

MULTIMOB Group
Internet-Draft
Intended status: Standards Track
Expires: January 14, 2014

T. Schmidt, Ed.
HAW Hamburg
S. Gao
H. Zhang
Beijing Jiaotong University
M. Waehlich
link-lab & FU Berlin
July 13, 2013

Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains
draft-ietf-multimob-pmipv6-source-04

Abstract

Multicast communication can be enabled in Proxy Mobile IPv6 domains via the Local Mobility Anchors by deploying MLD Proxy functions at Mobile Access Gateways, via a direct traffic distribution within an ISP's access network, or by selective route optimization schemes. This document describes the support of mobile multicast senders in Proxy Mobile IPv6 domains for all three scenarios. Protocol optimizations for synchronizing PMIPv6 with PIM, as well as a peering function for MLD Proxies defined. Mobile sources always remain agnostic of multicast mobility operations.

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

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 January 14, 2014.

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. Base Solution for Source Mobility and PMIPv6 Routing	4
3.1. Overview	4
3.2. Base Solution for Source Mobility: Details	7
3.2.1. Operations of the Mobile Node	7
3.2.2. Operations of the Mobile Access Gateway	7
3.2.3. Operations of the Local Mobility Anchor	7
3.2.4. IPv4 Support	8
3.2.5. Efficiency of the Distribution System	9
4. Direct Multicast Routing	10
4.1. Overview	10
4.2. MLD Proxies at MAGs	11
4.2.1. Considerations for PIM-SM on the Upstream	12
4.2.2. SSM Considerations	12
4.3. PIM-SM at MAGs	12
4.3.1. Routing Information Base for PIM-SM	12
4.3.2. Operations of PIM in Phase One	13
4.3.3. Operations of PIM in Phase Two	14
4.3.4. Operations of PIM in Phase Three	14
4.3.5. PIM-SSM Considerations	15
4.3.6. Handover Optimizations for PIM	15
4.4. BIDIR-PIM	16
4.4.1. Routing Information Base for BIDIR-PIM	16
4.4.2. Operations of BIDIR-PIM	16
5. Extended MLD Proxy Function for Optimized Source Mobility in PMIPv6	17
5.1. Peering Function for MLD Proxies	17
5.1.1. Operations at the Multicast Sender	18
5.1.2. Operations at the Multicast Listener	18
6. IANA Considerations	19

7. Security Considerations	19
8. Acknowledgements	20
9. References	20
9.1. Normative References	20
9.2. Informative References	21
Appendix A. Multiple Upstream Interface Proxy	21
A.1. Operations for Local Multicast Sources	22
A.2. Operations for Local Multicast Subscribers	22
Appendix B. Change Log	23
Authors' Addresses	24

1. Introduction

Proxy Mobile IPv6 (PMIPv6) [RFC5213] extends Mobile IPv6 (MIPv6) [RFC6275] 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). An MN connected to a PMIPv6 domain, which only operates according to the base specifications of [RFC5213], cannot participate in multicast communication, as MAGs will discard group packets.

Multicast support for mobile listeners can be enabled within a PMIPv6 domain by deploying MLD Proxy functions at Mobile Access Gateways, and multicast routing functions at Local Mobility Anchors [RFC6224]. This base deployment option is the simplest way to PMIPv6 multicast extensions in the sense that it follows the common PMIPv6 traffic model and neither requires new protocol operations nor additional infrastructure entities. Standard software functions need to be activated on PMIPv6 entities, only, at the price of possibly non-optimal multicast routing.

Alternate solutions leverage performance optimization by providing multicast routing at the access gateways directly, or by selective route optimization schemes. Such approaches (partially) follow the business model of providing multicast data services in parallel to PMIPv6 unicast routing.

Multicast listener support satisfies the needs of receptive use cases such as IPTV or server-centric gaming on mobiles. However, current trends in the Internet enfold towards user-centric, highly interactive group applications like user generated streaming, conferencing, collective mobile sensing, etc. Many of these popular applications create group content at end systems and can largely profit from a direct data transmission to a multicast-enabled network.

This document describes the support of mobile multicast senders in Proxy Mobile IPv6 domains subsequently for the base deployment scenario [RFC6224], for direct traffic distribution within an ISP's access network, as well as for selective route optimization schemes. The contribution of this work reflects the source mobility problem as discussed in [RFC5757]. Mobile Nodes in this setting remain agnostic of multicast mobility operations.


