

Internet Engineering Task Force (IETF)
Request for Comments: 7028
Category: Experimental
ISSN: 2070-1721

JC. Zuniga
InterDigital Communications, LLC
LM. Contreras
Telefonica I+D
CJ. Bernardos
UC3M
S. Jeon
Instituto de Telecomunicacoes
Y. Kim
Soongsil University
September 2013

Multicast Mobility Routing Optimizations for Proxy Mobile IPv6

Abstract

This document proposes some experimental enhancements to the base solution to support IP multicasting in a Proxy Mobile IPv6 (PMIPv6) domain. These enhancements include the use of a multicast tree mobility anchor as the topological anchor point for multicast traffic, as well as a direct routing option where the Mobile Access Gateway can provide access to multicast content in the local network. The goal of these enhancements is to provide benefits such as reducing multicast traffic replication and supporting different PMIPv6 deployment scenarios.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7028>.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
2. Terminology	4
3. Overview	5
3.1. MTMA/Direct Routing Mode Selection	5
3.2. Multicast Tree Mobility Anchor (Subscription via MTMA)	5
3.3. Direct Routing (Subscription via Direct Routing)	7
4. Mobile Access Gateway Operation	9
4.1. Extensions to Binding Update List Data Structure	9
4.2. MAG as MLD Proxy	9
4.2.1. MTMA Mode (Subscription via MTMA)	9
4.2.2. Direct Routing Mode (Subscription via Direct Routing)	11
5. Local Mobility Anchor Operation	14
5.1. Dynamic IP Multicast Selector Option	14
5.1.1. Option Application Rules	14
5.1.2. Option Format	14
6. Multicast Tree Mobility Anchor Operation	16
6.1. Conceptual Data Structures	17
7. Mobile Node Operation	17
8. IPv4 Support	17
9. IANA Considerations	18
10. Security Considerations	18
11. Contributors	19
12. References	20
12.1. Normative References	20
12.2. Informative References	21
Appendix A. MTMA Deployment Use Cases	22
A.1. PMIPv6 Domain with Ratio 1:1	22
A.2. PMIPv6 Domain with Ratio N:1	22
A.3. PMIPv6 Domain with Ratio 1:N	24
A.4. PMIPv6 Domain with H-LMA	26

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 performs 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. However, a PMIPv6 domain may handle data from both unicast and multicast sources.

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 Version 2 (MLDv2) [RFC3810] is used in a similar way by IPv6 routers to discover the presence of IPv6 multicast hosts. Also, the IGMP/MLD proxy specification [RFC4605] allows an intermediate (i.e., edge) node to appear as a multicast router to downstream hosts and as a host to upstream multicast routers. IGMP- and MLD-related protocols however were not originally designed to address the IP mobility of multicast listeners (i.e., IGMP and MLD protocols were originally designed for fixed networks).

A base solution to support both IPv4 and IPv6 multicast listener mobility in a PMIPv6 domain is specified in [RFC6224], which describes deployment options without modifying mobility and multicast protocol standards. PMIPv6 allows a mobile access gateway to establish multiple PMIPv6 tunnels with different local mobility anchors, e.g., up to one per mobile node. In the presence of multicast traffic, multiple instances of the same traffic can converge to the same MAG. Hence, when IP multicasting is applied into PMIPv6, it may lead to redundant traffic at a MAG. This is the tunnel convergence problem.

In order to address this issue, this document proposes an experimental solution, consisting of two complementary enhancements: multicast anchor and direct routing. The first enhancement makes use of a Multicast Tree Mobility Anchor (MTMA) as the topological anchor point for remotely delivering multicast traffic, while the second enhancement uses direct routing taking advantage of local multicast source availability, allowing a mobile access gateway to connect directly to a multicast router for simple access to local content. Neither of the two schemes has any impact on the mobile node to support IPv4 and IPv6 multicast listener mobility, nor on the wider Internet, as they only affect the PMIPv6 domains where they are deployed. Although references to "MLD proxy" are used in the document, it should be understood to also include "IGMP/MLD proxy" functionality (see Section 8 for details). The status of this

proposal is Experimental. The status of this proposal may be reconsidered in the future, once more implementation feedback and deployment experience is gathered, reporting on the performance of the two proposed schemes as well as operational feedback on scheme selection.

2. Terminology

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

This document uses the terminology defined in [RFC5213], [RFC6275], 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 [RFC5213]. 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.

In this document we refine the definition from the point of view of the kind of traffic served to the MN in the following way:

PMIPv6 unicast domain: PMIPv6 unicast domain refers to the network covered by one LMA for unicast service. This service supports mobility as the MN moves from one MAG to another one, both associated with the same LMA regarding the MN unicast traffic.

PMIPv6 multicast domain: PMIPv6 multicast domain refers to the network covered by one network element named MTMA (defined below) for multicast service in such a way that an MN using that service is not aware of mobility as it moves from one MAG to another.

From the definitions above, it can be stated that a PMIPv6 domain can have several PMIPv6 unicast domains and PMIPv6 multicast domains. Additionally, some other definitions are introduced, as follows.

MTMA or multicast tree mobility anchor: An entity working as topological anchor point for multicast traffic. It manages the multicast groups subscribed by all (or a subset of) the MAGs in a PMIPv6 multicast domain, on behalf of the MNs attached to them. Hence, an MTMA performs the functions of either a designated multicast router or an MLD proxy.

H-LMA or Hybrid-LMA: An entity that is dedicated to both unicast and multicast services and able to work as both LMA and MTMA simultaneously.

Direct routing: This scheme uses the native multicast infrastructure for retrieving multicast data. For an operator having its own local content, this technique also includes the case where the content source is directly connected to the MAG.

Subscription via MTMA: Multicast subscription mode in which the content is retrieved from the remote (or home) MTMA.

Subscription via direct routing: Multicast subscription mode in which the content is retrieved using direct routing from the local domain.

3. Overview

3.1. MTMA/Direct Routing Mode Selection

This specification describes two complementary operational modes that can be used to deliver multicast traffic in a PMIPv6 domain: multicast tree mobility anchor and direct routing. There are different approaches that can be followed to perform this operational mode selection, depending on the operator's preferences and PMIPv6 deployment characteristics. For example, the mode can be manually configured at the mobile access gateway, according to the multicast tree deployment in the PMIPv6 domain, following operator's configuration of the multicast distribution on it. Another option is the use of dynamic policies, conveyed in the PBU (Proxy Binding Update) / PBA (Proxy Binding Acknowledgement) signaling using the Dynamic IP Multicast Selector option described in Section 5.1. Next, each of the two operational modes is introduced.

