Internet Engineering Task Force (IETF)
Request for Comments: 7287
Category: Experimental
ISSN: 2070-1721

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

June 2014

Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains

Abstract

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

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/rfc7287.

Table of Contents

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 Local Mobility Anchor (LMAs) 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 [MULTI-EXT] or by using selective route optimization schemes [RFC7028]. Such approaches (partially) follow the model of providing multicast data services in parallel to PMIPv6 unicast routing [RFC7161].

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 develop 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 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 source mobility problem as discussed in [RFC5757] serves as a foundation of this document. 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 routing [RFC4601] and edge-related protocols [RFC3376], [RFC3810], and [RFC4605].

Throughout this document, we use the following acronyms:

HNP Home Network Prefix as defined in [RFC5213].

MAG Mobile Access Gateway as defined in [RFC5213].

MLD Multicast Listener Discovery as defined in [RFC2710] and [RFC3810].

PIM Protocol Independent Multicast as defined in [RFC4601].

PMIPv6 Proxy Mobile IPv6 as defined in [RFC5213].

2.1. 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].

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. It displays the general setting for source mobility -- mobile nodes (MNs) with Home Network Prefixes (HNPs) that receive services via tunnels, which are spanned between a Local Mobility Anchor Address (LMAA) and a Proxy Care-of-Address (Proxy-CoA) at a Mobility Access Gateway (MAG). 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.

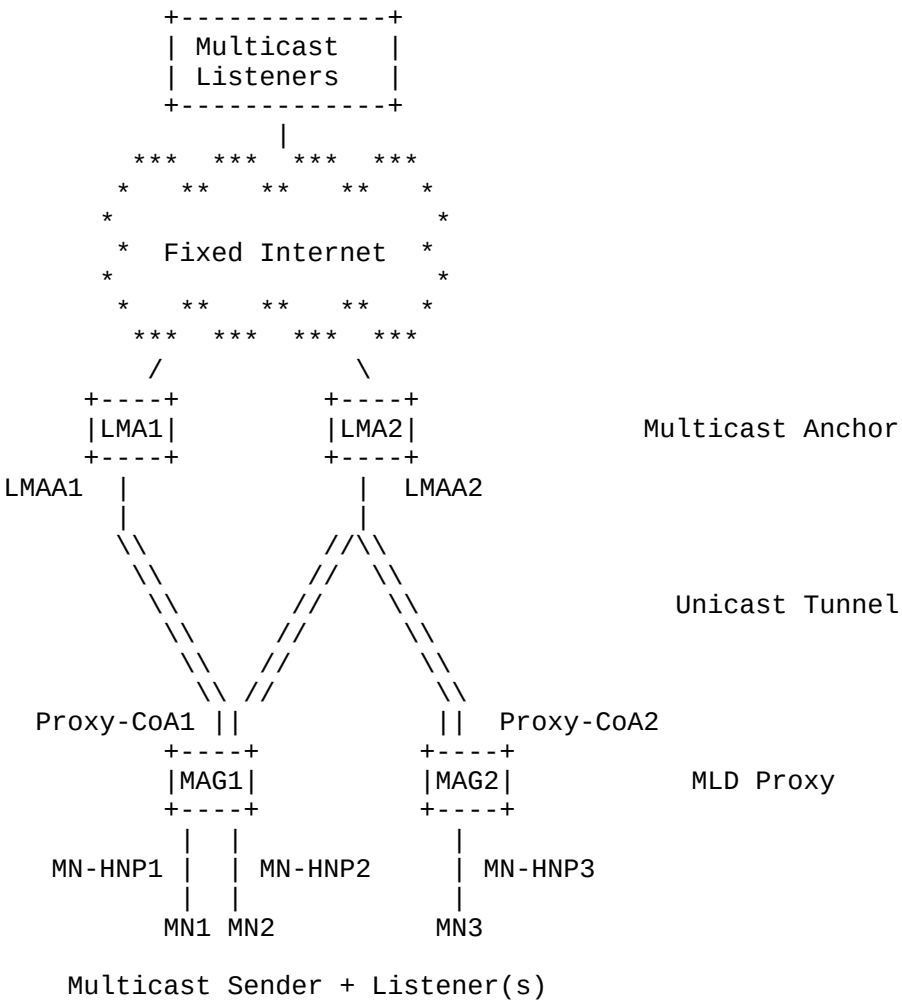

Figure 1: Reference Network for Multicast Deployment in PMIPv6

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 an 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 instructions for packet processing, i.e., 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 setup, the MLD proxy instance serving a mobile multicast source has configured its upstream interface at the tunnel towards the 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 (being unaware of IP mobility) can continue to send multicast packets as soon as network connectivity is re-established. 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 an uplink 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.

