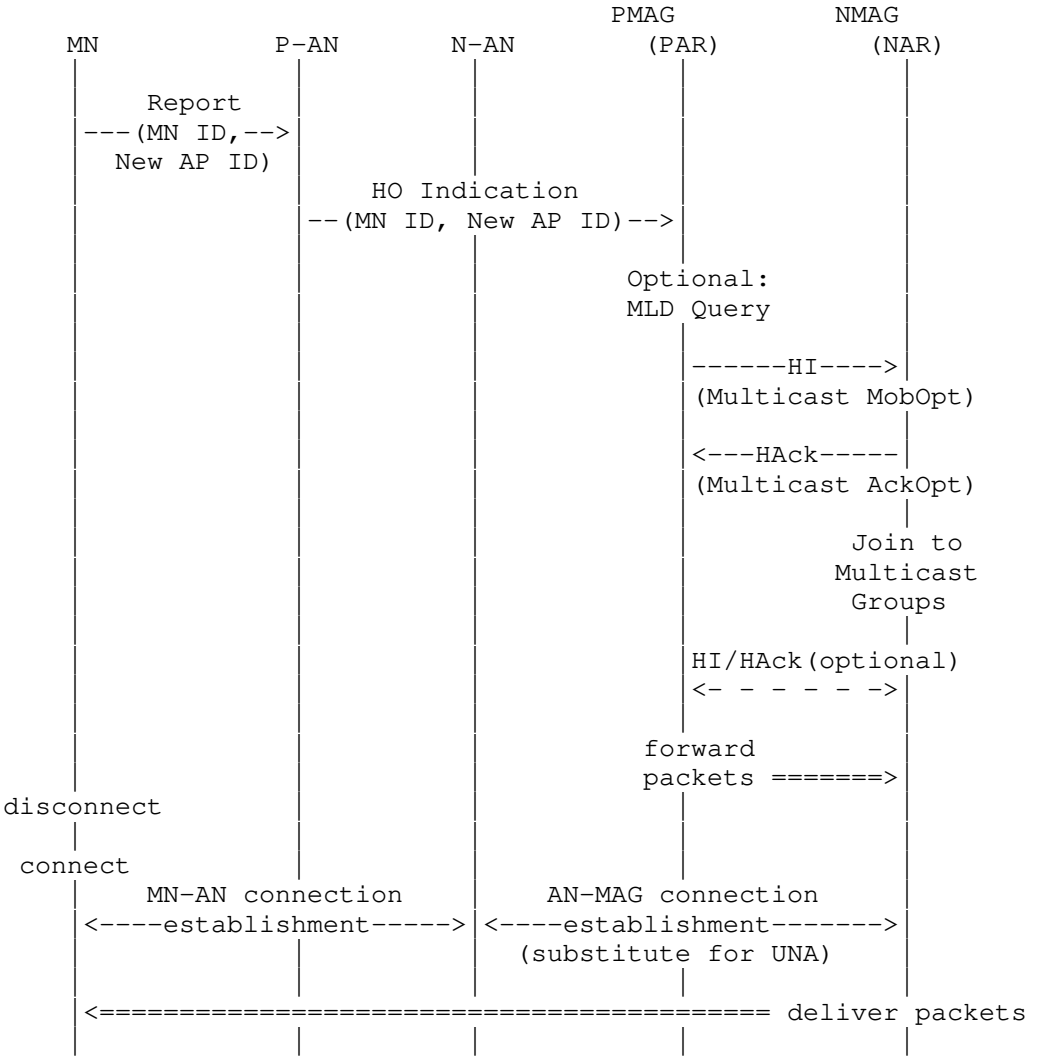

Figure 4: Predictive Multicast Handover for PFMIPv6

In reactive mode, the NMAG (NAR) will learn about MN's attachment to the N-AN and establish connectivity by means of PMIPv6 protocol operations. However, it will have no knowledge about multicast state at the MN. Triggered by a MN attachment, the NMAG will send a general MLD query and thereafter join the requested groups. In the case of a reactive handover, the binding is initiated by NMAG, and the HI/HACK message semantic is inverted (see [RFC5949]). For multicast context transfer, the NMAG attaches to its HI message those group identifiers it requests to be forwarded from PMAG. Using the identical syntax in its multicast mobility option headers as defined in Section 5.4, PMAG acknowledges the group forwarding request in its HACK answer. The corresponding call flow is displayed in Figure 5.

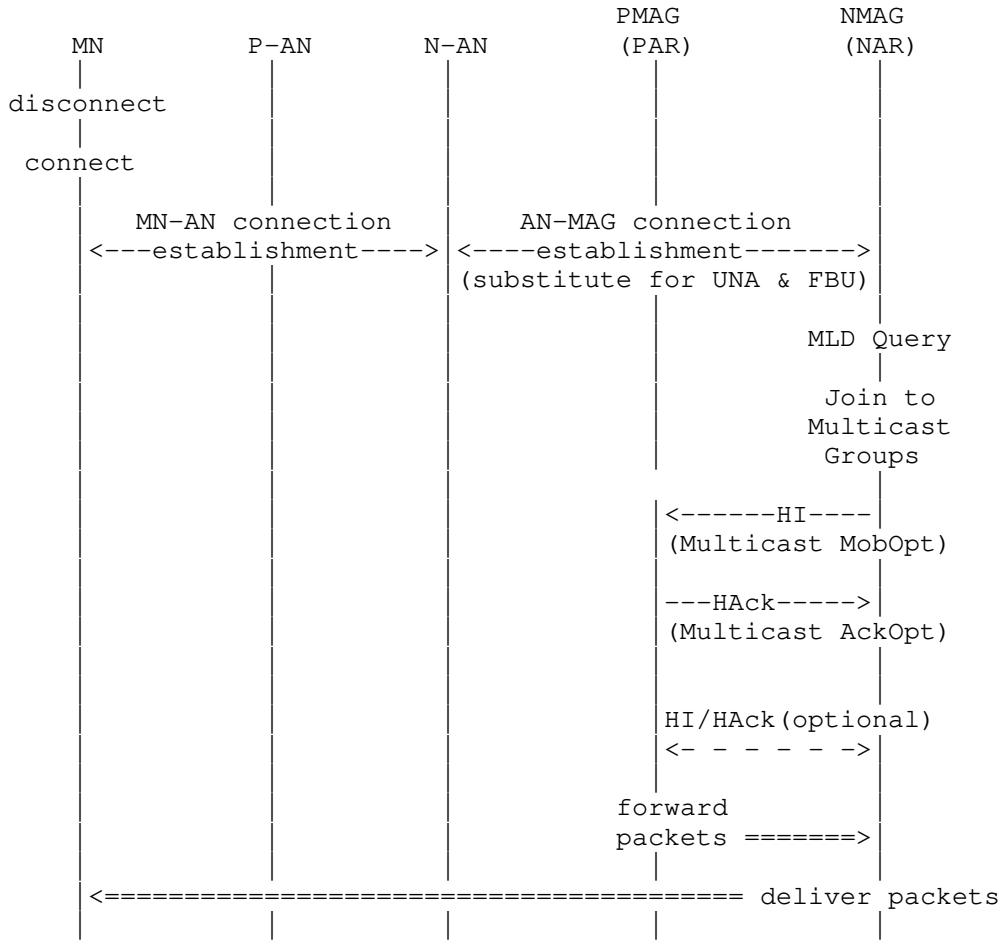

Figure 5: Reactive Multicast Handover for PFMIPv6

4. Protocol Details

4.1. Protocol Operations Specific to FMIPv6

4.1.1. Operations of the Mobile Node

A Mobile Node willing to manage multicast traffic within fast handover operations will inform about its MLD listener state records within handover signaling.

When sensing a handover in predictive mode, an MN will build a Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) multicast listener state and append it to the Fast Binding Update (FBU) prior to signaling with PAR. It will receive the Multicast Acknowledgement Option(s) within Fast Binding Acknowledge (FBack) (see Section 5.4) and learn about unsupported or prohibited groups at the NAR. The MN MAY take appropriate actions like home tunneling to bridge missing multicast services in the new access network. No multicast-specific operation is required by the MN when re-attaching in the new network besides standard FMIPv6 signaling.

In reactive mode, the MN appends an identical Multicast Mobility Option to FBU sent after its reconnect. In response, it will learn about the Multicast Acknowledgement Option(s) from FBACK and expect corresponding multicast data. Concurrently it joins all desired multicast groups (channels) directly on its newly established access link.

4.1.2. Operations of the Previous Access Router

A PAR will advertise its multicast support by setting the M-bit in PrRtAdv.

In predictive mode, a PAR will receive the multicast listener state of a MN prior to handover from the Multicast Mobility Option appended to the FBU. It will forward these records to NAR within HI messages and will expect Multicast Acknowledgement Option(s) in HACK, which itself is returned to the MN as an appendix to FBACK. In performing multicast context exchange, the AR is instructed to include the PAR-to-NAR tunnel obtained from unicast handover management in its multicast downstream interfaces and await MLD listener reports from NAR. In response to receiving multicast subscriptions, PAR will forward group data acting as a normal multicast router or proxy.

In reactive mode, PAR will receive the FBU augmented by the Multicast Mobility Option from the new network, but will continue with an identical multicast record exchange in the HI/HACk dialog. As in the predictive case, it will configure the PAR-to-NAR tunnel for multicast downstream and forward data according to MLD reports obtained from NAR.

In both modes, PAR will interpret the first of the two events, the departure of the MN or the reception of the Multicast Acknowledgement Option(s) as a multicast LEAVE message of the MN and react according to the signaling scheme deployed in the access network (i.e., MLD querying, explicit tracking).

4.1.3. Operations of the New Access Router

NAR will advertise its multicast support by setting the M-bit in PrRtAdv.

In predictive mode, a NAR will receive the multicast listener state of an expected MN from the Multicast Mobility Option appended to the HI message. It will extract the MLD/IGMP records from the message and intersect the request subscription with its multicast service offer. Further on it will adjoin the supported groups (channels) to the MLD listener state using loopback as downstream interface. This will lead to suitable regular subscriptions on its native multicast upstream interface without additional forwarding. Concurrently, NAR builds a Multicast Acknowledgement Option(s) (see Section 5.4) listing those groups (channels) unsupported on the new access link and returns them within HACK. As soon as the bidiretional tunnel from PAR to NAR is operational, NAR joins the groups desired for forwarding on the tunnel link.

In reactive mode, NAR will learn about the multicast listener state of a new MN from the Multicast Mobility Option appended to HI at a time, when the MN has already performed local subscriptions of the multicast service. Thus NAR solely determines the intersection of requested and supported groups (channels) and issues the join requests for group forwarding on the PAR-NAR tunnel interface.

In both modes, NAR MUST send a LEAVE message to the tunnel immediately after forwarding of a group (channel) becomes unneeded, e.g., after native multicast traffic arrives or group membership of the MN terminates.

4.2. Protocol Operations Specific to PFMIPv6

4.2.1. Operations of the Mobile Node

A Mobile Node willing to participate in multicast traffic will join, maintain and leave groups as if located in the fixed Internet. It will cooperate in handover indication as specified in [RFC5949] and required by its access link-layer technology. No multicast-specific mobility actions nor implementations are required at the MN in a PMIPv6 domain.

4.2.2. Operations of the Previous MAG

A MAG receiving a handover indication for one of its MNs follows the predictive fast handover mode as a PMAG. It MUST issue an MLD General Query immediately on its corresponding link unless it performs an explicit tracking on that link. After gaining knowledge of the multicast subscriptions of the MN, the PMAG builds a Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) multicast listener state. If not empty, this Mobility Option is appended to the regular fast handover HI messages, or - in the case of unicast HI message being submitted prior to multicast state detection - sent in an additional HI message to the NMAG. PMAG then waits for receiving the Multicast Acknowledgement Option(s) with HACK (see Section 5.4) and the creation of the bidirectional tunnel with NMAG. Thereafter PMAG will add the tunnel to its downstream interfaces in the multicast forwarding database. For those groups (channels) reported in the Multicast Acknowledgement Option(s), i.e., not supported in the new access network, PMAG takes appropriate actions (e.g., forwarding, termination) in concordance with the network policy. It SHOULD start forwarding traffic down the tunnel interface for those groups it receives an MLD listener report message from NMAG. After the departure of the MN and on the reception of LEAVE messages for groups/channels, PMAG MUST terminate forwarding of the specific groups and update its multicast forwarding database. Correspondingly it issues a group/channel LEAVE to its upstream link, if no more listeners are present on its downstream links.

A MAG receiving a HI message with Multicast Mobility Option for a currently attached node follows the reactive fast handover mode as a PMAG. It will return Multicast Acknowledgement Option(s) (see Section 5.4) within HACK listing those groups/channels unsupported at NMAG. It will add the bidirectional tunnel with NMAG to its downstream interfaces and will start forwarding multicast traffic for those groups it receives an MLD listener report message from NMAG. At the reception of LEAVE messages for groups (channels), PMAG MUST terminate forwarding of the specific groups and update its multicast forwarding database. According to its multicast forwarding states it MAY need to issue a group/channel LEAVE to its upstream link, if no more listeners are present on its downstream links.

In both modes, PMAG will interpret the departure of the MN as a multicast LEAVE message of the MN and react according to the signaling scheme deployed in the access network (i.e., MLD querying, explicit tracking).

4.2.3. Operations of the New MAG

A MAG receiving a HI message with Multicast Mobility Option for a currently unattached node follows the predictive fast handover mode as NMAG. It will decide on those multicast groups/channels it wants forwarded from the PMAG and builds a Multicast Acknowledgement Option (see Section 5.4) that enumerates only unwanted groups/channels. This Mobility Option is appended to the regular fast handover HACK messages, or – in the case of unicast HACK message being submitted prior to multicast state acknowledgement – sent in an additional HACK message to the PMAG. Immediately thereafter, NMAG SHOULD update its MLD listener state by the new groups/channels obtained from the Multicast Mobility Option. Until the MN re-attaches, NMAG uses its loopback interface for downstream and does not forward traffic to the potential link of the MN. NMAG SHOULD issue JOIN messages for those newly adopted groups to its regular multicast upstream interface. As soon as the bidirectional tunnel with PMAG is established, NMAG additionally joins those groups/channels on the tunnel interface that it wants to receive by forwarding from PMAG. NMAG MUST send a LEAVE message to the tunnel immediately after forwarding of a group/channel becomes unneeded, e.g., after native multicast traffic arrives or group membership of the MN terminates.

A MAG experiencing a connection request for a MN without prior reception of a corresponding Multicast Mobility Option is operating in the reactive fast handover mode as NMAG. Following the re-attachment, it immediately issues an MLD General Query to learn about multicast subscriptions of the newly arrived MN. Using standard multicast operations, NMAG joins the missing groups (channels) on its regular multicast upstream interface. Concurrently, it selects groups (channels) for forwarding from PMAG and builds a Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) multicast listener state. If not empty, this Mobility Option is appended to the regular fast handover HI messages with the F flag set, or – in the case of unicast HI message being submitted prior to multicast state detection – sent in an additional HI message to the PMAG. Upon reception of the Multicast Acknowledgement Option and upcoming of the bidirectional tunnel, NMAG additionally joins those groups/channels on the tunnel interface that it wants to receive by forwarding from PMAG. When multicast streams arrive, the NMAG forwards data to the appropriate downlink(s). NMAG MUST send a LEAVE message to the tunnel immediately after forwarding of a group/channel becomes unneeded, e.g., after native multicast traffic arrives or

group membership of the MN terminates.