3.2. Multicast Tree Mobility Anchor (Subscription via MTMA)

A multicast tree mobility anchor is used to serve as the mobility anchor for multicast traffic. The MTMA is either a designated multicast router or an MLD proxy. Typically, the MTMA will be used to get access to remote multicast content.

The multicast tree mobility anchor connects to the mobile access gateway, as described in [RFC6224], and it can reuse native PMIPv6 features such as tunnel establishment and security [RFC5213], heartbeat [RFC5847], etc. Unicast traffic will go normally to the local mobility anchors in the PMIPv6 domain as described in [RFC5213]. A MAG connecting to the MTMA acts as an MLD proxy.

This section describes how the MTMA works in scenarios of MN attachment and multicast mobility. It concentrates on the case of both LMA and MTMA defining a unique PMIPv6 domain. Some other deployment scenarios are presented in Appendix A.

Figure 1 shows an example of a PMIPv6 domain supporting multicast mobility. The local mobility anchor is dedicated to unicast traffic, and the multicast tree mobility anchor is dedicated to multicast traffic. The MTMA can be considered to be a form of upstream multicast router with tunnel interfaces allowing subscription via MTMA for the MNs.

As shown in Figure 1, MAG1 may connect to both unicast (LMA) and multicast (MTMA) entities. 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.

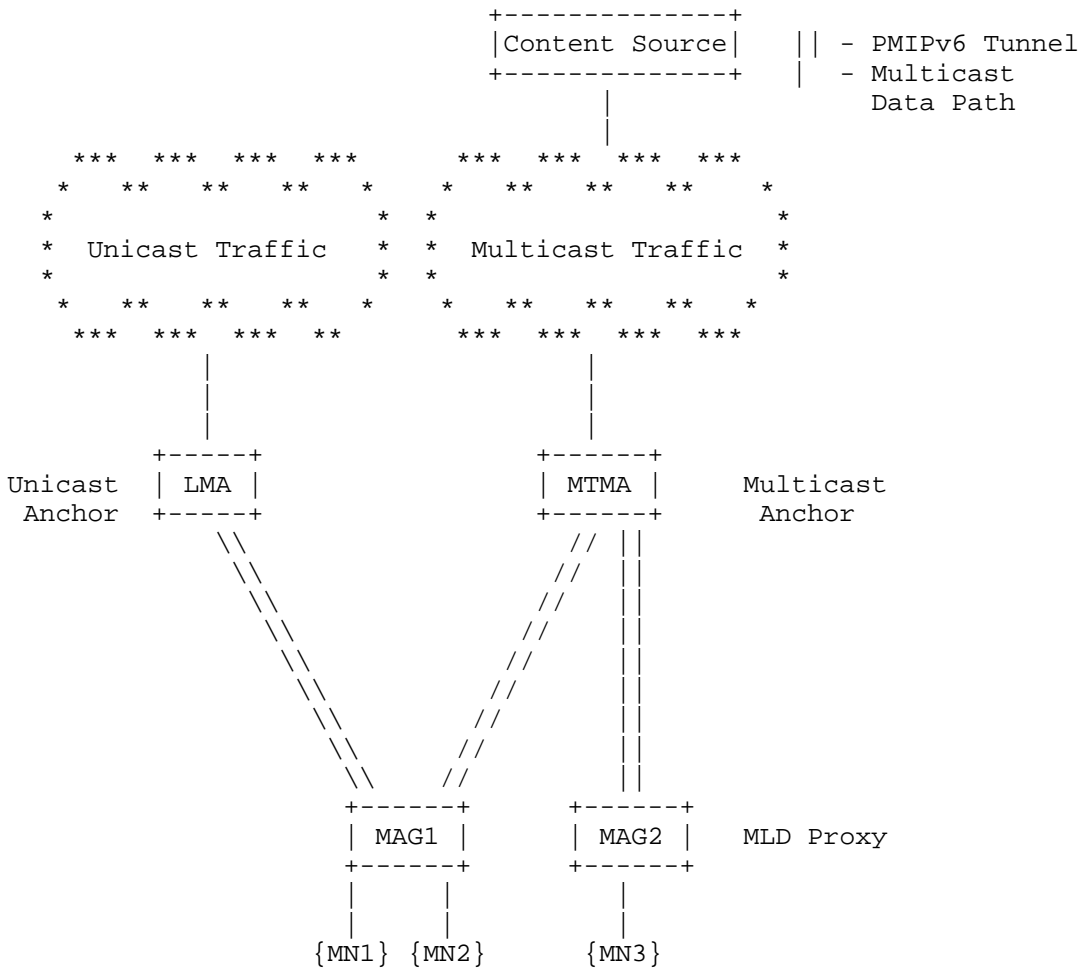


Figure 1: Architecture of Multicast Tree Mobility Anchor (MTMA)

3.3. Direct Routing (Subscription via Direct Routing)

Direct routing uses a native multicast infrastructure, allowing a mobile access gateway to directly connect to a multicast router (as next hop) in the PMIPv6 domain. A MAG acts as an MLD proxy.

The main purpose of direct routing is to provide optimal connectivity for local content. As a consequence, it replaces the MTMA of the channel management and data delivery of locally available content. Unicast traffic will go as normally to the LMAs in the PMIPv6 domain.

This section describes how the direct routing works in scenarios of MN attachment and multicast mobility.