2. Terminology

This document uses the terminology as defined for the mobility protocols [RFC6275], [RFC5213] and [RFC5844], as well as the multicast edge related protocols [RFC3376], [RFC3810] and [RFC4605].

3. Base Solution for Source Mobility and PMIPv6 Routing

3.1. Overview

The reference scenario for multicast deployment in Proxy Mobile IPv6 domains is illustrated in Figure 1. MAGs play the role of first-hop access routers that serve multiple MNs on the downstream while running an MLD/IGMP proxy instance for every LMA upstream tunnel.






Figure 1: Reference Network for Multicast Deployment in PMIPv6 with Source Mobility

An MN in a PMIPv6 domain will decide on multicast data transmission completely independent of its current mobility conditions. It will send packets as initiated by applications, using its source address with Home Network Prefix (HNP) and a multicast destination address chosen by application needs. Multicast packets will arrive at the currently active MAG via one of its downstream local (wireless) links. A multicast unaware MAG would simply discard these packets in the absence of a multicast routing information base (MRIB).

An MN can successfully distribute multicast data in PMIPv6, if MLD proxy functions are deployed at the MAG as described in [RFC6224]. In this set-up, the MLD proxy instance serving a mobile multicast source has configured its upstream interface at the tunnel towards MN's corresponding LMA. For each LMA, there will be a separate instance of an MLD proxy.

According to the specifications given in [RFC4605], multicast data arriving from a downstream interface of an MLD proxy will be forwarded to the upstream interface and to all but the incoming downstream interfaces that have appropriate forwarding states for this group. Thus multicast streams originating from an MN will arrive at the corresponding LMA and directly at all mobile receivers co-located at the same MAG and MLD Proxy instance. Serving as the designated multicast router or an additional MLD proxy, the LMA forwards data to the fixed Internet, whenever forwarding states are maintained by multicast routing. If the LMA is acting as another MLD proxy, it will forward the multicast data to its upstream interface, and to downstream interfaces with matching subscriptions, accordingly.

In case of a handover, the MN (unaware of IP mobility) can continue to send multicast packets as soon as network connectivity is reconfigured. At this time, the MAG has determined the corresponding LMA, and IPv6 unicast address configuration (including PMIPv6 bindings) has been completed. Still multicast packets arriving at

the MAG are discarded (if not buffered) until the MAG has completed the following steps.

1. The MAG has determined that the MN is admissible to multicast services.
2. The MAG has added the new downstream link to the MLD proxy instance with up-link to the corresponding LMA.

As soon as the MN's uplink is associated with the corresponding MLD proxy instance, multicast packets are forwarded again to the LMA and eventually to receivers within the PMIP domain (see the call flow in Figure 2). In this way, multicast source mobility is transparently enabled in PMIPv6 domains that deploy the base scenario for multicast.




Figure 2: Call Flow for Group Communication in Multicast-enabled PMIP

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]. IPv4 multicast is handled by an IGMP proxy function at the MAG in an analogous way.

Following these deployment steps, multicast traffic distribution transparently inter-operates with PMIPv6. It is worth noting that an MN - while being attached to the same MAG as the mobile source, but associated with a different LMA - cannot receive multicast traffic on a shortest path. Instead, multicast streams flow up to the LMA of the mobile source, are transferred to the LMA of the mobile listener and tunneled downwards to the MAG again (see Section 5 for further optimizations).

3.2. Base Solution for Source Mobility: Details

A support of multicast source mobility in PMIPv6 requires to deploy general multicast functions at PMIPv6 routers and to define their interaction with the PMIPv6 protocol in the following way.

3.2.1. Operations of the Mobile Node

A Mobile Node willing to send multicast data will proceed as if attached to the fixed Internet. No specific mobility or other multicast related functionalities are required at the MN.

3.2.2. Operations of the Mobile Access Gateway

A Mobile Access Gateway is required to have MLD proxy instances deployed, one for each tunnel to an LMA, which serves as its unique upstream link (cf., [RFC6224]). On the arrival of an MN, the MAG decides on the mapping of downstream links to a proxy instance and the upstream link to the LMA based on the regular Binding Update List as maintained by PMIPv6 standard operations. When multicast data is received from the MN, the MAG MUST identify the corresponding proxy instance from the incoming interface and forwards multicast data upstream according to [RFC4605].

The MAG MAY apply special admission control to enable multicast data transition from an MN. It is advisable to take special care that MLD proxy implementations do not redistribute multicast data to downstream interfaces without appropriate subscriptions in place.