4.2.4. IPv4 Support Considerations

An MN in a PMIPv6 domain may use an IPv4 address transparently for communication as specified in [RFC5844]. For this purpose, LMAs can register IPv4-Proxy-CoAs in its Binding Caches and MAGs can provide IPv4 support in access networks. Correspondingly, multicast membership management will be performed by the MN using IGMP. For multiprotocol multicast support on the network side, IGMPv3 router functions are required at both MAGs (see Section 5.6 for compatibility considerations with previous IGMP versions). Context transfer between MAGs can transparently proceed in HI/HACK message exchanges by encapsulating IGMP multicast state records within Multicast Mobility Options (see Section 5.3 and Section 5.4 for details on message formats.

It is worth mentioning the scenarios of a dual-stack IPv4/IPv6 access network, and the use of GRE tunneling as specified in[RFC5845]. Corresponding implications and operations are discussed in the PMIP Multicast Base Deployment document, cf., [I-D.ietf-multimob-pmipv6-base-solution].

5. Message Formats

5.1. Multicast Indicator for Proxy Router Advertisement (PrRtAdv)

An FMIPv6 AR will indicate its multicast support by activating the M-bit in its Proxy Router Advertisements (PrRtAdv). The message extension has the following format.

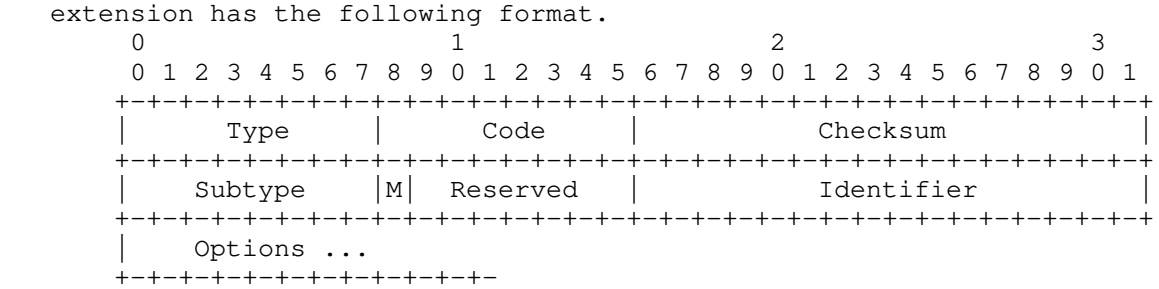

Figure 6: Multicast Indicator Bit for Proxy Router Advertisement (PrRtAdv) Message

5.2. Extensions to Existing Mobility Header Messages

The fast handover protocols use a new IPv6 header type called Mobility Header as defined in [RFC3775]. Mobility headers can carry variable Mobility Options.

Multicast listener context of an MN is transferred in fast handover operations from PAR/PMAG to NAR/NMAG within a new Multicast Mobility Option, and acknowledged by a corresponding Acknowledgement Option. Depending on the specific handover scenario and protocol in use, the corresponding option is included within the mobility option list of HI/HAck only (PFMIPv6), or of FBU/FBAck/HI/HAck (FMIPv6).

5.3. New Multicast Mobility Option

The Multicast Mobility Option contains the current listener state record of the MN obtained from the MLD Report message, and has the format displayed in Figure 7.

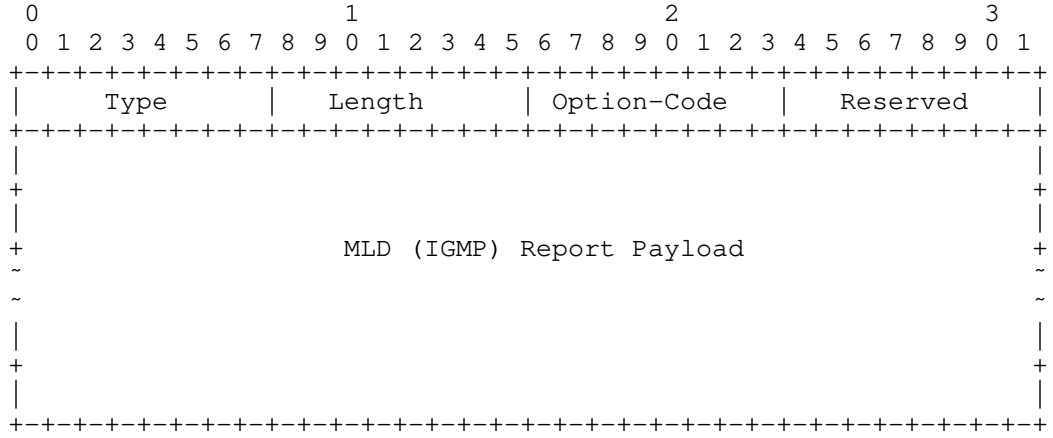

Figure 7: Mobility Header Multicast Option

Type: TBD

Length: 8-bit unsigned integer. The size of this option in 8 octets including the Type, Option-Code, and Length fields.

Option-Code:

1: IGMPv3 Payload Type

2: MLDv2 Payload Type

3: IGMPv3 Payload Type from IGMPv2 Compatibility Mode

4: MLDv2 Payload Type from MLDv1 Compatibility Mode

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

MLD (IGMP) Report Payload: this field is composed of the MLD (IGMP) Report message after stripping its ICMP header. Corresponding message formats are defined for MLDv2 in [RFC3810], and for IGMPv3 in [RFC3376].

Figure 8 shows the Report Payload for MLDv2, while the payload format for IGMPv3 is defined corresponding to the IGMPv3 payload format (see Section 5.2. of [RFC3810], or Section 4.2 of [RFC3376]) for the definition of Multicast Address Records).

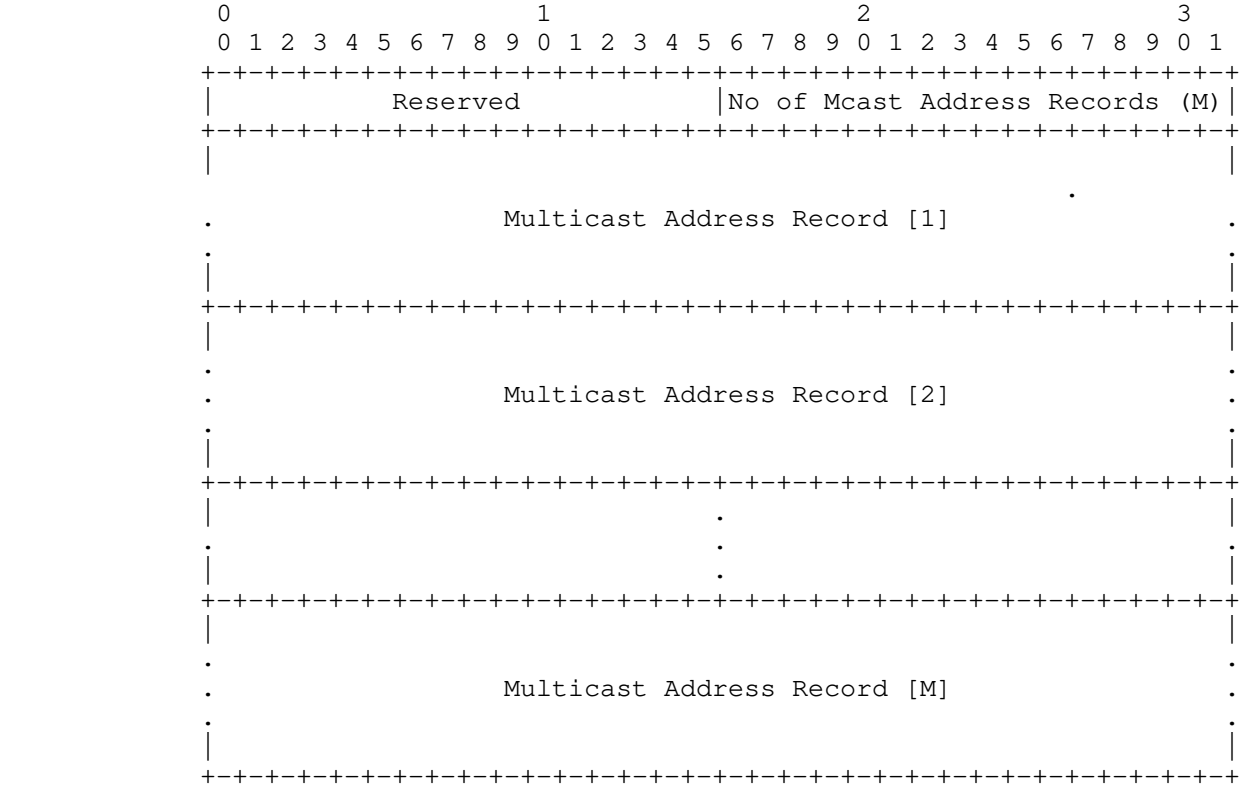

Figure 8: MLDv2 Report Payload

5.4. New Multicast Acknowledgement Option

The Multicast Acknowledgement Option reports the status of the context transfer and contains the list of state records that could not be successfully transferred to the next access network. It has the format displayed in Figure 9.

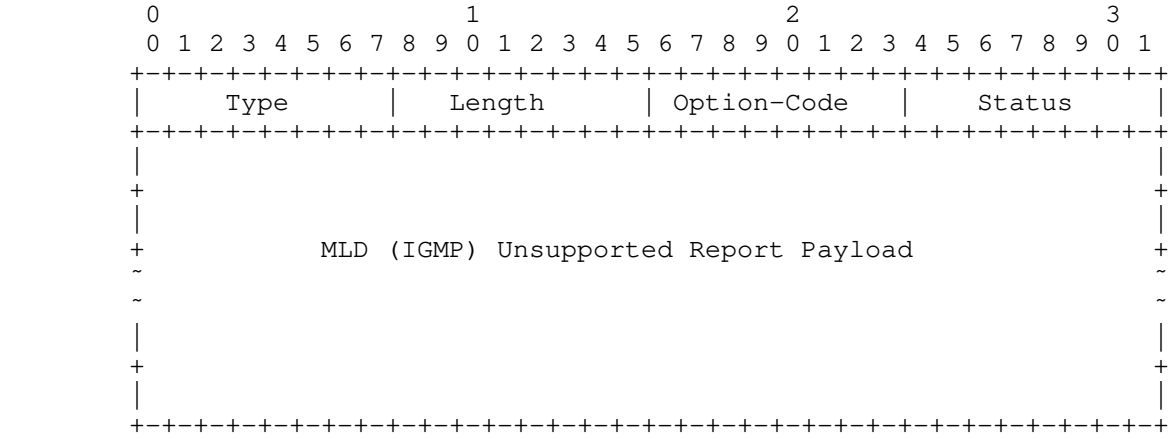

Figure 9: Mobility Header Multicast Acknowledgement Option

Type: TBD

Length: 8-bit unsigned integer. The size of this option in 8 octets. The length is 1 when the MLD (IGMP) Unsupported Report Payload field contains no Mcast Address Record.

Option-Code: 0

Status:

1: Report Payload type unsupported

2: Requested group service unsupported

3: Requested group service administratively prohibited

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

MLD (IGMP) Unsupported Report Payload: this field is syntactically identical to the MLD (IGMP) Report Payload field described in Section 5.3, but is only composed of those multicast address records that are not supported or prohibited in the new access network. This field MUST always contain the first header line (reserved field and

No of Mcast Address Records), but MUST NOT contain any Mcast Address Records, if the status code equals 1.

Note that group subscriptions to specific sources may be rejected at the destination network, and thus the composition of multicast address records may differ from initial requests within an MLD (IGMP) Report Payload option.

5.5. Length Considerations: Number of Records and Addresses

Mobility Header Messages exchanged in HI/HACK and FBU/FBACK dialogs impose length restrictions on multicast context records. The maximal payload length available in FBU/FBACK messages is the PATH-MTU - 40 octets (IPv6 Header) - 6 octets (Mobility Header) - 6 octets (FBU/ FBACK Header). For example, on an Ethernet link with an MTU of 1500 octets, not more than 72 Multicast Address Records of minimal length (without source states) may be exchanged in one message pair. In typical handover scenarios, this number reduces further according to unicast context and Binding Authorization data. A larger number of MLD Report Payloads MAY be sent within multiple HI/HACK or FBU/FBACK message pairs. In PFMIPv6, context information can be fragmented over several HI/HACK messages. However, a single MLDv2 Report Payload MUST NOT be fragmented. Hence, for a single Multicast Address Record on an Ethernet link, the number of source addresses is limited to 89.

5.6. MLD (IGMP) Compatibility Aspects

Access routers (MAGs) MUST support MLDv2 (IGMPv3). To enable multicast service for MLDv1 (IGMPv2) listeners, the routers MUST follow the interoperability rules defined in [RFC3810] ([RFC3376]) and appropriately set the Multicast Address Compatibility Mode. When the Multicast Address Compatibility Mode is MLDv1 (IGMPv2), a router internally translates the following MLDv1 (IGMPv2) messages for that multicast address to their MLDv2 (IGMPv2) equivalents and uses these messages in the context transfer. The current state of Compatibility Mode is translated into the code of the Multicast Mobility Option as defined in Section 5.3. A NAR (nMAG) receiving a Multicast Mobility Option during handover will switch to the minimum obtained from its previous and newly learned value of MLD (IGMP) Compatibility Mode for continued operation.

6. Security Considerations

Security vulnerabilities that exceed issues discussed in the base protocols of this document ([RFC5568], [RFC5949], [RFC3810], [RFC3376]) are identified as follows.

Multicast context transfer at predictive handovers implements group states at remote access routers and may lead to group subscriptions without further validation of the multicast service requests. Thereby a NAR (nMAG) is requested to cooperate in potentially complex multicast re-routing and may receive large volumes of traffic. Malicious or inadvertent multicast context transfers may result in a significant burden of route establishment and traffic management onto the backbone infrastructure and the access router itself. Rapid re-routing or traffic overload can be mitigated by a rate control at the AR that restricts the frequency of traffic redirects and the total number of subscriptions. In addition, the wireless access network remains protected from multicast data injection until the requesting MN attaches to the new location.

7. IANA Considerations

This document defines new flags and status codes in the HI and HAck messages as well as two new mobility options. The Type values for these mobility options are assigned from the same numbering space as allocated for the other mobility options defined in [RFC3775]. Those for the flags and status codes are assigned from the corresponding numbering space defined in [RFC5568], or [RFC5949] and requested to be created as new tables in the IANA registry (marked with asterisks). New values for these registries can be allocated by Standards Action or IESG approval [RFC5226].

8. Acknowledgments

Protocol extensions to support multicast in Fast Mobile IPv6 have been loosely discussed since several years. Repeated attempts have been taken to define corresponding protocol extensions. The first draft [fmcast-mip6] was presented by Suh, Kwon, Suh, and Park already in 2004.

This work was stimulated by many fruitful discussions in the MobOpts research group. We would like to thank all active members for constructive thoughts and contributions on the subject of multicast mobility.

9. References

9.1. Normative References

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

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

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

[RFC5568] Koodli, R., "Mobile IPv6 Fast Handovers", RFC 5568, July 2009.

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

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

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

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

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

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

9.2. Informative References

[RFC5757] Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey", RFC 5757, February 2010.

[fmcast-mip6]
Suh, K., Kwon, D., Suh, Y., and Y. Park, "Fast Multicast Protocol for Mobile IPv6 in the fast handovers environments", draft-suh-mipshop-fmcast-mip6-00 (work in progress), July 2004.

[FMIPv6-Analysis]
Schmidt, TC. and M. Waehlisch, "Predictive versus Reactive - Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility", Telecommunication

Systems Vol 33, No. 1-3, pp. 131-154, November 2005.

[I-D.ietf-multimob-pmipv6-base-solution]
Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-06 (work in progress), October 2010.

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

[RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", RFC 5845, June 2010.

Appendix A. Change Log

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-02.