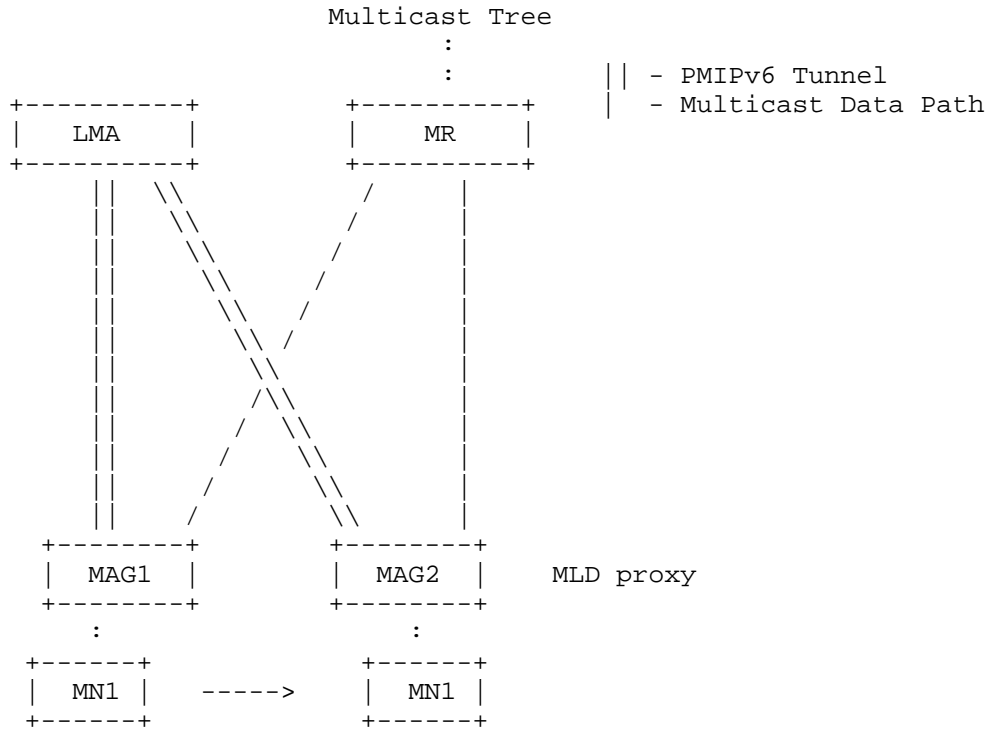


Figure 2: Architecture for Direct-Routing-Based PMIPv6 Multicasting

Figure 2 shows the architecture for the local routing case using native multicasting infrastructure [PMIPv6-REQ].

The local mobility anchor is dedicated to unicast traffic, and the multicast traffic is obtained from an upstream multicast router present in the PMIPv6 domain. Note that there can be multiple LMAs for unicast traffic (not shown in Figure 1 for simplicity) in a given PMIPv6 domain.

As shown in Figure 2, a mobile access gateway may connect to both unicast (LMA) and multicast routers (MRs). Thus, a given mobile node may simultaneously receive both unicast and multicast traffic.

As seen in Figure 2, each MAG has a direct connection (i.e., not using the PMIPv6 tunnel interface) with a multicast router. Depending on the multicast support on the visited network, different schemas can be used to provide this direct connection between the

MAGs and the multicast router(s), e.g., being connected to the same shared link or using a tunneling approach, such as Generic Routing Encapsulation (GRE) tunnels [RFC2784] or Automatic Multicast Tunneling (AMT) [AUTO]. To facilitate IGMP/MLD signaling and multicast traffic forwarding, an MLD proxy function defined in [RFC4605] SHOULD be implemented in the MAG. There SHOULD be direct connectivity between the MAG and the local multicast router (or additional MLD proxy).

4. Mobile Access Gateway Operation

This section describes the operation of the mobile access gateway, considering that the MAG incorporates MLD proxy functions as per [RFC4605].

4.1. Extensions to Binding Update List Data Structure

A Binding Update List (BUL) at the MAG, like the one specified in [RFC5213], MUST be maintained to handle the relationship between the serving entities (e.g., MTMA and LMA) and the mobile nodes for both unicast and multicast traffic.

4.2. MAG as MLD Proxy

4.2.1. MTMA Mode (Subscription via MTMA)

In case of subscription via MTMA, all MAGs that are connected to the MTMA must support the MLD proxy function [RFC4605]. Specifically in Figure 1, each of the MAG1-MTMA and MAG2-MTMA tunnel interfaces define an MLD proxy domain. The mobile nodes are considered to be on the downstream interface of the MLD proxy (of the MAG), and the MTMA is considered to be on the upstream interface (of the MAG) as per [RFC4605]. Note that the mobile access gateway could also be an IGMP proxy.

Figure 3 shows the procedure when MN1 attaches to a MAG, and establishes associations with the LMA (unicast) and the MTMA (multicast).

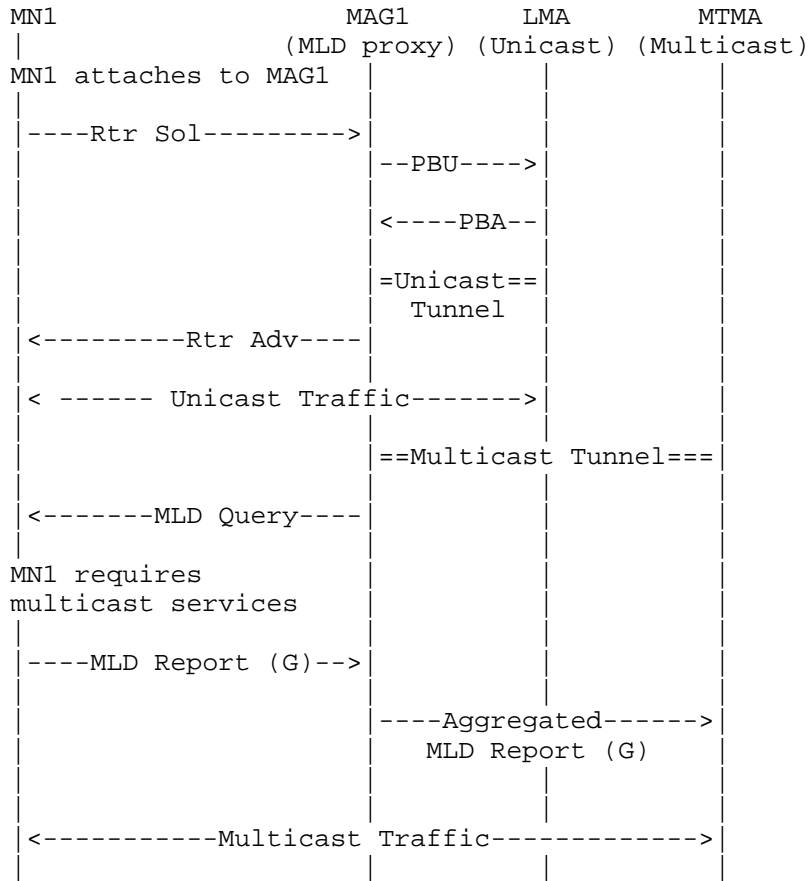


Figure 3: MN Attachment and Multicast Service Establishment for MTMA

In Figure 3, the MAG first establishes the PMIPv6 tunnel with LMA 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 LMA.