3.2.3. 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 upstream for the

corresponding MAG. It will manage and maintain a multicast forwarding information base for all group traffic arriving from its mobile sources. It SHOULD participate in multicast routing functions that enable traffic redistribution to all adjacent LMAs within the PMIPv6 domain and thereby ensure a continuous receptivity while the source is in motion.

3.2.3.1. Local Mobility Anchors Operating PIM

Local Mobility Anchors that operate the PIM-SM routing protocol [RFC4601] will require sources to be directly connected for sending PIM registers to the RP. This does not hold in a PMIPv6 domain, as MAGs are routers intermediate to MN and the LMA. In this sense, MNs are multicast sources external to the PIM-SM domain.

To mitigate this incompatibility common to all subsidiary MLD proxy domains, the LMA MUST act as a PIM Border Router and activate the Border-bit. In this case, the `DirectlyConnected(S)` is treated as being TRUE for mobile sources and the PIM-SM forwarding rule "`iif == RPF_interface(S)`" is relaxed to be TRUE, as the incoming tunnel interface from MAG to LMA is considered as not part of the PIM-SM component of the LMA (see A.1 of [RFC4601]).

In addition, an LMA serving as PIM Designated Router is connected to MLD proxies via individual IP-tunnel interfaces and will experience changing PIM source states on handover. As the incoming interface connects to a point-to-point link, PIM Assert contention is not active, and incoming interface validation is only performed by RPF checks. Consequently, a PIM DR SHOULD update incoming source states, as soon as RPF inspection succeeds, i.e., after PMIPv6 forwarding state update. Consequently, PIM routers SHOULD be able to manage these state changes, but some implementations are expected to incorrectly refuse packets until the previous state has timed out.

Notably, running BIDIR PIM [RFC5015] on LMAs remains robust with respect to source location and does not require special configurations or state management for sources.

3.2.4. IPv4 Support

An MN in a PMIPv6 domain may use an IPv4 address transparently for communication as specified in [RFC5844]. For this purpose, an LMA can register an IPv4-Proxy-CoA in its Binding Cache and the MAG can provide IPv4 support in its access network. 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.

Multicast streams from and to MNs arrive at a MAG on point-to-point links (identical to unicast). Multicast data transmission from the MAG to the corresponding LMA is link-local between the routers and routing/forwarding remains independent of any individual MN. So the MAG-proxy and the LMA SHOULD NOT use GRE key identifiers, but plain GRE encapsulation in multicast communication (including MLD queries and reports). Multicast traffic is transmitted as router-to-router forwarding via the MAG-to-LMA tunnels and according to the multicast routing information base (MRIB) of the MAG or the LMA. It remains independent of MN's unicast addresses, while the MAG proxy instance re-distributes multicast data down the point-to-point links (interfaces) according to its local subscription states, independent of IP addresses of the MN.

3.2.5. Efficiency of the Distribution System

The distribution system of the base solution directly follows PMIPv6 routing rules, and organizes multicast domains with respect to LMAs. Thus, no coordination between address spaces or services is required between the different instances, provided their associated LMAs belong to disjoint multicast domains. Routing is optimal for communication between MNs of the same domain, or stationary subscribers.

In the following, efficiency-related issues remain.

Multicast reception at LMA In the current deployment scenario, the LMA will receive all multicast traffic originating from its associated MNs. There is no mechanism to suppress upstream forwarding in the absence of receivers.

MNs on the same MAG using different LMAs For a mobile receiver and a source that use different LMAs, the traffic has to go up to one LMA, cross over to the other LMA, and then be tunneled back to the same MAG, causing redundant flows in the access network and at the MAG.


4. Direct Multicast Routing

There are deployment scenarios, where multicast services are available throughout the access network independent of the PMIPv6 routing system [I-D.ietf-multimob-pmipv6-ropt]. In these cases, the visited networks grant a local content distribution service (in contrast to LMA-based home subscription) with locally optimized traffic flows. It is also possible to deploy a mixed service model of local and LMA-based subscriptions, provided a unique way of service selection is implemented. For example, access routers (MAGs) could decide on service access based on the multicast address G or the SSM channel (S,G) under request (see Appendix A for further discussions).

4.1. Overview

Direct multicast access can be supported by

- o native multicast routing provided by one multicast router that is neighboring MLD proxies deployed at MAGs within a flat access network, or via tunnel uplinks,
o a multicast routing protocol such as PIM-SM [RFC4601] or BIDIR-PIM [RFC5015] deployed at the MAGs.