1. Detailed operations on PFMIPv6 entitities completed.

2. Some editorial improvements & clarifications.

3. References updated.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-01.

1. First detailed operations on PFMIPv6 added.

2. IPv4 support considerations for PFMIPv6 added.

3. Section on length considerations for multicast context records corrected.

4. Many editorial improvements & clarifications.

5. References updated.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-00.

1. Editorial improvements & clarifications.

2. Section on length considerations for multicast context records added.

3. Section on MLD/IGMP compatibility aspects added.

4. Security section added.

Authors' Addresses

Thomas C. Schmidt
HAW Hamburg
Dept. Informatik
Berliner Tor 7
Hamburg, D-20099
Germany

Email: schmidt@informatik.haw-hamburg.de

Matthias Waehlisch
link-lab & FU Berlin
Hoenower Str. 35
Berlin D-10318
Germany

Email: mw@link-lab.net

Rajeev Koodli
Cisco Systems
30 International Place
Xuanwu District,
Tewksbury MA 01876
USA

Email: rkoodli@cisco.com

Godred Fairhurst
University of Aberdeen
School of Engineering
Aberdeen AB24 3UE
UK

Email: gorry@erg.abdn.ac.uk

MULTIMOB Working Group
Internet-Draft
Intended status: Informational
Expires: December 9, 2010

D. von Hugo
Deutsche Telekom Laboratories
H. Asaeda
Keio University
B. Sarikaya
Huawei USA
P. Seite
France Telecom - Orange

June 8, 2010

Evaluation of further issues on Multicast Mobility: Potential future work for WG MultiMob
<draft-von-hugo-multimob-future-work-02.txt>

Abstract

The WG MultiMob aims at defining a basic mobile multicast solution leveraging on network localized mobility management, i.e. Proxy Mobile IPv6 protocol. The solution would be basically based on multicast group management, i.e. IGMP/MLD, proxying at the access gateway. If such a basic solution is essential from an operational point of view, challenges with efficient resource utilization and user perceived service quality still persist. These issues may prevent large scale deployments of mobile multicast applications.

This document attempts to identify topics for near future extension of work such as modifying multimob base solution, PMIPv6 and MLD/IGMP for optimal multicast support, and adaptation of Handover optimization. Far future items such as extending to and modifying of MIPv4/v6 and DSMIP, sender (source) mobility, consideration of multiple flows and multihoming will be dealt with in a future version.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at

http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on December 9, 2010.

Table of Contents

1. Introduction

Chartered work of WG MultiMob focuses on documentation of proper configuration and usage of existing (specified standard) protocols within both mobility and multicast related areas to enable and support mobility for multicast services and vice versa. The current WG document [I-D.ietf-multimob-pmipv6-base-solution] does not address specific optimizations and efficiency improvements of multicast routing for network-based mobility and thus the operation may be not resource efficient nor grant the service quality expected by the end user.

The described solution resolves the problem to ensure multicast reception in PMIPv6-enabled [RFC5213] networks without appropriate multicast support. However it neither automatically minimizes multicast forwarding delay to provide seamless and fast handovers for real-time services nor minimizes packet loss and reordering that result from multicast handover management as stated in [RFC5757]. Also Route Optimization is out of scope of the basic solution - an issue for reducing amount of transport resource usage and transmission delay. Thus possible enhancements and issues for solutions beyond a basic solution need to be described to enable current PMIPv6 protocols to fully support efficient mobile multicast services. Such extensions may include protocol modifications for both mobility and multicast related protocols to achieve optimizations for resource efficient and performance increasing multimob approaches. The document includes the case of mobile multicast senders using Any Source Multicast (ASM) and Source Specific Multicast (SSM) [RFC4607].

This document focuses on discussion work on multicast protocols such as IGMP/MLD operational tuning (e.g. as proposed in [I-D.asaeda-igmp-mld-optimization]) and enhancements of IGMP/MLD protocol behaviors and messages for optimal multicast support (proposed in [I-D.asaeda-igmp-mld-mobility-extension]).

An alternative approach proposes the addition of acknowledgement messages on group management ([I-D.liu-multimob-reliable-igmp-mld]) and changes the unreliable protocol concept.

Furthermore a modification of PMIPv6 by introducing a dedicated multicast tunnel and support of local routing is discussed in [I-D.asaeda-multimob-pmip6-extension]. Other performance improvements have been outlined in [I-D.schmidt-multimob-fmipv6-pfmipv6-multicast] where extensions to Mobile IPv6 Fast Handovers (FMIPv6) [RFC5568], and the corresponding extension for Proxy MIPv6 operation [I-D.ietf-mipshop-pfmipv6].

Another type of multimob work aims directly at enhancements of the current multimob base solution [I-D.ietf-multimob-pmipv6-base-solution] towards introduction of multicast traffic replication mechanisms and a reduction of the protocol complexity in terms of time consuming tunnel set-up by definition of pre- or post-configured tunnels (as provided by e.g. [I-D.zuniga-multimob-smspmip]). Further work within this topic deals with direct routing (e.g. [I-D.sijeon-multimob-mms-pmip6]) and with dynamic or automatic tunnel configuration (see e.g. [I-D.ietf-mboned-auto-multicast]).

A large field of additional investigations which are partly described in detail in [RFC5757] will be mentioned for completeness and may be subject of a later WG re-chartering.

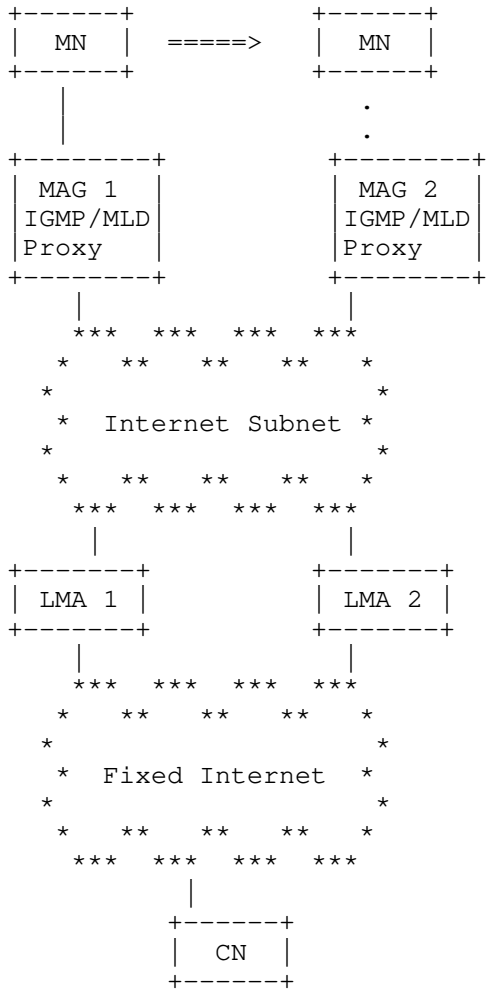

Figure 1: MultiMob Scenario for chartered PMIP6 issue

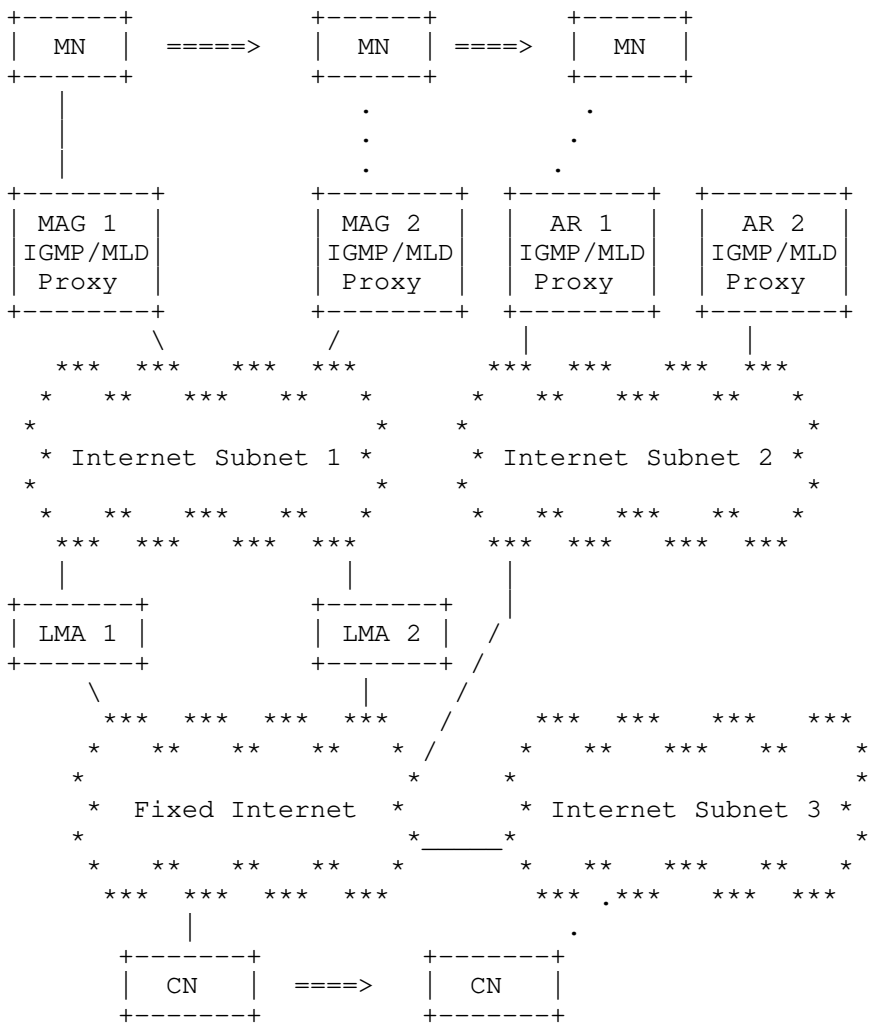

Figure 2: MultiMob scenario for extended MultiMob issues

Figure 1 illustrates the key components of the foreseen basic Multimob solution. The extended multicast mobility scenario, leading to above issues, is sketched in Figure 2.

In summary additional to a 'Single hop, link, flow' Proxy MIP mobility for listening MNs (scenario shown in Figure 1), future work towards a complete performance-optimized scenario of a 'Multi-hop, -homed, -flow' client mobility (i.e. including MIPv6 [RFC3775] and DSMIPv6 [RFC5555]) would cover a plurality of issues. For the near

future we see the following issues as most important:

- o Extension of multimob base solution
- o Modification of base PMIPv6 and MLD/IGMP for optimal multicast support.
- o Consideration of Handover optimization.

All further issues which would include extensions to and modifications of MIPv4/v6 and DSMIP using IGMP/MLD Proxy and the Foreign Agent/Access Router, consideration of sender (source) mobility, support of multiple flows on multihomed mobile nodes, multi-hop transmission, Routing optimization, and so forth will be topics for a potential next stage of future work extension.

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 BCP 14 [RFC2119].

This document uses the terminology defined in [RFC3775], [RFC3376], [RFC3810], [RFC5213], [RFC5757].

3. IGMP/MLD Proxy Architecture

Multimob basic solution is based on IGMPv3/MLDv2 Proxy support at the mobile access gateway (MAG) of Proxy Mobile IPv6 as shown in Figure 1. IGMPv3/MLDv2 proxy keeps multicast state on the subscriptions of the mobile nodes and only an aggregate state is kept at the local mobility anchor (LMA). When LMA receives multicast data it can forward it to the MAG without duplication because MAG takes of the packet duplication. This leads to solving the avalanche problem.

By keeping multicast state locally, IGMPv3/MLDv2 Proxy introduces mobility related problems such as possible packet loss when a mobile node does a handover to another MAG and its multicast state is not modified fast enough at the LMA.

IGMPv3/MLDv2 introduces tunnel convergence problem which occurs when a given MAG serves MNs that belong to different LMAs and MNs subscribe to the same multicast group. In that case MNs receive duplicate multicast data forwarded from more than one LMA.

It can be foreseen that mobile access gateways will serve both mobile and fixed terminals concurrently. The tuning of multicast-related

protocol parameters based on the terminal characteristics is needed. Parameters only applicable to mobile users need to be distinguished from the parameters applicable to fixed users. It should be also possible to distinguish between slow and fast movement and handover frequency to form corresponding tunnels for mobile users.

Based on the above observations we will state the problems next and then list the requirements on possible solutions.

4. Problem Description

The general issues of multicast mobility are extensively discussed and described in [RFC5757]. To reduce the complexity of the plethora of requirements listed in [RFC5757] and also in [I-D.deng-multimob-pmip6-requirement] this document summarises some lightweight solutions for multicast mobility which allow for easy deployment within realistic scenarios and architectures. Moreover we focus on approaches building directly on basic MultiMob solution [I-D.ietf-multimob-pmipv6-base-solution] which is based on IGMP/MLD Proxy functionality at the mobile access gateway, and for which already solution proposals have been described.

4.1. Modification of base PMIPv6 for optimal multicast support

Currently discussed aspects of multicast optimization for PMIPv6 include introduction of multicast tunnels and support of local routing such as described in [I-D.asaeda-multimob-pmip6-extension]. For a PMIPv6 domain the establishment of a dedicated multicast tunnel is proposed which may either be dynamically set up and released or be pre-configured in a static manner. Both mobility entities MAG and LMA may be operate as MLD proxy or multicast router.
Since further functional enhancements of PMIPv6 are currently under way in NETEXT WG, both the impact of new features on Mobile Multicast as well as such a Multicast-initiated proposal for PMIPv6 modification have to be considered in a continuous exchange process between MultiMob and NETEXT WGs.

4.2. Modification of MLD/IGMP for optimal multicast support

Potential approaches for enhancement of group management as specified e.g. by MLDv2 [RFC3810] include operational improvements such as proper tuning in terms of default timer value modification, specific query message introduction, and standard (query) reaction suppression, beside introducing multicast router attendance control in terms of e.g. specification of a Listener Hold message as proposed in [I-D.asaeda-multimob-igmp-mld-mobility-extensions].

4.3. Consideration of Handover Optimization

Ideally the customer experience while using multicast services should not be affected by transmission issues whether the terminal is operated in a fixed or a mobile environment. This implies not only that the terminal should be unaware of changes at network layer connectivity (seamless communication) as is typically the case in a PMIPv6 domain, but also that any impact of connectivity changes (handover) should be minimized. In the framework of Multimob this relates to reduction of delay, packet loss, and packet reordering effort for mobile multicast by applying fast handover mechanisms, which have originally been developed for unicast traffic to multicast group management. [I-D.schmidt-multimob-fmipv6-pfmipv6-multicast] works on specification of extension of the Mobile IPv6 Fast Handovers (FMIPv6) [RFC5568] and the Fast Handovers for Proxy Mobile IPv6 (PFMIPv6) [I-D.ietf-mipshop-pfmipv6] protocols to include multicast traffic management in fast handover operations. Issues for further work are details of including multicast group messaging in context transfer, for both predictive and reactive handover mode, as well as details of corresponding message exchange protocols and message design.

4.4. Specific PMIP deployment issues

Currently several proposals are under work which describe extensions of the base protocol WG draft [I-D.ietf-multimob-pmipv6-base-solution]. While MAG operation will remain that of an MLD proxy additional LMA functionalities are described in [I-D.zuniga-multimob-smspmip] which allow for replication of multicast traffic and solution of the tunnel convergence problem. The dedicated multicast LMA may either set up dedicated multicast tunnels dynamically or a-priory via pre-configuration or a delayed release.

Another solution on dynamic and/or automatic tunnel configuration is proposed within multicast WG MBONED [I-D.ietf-mboned-auto-multicast].

A direct or local routing approach is described in [I-D.sijeon-multimob-mms-pmip6]. This scenario may hold for short term deployment focusing on an architecture where multicast traffic is provided via the home network. However, depending on the network topology, namely the location of the content delivery network, the LMA may not be on the optimal multicast service delivery path. This enables mobile nodes to access locally available multicast services such as local channels.