For multicast traffic, a multicast tunnel may have been pre-configured between MAG and MTMA, or 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 MAG (generated as defined by [RFC6224] upon handover). The MAG, acting as an MLD proxy defined in [RFC4605], will then send an Aggregated MLD Report to the multicast anchor, MTMA (assuming that this is a new multicast group that the MAG had not previously

subscribed to). Multicast traffic will then flow from the MTMA towards MN1. The MTMA acts as an MLD Querier, so it will periodically query each mobile access gateway about the subscriptions it maintains (not shown in Figure 3).

We next consider a mobility scenario in which MN1 with an ongoing multicast subscription moves from one MAG to another MAG. According to the baseline solution signaling method described in [RFC6224], after MN1 mobility, the new mobile access gateway acting in its role of MLD proxy will send an MLD Query to the newly observed mobile node on its downlink. Assuming that the subsequent MLD Report from MN1 requests membership for a new multicast group (from the new MAG's point of view), this will then result in an Aggregated MLD Report being sent to the MTMA from the new mobile access gateway. This message will be sent through a multicast tunnel between the new MAG and MTMA (pre-established or dynamically established).

When MN1 detaches, the old MAG may keep the multicast tunnel with the multicast MTMA if there are still other MNs using the multicast tunnel. Even if there are no mobile nodes currently on the multicast tunnel, the old MAG may decide to keep the multicast tunnel temporarily for potential future use.

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

4.2.2. Direct Routing Mode (Subscription via Direct Routing)

In this case, the MLD proxy instance is configured to obtain the multicast traffic locally. Figure 4 shows an example of multicast service establishment. The mobile access gateway first establishes the PMIPv6 tunnel with the local mobility anchor for unicast traffic as defined in [RFC5213] after being triggered by the Router Solicitation message from the mobile node. Unicast traffic will then flow between the MN and LMA.

For multicast traffic, it is assumed that the upstream interface of the MLD proxy instance has been configured pointing to a multicast router internal to the PMIPv6 domain (or towards an additional MLD proxy node in the domain), for all the multicast channels (which, in consequence, have to be local). There should be direct connectivity between the MAG and the local multicast router (or additional MLD proxy).

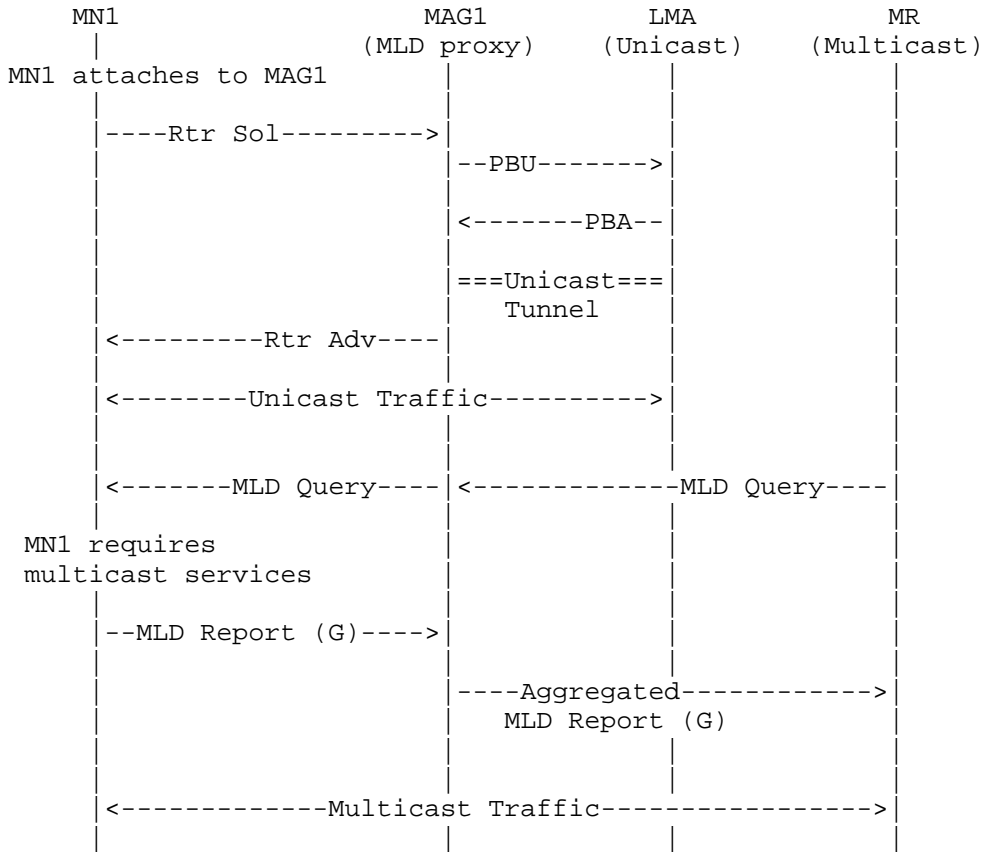


Figure 4: Multicast Service Establishment for Direct Routing

Upon detecting node attachment from an incoming interface, the MAG adds each downstream interface to the MLD proxy instance with an upstream link to an MR according to the standard MLD proxy operations [RFC4605] and sends an MLD Query message towards the MN. The mobile node sends the MLD report message (when required by its upper-layer applications) in response to an MLD Query from the MAG. Upon receiving the MLD Report message from each incoming interface, the MAG checks the MLD proxy instance associated with the downstream interface and then the MLD Report messages will be aggregated and forwarded to the upstream link associated with the MR (assuming that this is a new multicast group that the MAG had not previously subscribed to). Multicast traffic will then flow from the local multicast router towards the mobile node.

