

MULTIMOB Working Group
Internet Draft
Intended status: Experimental
Expires: April 16, 2010

Luis M. Contreras
Carlos J. Bernardos
Ignacio Soto
UC3M
October 15, 2009

Multicast service delivery in PMIPv6 domains
draft-contreras-multimob-msd-00.txt

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 April 16, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

A new working group, Multimob, has been chartered for providing an implementation of multicast service distribution to multicast listeners as they move in currently defined Proxy Mobile IPv6 (PMIPv6) domains according to [RFC 5213](#) [1]. No protocol modifications or extensions are considered at this stage.

Several proposals [2, 3, 4] have been already presented to face this issue. This contribution sums up some of the previous ideas extending them when necessary to overcome some limitations existing on such solutions.

Table of Contents

1.	Introduction.....	2		
2.	Conventions and terminology.....	3		
3.	Overview.....	3		
4.	Unicast and multicast service delivery integration.....	4		
5.	Local Mobility Anchor (LMA) operation.....	5		
6.	Mobile Access Gateway (MAG) operation.....	5		
	6.1.	MAG operation as MLD proxy.....	5	
	6.2.	MAG operation as multicast router.....	5	
7.	Multicast traffic hand-off process.....	6		
	7.1.	Newly-attached MN interrogation.....	6	
	7.2.	Multicast context transfer among MAGs.....	7	
	7.3.	pMAG-assisted handover.....	8	
		7.3.1.	Multicast flow encapsulation.....	11
		7.3.2.	Multicast flow delivery to the MN from the nMAG...	12
8.	Security Considerations.....	12		
9.	IANA Considerations.....	12		
10.	References.....	12		
	10.1.	Normative References.....	12	
	10.2.	Informative References.....	13	
11.	Acknowledgments.....	13		

1. Introduction

A new working group, Multimob, has been chartered for providing an implementation of multicast service distribution to multicast listeners as they move in currently defined Proxy Mobile IPv6 (PMIPv6) domains according to [RFC 5213](#) [1]. No protocol modifications or extensions are considered at this stage.

This draft takes as starting point some existing contributions [3, 4, 5] to define a comprehensive architecture for multicast traffic delivery in a PMIPv6 domain, introducing a new mechanism to allow fast handover of the multicast flow, thus minimising multicast service disruption when the multicast mobile listener moves across the network.

2. Conventions and terminology

This document uses the terminology referring to PMIPv6 components as defined in [1].

3. Overview

In currently defined PMIPv6 solution [1] two main entities, named Local Mobility Anchor (LMA) and Mobile Access Gateway (MAG), are in charge of managing both Mobile Node (MN) mobility and traffic delivery from/to the MN.

The LMA plays a central role in the unicast traffic delivery from/to the MN due that it is responsible of guaranteeing that the MN is reachable through the PMIPv6 domain by means of announcing an available route towards the MN, and forwarding traffic to it. By contrast, it plays a passive role in the mobility process, just pointing out to the MAG which acts as next-hop router to reach the MN.

On the other hand, the MAG's main role is to manage the MN node mobility within the PMIPv6 domain, notifying to the LMA the MN's current location as it moves within the PMIPv6 domain. Respect to the unicast traffic delivery, the MAG behaves as a passive element, just forwarding the MN's terminating traffic on the point-to-point link connecting both the MAG and the MN, or forwarding the traffic originated in the MN to the LMA as default router.

The case for multicast traffic is slightly different to the unicast one. The IP destination address of multicast packets does not identify the receiver, so the LMA would need additional information to be able to forward the multicast flow to the MNs requiring such content. In order to manage multicast traffic terminated on an MN, the LMA would need to keep the multicast status for every MN subscribing to a content. Nevertheless, the MAG is the element placed as first-hop router for the MNs, and thus, the MAG is the element

which directly receives the MLD signalling messages coming from the MNs. It is the element better positioned in the network to maintain the MN multicast status.

As consequence of that, in this draft the LMA is considered to manage unicast traffic terminating on the MN, whereas the MAG is considered to manage the multicast one, either if it plays the role of an MLD proxy or a fully-enabled multicast router.

This draft takes as starting point some existing contributions [3, 4, 5] to define a comprehensive architecture for multicast traffic delivery in a PMIPv6 domain, extending them in three ways:

1. An integrated vision of multicast service delivery together with unicast standard process.
2. The extension of the MAG role to also consider it to be a fully enabled multicast router.
3. A new proposal for minimising multicast service disruption on the handover event between MAGs.

4. Unicast and multicast service delivery integration

This draft considers that every MN demanding multicast services is previously registered in the PMIPv6 domain based on a unicast IP address as stated in [1]. This registration can be required for several purposes such as remote management, billing, etc.