Figure 3 illustrates the use-case for local routing.

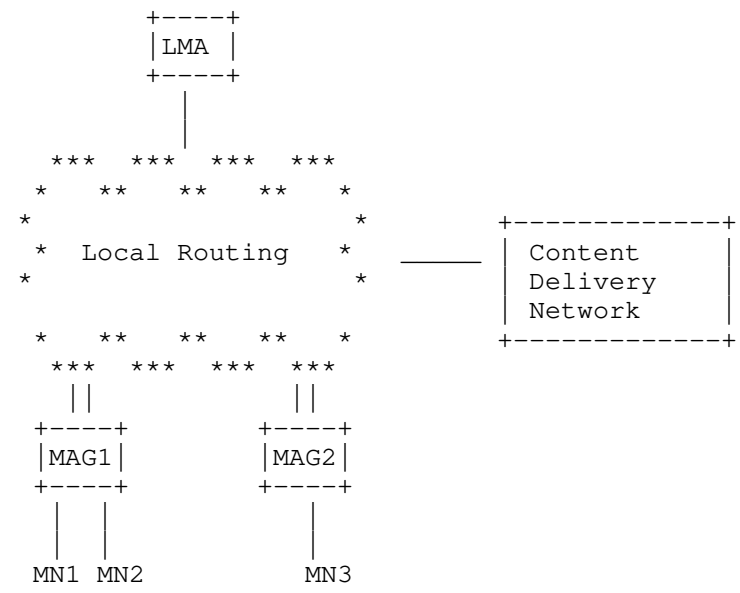

Figure 3: local Multicast routing

In such a case, the MAG should act as a multicast router to construct the optimal multicast delivery path. If the MAG also supports MLD proxy function issue raises up on the dual mode behaviour. In such a case, a pragmatic approach could be to leverage only on multicast routing at the MAG in the PMIP domain.

Whatever is the MAG operation mode, the multicast state is locally kept at the access gateway, so unknown from the mobility anchor. In other words, the multicast service is independent from the mobility service that the mobile node is receiving from the network in the form of PMIPv6 or DSMIPv6. However, handover support is still desirable but cannot be provided by the mobility anchor (i.e. HA or LMA). In such a case mobility support for locally available multicast should be provided by extending multicast protocols of IGMP or MLD.

5. Requirements on Solutions

This section tries to identify requirements from the issues discussed in previous section.

- o Seamless handover (low latency and during the handover).
- o Similar packet loss to unicast service.
- o Multiple LMAs architecture.
- o Agnostic mobile host re-subscription. So, MAGs must be able to retrieve multicast contexts of the mobile nodes.
- o Solution address IPv6, IPv4 only and dual stack nodes.
- o Supports sender (source) mobility.
- o Optimal local routing.
- o To be completed...

6. Security Considerations

This draft introduces no additional messages. Compared to [RFC3376], [RFC3810], [RFC3775], and [RFC5213] there have no additional threats been introduced.

7. IANA Considerations

Whereas this document does not explicitly introduce requests to IANA some of the proposals referenced above (such as [I-D.asaeda-multimob-pmip6-extension] and [I-D.schmidt-multimob-fmipv6-pfmipv6-multicast]) specify flags for mobility messages or options. For details please see those documents.

8. Acknowledgements

The authors would thank all active members of MultiMob WG, especially (in no specific order) Gorry Fairhurst, Jouni Korhonen, Thomas Schmidt, Suresh Krishnan and Matthias Waehlisch for providing continuous support and helpful comments.

9. References

9.1. Normative References

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

[RFC3314] Wasserman, M., "Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards", RFC 3314, September 2002.

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

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

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

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

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

[RFC5555] Soliman, H., "Mobile IPv6 Support for Dual Stack Hosts and Routers", RFC 5555, June 2009.

9.2. Informative References

[23246] "3GPP TS 23.246 V8.2.0, Multimedia Broadcast/Multicast Service (MBMS); Architecture and functional description (Release 8).", 2008.

[23401] "3GPP TS 23.401 V8.2.0, General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 8).", 2008.

[23402] "3GPP TS 23.402 V8.4.1, Architecture enhancements for non-3GPP accesses (Release 8).", 2009.

[I-D.asaeda-multimob-igmp-mld-mobility-extensions] Asaeda, H. and T. Schmidt, "IGMP and MLD Hold and Release Extensions for Mobility", draft-asaeda-multimob-igmp-mld-mobility-extensions-03 (work in progress), July 2009.

[I-D.asaeda-multimob-igmp-mld-optimization] Asaeda, H. and S. Venaas, "Tuning the Behavior of IGMP and MLD for Mobile Hosts and Routers", draft-asaeda-multimob-igmp-mld-optimization-02 (work in progress), March 2010.

[I-D.asaeda-multimob-pmip6-extension] Asaeda, H., Seite, P., and J. Xia, "PMIPv6 Extensions for Multicast", draft-asaeda-multimob-pmip6-extension-02 (work in progress), July 2009.

[I-D.deng-multimob-pmip6-requirement] Deng, H., Chen, G., Schmidt, T., Seite, P., and P. Yang, "Multicast Support Requirements for Proxy Mobile IPv6", draft-deng-multimob-pmip6-requirement-02 (work in progress), July 2009.

[I-D.liu-multimob-reliable-igmp-mld]
Liu, H. and Q. Wu, "Reliable IGMP and MLD Protocols in Wireless Environment", draft-liu-multimob-reliable-igmp-mld-00 (work in progress), March 2010.

[I-D.schmidt-multimob-fmipv6-pfmipv6-multicast]
Schmidt, T., Waehlisch, M., Koodli, R., and G. Fairhurst, "Multicast Listener Extensions for MIPv6 and PMIPv6 Fast Handovers", draft-schmidt-multimob-fmipv6-pfmipv6-multicast-01 (work in progress), March 2010.

[I-D.sijeon-multimob-mms-pmip6]
Jeon, S. and Y. Kim, "Mobile Multicasting Support in Proxy Mobile IPv6", draft-sijeon-multimob-mms-pmip6-02 (work in progress), March 2010

[I-D.zuniga-multimob-smspmip]
Zuniga, J., Lu, G., and A. Rahman, "Support Multicast Services Using Proxy Mobile IPv6", draft-zuniga-multimob-smspmip-02 (work in progress), June 2010.

[I-D.ietf-mboned-auto-multicast]
Thaler, D., Talwar, M., Aggarwal, A., Vicisano, L., and T. Pusateri, "Automatic IP Multicast Without Explicit Tunnels (AMT)", draft-ietf-mboned-auto-multicast-10 (work in progress), March 2010

[I-D.ietf-16ng-ipv4-over-802-dot-16-ipcs]
Madanapalli, S., Park, S., Chakrabarti, S., and G. Montenegro, "Transmission of IPv4 packets over IEEE 802.16's IP Convergence Sublayer", draft-ietf-16ng-ipv4-over-802-dot-16-ipcs-07 (work in progress), June 2010.

[I-D.ietf-manet-smf]
Macker, J. (editor), "Simplified Multicast Forwarding", draft-ietf-manet-smf-10 (work in progress), March 2010.

[I-D.ietf-mipshop-pfmipv6]
Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", draft-ietf-mipshop-pfmipv6-14 (work in progress), May 2010

[I-D.ietf-multimob-pmipv6-base-solution]
Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-02 (work in progress), May 2010.

[RFC5757] Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast Mobility in MIPv6: Problem Statement and Brief Survey", RFC 5757, June 2010.

[RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.

[RFC5121] Patil, B., Xia, F., Sarikaya, B., Choi, JH., and S. Madanapalli, "Transmission of IPv6 via the IPv6 Convergence Sublayer over IEEE 802.16 Networks", RFC 5121, February 2008.

Authors' Addresses

Dirk von Hugo
Deutsche Telekom Laboratories
Deutsche-Telekom-Allee 7
64295 Darmstadt, Germany

Email: dirk.von-hugo@telekom.de

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-8520
Japan

Email: asaeda@wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~asaeda/

Behcet Sarikaya
Huawei USA
1700 Alma Dr. Suite 500
Plano, TX 75075

Email: sarikaya@ieee.org

Pierrick Seite
France Telecom - Orange
4, rue du Clos Courtel
BP 91226
Cesson-Sevigne, BZH 35512
France

Email: pierrick.seite@orange-ftgroup.com

Network working group
Internet Draft
Category: Informational
Created: October 25, 2010
Expires: April 2011

Q. Wu
H. Liu
Huawei

Proposal for Tuning IGMPv3/MLDv2 Protocol Behavior in Wireless and Mobile networks

draft-wu-multimob-igmp-mld-tuning-03

Abstract

This document proposes a variety of optimization approaches for tuning IGMPv3 and MLDv2 protocols. It aims to provide useful guideline to allow efficient multicast communication in wireless and mobile networks using the current IGMP/MLD protocols.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Status of this Memo

at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on August 15, 2009.

Table of Contents

1. Introduction

Multicasting is more efficient a method of supporting group communication than unicasting. With the wide deployment of different wireless networks, multicast communication over wireless network comes to attract more and more interests from content and service providers, but still faces great challenges when considering dynamic group membership and constant update of delivery path due to node movement, which is highly required in the wireless or mobile network. On the other hand, unlike wired network, some of wireless networks often offer limited reliability, consume more power and cost more transmission overhead, thus in worse case are more prone to loss and congestion.

Multicast network is generally constructed by IGMP/MLD group management protocol to track valid receivers and by multicast routing protocol to build multicast delivery paths. This document focuses only on IGMP/MLD protocols, which are used by a mobile user to subscribe a multicast group and are most possibly to be exposed to wireless link to support terminal mobility. As IGMP and MLD are designed for fixed users using wired link, they does not work perfectly for wireless link types. They should be enhanced or tuned to adapt to wireless and mobile environment to meet the reliability and efficiency requirements in the scenarios described in [REQUIRE][RFC 5757].

This memo proposes a variety of optimization approaches for tuning IGMP/MLD protocols in wireless or mobile communication environment. It aims to make the minimum tuning on the protocol behavior without introducing interoperability issues, and to improve the performance of wireless and mobile multicast networks. These solutions can also be used in wired network when efficiency and reliability are required. They are discussed in detail in Section 4.

2. Impact of wireless and mobility on IGMP/MLD

This section analyzes the impact of wireless or mobility on IGMP/MLD by comparing wireless multicast with wired multicast and comparing different wireless link models. It then gives the requirements of

wireless and mobile multicast on IGMP/MLD protocols according to the analysis.

2.1. Comparison analysis between wired and wireless multicast

Existing multicast support for fixed user can be extended to mobile users in wireless environments. However applying such support to wireless multicast is difficult for the following five reasons.

- O Limited Bandwidth: In contrast with wired link, wireless link usually has limited bandwidth. This situation will be made even worse if wireless link has to carry high volume video multicast data. Also the bandwidth available in upstream direction and downstream direction may not be equal.

- O Large packets Loss: In contrast with wired multicast, wireless multicast has packet loss that range between 1% and 30%, based on the links types and conditions. And when packets have to travel between home and access networks e.g. through tunnel, the packets are prone to be lost if the distance between the two networks is long.

- O Frequent Membership change: In fixed multicast, membership change only happens when a user leave or joins a group while in the mobile multicast, membership changes may also occur when a user changes its location.

- O Prone to performance degradation: Due to possible unwanted interaction of protocols across layers and user movement, the wireless network may be overwhelmed with more excessive traffic than wired network. In worse case, this may lead to network performance degrading and network connection complete loss.

- O Increased Leave Latency: Unlike fixed multicast, the leave latency in the mobile multicast will be increased due to user movement. And if the traffic has to be transmitted between access network and the home network, or if the handshake is required between these two networks, the Leave Latency will be increased further more.

Figure 1 shows the details for the difference between wired/fixed multicast and wireless/mobile multicast.

Issues	Wired or fixed Multicast	Wireless/mobile multicast
Bandwidth	Plentiful	Limited and variable possibly asymmetric
Loss of Packets	Infrequent(<1%)	Frequent and variable (1%-30% based on links)
Membership Changes	Only when a user leaves and joins a group	Also when a user moves to another location
Reliability	Possible use of a transport-layer protocol(such as the Multicast File Transfer Protocol)	More complex due to wireless links and user mobility; possible unwanted interaction of protocols at transport and link layers
Leave Latency	not changed by user movement	Increased due to user movement and lost packet

Figure 1. Comparison between wired/fixed multicast and wireless/mobile multicast

2.2. Link models analysis for wireless multicast

There are various types of wireless links, each with different feature and performance. In this document, we according to the transmission mode categorize the wireless link type into three typical link models:

- O Point To Point (PTP) link model
- O Point To Multipoint (PTMP) link model
- O Broadcast link model

PTP link model is the model with one dedicated link that connects exactly two communication facilities. For multicast transmission, each PTP link has only one receiver and the bandwidth is dedicated

for each receiver. Also one unique prefix or set of unique prefixes will be assigned to each receiver. Such link model can be accomplished by running PPP on the link or having separate VLAN for each receiver.

PTMP link model is the model with multipoint link which consists of a series of receivers and one centralized transmitter. Unlike P2P link model, PTMP provide downlink common channels and dedicated uplink channel for each user. Bandwidth and prefix in this model are shared by all the receivers on the same link. Therefore Duplicate Address Detection (DAD) should be performed to check whether the assigned address is used by other receivers.

Broadcast link model is the model with the link connecting two or more nodes and supporting broadcast transmission. Such link model is quite similar to fixed Ethernet link model and its link resource is shared in both uplink and downlink directions. The bandwidth and prefix are shared by all the receivers and DAD is required to avoid address collision.

Figure 2 shows the details for the difference between different wireless link models.

Features	PTP link model	PTMP link model	Broadcast link model
Shared link/ Dedicated link	Dedicated uplink and downlink channels for each user	Common downlink channels and dedicated uplink channels for each user	common downlink Channel for each user
Shared Prefix /Dedicated Prefix	Per Prefix for each receiver No need DAD	Prefix shared by all receivers DAD is required	Prefix shared by all receivers DAD is required
Shared Service Support	Not Support	Support	Support
link layer Broadcast Multicast Support	Only one node On the link Forward multicast packets to the only receiver on the link	Link Layer Multicast Support using Backend (e.g.,AR) IGMP/MLD Snooping at AR	Broadcast Support at L2 using switch IGMP/MLD Snooping at switch
Ethernet link Support	Not support	Not support	Ethernet Support By Implementing Bridge

Figure 2. Wireless Link Models Analysis

2.3. Requirements of wireless and mobile multicast on IGMP/MLD

Due to the characteristics of wireless and mobile multicast described in the section 2.1 and 2.2, it is desirable for IGMP and MLD to have the following characteristics when used in wireless and mobile networks [REQUIRE]:

o Adaptive to different link characteristics: IGMP and MLD are originally designed for wired multicast and some of their processing is not applicable to wireless multicast for its asymmetrical link, limited bandwidth, larger packet loss rate, increased leave latency, and etc. Also Wireless network has various link types, each of them has different bandwidth and performance. These require IGMP/MLD protocol behavior should be tuned to adapt to different link model and link conditions.

o Minimal Join and Leave Latency: Fast join and leave of a subscriber helps to improve the user's experience during channel join and channel zapping. Fast leave also facilitates releasing of unused network resources quickly. Besides, mobility and handover may cause a user to join and leave a multicast group frequently, which also require fast join and leave to accelerate service activation and to optimize resource usages.

o Robustness to packet loss: Wireless link has the characteristic that packet transmission is unreliable due to instable link conditions and limited bandwidth. For mobile IP network, packets sometimes have to travel between home network and foreign network and have the possibility of being lost due to long distance transmission. These network scenarios have more strict robustness requirement on delivery of IGMP and MLD protocol messages.

o Minimum packet transmission: Wireless link resources are usually more precious and limited compared to their wired counterpart, and are prone to be congested when carrying high volume multicast stream. Minimizing packet exchange without degrading general protocol performance should also be emphasized to improve efficiency and make good use of network capacity and processing capability.

o Avoiding packet burst: Large number of packets generated within a short time interval may have the tendency to deteriorate wireless network conditions. IGMP and MLD when using in wireless and mobile networks should be optimized if their protocol message generation has the potential of introducing packet burst.