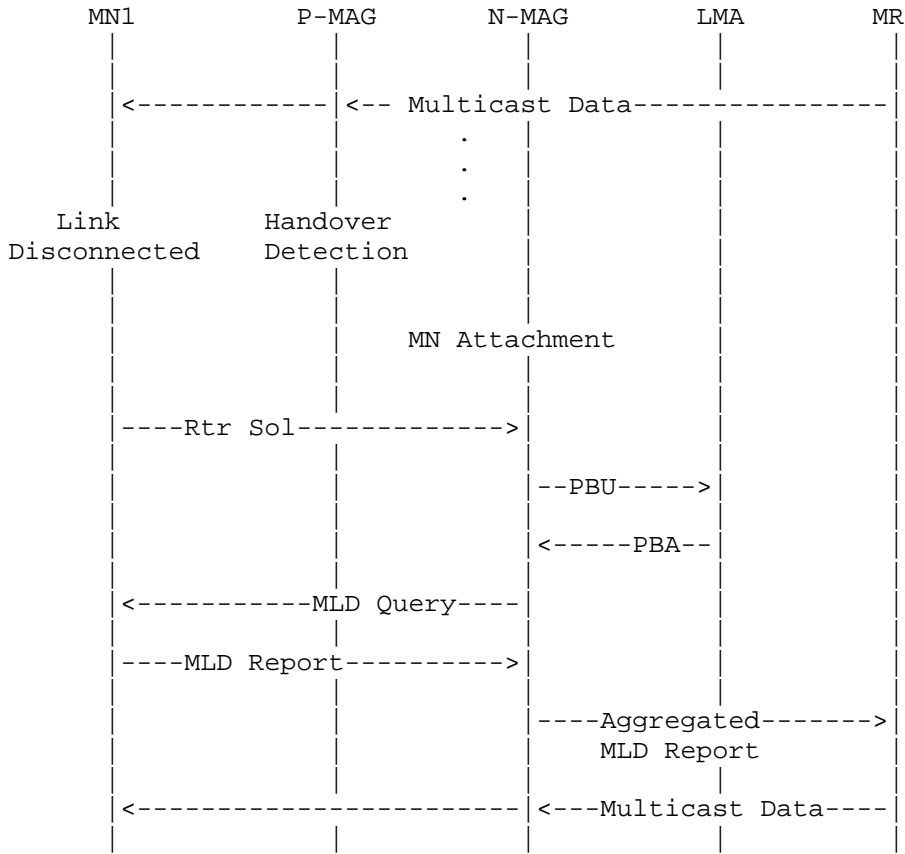


Figure 5: Multicast Mobility Signaling for Direct Routing

Figure 5 shows the handover operation procedure for the direct routing operation mode. When MN1 hands off to the next MAG (N-MAG) from the previous MAG (P-MAG), the N-MAG detects the newly arrived attached mobile node and performs binding update procedure by exchanging PBU/PBA signaling messages with LMA. At the same time, an MLD proxy instance detecting MN1 transmits an MLD query message to the mobile node. After receiving the MLD query message, MN1 sends an MLD report message that includes the multicast group information. The N-MAG then sends an aggregated MLD report message to the upstream link associated with the MR. An upstream interface of MLD proxy instance is chosen towards certain multicast router. The upstream interface selection can be done according to dynamic policies conveyed in the Dynamic IP Multicast Selector option (as described in Section 5.1) or according to manually configured policies. Note that in the base solution defined in [RFC6224], the interface selection is determined for each MN based on the Binding Update List. When the

N-MAG receives the multicast packets from the MR, it then simply forwards them without tunnel encapsulation. The N-MAG updates MN1's location information to the LMA by exchanging PBU/PBA signaling messages.

5. Local Mobility Anchor Operation

This section includes a new mobility option to support dynamic policies on subscription via MTMA/direct routing based on the local mobility anchor conveying the required info to the mobile access gateway in the proxy binding acknowledgement message.

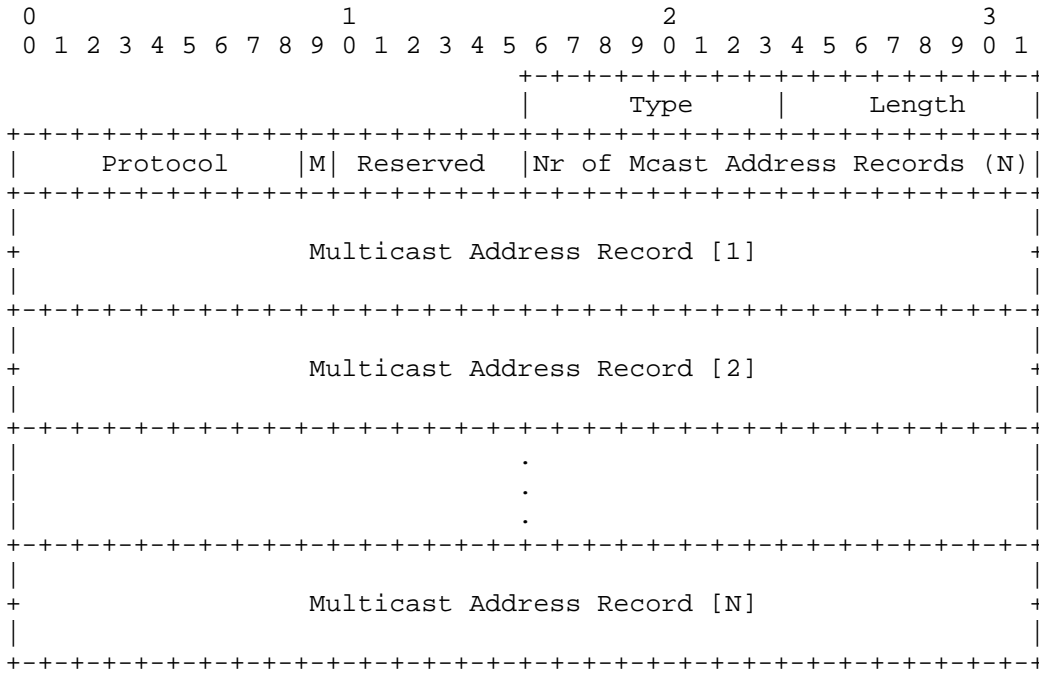
5.1. Dynamic IP Multicast Selector Option

5.1.1. Option Application Rules

A new TLV-encoded mobility option, the Dynamic IP Multicast Selector option, is defined for use with the proxy binding acknowledgement message exchanged between an LMA and a MAG to convey dynamic policies on subscription via MTMA/direct routing. This option is used for exchanging the IP addresses of both the group subscribed to by the MN, and the source(s) delivering it, as well as the applicable filter mode. This information is carried by using directly the Multicast Address Record format defined in [RFC3810]. There can be multiple "Dynamic IP Multicast Selector" options present in the message, up to one for each active subscription maintained by the MN.