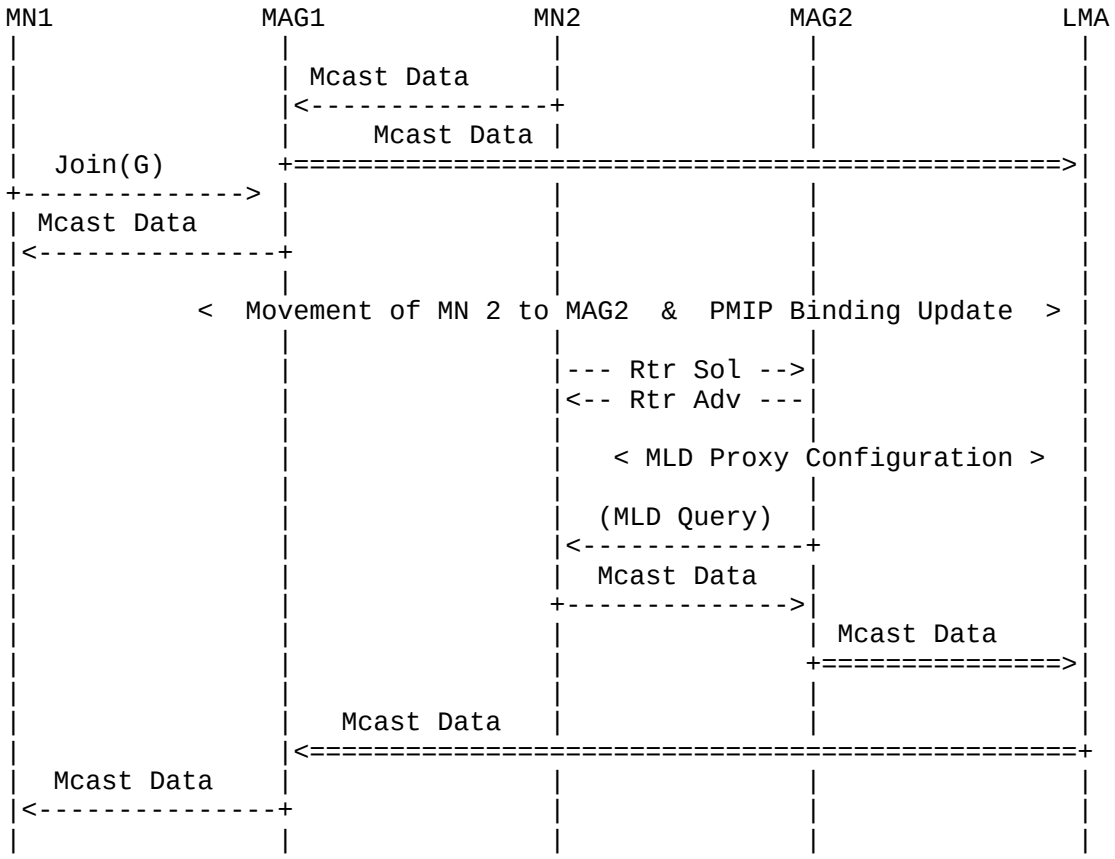

Figure 2: Call Flow for Group Communication in Multicast-Enabled PMIP

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 interoperates 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 are tunneled downwards to the MAG again. (See Section 5 for further optimizations.)

3.2. Base Solution for Source Mobility: Details

Support of multicast source mobility in PMIPv6 requires that general multicast functions be deployed at PMIPv6 routers and that their interactions with the PMIPv6 protocol be defined as follows.

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 transmission 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 Protocol Independent Multicast - Sparse Mode (PIM-SM) routing protocol [RFC4601] will require sources to be directly connected for sending PIM registers to the Rendezvous Point (RP). This does not hold in a PMIPv6 domain, as MAGs are routers intermediate to the 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 not considered part of the PIM-SM component of the LMA (see Appendix A.1 of [RFC4601]).

In addition, an LMA serving as the PIM Designated Router (DR) 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 Reverse Path Forwarding (RPF) checks. Consequently, a PIM DR SHOULD update incoming source states, as soon as RPF inspection succeeds, i.e., after the 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 Bidirectional PIM (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 behavior 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 they 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.