According to these requirements, in the following parts of the document, current versions of IGMP/MLD protocols are evaluated whether their various protocol aspects are applicable to wireless and mobile multicast communications. They will be optimized to meet these requirements without new features introduced on the wire or link, without new message type defined, and without interoperability issues introduced, which is referred to as "tuning" of IGMP/MLD protocols.

3. Evaluation of IGMP/MLD on wireless and mobile multicast

This section analyzes the applicability of IGMP and MLD to wireless communication in the following aspects:

- O General evaluation of different versions: IGMPv2 [RFC2236] and MLDv1 [RFC2710] only support ASM communication mode. They do not support SSM subscription and explicit tracking. IGMPv3 [RFC3376] and MLDv2 [RFC3810] and their lightweight version LW-IGMPv3/LW-MLDv2 [RFC5760] support all the features of ASM/SSM communication modes and explicit tracking. Because SSM is more efficient and secure than ASM for IPTV application, and explicit tracking enables faster channel zapping and better manageability capability, IGMPv3/MLDv2 and LW-IGMPv3/MLDv2 are more promising to be deployed widely than IGMPv2 and MLDv1.

- O Robustness: IGMP/MLD actively sends unsolicited Report or Leave message to join or leave a group, and solicited Report to respond to Queries. Unsolicited Report and Leave messages are more important for ensuring satisfactory user experience and should be guaranteed to improve service performance. Current IGMP and MLD provide the reliability for these messages by non responsive retransmission, which is not adequate from both the robustness and efficiency aspects when they are used on unreliable wireless link or have to be exchanged over the tunnel between home network and access network separated by long distance [ROBUST][ACK]. For IGMPv3/MLDv2, because unsolicited report and leave messages will not be suppressed by report from other host, it is possible to adopt acknowledgement-retransmission to improve reliability and reduce superfluous packet transmission [IGMP-ACK].

  Besides, for IGMPv3/MLDv2, because the router could by explicit tracking establishes membership database recording each valid receiver, it is possible to deduce the possible loss of some protocol messages according to the feedback after their transmission, and to take some remedies (e.g. by retransmission)

to enable more reliable transmission of these messages in bad conditions.

O Efficiency: IGMPv2 and MLDv1 use host suppression to suppress duplicated membership reports on the link. In IGMPv3 and MLDv2, because host suppression is not adopted, the report count will be numerous if the number of valid receivers on the network is large. IGMPv3 and MLDv2 should be optimized to try to minimize unnecessary packet transmission to compensate this drawback. As an example, because an IGMPv3/MLDv2 router has record of each user in its state database by explicit tracking, it is possible to eliminate the need for query timeouts when receiving leave messages and to improve the efficiency by reducing both the unnecessary Queries and reports generated on a network.

And as described in [REQUIRE] and [RFC5757], the default timer values and counter values specified in IGMP and MLD were not designed for the mobility context. This may result in a slow reaction following a client join or leave, in possible packet loss under worse conditions, or in overburdening the wireless link by excessive packets exchange than necessary. These issues can be addressed by tuning these parameters for the expected packet loss on a link to optimize service performance and resource usage.

The comparison between IGMPv2/MLDv1 and IGMPv3/MLDv2 is illustrated in figure 3. In summary, it is desirable to choose IGMPv3/MLDv2 or LW-IGMPv3/MLDv2 as the group management protocol for wireless or mobile multicast. They should be optimized to adapt to wireless and mobile networks to meet the efficiency and reliability requirement for these networks. These optimizations range from the tuning of the parameters (e.g. the Query Interval and other variables), to the tuning of protocol behavior without introducing interoperability issues. Considering an enhancement in one direction might introduce side effects in another one, balances should be taken carefully to avoid defects and improve protocol performance as a whole.

Issues	IGMPv2/MLDv1	IGMPv3/MLDv2
Default Timer and Robustness Variable	Not designed for Mobility context Need to be tuned	Not designed for Mobility context Need to be tuned
Explicit Tracking	Not Support	Support
ASM and SSM Subscription	Only Support ASM Subscription	Both Support
Explicit Join and Leave	Support	Support
Host Suppression	Support	Not Support

Figure 3. Comparison between IGMPv2/MLDv1 and IGMPv3/MLDv2

4. IGMP/MLD tuning optimization for Wireless or Mobile Network

As mentioned in section 2, IGMPv3/MLDv2 or LW-IGMPv3/MLDv2 is recommended to be used as the basis for optimization of IGMP/MLD to adapt to wireless and mobile networks. In this section, taking these characteristics requirement into account, we will discuss several optimization approaches for tuning of IGMPv3 and MLDv2 in wireless environment. The optimizations try to minimize the packet transmission for both the Reports and Queries, and at the meanwhile take the factor of improving reliability into account, with minimum cost. Different link types are also considered for the tuning behavior.

4.1. Explicit Tracking and Query Suppression

In IGMPv2/MLDv1, the member reports are suppressed if the same report has already been sent by another host in the network which is also referred to as host suppression. As described in the A.2 of [RFC3810], the suppression of multicast listener reports has been removed in MLDv2 due to the following reasons:

O Routers may want to track per-host multicast listener status on an interface. This enables the router to track each individual host that is joined to a particular group or channel and allow minimal leave latencies when a host leaves a multicast group or channel.

o Multicast Listener Report suppression does not work well on bridged LANs. Many bridges and Layer2/Layer3 switches that implement MLD snooping do not forward MLD messages across LAN segments in order to prevent multicast listener report suppression.

o By eliminating multicast listener report suppression, hosts have fewer messages to process; this leads to a simpler state machine implementation.

o In MLDv2, a single multicast listener report now bundles multiple multicast address records to decrease the number of packets sent. In comparison, the previous version of MLD required that each multicast address be reported in a separate message.

Without host suppression, it is possible to enable explicit tracking on a router by which the local replication can be used by the router to inspect incoming join and leave requests, record or refresh the membership state for each host on the interface, and take appropriate action to each received report. In the meanwhile, the router builds a table to track which channel being forwarded to each port. If the channel being requested to view is already being received at the router, it can replicate the stream and forward to this new requester which ensure good response time.

By using the tracking table mentioned above, the router has the capability to learn if a particular multicast address has any members on an attached link or if any of the sources from the specified list for the particular multicast address has any members on an attached link or not. Such capability makes Group specific Query or Source-and-Group Specific Queries, which are sent to query other members when a member leaves, unnecessary to be used because the router has already known who are active on the interface using explicit tracking. Therefore it is desirable that these two Queries are eliminated when explicit tracking is used. But General periodical Query by a router to solicit current state reports to refresh existing membership state database should still be used to prevent incorrectness of the database due to the possible loss of explicit join and leave message in some cases.

The main benefits of using explicit tracking without Group specific Query or Source-and-Group Specific Queries are that it provides:

- O minimizing packet number and packet burst: Elimination of Group and Source-Group specific Queries when a member leaves a group will reduce the number of transmitted Group Specific Queries. And finally the total number of Reports in response to Group Specific Queries can be drastically reduced.

- O Minimal leave latencies: an IGMPv3/MLDv2 router configured with explicit tracking can immediately stop forwarding traffic if the last host to request to receive traffic from the router indicates its leave from the group.

- O Faster channel changing: The channel change time of the receiver application depends on the leave latency, that is to say, single host can not receive the new multicast stream before forwarding of the old stream has stopped.

- O Reducing Power consumption: Due to elimination of the suppression of membership reports, the host does not need to spend processing power to hear and determine if the same report has already been sent by another host in the network, which is beneficial to mobile hosts that do not have enough battery power.

4.2. Report Suppression for the hosts

The large number of Reports and bad link condition may result in packets burst. This packet burst can be mitigated by having the router aggregate the responses (membership reports) from multiple clients. The router can intercept IGMP/MLD reports coming from hosts, and forwards a summarized version to the upstream router only when necessary. Typically this means that the router will forward IGMP/MLD membership reports as follows:

- Unsolicited membership reports (channel change requests) are forwarded only when the first subscriber joins a multicast group, or the last subscriber leaves a multicast group. This tells the upstream router to begin or stop sending this channel to this router.

- Solicited membership reports (sent in response to a query) are forwarded once per multicast group. The router may also aggregate multiple responses together into a single membership report.

4.3. Query Suppression for the routers

The large number of Queries and bad link condition may result in packets burst. This packet burst can be mitigated by having the downstream router stop forwarding IGMP/MLD Queries packets sent to

the hosts and respond with report as proxy to the upstream router. Typically this means that the router will:

- Never send a specific query to any client, and

- Send general queries only to those clients receiving at least one multicast group

4.4. Minimizing Query Frequency by increasing interval each time

In IGMPv3/MLDv2, Group Specific Queries and Source and Group specific Queries are sent for [Last Member Query Count] times with short fixed [Last Member Query Interval], to learn whether there are valid members from an attached link. If the network is undergoing congestion, the multiple transmissions of the queries may further deteriorate the bad conditions. To eliminate the bad effects for this, these Queries can be slowed down when a router can not collect successfully expected members' report responses in the mean while it detects the network congestion is going to happen. The slowing down process of the Queries could be arranged in a prolonged time interval as described in [ADAPTIVE].

The slow down behavior is: a router after sending a Query, if acquires the expected responses from the receivers, refreshes its state database and stop the querying retransmission process, or if after a time interval fails to get the expected report responses, resends a Query with an increased (e.g. double) interval. This process can be repeated, for each time the retransmission is arranged in a prolonged time interval, till the router receives the expected responses, or determines the receiver is unreachable and then stops the sending of the Query ultimately. The router can make judgment on not getting expected response from the Queries in the following cases:

O When Group Specific Query and Source and Group Specific Queries are used to track other numbers, the router can not collect any response from the link.

O When all group members leave the group or move out of scope, the General Query sent by the router can not solicit any responses from the link, as mentioned in section 4.9.

O When General Query is retransmitted due to possible loss deducing from no responses from valid members in the database.

- O When General Query is retransmitted by a router on startup [RFC3376][RFC3810], it gets no membership response from the interface.

- O When unicast Query is sent to solicit a particular receiver, if the router can not get responses from the receiver, as described in section 4.5 and 4.6.

In the above cases, if the router fails to get expected response from the network, and if the link condition is bad or in congestion, the router could retransmit the Queries in increased interval. This query retransmission with incremental interval enables the router to reduce the total packet retransmission times in the same time period comparing with retransmission for multiple times with fixed interval, and at the mean time gain some degree of reliability. The variable time interval and the termination condition should be configurable and could be set according to actual network condition, which is out the scope of this document.

4.5. Switching Between Unicast Query and Multicast Query

IGMP/MLD protocols define the use of multicast Queries whose destination addresses are multicast addresses and also allow use of unicast Queries with unicast destination. The unicast Query is sent only for one destination and has the advantages of not affecting other host on the same link. This is especially desirable for wireless communication because the mobile terminal often has limited battery power. But if the number of valid receivers is large, using unicast Query instead of multicast Query will introduce large number of Queries because each Query will be generated for each member, which will not be an efficient use of link resources. In this case the normal multicast Query will be a good choice because only one Query needs to be sent. On the other hand of the number of receivers to be queried is small, the unicast Query is advantageous over multicast one.

The router can choose to switch between unicast and multicast Query according to the practical network conditions. For example, if the receiver number is small, the router could send unicast Queries respectively to each receiver to solicit their membership states, without arousing other host which is in the dormant state. When the receiver number reaches a predefined level, the router could change to use multicast Queries. The router could make the switching flexibly according to practical conditions to improve the efficiency.

4.6. Using General Query with Unicast Query

Unicast Query also can be used in addition to General Query to improve the robustness of solicited reports when General Query fails to collect its valid members. It requires the explicit tracking to be enabled on the router. Its basic behavior is: a router after sending a periodical Query collects successfully all the members' report responses except for one or two which are currently still valid in its database. This may be because the non-respondent ones silently leave the network without any notification, or because their reports are lost due to some unknown reason. The router in this case could choose to unicast a Query respectively to each non-respondent receiver to check whether they are still alive for the multicast reception, without affecting the majority of receivers that have already responded. Unicast Queries under this condition could be sent for [Last Member Query Count] times, following the same rule of [3376] or [3810], or could be resent in incremental interval, as described in section 4.4.

4.7. Retransmission of General Queries

In IGMPv3 and MLDv2, apart from the continuously periodical transmission, General Query is also transmitted during a router's startup. It will be transmitted for [Startup Query Count] times with [Startup Query Interval], to improve reliability of General Query during startup. There are some other cases where retransmission of General Query is beneficial which are not covered by current IGMPv3/MLDv2 protocols as shown in the following.

For example, a router which keeps track of all its active receivers, if after sending a General Query, may fail to get any response from the receivers which are still valid in its membership database. This may be because all the valid receivers leaves the groups or moves out of the range of the link at the moment, or because all the responses of the receivers are lost, or because the sent Query does not arrive at the other side of the link. If current database indicates the number of the valid receiver is not small, the router could choose to compensate this situation by retransmitting the General Query to solicit its active members.

This compensating General Query could be sent several times, if the router can not get any feedback from the receivers which are previous in the database. The repetition of the transmission could in fixed

interval such as [Last Member Query Interval], or could in prolonged interval if the link condition is not good.

4.8. General Query Suppression with no receiver

In IGMPv3 and MLDv2, General Query is multicast sent periodically and continuously without any limitations. It helps solicit the state of current valid member but has influence on all terminals, whether they are valid multicast receivers or not. When there is no receiver on the link, the transmission of the General Query is a waste of resources for both terminals and the router.

The IGMPv3/MLDv2 router could suppress its transmission of General Query if there is no valid multicast receiver on the link, e.g. in the following cases:

O If the last member reports its leave for a group. This could be judged by an explicit tracking router checking its membership database, or by a non explicit tracking router sending Group and Source Group Specific Queries;

O If the only member on a PTP link reports its leaving;

O If the router after retransmission of General Queries on startup fails to get any response from any member;

O If the router previously has valid members but fails to get any response from any member after several rounds of General Queries or Unicast Queries;

In these cases the router could make a decision that no member is on this link and totally stop its transmission of periodical General Queries. If afterwards there is valid multicast receiver joins a group, the router could resume the original cycle of transmission of General Queries. Because General Query has influences on all the terminals on the link, suppressing it when it is not needed is beneficial for both the link efficiency and terminal power saving.

4.9. Tuning Response Delay according to link type and status

IGMPv3 and MLDv2 use delayed response mechanism to spread Report messages from different hosts over a longer interval which can greatly reduce possibility of packet burstiness. This is implemented by the host responding to a Query in a specific time randomly chosen between 0 and [Maximum Response Delay]. The value of [Maximum Response Delay] parameter is determined by the router and is carried

in Query messages to inform the valid hosts to make the selection. A long delay will lessen the burstiness but will increase leave latency (the time between when the last listener stops listening to a source or multicast address and when the traffic stops flowing).

In order to avoid burstiness of MLD messages and reduce leave latency, explicit tracking with Group Specific Query eliminated is recommended to be used first to reduce leave latency. Then the Response Delay may be dynamically calculated based on the expected number of Reporters for each Query and link type and link status.