5.1.2. Option Format

The format of this new option is as follows:



Type:

54

Length:

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

Protocol:

Field used to identify the multicast membership protocol in use, and the corresponding format of the next Multicast Address Record. This field maps the type codification used in the original MLD specifications for the Report message, namely for MLDv2 [RFC3810] the Protocol value MUST be 143, whereas for MLDv1 [RFC2710] the Protocol value MUST be 131.

Dynamic IP Multicast Selector Mode Flag (M-bit):

This field indicates the subscription via MTMA/direct routing mode. If the (M) flag value is set to a value of (1), it is an indication that the IP multicast traffic associated with the multicast group(s) identified by the Multicast Address Record(s)

in this mobility option SHOULD be routed locally (subscription via direct routing mode). If the (M) flag value is set to a value of (0), it is an indication that IP multicast traffic associated with the multicast group(s) identified by the Multicast Address Record in this mobility option(s) SHOULD be routed to the home network, via the MTMA (subscription via MTMA mode). The mobile access gateway MAY also choose to use static pre-established policies instead of following the indications provided by the local mobility anchor. All other IP traffic associated with the mobile node is managed according to a default policy configured at the PMIPv6 multicast domain.

Reserved:

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

Nr of Mcast Address Records (N)

16-bit unsigned integer indicating the number of Mcast Address Records (N) present in this option.

Multicast Address Record:

Multicast subscription information corresponding to a single multicast address as defined in [RFC3810], or as defined in [RFC2710] for MLDv1.

6. Multicast Tree Mobility Anchor Operation

The MTMA provides connectivity to the multicast infrastructure out of the PMIPv6 domain. The MTMA itself either could act as an additional MLD proxy (only in the case where all the connected mobile access gateways act also as MLD proxies), reporting to a further node an aggregated view of the subscriptions in a PMIPv6 multicast domain, or can act as a designated multicast router for all the MAGs in a PMIPv6 multicast domain. The multicast tree mobility anchor will then request the multicast content on behalf of the MAGs (and mobile nodes behind them). In addition, the MTMA will create and maintain the corresponding multicast forwarding states per each tunnel interface towards the MAGs. Whatever the role played, when the MAGs act as MLD proxy, the MTMA becomes the MLD querier of the MLD proxy instance located in each MAG.

6.1. Conceptual Data Structures

The multicast tree mobility anchor does not directly interact with the mobile nodes attached to any of the mobile access gateways. The MTMA only manages the multicast groups subscribed per MAG on behalf of the MNs attached to it. Having this in mind, the relevant information to be stored in the MTMA should be the tunnel interface identifier (tunnel-if-id) of the bidirectional tunnel for multicast between the MTMA and every MAG (e.g., similar to what is stated in [RFC5213] for the unicast case), the IP addresses of the multicast group delivered per tunnel to each of the MAGs, and the IP addresses of the sources injecting the multicast traffic per tunnel to the multicast domain defined by the MTMA.

7. Mobile Node Operation

The mobile node operation is not impacted by the existence of an MTMA as anchor for the multicast traffic being subscribed or the use of direct routing. The MN will act according to the stated operations in [RFC5213] and [RFC6224].

This document considers that every mobile node requesting multicast-only services is previously registered in a PMIPv6 unicast domain to get a unicast IP address. The registration can also be required for several purposes such as remote management, billing, multicast configuration, etc.

A given mobile node's policy profile information must be updated to be able to store the IPv6 addresses of both the local mobility anchor and multicast tree mobility anchor, the later for the subscription via MTMA case.

8. IPv4 Support

This document does not introduce any IPv4-specific issue regarding [RFC5844]. In order for the solution to support IPv4, all the described network elements (i.e., MAG, MTMA, and MR) must support IGMP. In this case, the functionalities of the MAG and MTMA would be as described in [RFC6224], with the MTMA replicating the requirements described for the LMA. For the case of the MR, it must also be dual-stack (i.e., IPv6/IPv4) enabled.

Although references to "MLD proxy" have been used in the document, it should be understood to also include "IGMP/MLD proxy" functionality.

Regarding the Dynamic IP Multicast Selector Option format, it SHOULD consider IPv4 compatibility in the following way:

Protocol field:

For IPv4, this field maps the type codification used in the original IGMP specifications for the Report message, in the following way:

It MUST be 0x12 in case of using IGMPv1.

It MUST be 0x16 in case of using IGMPv2.

It MUST be 0x22 in case of using IGMPv3.

Multicast Address Record field:

This field takes different formats depending on the IGMP version being used by the MN, as follows:

- * For IGMPv1, it takes the format given by the Group Address in [RFC1112].
- * For IGMPv2, it takes the format given by the Group Address in [RFC2236].
- * For IGMPv3, it takes the format given by the Group Record in [RFC3376].

9. IANA Considerations

This document defines a new mobility option, the Dynamic IP Multicast Selector, which has been assigned the Type 54 by IANA. The Type value for these options has been assigned from the same numbering space as allocated for the other mobility options, as defined in [RFC6275]: <http://www.iana.org/assignments/mobility-parameters>.

10. Security Considerations

This document describes two complementary operational modes that can be used to deliver multicast traffic in a PMIPv6 domain: multicast anchor and direct routing. Different approaches are described in the document to decide which operational mode is selected: i) the use of pre-configured/pre-provisioned policies at the mobile access gateway, or ii) the use of dynamic policies. Approach ii) could introduce a potential security issue if the protocol signaling is not properly secured. The use of the Dynamic IP Multicast Selector option described in the document requires message integrity protection and source authentication. Hence, the IPsec security mechanism

recommended by Proxy Mobile IPv6 [RFC5213] MUST be used to secure the Dynamic IP Multicast Selector option conveyed in the PBA (Proxy Binding Acknowledgement).

This document does not introduce any additional security threats beyond the current security considerations of PMIPv6 [RFC5213], MLD [RFC3810], IGMP [RFC3376], and IGMP/MLD Proxying [RFC4605].