Figure 3: Reference Networks for (a) Proxy-assisted Direct Multicast Access and (b) Dynamic Multicast Routing at MAGs

Figure 3 displays the corresponding deployment scenarios, which separate multicast from PMIPv6 unicast routing. It is assumed throughout these scenarios that all MAGs (MLD proxies) are linked to a single multicast routing domain. Notably, this scenario requires coordination of multicast address utilization and service bindings.

Multicast traffic distribution can be simplified in these scenarios. A single proxy instance at MAGs with up-link into the multicast domain will serve as a first hop multicast gateway and avoid traffic duplication or detour routing. Multicast routing functions at MAGs will seamlessly embed access gateways within a multicast cloud. However, mobility of the multicast source in this scenario will require some multicast routing protocols to rebuild distribution trees. This can cause significant service disruptions or delays (see [RFC5757] for further aspects). Deployment details are specific to the multicast routing protocol in use, in the following described for common protocols.

4.2. MLD Proxies at MAGs

In a PMIPv6 domain, single MLD proxy instances can be deployed at each MAG that enable multicast service at the access via an uplink to a multicast service infrastructure (see Figure 3 (a)). To avoid service disruptions on handovers, the uplinks of all proxies SHOULD be adjacent to the same next-hop multicast router. This can either be achieved by arranging proxies within a flat access network, or by upstream tunnels that terminate at a common multicast router.

Multicast data submitted by a mobile source will reach the MLD proxy at the MAG that subsequently forwards flows to the upstream, and all downstream interfaces with appropriate subscriptions. Traversing the upstream will lead traffic into the multicast infrastructure (e.g., to a PIM Designated Router) which will route packets to all local MAGs that have joined the group, as well as further upstream according to protocol procedures and forwarding states.

On handover, a mobile source will reattach to a new MAG and can continue to send multicast packets as soon as PMIPv6 unicast configurations have completed. Like at the previous MAG, the new MLD proxy will forward data upstream and downstream to subscribers. Listeners local to the previous MAG will continue to receive group traffic via the local multicast distribution infrastructure following aggregated listener reports of the previous proxy. In general, traffic from the mobile source continues to be transmitted via the same next-hop router using the same source address and thus remains unchanged when seen from the wider multicast infrastructure.

4.2.1. Considerations for PIM-SM on the Upstream

A mobile source that transmits data via an MLD proxy will not be directly connected to a PIM Designated Router as discussed in Section 3.2.3.1. Countermeasures apply correspondingly.

A PIM Designated Router that is connected to MLD proxies via individual IP-tunnel interfaces will experience invalid PIM source states on handover. In some implementations of PIM-SM this could lead to an interim packet loss (see Section Section 3.2.3.1). This problem can be mitigated by aggregating proxies on a lower layer.

4.2.2. SSM Considerations

Source-specific subscriptions invalidate with routes, whenever the source moves from or to the MAG/proxy of a subscriber. Multicast forwarding states will rebuild with unicast route changes. However, this may lead to noticeable service disruptions for locally subscribed nodes.

4.3. PIM-SM at MAGs

The full-featured multicast routing protocol PIM-SM MAY be deployed in the access network for providing multicast services in parallel to unicast routes. Throughout this section, it is assumed that the PMIPv6 mobility domain is part of a single PIM-SM multicast routing domain with PIM-SM routing functions present at all MAGs and all LMAs. The PIM routing instance at a MAG SHALL then serve as the Designated Router (DR) for all directly attached Mobile Nodes. For expediting handover operations, it is advisable to position PIM Rendezvous Points (RPs) in the core of the PMIPv6 network domain. However, regular IP routing tables need not be present in a PMIPv6 deployment, and additional effort is required to establish reverse path forwarding rules as required by PIM-SM.

4.3.1. Routing Information Base for PIM-SM

In this scenario, PIM-SM will rely on a Multicast Routing Information Base (MRIB) that is generated independently of the policy-based routing rules of PMIPv6. The granularity of mobility-related routing locators required in PIM depends on the complexity (phases) of its deployment.

The following information is needed for all phases of PIM.

- o All routes to networks and nodes (including RPs) that are not mobile members of the PMIPv6 domain MUST be defined consistently among PIM routers and MUST remain unaffected by node mobility. The setup of these general routes is expected to follow the topology of the operator network and is beyond the scope of this document.

The following route entries are required at a PIM-operating MAG when phases two or three of PIM, or PIM-SSM are in operation.