O If the expected number of Reporters is large and link condition is bad, the system administrator MUST choose the longer Maximum Response Delay; if the expected number of Reporters is small and the link condition is good, the administrator may choose the smaller Maximum response Delay. In this case, the IGMP/MLD packet burstiness can be reduced.

o Another case is if the link type is PTP which means the resource is dedicated for one receiver on each link, then the Maximum Response Delay can be chosen smaller, if the link type is shared medium link or P2MP, then the Maximum Response Delay can be configured larger.

The Maximum Response Delay can be configured by the administrator as mentioned above, or be calculated automatically by software tool implemented according to experiential model on different link modes. As the router arrives at a value appropriate for current link type and conditions, it will encode the value in Query messages to inform the host to make the response. The determination of the instant Maximum Response Delay value is out of this document's scope.

4.10. Triggering reports and queries quickly during handover

As a mobile terminal is moving from one network to another, if it is a multicast receiver from a group, its new access network should try to deliver the content to the receiver without disruption or performance deterioration. For the smooth switching between networks, the terminal's membership should be acquired as quickly as possible by the new access network.

For the access router, it could trigger a Query to the terminal as soon as it detects a new terminal on its link. This could be a General Query if the router does not know whether or not the terminal is a valid receiver or if the number of the entering terminals is not small. Or this Query could also be a unicast Query

for only a small quantity of terminals to prevent unnecessary action of other terminals in the switching area.

For the terminal, it could trigger a report if it is currently in the multicast reception state. This helps establish more quickly the membership states and enable faster multicast stream injection because active report from the host does not requires the router to wait for the query-response round in the passive reporting cases.

5. Security Considerations

They will be described in the later version of this draft.

6. Acknowledgement

The authors would like to thank Stig,Venaas, Gorry Fairhurst, Thomas C. Schmidt, Marshall Eubanks, Suresh Krishnan, J.William Atwood, WeeSan Lee, Imed Romdhani, Hitoshi Asaeda, Liu Yisong and Wei Yong for their valuable comments and suggestions on this document.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirement levels", RFC 2119, March 1997.

[RFC1112] Deering, S. "Host Extensions for IP Multicasting", RFC1112, August 1989.

[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.

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight IGMPv3 and MLDv2 Protocols", RFC5790, February 2010.

7.2. Informative Referencess

[REQUIRE] H. Liu, Q. Wu, H. Asaeda and TM. Eubanks, "Mobile and Wireless Multicast Requirements on IGMP/MLD Protocols", draft-liu-multimob-igmp-mld-mobility-req-03.txt, March 2010.

[ROBUST] A. Sen Mazumder, "Facilitating Robust Multicast Group Management", NOSSDAV'05, June 13-14, 2005, Stevenson, Washington, USA.

[ACK] Nikaein, N. and Bonnet, C. "Wireless multicasting in an IP environment" In Proceedings of the 5th International Workshop on Mobile Multimedia Communication MoMuc'98 (Berlin, Germany, Oct. 12-14). IEEE Computer Society Press, 1998.

[IGMP-ACK] H. Liu, Q, Wu, "Reliable IGMP and MLD Protocols in Wireless Environment", draft-liu-multimob-reliable-igmp-mld-00.txt, February 2010.

[ADAPTIVE] I. Romdhani, J. Munoz, H. Bettahar, and A. Bouabdallah, "Adaptive Multicast Membership Management for Mobile Multicast Receivers", IEEE, 2006.

[RFC5757] Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey", RFC 5757, February 2010.

Authors' Addresses

Qin Wu
Huawei Technologies Co., Ltd.
Site B, Floor 12, Huihong Mansion,No.91 Baixia Rd.
Nanjing, Jiangsu 21001
China
Phone: +86-25-84565892

EMail: sunseawq@huawei.com

Hui Liu
Huawei Technologies Co., Ltd.
Huawei Bld., No.3 Xinxi Rd.
Shang-Di Information Industry Base
Hai-Dian Distinct, Beijing 100085
China

EMail: Liuhui47967@huawei.com

MULTIMOB Group
INTERNET-DRAFT
Intended Status: Standards Track
Expires: April 2011

J.C. Zuniga
A. Rahman
InterDigital Communications, LLC
L.M. Contreras
C.J. Bernardos
Universidad Carlos III de Madrid
I. Soto
Universidad Politecnica de Madrid

October 25, 2010

Support Multicast Services Using Proxy Mobile IPv6
draft-zuniga-multimob-smspmip-04.txt

Abstract

The MULTIMOB group has specified a base solution to support IP multicasting in a PMIPv6 domain [I-D.draft-ietf-multimob-pmipv6-base-solution]. In this document, an enhancement is proposed to the base solution to use a dedicated multicast LMA as the topological anchor point for multicast traffic, while the MAG remains as an IGMP/MLD proxy. This enhancement provides benefits such as reducing multicast traffic replication and supporting different PMIPv6 deployments scenarios.

Table of Contents

1 Introduction

Proxy Mobile IPv6 [RFC5213] is a network-based approach to solving the IP mobility problem. In a Proxy Mobile IPv6 (PMIPv6) domain, the Mobile Access Gateway (MAG) behaves as a proxy mobility agent in the network and does the mobility management on behalf of the Mobile Node (MN). The Local Mobility Anchor (LMA) is the home agent for the MN and the topological anchor point. PMIPv6 was originally designed for unicast traffic.

The Internet Group Management Protocol (IGMPv3) [RFC3376] is used by IPv4 hosts to report their IP multicast group memberships to neighboring multicast routers. Multicast Listener Discovery (MLDv2) [RFC3810] is used in a similar way by IPv6 routers to discover the presence of IPv6 multicast hosts. Also, the IGMP/MLD proxy [RFC4605] allows an intermediate (edge) node to appear as a multicast router to downstream hosts, and as a host to upstream multicast routers. IGMP and MLD related protocols were not originally designed to address IP mobility of multicast listeners (i.e. IGMP and MLD protocols were originally designed for fixed networks).

The MULTIMOB group has specified a base solution to support IP multicast listener mobility in a PMIPv6 domain [I-D.draft-ietf-multimob-pmipv6-base-solution]. In this document, an enhancement is proposed to the base solution to use a dedicated multicast LMA (M-LMA) as the topological anchor point for multicast traffic, while the MAG remains as an IGMP/MLD proxy. This enhancement allows different PMIPv6 deployment scenarios. It also eliminates the so called "Tunnel Convergence problem" where the MAG may receive the same multicast packet from several LMAs. There are no impacts to the MN to support multicast listener mobility from this document.

2 Conventions and 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 RFC 2119 [RFC2119].

This document uses the terminology defined in [RFC5213], [RFC3775], and [RFC3810]. Specifically, the definition of PMIPv6 domain is reused from [RFC5213] and reproduced here for completeness.

- Proxy Mobile IPv6 Domain (PMIPv6-Domain): Proxy Mobile IPv6 domain refers to the network where the mobility management of a mobile node is handled using the Proxy Mobile IPv6 protocol as defined in this specification. The Proxy Mobile IPv6 domain includes local mobility anchors and mobile access gateways between

which security associations can be set up and authorization for sending Proxy Binding Updates on behalf of the mobile nodes can be ensured.

Additionally, some definitions are introduced, as follows.

- U-LMA or Unicast-LMA: LMA entity dedicated to unicast service exclusively.

- M-LMA or Multicast-LMA: LMA entity dedicated to multicast service exclusively.

- H-LMA or Hybrid-LMA: LMA entity dedicated to both unicast and multicast services.

3 Solution

A PMIPv6 domain may handle data from both unicast and multicast sources. A dedicated multicast LMA can be used to serve as the mobility anchor for multicast traffic. Unicast traffic will go normally to the other LMAs in the PMIPv6 domain. This section describes how the multicast LMA works in scenarios of MN attachment and multicast mobility. We first concentrate on the case of both LMAs (multicast and unicast) defining a unique PMIPv6 domain, and then different deployment scenarios are presented.

3.1 Architecture

Figure 1 shows an example of a PMIPv6 domain supporting multicast mobility. LMA1 is dedicated to unicast traffic, and LMA2 is dedicated to multicast traffic. The multicast traffic LMA (LMA2) can be considered to be a form of upstream multicast router with tunnel interfaces allowing remote subscription for the MNs. Note that there can be multiple LMAs for unicast traffic (not shown in Figure 1) in a given PMIPv6 domain. Similarly, more than one multicast dedicated LMA can be deployed by the operator (not shown in Figure 1).

Also in this architecture, all MAGs that are connected to the multicast LMA must support the MLD proxy [RFC4605] function. Specifically in Figure 1, each of the MAG1-LMA2 and MAG2-LMA2 tunnel interfaces defines an MLD proxy domain. The MNs are considered to be on the downstream interface of the MLD proxy (in the MAG), and LMA2 is considered to be on the upstream interface (of the MAG) as per [RFC4605]. Note that MAG could also be an IGMP proxy. For brevity this document will refer primarily to MLD proxy, but all references to "MLD proxy" should be understood to also include "IGMP/MLD proxy" functionality.

As shown in Figure 1, MAG1 may connect to both unicast and multicast LMAs. Thus, a given MN may simultaneously receive both unicast and multicast traffic. In Figure 1, MN1 and MN2 receive unicast traffic, multicast traffic, or both, whereas MN3 receives multicast traffic only.

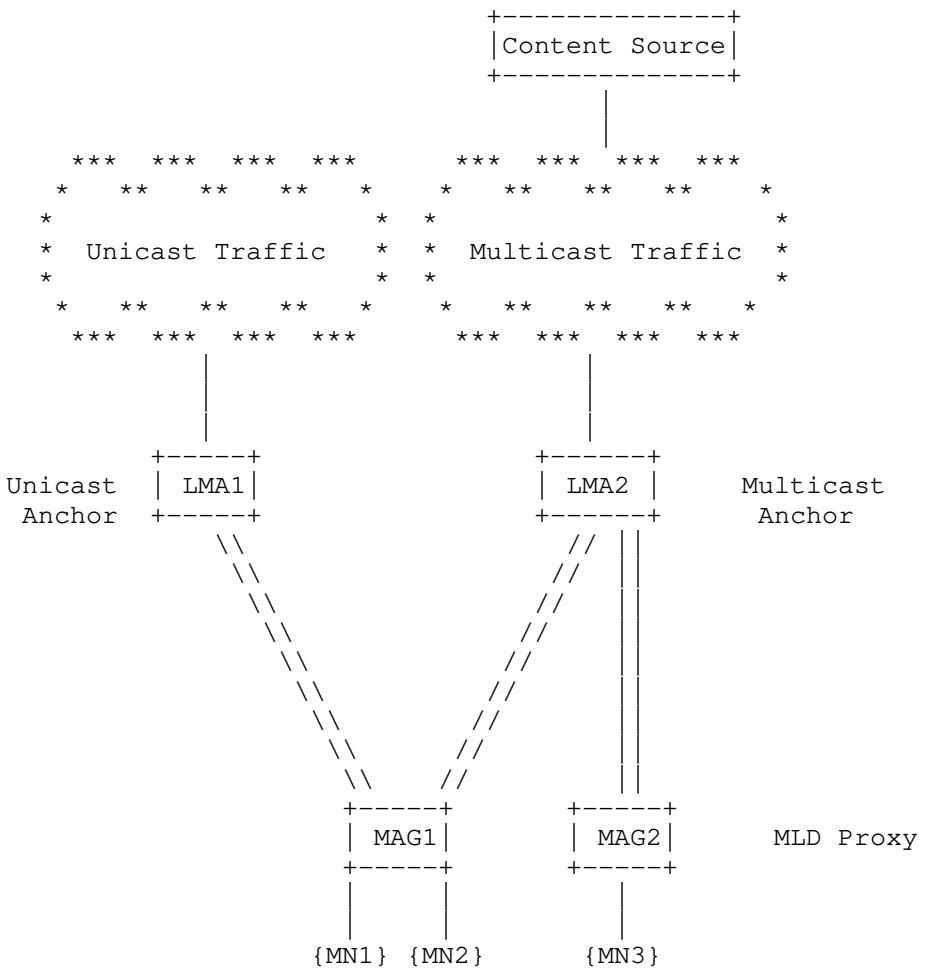

Figure 1. Architecture of Dedicated LMA as Multicast Anchor

3.2 Deployment Scenarios

From the network architecture point of view, there are a several options when considering the dedicated multicast LMA (M-LMA) approach. These options can be distinguished in terms of the number

of unicast and multicast LMAs present in a PMIPv6 domain and the service relationship that a set of MN get from them, in the form of a "U-LMA : M-LMA" ratio. According to that, it is possible to differentiate the following approaches:

- A set of MNs is served in a PMIPv6 domain by two LMAs, one for multicast service, the other one for unicast, in such a way that the ratio is 1:1.

- A set of MNs is served in a PMIPv6 domain by several LMAs, one for multicast service, while the rest for unicast, in such a way that the ratio is N:1.

- A set of MNs is served in a PMIPv6 domain by several LMAs, one for unicast, while the rest are devoted to multicast service, in such a way that the ratio is 1:N.

Scenarios with an N:M ratio are considered to be a combination of the previous ones.

3.2.1 PMIPv6 domain with ratio 1:1

This approach basically refers to the architecture presented in figure 1. Within this approach, a common set of MNs is served by a couple of LMAs, one for unicast and the other one for multicast. All the MNs of the set are served by these two LMAs as they move in the PMIPv6 domain.

3.2.2 PMIPv6 domain with ratio N:1

This approach basically refers to the situation where a common set of MNs is served by a unique LMA for multicast service, but simultaneously there are subsets from that group of MNs which are served by distinct LMAs for unicast service as they move in the PMIPv6 domain. Each particular MN association with the LMAs (unicast and multicast) remains always the same as it moves in the PMIPv6 domain.

Figure 2 shows the scenario here described.

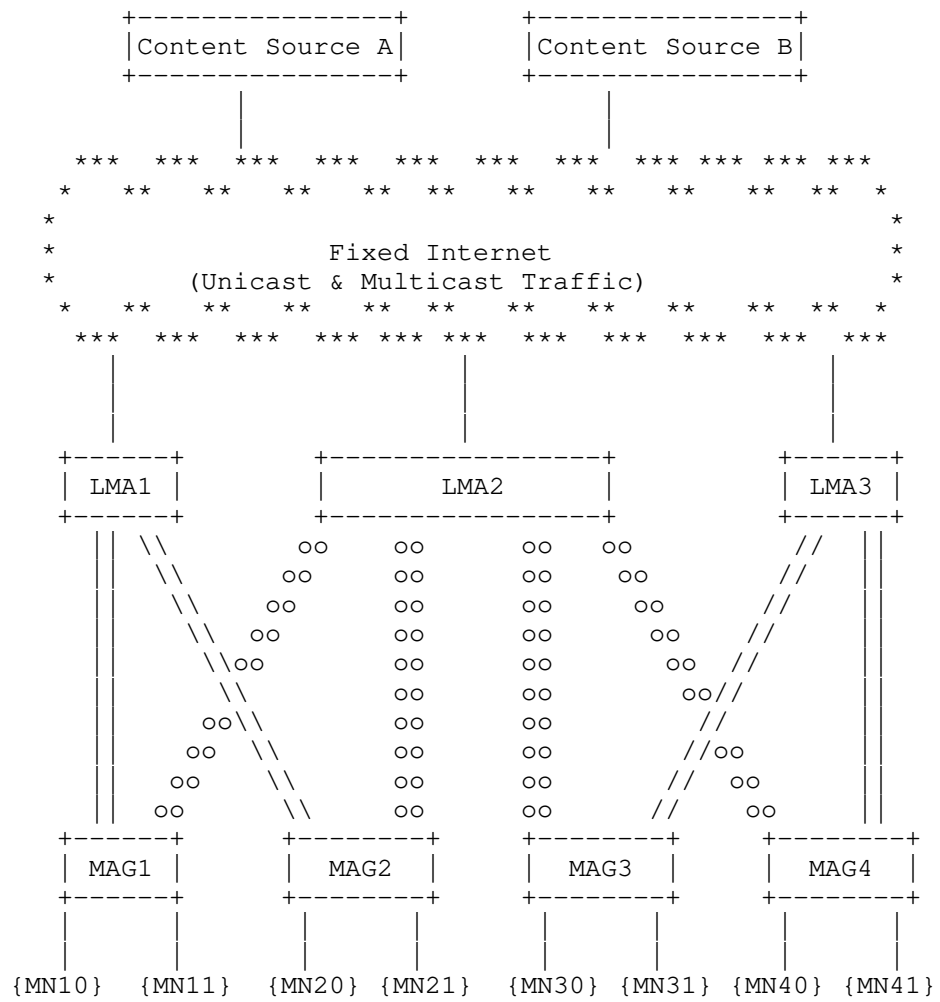

Figure 2. PMIPv6 domain with ratio N:1

The figure 2 proposes an architecture where there are two LMAs, LMA1 and LMA3, acting as U-LMAs, while there is another one, the LMA2, working as dedicated M-LMA. The tunnels among MAGs and LMAs represented by lines ("||") indicate a tunnel transporting unicast traffic, while the tunnels depicted with circles ("o") show a tunnel transporting multicast traffic.