The following points are worth noting about the scenario in [RFC5845] in which overlapping private address spaces of different operators can be hosted in a PMIP domain by using Generic Routing Encapsulation (GRE) 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 in multicast communication (including MLD queries and reports). Multicast traffic is transmitted using router-to-router forwarding via the MAG-to-LMA tunnels and according to the MRIB of the MAG or the LMA. It remains independent of MN’s unicast addresses, while the MAG proxy instance redistributes 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 situations, 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.

These remaining deficits in routing efficiency can be resolved by adding peering functions to MLD proxies as described in Section 5.

4. Direct Multicast Routing

There are deployment scenarios, where multicast services are available throughout the access network independent of the PMIPv6 routing system [RFC7028]. 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 that 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 source-specific multicast (SSM) channel (S,G) under request. (See Appendix A for further discussions.)

4.1. Overview

Direct multicast access can be supported by

- 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,

- a multicast routing protocol such as PIM-SM [RFC4601] or BIDIR-PIM [RFC5015] deployed at the MAGs.

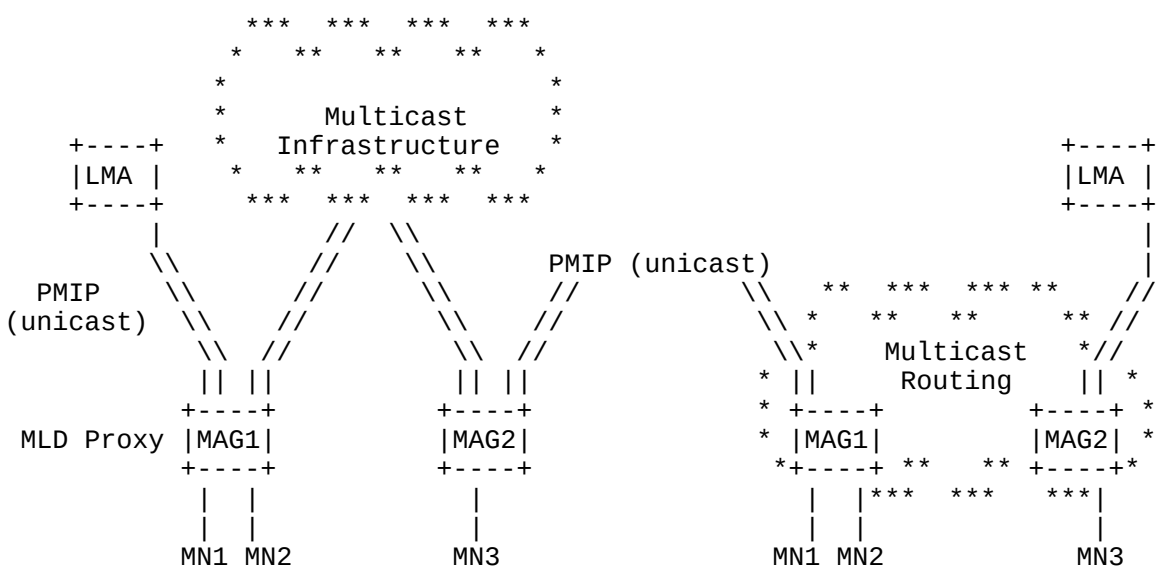

(a) Multicast Access at Proxy Uplink (b) Multicast Routing at MAG

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 that 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 uplinks 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; this is described below 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 using 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 to all downstream interfaces with appropriate subscriptions. Traversing the upstream will transfer traffic into the multicast infrastructure (e.g., to a PIM Designated Router) that 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 been 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 multicast 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 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 (see Figure 3(b)). 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 (specific phase) of its deployment.

For all three phases in the operation of PIM (see [RFC4601]), the following information is needed.

- 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 phase two or three of PIM or PIM-SSM is in operation.

- 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.

- 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 they MAY be eventually expressed by an appropriate default entry.

4.3.2. Operations of PIM in Phase One (RP Tree)

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 that introduces little overhead. In a PMIPv6 deployment, PIM RPs MAY be configured to uninitiated (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 (Register-Stop)

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 lower routing efficiency than the PIM source register tunnel established in phase one.