- o Local routes to the Home Network Prefixes (HNPs) of all MNs associated with their corresponding point-to-point attachments that MUST be included in the local MRIB.
- o All routes to MNs that are attached to distant MAGs of the PMIPv6 domain point towards their corresponding LMAs. These routes MUST be made available in the MRIB of all PIM routers (except for the local MAG of attachment), but MAY be eventually expressed by an appropriate default entry.

4.3.2. Operations of PIM in Phase One

A new mobile source S will transmit multicast data of group G towards its MAG of attachment. Acting as a PIM DR, the access gateway will unicast-encapsulate the multicast packets and forward the data to the Virtual Interface (VI) with encapsulation target RP(G), a process known as PIM source registering. The RP will decapsulate and natively forward the packets down the RP-based distribution tree towards (mobile and stationary) subscribers.

On handover, the point-to-point link connecting the mobile source to the old MAG will go down and all (S,*) flows terminate. In response, the previous DR (MAG) deactivates the data encapsulation channels for the transient source (i.e., all `DownstreamJPState(S,*,VI)` are set to NoInfo state). After reattaching and completing unicast handover negotiations, the mobile source can continue to transmit multicast packets, while being treated as a new source at its new DR (MAG). Source register encapsulation will be immediately initiated, and (S,G) data continue to flow natively down the (*,G) RP-based tree.

Source handover management in PIM phase one admits low complexity and remains transparent to receivers. In addition, the source register tunnel management of PIM is a fast protocol operation and little overhead is induced thereof. In a PMIPv6 deployment, PIM RPs MAY be configured to not initiate (S,G) shortest path trees for mobile sources, and thus remain in phase one of the protocol. The price to pay for such simplified deployment lies in possible routing detours by an overall RP-based packet distribution.

4.3.3. Operations of PIM in Phase Two

After receiving source register packets, a PIM RP eventually will initiate a source-specific Join for creating a shortest path tree to the (mobile) source S, and issue a source register stop at the native arrival of data from S. For initiating an (S,G) tree, the RP, as well as all intermediate routers, require route entries for the HNP of the MN that - unless the RP coincides with the MAG of S - point towards the corresponding LMA of S. Consequently, the (S,G) tree will proceed from the RP via the (stable) LMA, down the LMA-MAG tunnel to the mobile source. This tree can be of lesser routing efficiency than the PIM source register tunnel established in phase one, but provides the advantage of immediate data delivery to receivers that share a MAG with S.

On handover, the mobile source reattaches to a new MAG (DR), and PMIPv6 unicast management will transfer the LMA-MAG tunnel to the new point of attachment. However, in the absence of a corresponding multicast forwarding state, the new DR will treat S as a new source and initiate a source registering of PIM phase one with the RP. In response, the PIM RP will recognize the known source at a new (tunnel) interface immediately responds with a register stop. As the RP had joined the shortest path tree to receive from the source via the LMA, it will see an RPF change when data arrives at a new interface. Implementation-dependent, this can trigger an update of the PIM MRIB and trigger a new PIM Join message. Otherwise, the tree is periodically updated by Joins transmitted towards the new MAG on a path via the LMA. In proceeding this way, a quick recovery of PIM transition from phase one to two will be performed per handover.

4.3.4. Operations of PIM in Phase Three

In response to an exceeded threshold of packet transmission, DRs of receivers eventually will initiate a source-specific Join for creating a shortest path tree to the (mobile) source S, thereby transitioning PIM into the final short-cut phase three. For all receivers not sharing a MAG with S, this (S,G) tree will range from the receiving DR via the (stable) LMA, the LMA-MAG tunnel to the mobile source. This tree is of higher routing efficiency than

established in the previous phase two, but need not outperform the PIM source register tunnel established in phase one. It provides the advantage of immediate data delivery to receivers that share a MAG with S.

On handover, the mobile source reattaches to a new MAG (DR), and PMIPv6 unicast management will transfer the LMA-MAG tunnel to the new point of attachment. However, in the absence of a corresponding multicast forwarding state, the new DR will treat S as a new source and initiate a source registering of PIM phase one. A PIM implementation compliant with this change can recover phase three states in the following way. First, the RP recovers to phase two as described in the previous section, and will not forward data arriving via the source register tunnel. Tree maintenance eventually triggered by the RPF change (see Section 4.3.3) will generate proper states for a native forwarding from the new MAG via the LMA. Thereafter, packets arriving at the LMA without source register encapsulation are forwarded natively along the shortest path tree towards receivers.