In the figure it can be observed that all the MNs are served by LMA2 for the incoming multicast traffic from sources A or B. However, there are different subsets regarding unicast traffic which maintain distinct associations within the PMIPv6 domain. For instance, the

subset formed by MN10, MN11, MN20 and MN21 is served by LMA1 for unicast, and the rest of MNs are being served by LMA3. For the scenario described above, the association between each MN and the corresponding U-LMA and M-LMA is permanently maintained.

3.2.3 PMIPv6 domain with ratio 1:N

This approach is related to an scenario where a common group of MNs is served by a unique LMA for unicast service, but simultaneously there are subsets from that group of MNs which are served by distinct LMAs for multicast service as they move in the PMIPv6 domain. Each particular MN association with the LMAs (unicast and multicast) remains always the same as it moves in the PMIPv6 domain.

Figure 3 shows the scenario here described.

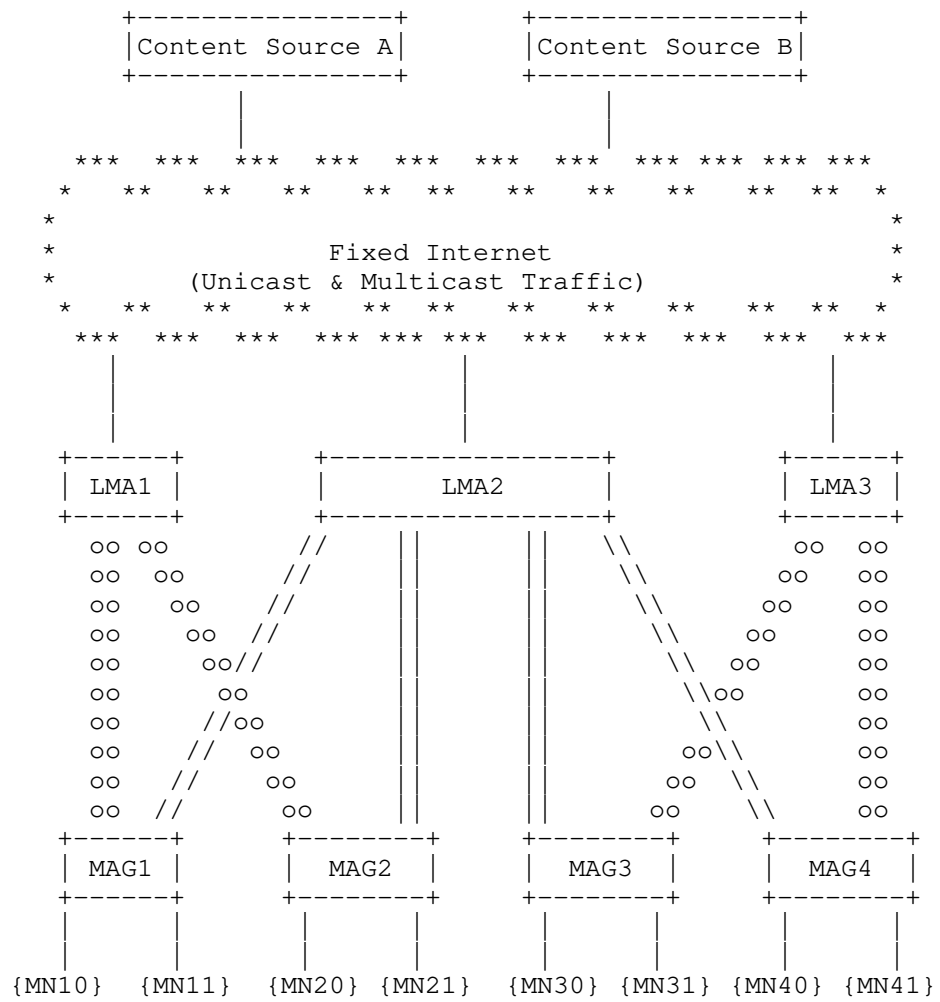

Figure 3. PMIPv6 domain with ratio 1:N

The figure 3 proposes an architecture where the LMA2 is the unique U-LMA for a certain group of MNs, while there are two others LMAs, LMA1 and LMA3, act as M-LMAs for different subsets of MNs of the same group. Each M-LMA could be devoted to carry on a different content (for instance, LMA1 for source A and LMA3 for source B) or not. Looking at the picture, the subset formed by MN10, MN11, MN20 and MN21 is served by LMA1 for multicast. The rest of MNs are being served by LMA3 also for multicast. Finally, all of them are served by LMA2 for unicast. For the scenario described above, the association between each MN and the corresponding U-LMA and M-LMA is permanently maintained.

3.2.4 PMIPv6 domain with H-LMA

The H-LMA is defined as an LMA which simultaneously transports unicast and multicast service. In the context of the dedicated M-LMA solution, an H-LMA can play the role of M-LMA for an entire group of MNs in a PMIPv6 domain, while acting simultaneously as U-LMA for a subset of them. The figure 4 adapts the PMIPv6 domain with ratio 1:N scenario of figure 3 to the case where LMA2 is an H-LMA, which serves multicast traffic to all the MNs in the picture, and simultaneously, it is able to serve unicast traffic to the subset formed by MN30, MN40 and MN41.

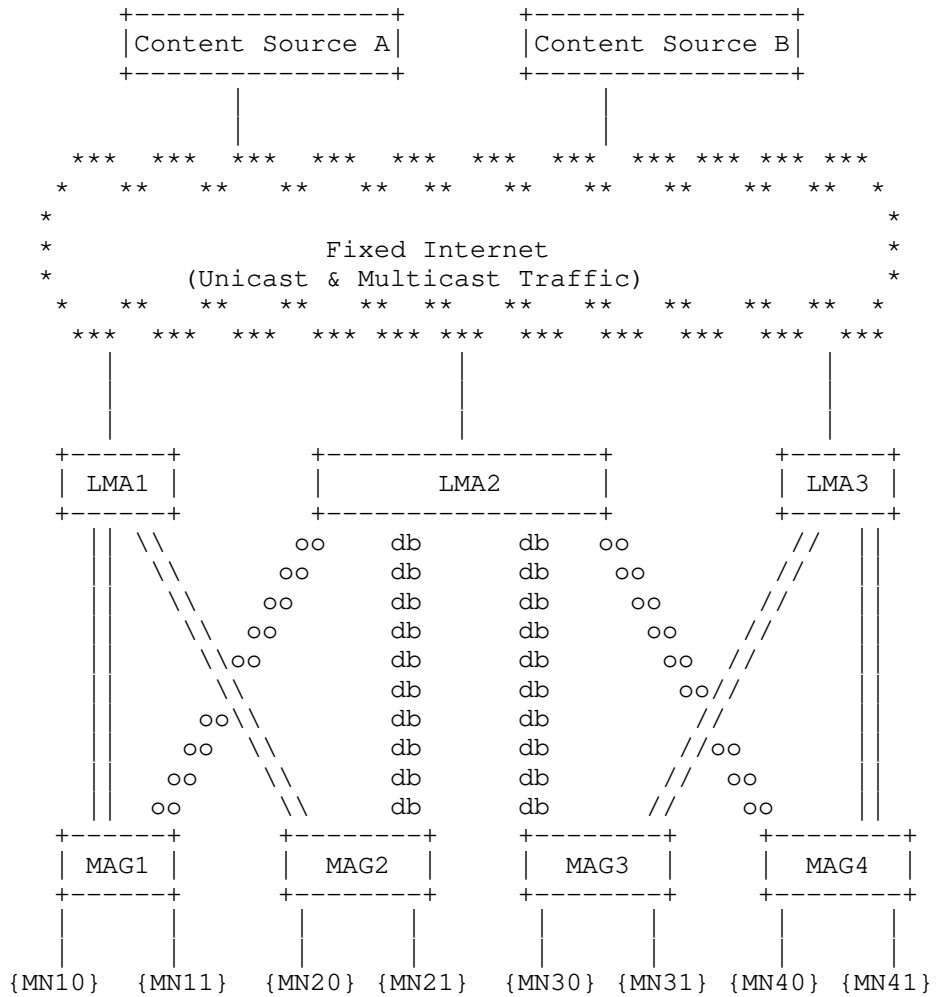

Figure 4. PMIPv6 domain with H-LMA

The figure 4 presents a PMIPv6 network where there are two pure unicast LMAs, LMA1 and LMA3, and a hybrid LMA, the LMA2. The LMA2 is a dedicated M-LMA from the perspective of MAG1 and MAG4. The tunnels among MAGs and LMAs represented by lines ("||") indicate a tunnel transporting exclusively unicast traffic, the tunnels depicted with circles ("o") show a tunnel transporting exclusively multicast traffic, and the tunnels with mixed lines and circles ("db") describe a tunnel transporting both types of traffic simultaneously.

All of the MNs in the figure receive the multicast traffic from LMA2, but it is possible to distinguish three subsets from the unicast service perspective. The first subset is the one formed by MN10, MN11 and MN 20, which receives unicast traffic from LMA1. A second subset is the one formed by MN21 and MN30, which receives unicast traffic from LMA2. And finally, a third subset is built on MN31, MN40 and MN41, which receives unicast traffic from LMA3. For the scenario described above, the association between each MN and the corresponding U-LMA and M-LMA is permanently maintained.

3.3 Multicast Establishment

Figure 5 shows the procedure when MN1 attaches to MAG1, and establishes associations with LMA1 (unicast) and LMA2 (multicast).

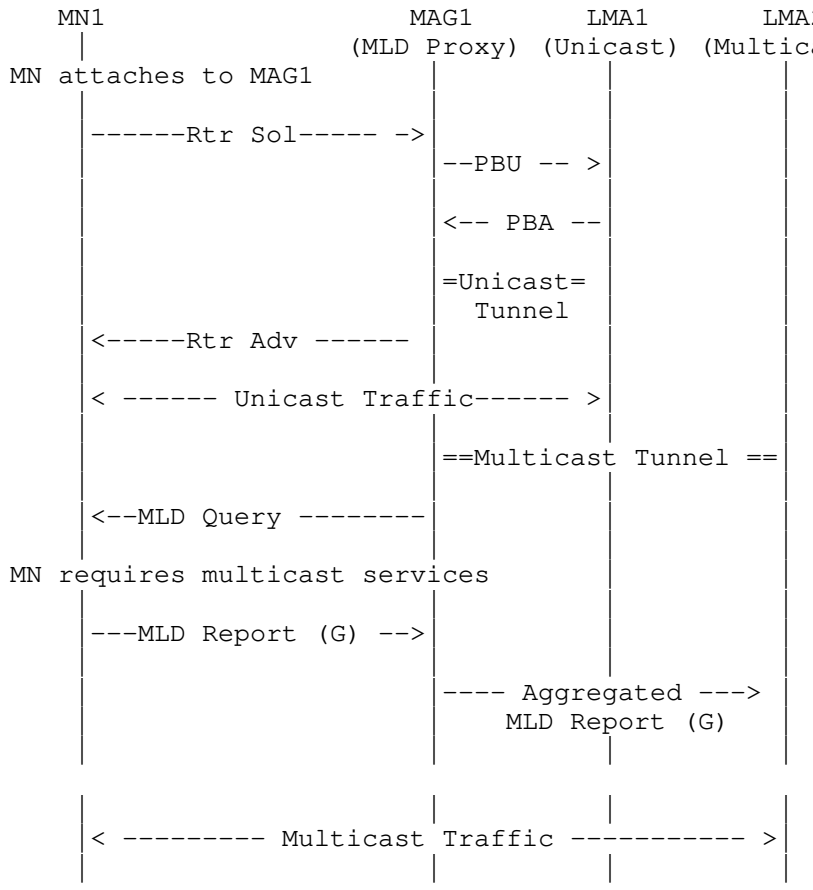

Figure 5. MN Attachment and Multicast Service Establishment

In Figure 5, MAG1 first establishes the PMIPv6 tunnel with LMA1 for unicast traffic as defined in [RFC5213] after being triggered by the Router Solicitation message from MN1. Unicast traffic will then flow between MN1 and LMA1.

For multicast traffic, a multicast tunnel may have been pre-configured between MAG1 and the multicast LMA (LMA2). Or the multicast tunnel may be dynamically established when the first MN appears at the MAG.

MN1 sends the MLD report message (when required by its upper layer applications) as defined in [RFC3810] in response to an MLD Query from MAG1. MAG1 acting as a MLD Proxy as defined in [RFC4605] will then send an Aggregated MLD Report to the multicast anchor, LMA2

(assuming that this is a new multicast group which MAG1 had not previously subscribed to). Multicast traffic will then flow from LMA2 towards MN1.

3.4 Multicast Mobility

Figure 6 illustrates the mobility scenario for multicast traffic. Specifically, MN2 with ongoing multicast subscription moves from MAG1 to MAG2. Note that, for simplicity, in this scenario MAG2 is connected only to LMA2 (multicast) and does not receive unicast traffic. Of course, if it was desired to support unicast traffic, the architecture will easily allow MAG2 to also connect to LMA1 to support unicast traffic.

After MN2 mobility, MAG2 acting in its role of MLD proxy will send an MLD Query to the newly observed MN on its downlink. Assuming that the subsequent MLD Report from MN2 requests membership of a new multicast group (from MAG2's point of view), this will then result in an Aggregated MLD Report being sent to LMA2 from MAG2. This message will be sent through a pre-established (or dynamically established) multicast tunnel between MAG2 and LMA2.

When MN2 detaches, MAG1 may keep the multicast tunnel with the multicast LMA2 if there are still other MNs using the multicast tunnel. Even if there are no MNs currently on the multicast tunnel, MAG1 may decide to keep the multicast tunnel for potential future use.

As discussed above, existing MLD (and Proxy MLD) signaling will handle a large part of the multicast mobility management for the MN.