As the registration is required by the network before any multicast service is provided, any MLD message originated by the MN will be discarded until the registration process is finished. This fact also prevents the network from accepting any subscription request or leave sent by the MN during the handover process.

In the same way, any subscription request or leave sent by the MN after a detachment procedure is initiated, will be discarded by the network, not being taken into account for multicast traffic management purposes.

5. Local Mobility Anchor (LMA) operation

As stated above, the LMA will not play any specific role in the multicast traffic management.

6. Mobile Access Gateway (MAG) operation

As stated above, the MAG will play a central role in the multicast traffic management. The MAG will keep the multicast status of every MN connected to it, and it will handle the multicast traffic accordingly to the MLD messages received on each point-to-point link.

Depending on the functionality deployed on the MAG, it is possible to consider the MAG as an MLD proxy or as a multicast router.

6.1. MAG operation as MLD proxy

When the MAG operates as MLD proxy it is in charge of summarizing the MN's subscription requests and delivering them towards the multicast router that can be reached through the configured proxy's upstream interface.

Conceptually, this scenario corresponds to the one described in Figure 1 of [5].

In this draft, just one MLD proxy instance is considered to be running in the MAG, and as consequence, just one interface can be defined as upstream interface.

The multicast router connected to the MLD proxy upstream interface could be any router in the network, even one of the routers acting as LMA. In such case, the upstream interface could be defined over the tunnel between the MAG and the LMA, once that tunnel has been set up.

6.2. MAG operation as multicast router

When a MAG operates as multicast router, it will send PIM messages to build the multicast tree that distributes the content required by a certain MN, if such content is not being received yet on the MAG.

Some content (the most popular one according to audience metrics or estimations) could be permanently subscribed at MAG to minimize the signalling involved in the establishment and pruning processes of the multicast tree leafs reaching the MAG.

One of the routers participating in the multicast distribution tree can be one of the routers that acts as LMA in the PMIPv6 domain. In that case, the multicast distribution tree could be extended over the tunnel between MAG and LMA, once that tunnel has been set up.

7. Multicast traffic hand-off process

As the MN moves and connects to different access networks with the same or different wireless access technology, the network should be able to guarantee the delivery of the traffic following the MN movement.

In the case of the multicast traffic, as seen above, the MAG entities have the knowledge about the multicast status (group subscription) of the MNs attached to them. MAG nodes contain enough information to support the handover process minimising multicast service disruption.

We describe below some alternatives [[4](#), [5](#)] to facilitate multicast service handover, including a new proposal.

7.1. Newly-attached MN interrogation

This method, proposed in [[4](#)], is based on the fact that MAG entities running an MLD instance have the capability of interrogating mobile hosts to obtain multicast memberships. In this case, an MN entering a new access network under responsibility of a new MAG will be interrogated by the new MAG, which sends an MLD Query towards the MN. Once the MLD Query is received, the MN will provide information about active memberships it has, if any. The MAG will get in this way knowledge of the multicast status for this newly-attached MN.

A number of pros and cons can be identified for this method.

Pros:

- o The MAG entity does not require any additional functionality as it already implements MLD capabilities.

- o The application of this method can be easily integrated within the MN registration process.

Cons:

- o Any new MN attaching the nMAG is interrogated, independently if it maintains or not an active multicast session at that time.
- o The signalling process wastes resources over the air interface.

7.2. Multicast context transfer among MAGs

This method, introduced in [5], proposes to trigger a context transfer process of the MN multicast status from the previous MAG (pMAG) to the new MAG (nMAG). As a previous step, the pMAG should predict the path the MN follows as it moves, to correctly identify the candidate nMAG to support the MN connection.

A number of pros and cons can be identified for this method.

Pros:

- o There is no additional signalling process over the air interface.
- o The context transfer method is an standardized process as per [2].

Cons:

- o The MAG entity requires new functionality to support multicast traffic handovers.
- o As predictive method, there is no total guarantee of success (the multicast context could not reach nMAG in all cases).
- o Coordination problems could arise in case of vertical handovers.

7.3. pMAG-assisted handover

This new proposal considers the fact that the MN keeps the same IP address in the PMIPv6 domain, thus the MN will be always reachable and traceable from the LMA as it moves from one wireless point of attachment to another.

The pMAG maintains the multicast status of an MN reconnecting to a nMAG. Before the detachment process, the multicast traffic reaching the MN passes through the pMAG. So, once a pMAG receives a Proxy Binding Ack (PBA) message from the LMA confirming that the MN has been detached from its point-to-point link, the pMAG can be able to temporally forward via the LMA a copy of the multicast content within an unicast flow towards the new location of the MN, once it is registered again in some point of the network. In case the MN is not registered again, the copy will be silently discarded in the LMA until timer expiration.