11. Contributors

The following individuals made significant contributions to this document.

Akbar Rahman
InterDigital Communications, LLC
EMail: akbar.rahman@interdigital.com

Ignacio Soto
Universidad Carlos III de Madrid
EMail: isoto@it.uc3m.es

12. References

12.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [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.
- [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.
- [RFC5847] Devarapalli, V., Koodli, R., Lim, H., Kant, N., Krishnan, S., and J. Laganier, "Heartbeat Mechanism for Proxy Mobile IPv6", RFC 5847, June 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

12.2. Informative References

- [AUTO] Bumgardner, G., "Automatic Multicast Tunneling", Work in Progress, July 2013.
- [MLDPROXY] Asaeda, H. and S. Jeon, "Multiple Upstream Interface Support for IGMP/MLD Proxy", Work in Progress, February 2013.
- [MUIIMP] Zhang, H. and T. Schmidt, "Multi-Upstream Interfaces IGMP/MLD Proxy", Work in Progress, July 2013.
- [MULTIMOB] Schmidt, T., Gao, S., Zhang, H., and M. Waehlich, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains", Work in Progress, July 2013.
- [PMIPv6-REQ] Deng, H., Chen, G., Schmidt, T., Seite, P., and P. Yang, "Multicast Support Requirements for Proxy Mobile IPv6", Work in Progress, July 2009.
- [RFC6224] Schmidt, T., Waehlich, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.
- [UPSTREAM] Contreras, LM., Bernardos, CJ., and JC. Zuniga, "Extension of the MLD proxy functionality to support multiple upstream interfaces", Work in Progress, February 2013.

Appendix A. MTMA Deployment Use Cases

This informative appendix describes, from the network architecture point of view, several deployment options considering the MTMA.

These options can be distinguished in terms of the number of LMAs and MTMAs present in a PMIPv6 domain and the service relationship that a set of MNs gets from them, in the form of a "LMA : MTMA" ratio. According to that, it is possible to differentiate the following approaches:

- o A set of MNs is served in a PMIPv6 domain by two entities, one MTMA for multicast service, and one LMA for unicast, in such a way that the ratio is 1:1 (one common PMIPv6 unicast and multicast domain).
- o A set of MNs is served in a PMIPv6 domain by several entities, one MTMA for multicast service, while the others (LMAs) for unicast, in such a way that the ratio is N:1 (N PMIPv6 unicast domains coexist with a unique multicast domain).
- o A set of MNs is served in a PMIPv6 domain by several entities, one LMA for unicast, while the others (MTMAs) are devoted to multicast service, in such a way that the ratio is 1:N (one single PMIPv6 unicast domain coexists with multiple multicast domains).

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

A.1. PMIPv6 Domain with Ratio 1:1

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

A.2. PMIPv6 Domain with Ratio N:1

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

Figure 6 shows the scenario here described.

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 served by LMA3. For the scenario described above, the association between each MN and the corresponding LMA and MTMA is permanently maintained.

A.3. PMIPv6 Domain with Ratio 1:N

This approach is related to a 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 that are served by distinct MTMAs for multicast service as they move in the PMIPv6 domain. Different MTMAs might be associated with serving different multicast groups. These associations remain the same even if the MNs move within the PMIPv6 domain.

Figure 7 shows the scenario here described.

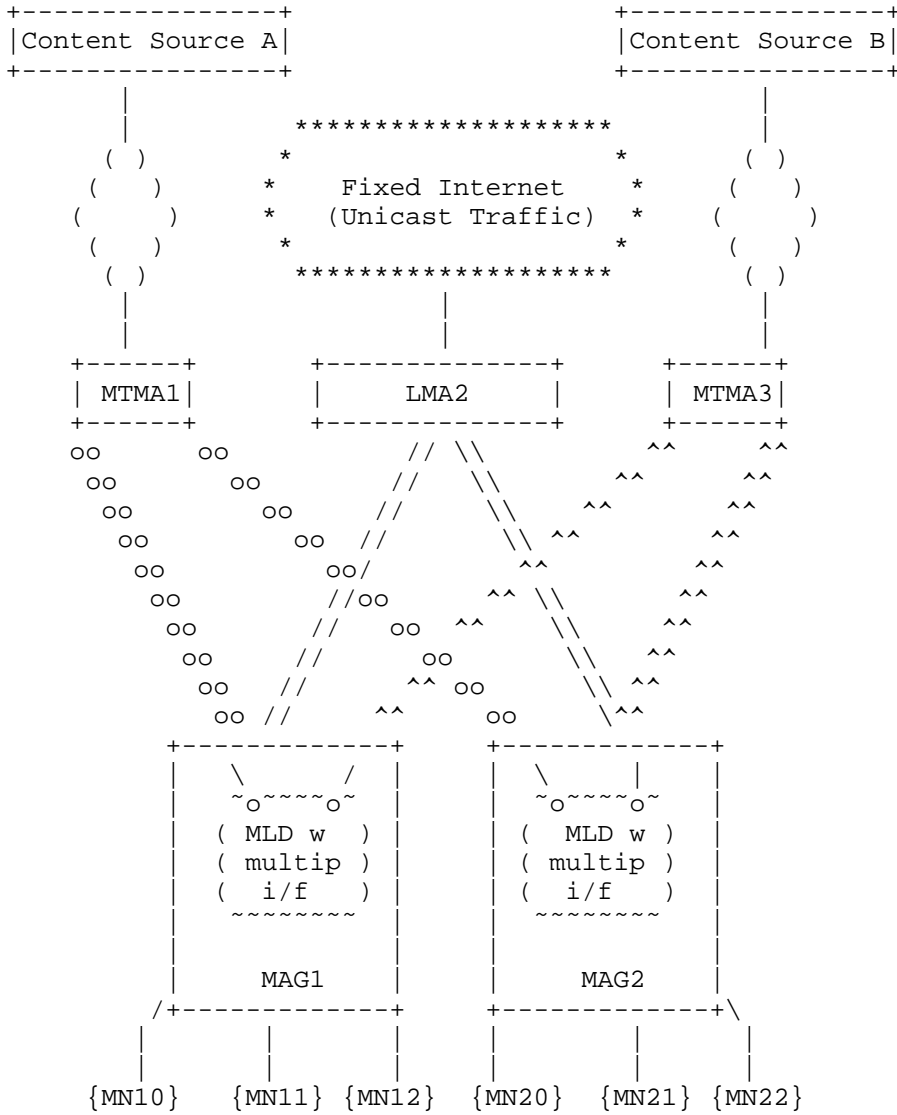


Figure 7: PMIPv6 Domain with Ratio 1:N

Figure 7 proposes an architecture where the LMA2 is the unique LMA for a certain group of MNs, while there are two other entities, MTMA1 and MTMA3, acting as MTMAs for different subsets of multicast content. MTMA1 and MTMA3 constitute two distinct multicast domains, whereas LMA2 forms a single unicast domain. Each MTMA could be devoted to carry on a different content (for instance, MTMA1 for source A and MTMA3 for source B). Looking at the figure, all MNs are