In consequence, the PIM transition from phase one to two and three will be quickly recovered per handover, but still leads to an enhanced signaling load and intermediate packet loss.

4.3.5. PIM-SSM Considerations

Source-specific Joins of receivers will guide PIM to operate in SSM mode and lead to an immediate establishment of source-specific shortest path trees. Such (S,G) trees will equal the distribution system of PIM's final phase three (see Section 4.3.4). However, on handover and in the absence of RP-based data distribution, SSM data delivery cannot be resumed via source registering as in PIM phase one. Consequently, data packets transmitted after a handover will be discarded at the MAG until regular tree maintenance has re-established the (S,G) forwarding state at the new MAG.

4.3.6. Handover Optimizations for PIM

Source-specific shortest path trees are constructed in PIM-SM (phase two and three), and in PIM-SSM that follow LMA-MAG tunnels towards a source. As PIM remains unaware of source mobility management, these trees invalidate under handovers with each tunnel re-establishment at a new MAG. Regular tree maintenance of PIM will recover the states, but remains unsynchronized and too slow to seamlessly preserve PIM data dissemination.

A method to quickly recover PIM (S,G) trees under handover SHOULD synchronize multicast state maintenance with unicast handover operations and MAY proceed as follows. On handover, an LMA reads all

(S,G) Join states from its corresponding tunnel interface and identifies those source addresses S_i that match moving HNPs. After re-establishing the new tunnel, it SHOULD associate the $(S_i,*)$ Join states with the new tunnel endpoint and immediately trigger a state maintenance (PIM Join) message. In proceeding this way, the source-specific PIM states are transferred to the new tunnel end point and propagated to the new MAG in synchrony with unicast handover procedures.

4.4. BIDIR-PIM

BIDIR-PIM MAY be deployed in the access network for providing multicast services in parallel to unicast routes. Throughout this section, it is assumed that the PMIPv6 mobility domain is part of a single BIDIR-PIM multicast routing domain with BIDIR-PIM routing functions present at all MAGs and all LMAs. The PIM routing instance at a MAG SHALL then serve as the Designated Forwarder (DF) for all directly attached Mobile Nodes. For expediting handover operations, it is advisable to position BIDIR-PIM Rendezvous Point Addresses (RPAs) in the core of the PMIPv6 network domain. As regular IP routing tables need not be present in a PMIPv6 deployment, reverse path forwarding rules as required by BIDIR-PIM need to be established.

4.4.1. Routing Information Base for BIDIR-PIM

In this scenario, BIDIR-PIM will rely on a Multicast Routing Information Base (MRIB) that is generated independently of the policy-based routing rules of PMIPv6. The following information is needed.

- o All routes to networks and nodes (including RPAs) that are not mobile members of the PMIPv6 domain MUST be defined consistently among BIDIR-PIM routers and remain unaffected by node mobility. The setup of these general routes is expected to follow the topology of the operator network and is beyond the scope of this document.

4.4.2. Operations of BIDIR-PIM

BIDIR-PIM will establish spanning trees across its network domain in conformance to its preconfigured RPAs and the routing information provided. Multicast data transmitted by a mobile source will immediately be forwarded by its DF (MAG) onto the spanning group tree without further protocol operations.

On handover, the mobile source re-attaches to a new MAG (DF), which completes unicast network configurations. Thereafter, the source can

immediately proceed with multicast packet transmission onto the pre-established distribution tree. BIDIR-PIM does neither require protocol signaling nor additional reconfiguration delays to adapt to source mobility and can be considered the protocol of choice for mobile multicast operations in the access. As multicast streams always flow up to the Rendezvous Point Link, some care should be taken to configure RPAs compliant with network capacities.

5. Extended MLD Proxy Function for Optimized Source Mobility in PMIPv6

A deployment of MLD Proxies (see [RFC4605]) at MAGs has proven a useful and appropriate approach to multicast in PMIPv6, see [RFC6224], [I-D.ietf-multimob-pmipv6-ropt]. However, deploying unmodified standard proxies can go along with significant performance degradation for mobile senders as discussed along the lines of this document. To overcome these deficits, an optimized approach to multicast source mobility based on extended peering functions among proxies is introduced in this section. Prior to presenting the solution, we will sketch the relevant requirements.

Solutions that extend MLD Proxies by additional uplinking functions need to comply to the following requirements.

Prevention of Routing Loops In the absence of a full-featured routing logic at an MLD Proxy, simple and locally decidable rules need to prevent source traffic from traversing the network in loops as potentially enabled by multiple uplinks.