The timer duration has to be long enough to cover the registration process and the content subscription at the nMAG (triggered by periodic MLD reports from the MN). Once the timer expires, the pMAG deletes the MN multicast status state and stops the unicast flow that encapsulates the multicast content.

Figure 1 presents the interaction between pMAG, LMA and nMAG.

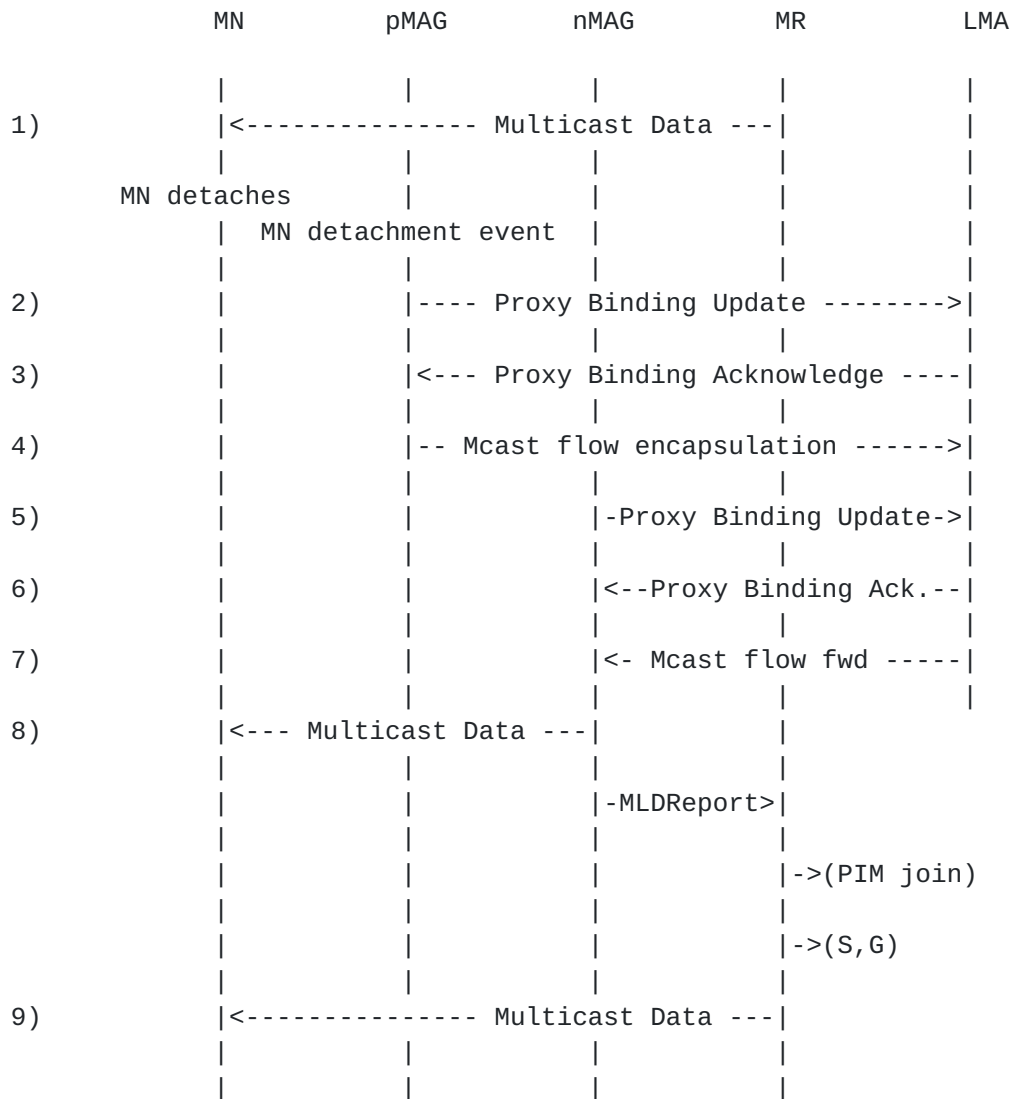


Figure 1. pMAG-assisted handover

Each step of the process is described as follows:

- 1) A registered MN is receiving a multicast content which has been previously subscribed by standard MLD request from the MN to the currently serving MAG, pMAG. The pMAG keeps the multicast status state of 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 a Proxy Binding Update message. From this moment, pMAG would discard any new MLD Report message coming from the MN requesting new content subscription.