served by LMA2 for unicast, while they might be simultaneously served by MTMA1 and MTMA3, depending on the multicast content. For the scenario described above, the association between multicast content and MTMA is permanently maintained. Note that this scenario would require support for MLD proxy with multiple interfaces [MULTIMOB], [UPSTREAM], [MLDPROXY], [MUIIMP] at the MAGs.

A.4. PMIPv6 Domain with H-LMA

The H-LMA is defined as an entity that simultaneously transports unicast and multicast service, that is, it simultaneously works as LMA and MTMA. In the context of the MTMA solution, an H-LMA can play the role of MTMA for an entire group of MNs in a PMIPv6 domain, while acting simultaneously as LMA for a subset of them. Figure 8 adapts the PMIPv6 domain with ratio N:1 scenario of Figure 6 to the case where MTMA2 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 MN21 and MN30.

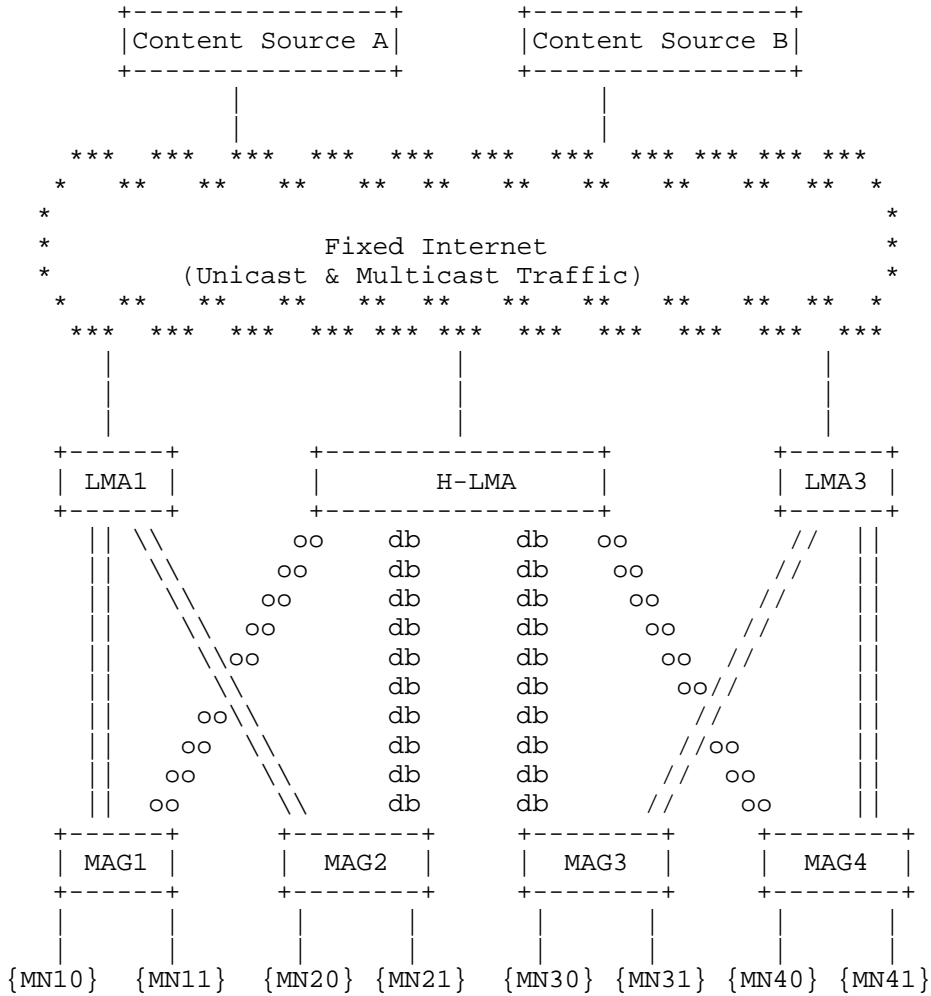


Figure 8: PMIPv6 Domain with H-LMA

Figure 8 presents a PMIPv6 network where there are two pure unicast LMAs, LMA1, and LMA3, and a hybrid LMA, labeled as H-LMA in the figure. The H-LMA is an MTMA 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 H-LMA (one single multicast domain), but it is possible to distinguish three subsets from the unicast service perspective (that is, three unicast domains). The first subset is the one formed by MN10, MN11, and MN20, which receives unicast traffic from LMA1. A second subset is the one formed by MN21 and MN30, which receives unicast traffic from H-LMA. 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 LMA and H-LMA is permanently maintained.

Authors' Addresses

Juan Carlos Zuniga
InterDigital Communications, LLC
1000 Sherbrooke Street West, 10th floor
Montreal, Quebec H3A 3G4
Canada
EMail: JuanCarlos.Zuniga@InterDigital.com
URI: <http://www.InterDigital.com/>

Luis M. Contreras
Telefonica I+D
Don Ramon de la Cruz, 82-84
Madrid 28006
Spain
EMail: lmcm@tid.es

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain
Phone: +34 91624 6236
EMail: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

Seil Jeon
Instituto de Telecomunicacoes
Campus Universitario de Santiago
Aveiro 3810-193
Portugal
EMail: seiljeon@av.it.pt
URI: <https://atnog.av.it.pt/~sjeon/>

Younghan Kim
Soongsil University
Sangdo-dong, Dongjak-gu
Seoul 511
Republic of Korea
EMail: yhkim@dcn.ssu.ac.kr
URI: <http://dcnlab.ssu.ac.kr/>