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 and will immediately respond with a register stop. As the RP had previously joined the shortest path tree towards the source via the LMA, it will see an RPF change when data arrives at a new interface. This is implementation dependent and can trigger an update of the PIM MRIB as well as a new PIM Join message that will install the multicast forwarding state missing at the new MAG. 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 (Shortest-Path Tree)

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 shortcut 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, and the serving MAG to the mobile source. This tree is of higher routing efficiency than that established in the previous phase two, but it 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 transitions from phase one to two to three will be quickly recovered per handover but still lead 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 reestablished 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. These RPF-trees traverse 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 it remains unsynchronized and too slow to seamlessly preserve PIM data distribution services.

A method to quickly recover PIM (S,G) trees under handover SHOULD synchronize multicast state maintenance with unicast handover operations and can 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 endpoint 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.

- 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 pre-configured RPAs and the routing information provided. Multicast data transmitted by a mobile source will immediately be forwarded by its DF (MAG) onto the spanning tree for the multicast group without further protocol operations.

On handover, the mobile source reattaches 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 not require protocol signaling or additional reconfiguration delays to adapt to source mobility, and it can be considered the protocol of choice for mobile multicast operations in the access network. As multicast streams always flow up to the Rendezvous Point Link, some care should be taken to configure RPAs compliant with network capacities.

5. MLD Proxy Peering 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] and [RFC7028]. However, deploying unmodified standard proxies can go along with significant performance degradation for mobile senders as discussed in this document. To overcome these deficits, an optimized approach to multicast source mobility based on extended peering functions among proxies is defined in this section. Based on such direct data exchange between proxy instances at MAGs, triangular routing is avoided and multicast streams can be disseminated directly within a PMIPv6 access network, and in particular within MAG routing machines. Prior to presenting the solution, we will summarize the relevant requirements.

5.1. 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 that would be 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 local filtering techniques, these requirements are met in the following solution.

5.2. Overview

A peering interface for MLD proxies 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 deployed). Peerings remain as silent virtual links in regular proxy operations. Data is exchanged on such links only in cases where one peering proxy on its downstream directly connects to a source of multicast traffic to which the other peering proxy actively subscribes. In such cases, the proxy connected to the source will receive a listener report on its peering interface and will forward traffic from its local source accordingly. It is worth noting that multicast traffic distribution on peering links does not follow reverse unicast paths to sources. In the following, operations are defined for Any-Source Multicast (ASM) and SSM, but they provide superior performance in the presence of source-specific signaling (IGMPv3/MLDv2) [RFC4604].

5.3. Operations in Support of Multicast Senders

An MLD proxy with 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 source that is locally attached at one of its downstream interfaces.

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 (e.g., due to mobility). Multicast packets that arrive from the upstream interface of the proxy are thus prevented from traversing any peering link, but they are only forwarded to regular downstream interfaces with appropriate subscription states. In this way, multihop forwarding on peering links is prevented.

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

5.4. Operations in Support of Multicast Listeners

On the listener side, peering interfaces appear as preferred upstream links. The multicast proxy will attempt to receive multicast services on peering links for as many groups (channels) as possible. The general upstream interface configured according to [RFC4605] will be used only for retrieving those groups (channels) that remain unavailable from peerings. From this extension of [RFC4605], an MLD proxy with peering interconnects will exhibit several interfaces for pulling remote traffic: the regular upstream and the peerings. Traffic available 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

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

In detail, an MLD proxy instance at a MAG first issues listener reports (in parallel) to all of its peering links. These links span at most one (virtual) hop. Whenever certain group traffic (SSM channels) does not arrive from the peerings after a waiting time (default: 10 ms (node-local) and 25 ms (remote)), additional reports (complementary reports, in the case of SSM) are sent to the standard upstream interface.

Whenever traffic from a peering link arrives, an MLD proxy MUST install source filters at its upstream interfaces (as described in RFC 4605) in the following way.

ASM with IGMPv2/MLDv1: In the presence of ASM 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 on the peering links. It is noteworthy that unwanted traffic is still replicated to the proxy via the (wired) provider backbone, but it is not forwarded into the wireless 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 forwarding of duplicate traffic 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.

MLD proxies will install data-driven timers (source-timeout) for each source but common to all peering interfaces to detect interruptions of data services from individual sources at proxy peers. Termination of source-specific flows may be application specific, but also may be due to a source handover or a transmission failure. After a handover, a mobile source may reattach at another MLD proxy with a peering relation to the listener, or at a proxy that does not peer. While in the first case, traffic reappears on another peering link; in the second case, data can only be retrieved via the regular upstream. To account for the latter, the MLD proxy revokes the source-specific filter(s) at its upstream interface, after the source-timeout expires (default: 50 ms). Corresponding traffic will then be pulled from the regular upstream interface. Source-specific filters MUST be reinstalled whenever traffic of this source arrives at any peering interface.