- 3) The LMA confirms the reception of the detach process notification by sending a Proxy Binding Acknowledge message to the pMAG.
- 4) After reception of the PBA message, the pMAG encapsulates the multicast flow being served to the MN prior to the detachment event, and temporally forwards it to the LMA. The reason for this is that the MN keeps its IP address through the PMIPv6 domain, and it will be always reachable through the LMA, as soon as the MN is registered again by the nMAG. During the time that the MN is not reachable by the LMA, the LMA will silently discard the encapsulated multicast flow sent by the pMAG.

The encapsulation method is described later in this draft.

A timer is set up by pMAG to control the duration of multicast flow forwarding via the LMA. The timer value must be at least the maximum time to cover the MN registration process at the MAG and the refreshment of the group membership by the MLD instance at the MN.

- 5) The MN triggers an attachment process in a different MAG, nMAG. This nMAG signals the LMA with the new mobility status of the MN.
- 6) The LMA confirms the new attachment, and configures the routing table accordingly to forward incoming traffic to the MN through the nMAG. At this time, any MLD messages coming from the MN are accepted by the nMAG.
- 7) Once the MN is reachable again, the LMA forwards the encapsulated multicast traffic originally subscribed by the MN to the nMAG.
- 8) The nMAG decapsulates the originally subscribed multicast flow and sends it over the point-to-point link connecting both the nMAG and the MN. With the information contained in the flow, the nMAG is able to build the multicast status of the MN. Additionally, in case the nMAG currently does not have a subscription to the required content, it can trigger a content subscription on behalf of the MN to set up the delivery tree to the nMAG.

If the MN sends MLD messages to leave originally subscribed multicast content, the encapsulated multicast flow coming from the pMAG will be silently discarded until timer expiration.

- 9) The delivery tree reaches nMAG. The multicast content is directly served by the multicast tree. The encapsulated multicast flow coming from the pMAG will be silently discarded until timer expiration.

A number of pros and cons can be identified for this method.

Pros:

- o There is no additional signalling process over the air interface.
- o The multicast traffic is sent towards the MN immediately after the new registration.

Cons:

- o The MAG entity require new functionality to support multicast traffic handovers.
- o The network temporally supports more traffic.

7.3.1. Multicast flow encapsulation

The originally subscribed multicast traffic has to be encapsulated from the pMAG to the LMA to forward it to the MN as it moves. The format of the tunnelled flow considered in this draft is the one already described in Figure 1 of reference [3], but with different addressing considerations. Briefly, the format proposed here is as follows:

1. One external tunnel, with source address pMAG's Proxy-CoA, and destination address LMA Address (LMAA).
2. One internal tunnel, with source address pMAG's Proxy-CoA, and destination address MN-HoA.
3. The multicast flow being received by the MN, with source address the IP address of the source injecting the multicast content to the network, and destination address the multicast IP address of the group subscribed by the MN.

Once these packets reach the LMA, if a routing entry already exists pointing out to the MN through the nMAG, the external tunnel addressing is modified accordingly: the source address is the LMAA, whereas the destination address is the nMAG's Proxy-CoA.

7.3.2. Multicast flow delivery to the MN from the nMAG

Here the same behaviour described in [section 5](#) of reference [3] is proposed for the nMAG.

After receiving the forwarded original multicast flow from pMAG, the nMAG will analyze the internal tunnel header. In case the source address is a Proxy-CoA address of any neighbouring MAG in the domain, and the destination address is the one of the MN, MN-HoA, the nMAG will determine that a second tunnel exists. The nMAG will remove the internal header and will deliver the original multicast content as is (with multicast source and group addresses) to the MN on the point-to-point link.

In this way, the MN is guaranteed to receive the original subscribed multicast flow as before the handover process.

This mechanism implies that any MAG within a PMIPv6 domain has to be aware of the Proxy-CoA addresses of the neighbouring MAGs in the domain, to be able to determine the existence of the second (internal) tunnel. The maximum number of IP addresses to take into account would be then equal to the number of MAGs in the domain.

8. Security Considerations

None.

9. IANA Considerations

This document makes no request of IANA.

10. References

10.1. Normative References

- [1] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), August 2008.

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

10.2. Informative References

- [3] S. Krishnan, B. Sarikaya and TC. Schmidt, "Proxy Mobile IPv6 Basic Multicast Support Solution", [draft-krishnan-multimob-pmip6basicmcast-solution-00](#), (work in progress), July 2009.
- [4] TC. Schmidt, M. Waehlich, B. Sarikaya, and S. Krishnan, "A Minimal Deployment Option for Multicast Listeners in PMIPv6 Domains", [draft-schmidt-multimob-pmip6-mcast-deployment-01](#), (work in progress), June 2009.
- [5] S. Jeon, Y. Kim, and J. Lee, "Mobile Multicasting Support in Proxy Mobile IPv6", [draft-sijeon-multimob-mms-pmip6-00](#), (work in progress), July 2009.

11. 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).

Authors' 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