Unique coverage of receivers Listener functions at Proxies require simple, locally decidable rules to initiate a unique delivery of multicast packets to all receivers.

Following different techniques, these requirements are met in the following solutions.

5.1. Peering Function for MLD Proxies

In this section, we define a peering interface for MLD proxies that allows for a direct data exchange of locally attached multicast sources. Such peering interfaces can be configured - as a direct link or a bidirectional tunnel - between any two proxy instances (locally deployed as in [RFC6224] or remotely) and remain as silent virtual links in regular proxy operations. Data on such link is exchanged only in cases, where one peering proxy directly connects on the downstream to a source of multicast traffic, which the other peering proxy actively subscribes to. Operations are defined for ASM and SSM, but provide superior performance in the presence of source-specific signaling (IGMPv3/MLDv2).

5.1.1. Operations at the Multicast Sender

An MLD Proxy in the perspective of a sender will see peering interfaces as restricted downstream interfaces. It will install and maintain source filters at its peering links that will restrict data transmission to those packets that originate from a locally attached source at the downstream. In detail, a proxy will extract from its configuration the network prefixes attached to its downstream interfaces and MUST implement a source filter base at its peering interfaces that restricts data transmission to IP source addresses from its local prefixes. This filter base MUST be updated, if and only if the downstream configuration changes. In this way, a multihop forwarding on peering links is prevented. Multicast packets that arrive from the upstream interface of the proxy are thus only forwarded to regular downstream interfaces with appropriate subscription states.

Multicast traffic arriving from a locally attached source will be forwarded to the regular upstream interface and all downstreams with appropriate subscription states (i.e., regular Proxy operations). In addition, local multicast packets are transferred to those peering interfaces with appropriate subscription states.

5.1.2. Operations at the Multicast Listener

From the listener side, peering interfaces appear as preferred upstream links. Thus an MLD proxy with peering interconnects will offer several interfaces for pulling remote traffic: the regular upstream and the peerings. Traffic arriving from any of the peering links will be mutually disjoint, but normally also available from the upstream. To prevent duplicate traffic from arriving at the listener side, the proxy

- o MAY delay aggregated reports to the upstream, and
- o MUST apply appropriate filters to exclude duplicate streams.

In detail, it first issues listener reports (in parallel) to its peering links, which only span one (virtual) hop. Whenever the expected traffic (e.g., SSM channels) does not completely arrive from the peerings after a waiting time (default: 10 ms), additional (complementary, in the case of SSM) reports are sent to the standard upstream interface.

After the arrival of traffic from peering links, an MLD proxy MUST install source filters at the upstream in the following way.

ASM with IGMPv2/MLDv1 In the presence of Any Source Multicast using IGMPv2/MLDv1, only, the proxy cannot signal source filtering to its upstream. Correspondingly, it applies (S,*) ingress filters at its upstream interface for all sources S seen in traffic of the peering links. It is noteworthy that unwanted traffic is still replicated to the proxy via the access network.

ASM with IGMPv3/MLDv2 In the presence of source-specific signaling (IGMPv3/MLDv2), the upstream interface is set to (S,*) exclude mode for all sources S seen in traffic of the peering links. The corresponding source-specific signaling will prevent duplicate traffic forwarding throughout the access network.

SSM In the presence of Source Specific Multicast, the proxy will subscribe on its uplink interface to those (S,G) channels, only, that do not arrive via the peering links.

In proceeding this way, multicast group data arrive from peering interfaces first, while only peer-wise unavailable traffic is retrieved from the regular upstream interface.

6. IANA Considerations

TODO.

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

7. Security Considerations

TODO

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

The authors would like to thank (in alphabetical order) Luis M. Contreras, Muhamma Omer Farooq, Bohao Feng, Dirk von Hugo, Ning Kong, Jouni Korhonen, He-Wu Li, Akbar Rahman, Stig Venaas, Li-Li Wang, Qian Wu, Zhi-Wei Yan for advice, help and reviews of the document. Funding by the German Federal Ministry of Education and Research within the G-LAB Initiative (projects HAMcast, Mindstone and SAFEST) 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.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [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.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, October 2007.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

9.2. Informative References

- [I-D.ietf-multimob-pmipv6-ropt]
Zuniga, J., Contreras, L., Bernardos, C., Jeon, S., and Y. Kim, "Multicast Mobility Routing Optimizations for Proxy Mobile IPv6", draft-ietf-multimob-pmipv6-ropt-07 (work in progress), July 2013.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC5757] Schmidt, T., Waehlich, 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.
- [RFC6224] Schmidt, T., Waehlich, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