There is a noteworthy trade-off between traffic minimization and available traffic at the upstream that is locally filtered at the proxy. Implementors can use this relation to optimize for service-specific requirements.

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

6. Security Considerations

This document defines multicast sender mobility based on PMIPv6 and common multicast routing protocols. Consequently, threats identified as security concerns in [RFC2236], [RFC2710], [RFC3810], [RFC4605], [RFC5213], and [RFC5844] are inherited by this document.

In addition, 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 arise from rapid state changes, as well as from high-volume data streams routed into access networks of limited capacities. In cases of PIM-SM deployment, handover operations of the MNs include re-registering sources at the Rendezvous Points at possibly high frequency. In addition to proper authorization checks of MNs, rate controls at routing agents and replicators may be needed to protect the agents and the downstream networks. In particular, MLD proxy implementations at MAGs SHOULD automatically erase multicast state on the departure of MNs, as mobile multicast listeners in the PMIPv6 domain will in general not actively terminate group membership prior to departure.

The deployment of IGMP/MLD proxies for multicast routing requires particular care, as routing loops on the upstream are not automatically detected. Peering functions between proxies extend this threat in the following way. Routing loops among peering and upstream interfaces are prevented by filters on local sources. Such filtering can fail whenever prefix configurations for downstream interfaces at a proxy are incorrect or inconsistent. Consequently, implementations of peering-enabled proxies SHOULD take particular care on keeping IP configurations consistent at the downstream in a reliable and timely manner. (See [RFC6224] for requirements on PMIPv6-compliant implementations of MLD proxies.)

7. Acknowledgements

The authors would like to thank (in alphabetical order) David Black, Luis M. Contreras, Spencer Dawkins, Muhamma Omer Farooq, Bohao Feng, Sri Gundavelli, Dirk von Hugo, Ning Kong, Jouni Korhonen, He-Wu Li, Cong Liu, Radia Perlman, Akbar Rahman, Behcet Sarikaya, Stig Venaas, Li-Li Wang, Sebastian Woelke, Qian Wu, and 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.

8. References

8.1. 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.

8.2. Informative References

[MULTI-EXT] Schmidt, T., Ed., Waehlisch, M., Koodli, R., Fairhurst, G., and D. Liu, "Multicast Listener Extensions for MIPv6 and PMIPv6 Fast Handovers", Work in Progress, May 2014.

[PEERING-ANALYSIS]
Schmidt, TC., Woelke, S., and M. Waehlisch, "Peer my Proxy - A Performance Study of Peering Extensions for Multicast in Proxy Mobile IP Domains", Proc. of 7th IFIP Wireless and Mobile Networking Conference (WMNC 2014), IEEE Press, May 2014.

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

[RFC4604] Holbrook, H., Cain, B., and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", RFC 4604, August 2006.

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

[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., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

[RFC7028] Zuniga, JC., Contreras, LM., Bernardos, CJ., Jeon, S., and Y. Kim, "Multicast Mobility Routing Optimizations for Proxy Mobile IPv6", RFC 7028, September 2013.

[RFC7161] Contreras, LM., Bernardos, CJ., and I. Soto, "Proxy Mobile IPv6 (PMIPv6) Multicast Handover Optimization by the Subscription Information Acquisition through the LMA (SIAL)", RFC 7161, March 2014.

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 interface of a single proxy instance can interconnect to one of the LMAs. With such ambiguous upstream options, appropriate forwarding rules MUST be supplied to

- unambiguously guide traffic forwarding from directly attached mobile sources, and
- 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 part 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 a 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) are 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

- 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. Implementation

An implementation of the extended IGMP/MLD proxy has been provided within the MCPROXY project (http://mcproxy.realmv6.org/). This open-source software is written in C++ and uses forwarding capabilities of the Linux kernel. It supports all regular operations according to [RFC4605] and allows for multiple proxy instances on one node, dynamically changing downstream links, proxy-to-proxy peerings, and multiple upstream links with individual configurations. The software can be downloaded from GitHub at <https://github.com/mcproxy/mcproxy>. Based on this software, an experimental performance evaluation of the proxy peering function has been reported in [PEERING-ANALYSIS].

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

EMail: mw@link-lab.net