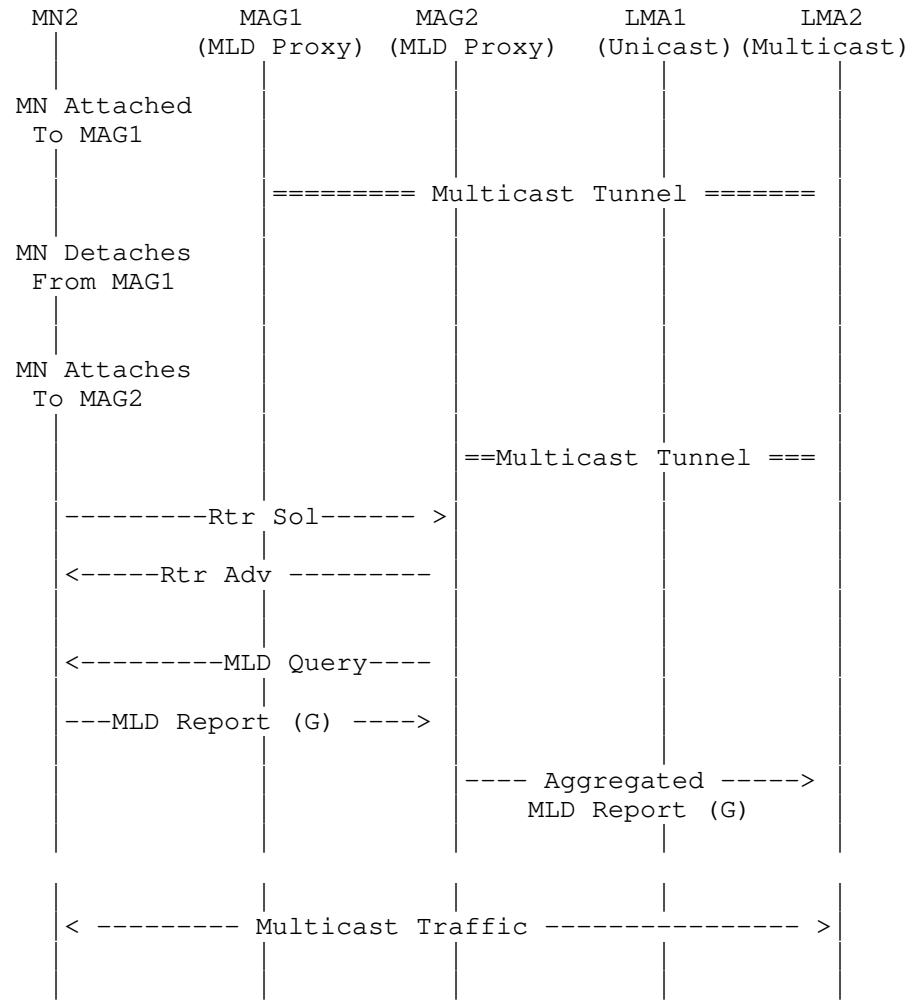

Figure 6. Multicast Mobility Signaling

3.5 PMIPv6 enhancements

This section describes the enhancements to the Proxy Mobile IPv6 [RFC5213] protocol required to support the M-LMA architecture.

3.5.1 New Binding Update List in MAG

The Binding Update List in the MAG must be updated to be able to

handle the fact that more than one LMA (i.e. U-LMA and M-LMA) may be serving the mobile node.

3.5.2 Policy Profile Information with Multicast Parameters

A given mobile node's policy profile information must be updated to be able to store the IPv6 addresses of both the U-LMA and M-LMA.

3.5.3 MAG to M-LMA attach requirements

The MAG procedures must be updated to be able to handle simultaneous attach for a given mobile node to both the U-LMA and M-LMA. For example, packets coming from a given mobile node must be screened to determine if it should be sent to the U-LMA or to the M-LMA.

3.6 Advantages

An advantage of the proposed dedicated multicast LMA (M-LMA) architecture is that it allows a PMIPv6 domain to closely follow a simple multicast tree topology for Proxy MLD forwarding (cf., sections 1.1 and 1.2 of [RFC4605]). In contrast, the combined unicast/multicast LMA as proposed in [I-D.draft-ietf-multimob-pmipv6-base-solution] will be a more complex set of trees.

Another advantage of the proposed dedicated multicast solution is that it allows a gradual network upgrade of a PMIPv6 domain to support multicast functionality. This is because the operator does not have to upgrade all the LMAs in the network to support multicast functionality. Only certain LMAs, dedicated to multicast support, will have to be upgraded to support the new multicast functionality. Also, multiple deployment scenarios are supported as required by the operator for expected traffic distributions.

A final advantage is that a dedicated multicast LMA minimizes replication of multicast packets (the Tunnel Convergence problem), in certain scenarios, compared to [I-D.draft-ietf-multimob-pmipv6-base-solution]. Figures 7 and 8 illustrate this point visually. For this simple scenario, it can be observed that the dedicated multicast LMA topology (Figure 7) generates 6 packets for one input multicast packet. In comparison, the combined unicast/multicast LMA topology (Figure 8) generates 8 packets for one input multicast packet.

In general, it can be seen that the extra multiplication of packets in the combined unicast/multicast LMA topology will be proportional to the number of LMAs, and the number of MNs (in a given MAG)

associated to different LMAs, for a given multicast group. The packet multiplication problem aggravates as more MNs associated to different LMAs receive the same multicast traffic when attached to the same MAG. Hence, the dedicated multicast architecture significantly decreases the network capacity requirements in this scenario.

(Note that in Figure 7, it is assumed that MN1 and MN2 are associated with MAG1-LMA1, and MN3 is associated with MAG2-LMA2 for multicast traffic. In Figure 8, it is assumed that MN1 is associated with MAG1-LMA1, MN2 is associated with MAG1-LMA2, and MN3 is associated with MAG2-LMA2 for multicast traffic. In both Figures 7 and 8, it is assumed that the packets are transmitted point to point on the last hop wireless link.)

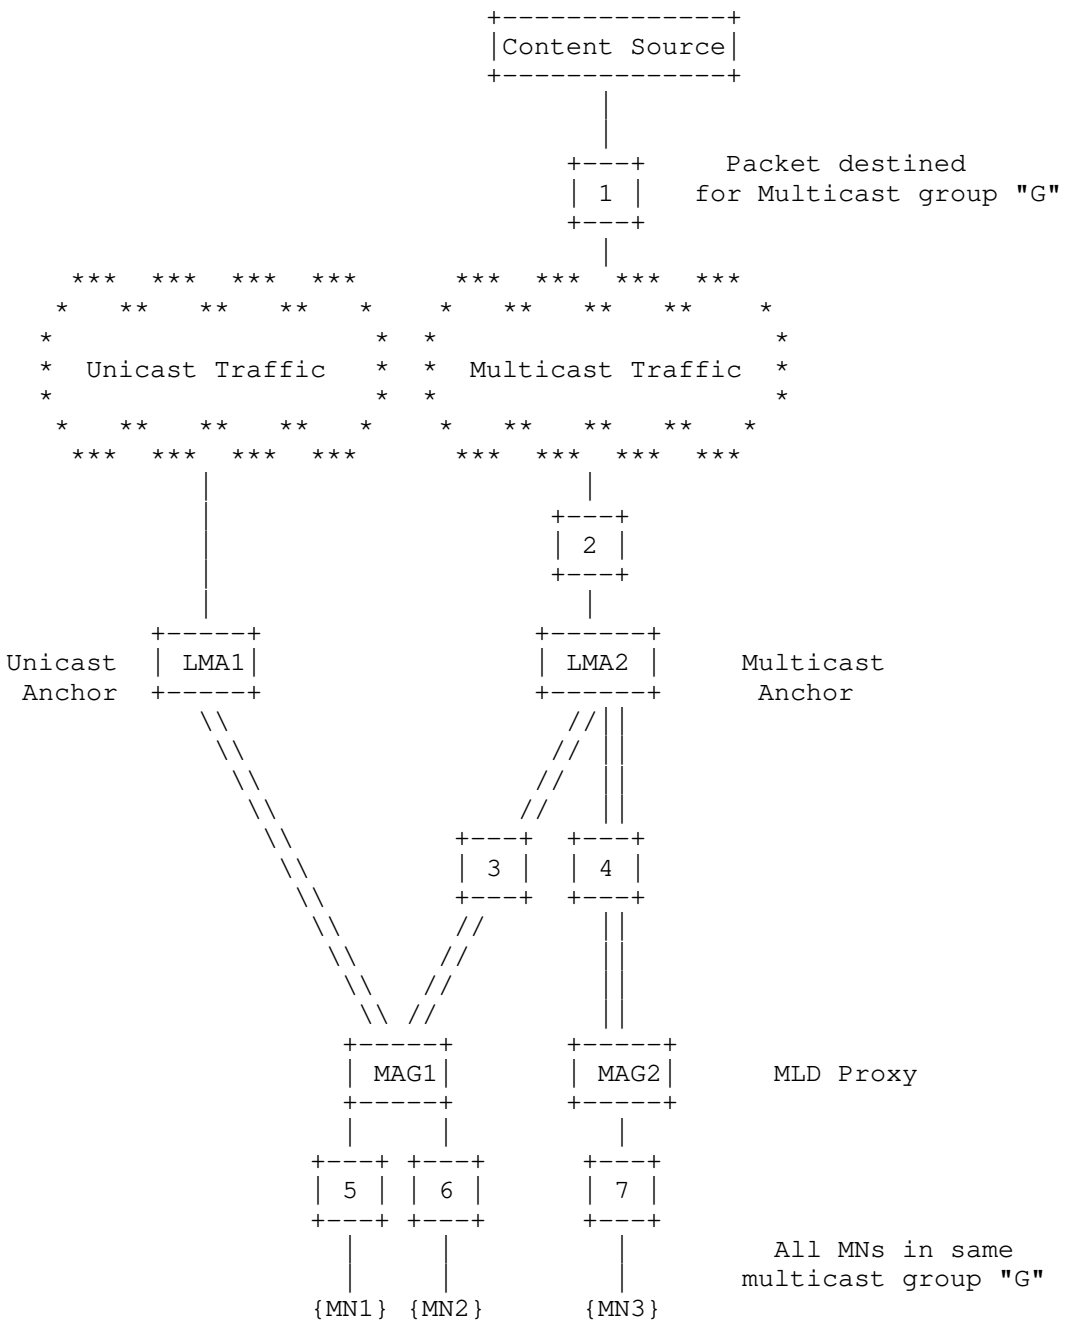

Figure 7. Packet Flow in a Dedicated Multicast LMA

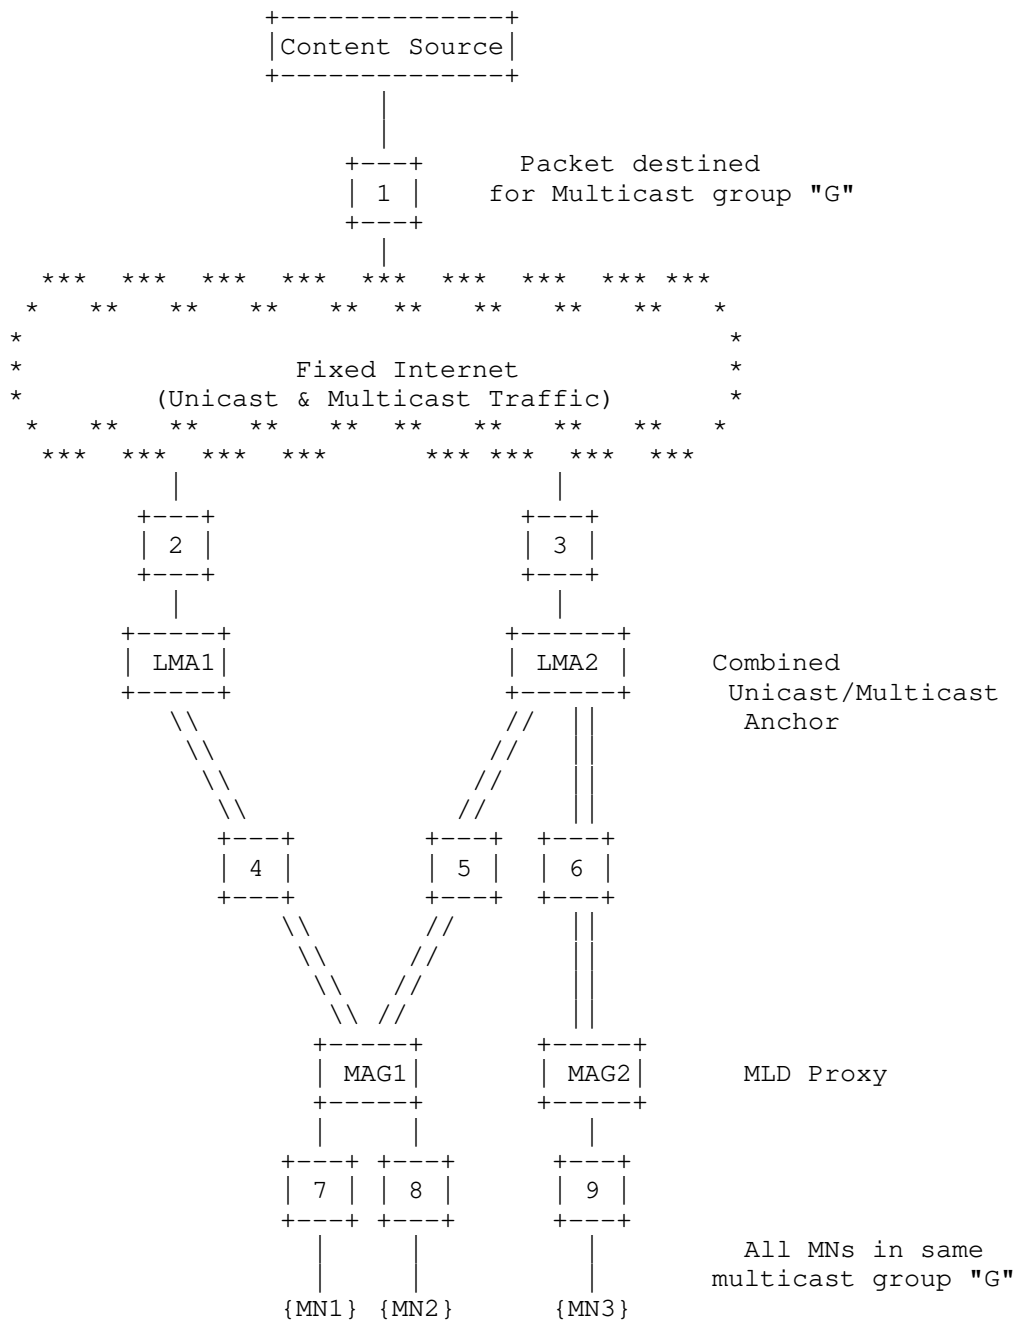

Figure 8. Packet Flow in a Combined Unicast/Multicast LMA

4 Security Considerations

This draft discusses the operations of existing protocols without modifications. It does not introduce new security threats beyond the current security considerations of PMIPv6 [RFC5213], MLD [RFC3810], IGMP [RFC3376] and IGMP/MLD Proxying [RFC4605].

5 IANA Considerations

This document makes no request of IANA.

6 References

6.1 Normative References

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

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

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

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

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

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP)/ Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

6.2 Informative References

[I-D.draft-ietf-multimob-pmipv6-base-solution] Schmidt, T.C., Waehlisch, M., and S.Krishnan, "Base Deployment for Multicast Listener Support in PMIPv6 Domains", draft-ietf-multimob-pmipv6-base-solution-05 (work in progress), July 28, 2010.

Author's Addresses

Juan Carlos Zuniga
InterDigital Communications, LLC
Email: JuanCarlos.Zuniga@InterDigital.com

Akbar Rahman
InterDigital Communications, LLC
Email: Akbar.Rahman@InterDigital.com

Luis M. Contreras
Universidad Carlos III de Madrid
Email: luisc@it.uc3m.es

Carlos J. Bernardos
Universidad Carlos III de Madrid
Email: cjbc@it.uc3m.es

Ignacio Soto
Universidad Politecnica de Madrid
Email: isoto@dit.upm.es