Appendix A. Multiple Upstream Interface Proxy

In this section, we document upstream extensions for an MLD proxy that were originally developed during the work on this document. Multiple proxy instances deployed at a single MAG (see Section 3) can be avoided by adding multiple upstream interfaces to a single MLD Proxy. In a typical PMIPv6 deployment, each upstream of a single proxy instance can interconnect to one of the LMAs. With such ambiguous upstream options, appropriate forwarding rules MUST be supplied to

- o unambiguously guide traffic forwarding from directly attached mobile sources, and
- o lead listener reports to initiating unique traffic subscriptions.

This can be achieved by a complete set of source- and group-specific filter rules (e.g., (S,*), (*,G)) installed at proxy interfaces. These filters MAY be derived in parts from PMIPv6 routing policies, and can include a default behavior (e.g., (*,*)).

A.1. Operations for Local Multicast Sources

Packets from a locally attached multicast source will be forwarded to all downstream interfaces with appropriate subscriptions, as well as up the interface with the matching source-specific filter.

Typically, the upstream interface for a mobile multicast source is chosen based on the policy routing (e.g., the MAG-LMA tunnel interface for LMA-based routing or the interface towards the multicast router for direct routing), but alternate configurations MAY be applied. Packets from a locally attached multicast source will be forwarded to the corresponding upstream interface with the matching source-specific filter, as well as all the downstream interfaces with appropriate subscriptions.

A.2. Operations for Local Multicast Subscribers

Multicast listener reports are group-wise aggregated by the MLD proxy. The aggregated report is issued to the upstream interface with matching group/channel-specific filter. The choice of the corresponding upstream interface for aggregated group membership reports MAY be additionally based on some administrative scoping rules for scoped multicast group addresses.

In detail, a Multiple Upstream Interface Proxy will provide and maintain a Multicast Subscription Filter Table that maps source- and group-specific filters to upstream interfaces. The forwarding decision for an aggregated MLD listener report is based on the first matching entry from this table, with the understanding that for IGMPv3/MLDv2 the MLD Proxy performs a state decomposition, if needed (i.e., a (*,G) subscription is split into (S,G) and (* \ S,G) in the presence of (*,G) after (S,G) interface entries), and that (S,*)-filters are always false in the absence of source-specific signaling, i.e. in IGMPv2/MLDv1 only domains.

In typical deployment scenarios, specific group services (channels) could be either associated with selected uplinks to remote LMAs, while a (*,*) default subscription entry (in the last table line) is bound to a local routing interface, or selected groups are configured as local services first, while a (*,*) default entry (in the last table line) points to a remote uplink that provides the general multicast support.

Appendix B. Change Log

The following changes have been made from version draft-ietf-multimob-pmipv6-source-03:

1. Fixed issues in Section Section 4.3 (PIM phase two and three transistion) according to WG feedback.
2. Editorial improvements, resolved nits.
3. Updated references.

The following changes have been made from version draft-ietf-multimob-pmipv6-source-02:

1. Added clarifications and details as requested by the working group, resolved nits.
2. Moved Multiple Upstream MLD Proxy to Appendix in response to WG desire.
3. Updated references.

The following changes have been made from version draft-ietf-multimob-pmipv6-source-01:

1. Added clarifications and details as requested by the working group, resolved nits.
2. Detailed out operations of Multiple Upstream MLD Proxies.
3. Clarified operations of MLD proxies with peering links.
4. Many editorial improvements.
5. Updated references.

The following changes have been made from version draft-ietf-multimob-pmipv6-source-00:

1. Direct routing with PIM-SM and PIM-SSM has been added.
2. PMIP synchronization with PIM added for improved handover.
3. Direct routing with BIDIR-PIM has been added.
4. MLD Proxy extensions requirements added.

5. Peering of MLD Proxies added.
6. First sketch of multiple upstream proxy added.
7. Editorial improvements.
8. Updated references.

Authors' Addresses

Thomas C. Schmidt (editor)
HAW Hamburg
Berliner Tor 7
Hamburg 20099
Germany

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

Shuai Gao
Beijing Jiaotong University
Beijing
China

Email: shgao@bjtu.edu.cn

Hong-Ke Zhang
Beijing Jiaotong University
Beijing
China

Email: hkzhang@bjtu.edu.cn

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

Email: mw@link-lab.net