                Base Deployment for Multicast Listener Support
                   in Proxy Mobile IPv6 (PMIPv6) Domains

Abstract

   This document describes deployment options for activating multicast
   listener functions in Proxy Mobile IPv6 domains without modifying
   mobility and multicast protocol standards. Similar to home agents in
   Mobile IPv6, Local Mobility Anchors of Proxy Mobile IPv6 serve as
   multicast subscription anchor points, while Mobile Access Gateways
   provide Multicast Listener Discovery (MLD) proxy functions. In this
   scenario, mobile nodes remain agnostic of multicast mobility
   operations. Support for mobile multicast senders is outside the
   scope of this document.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

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

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

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

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

2. Terminology

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

3. Overview

   The reference scenario for multicast deployment in Proxy Mobile IPv6
   domains is illustrated in Figure 1. Below, LMAA and MN-HNP are the
   LMA Address and Mobile Node's Home Network Prefix as defined in
   [RFC5213].

```
                        +-------------+
                        | Content     |
                        | Source      |
                        +-------------+
                               |
                     ***  ***  ***  ***
                    *    **   **   **    *
                    *                     *
                    *  Fixed Internet  *
                    *                     *
                    *    **   **   **    *
                     ***  ***  ***  ***
                     /                \
                 +----+            +----+
                 |LMA1|            |LMA2|            Multicast Anchor
                 +----+            +----+
             LMAA1  |                |  LMAA2
                    |                |
                    \\            //\\
                     \\          //  \\
                      \\       //     \\          Unicast Tunnel
                       \\     //       \\
                        \\   //         \\
                         \\ //           \\
             Proxy-CoA1 ||              ||  Proxy-CoA2
                     +----+            +----+
                     |MAG1|            |MAG2|            MLD Proxy
                     +----+            +----+
                      |  |              |
             MN-HNP1  |  | MN-HNP2      | MN-HNP3
                     MN1 MN2           MN3
```

   Figure 1: Reference Network for Multicast Deployment in PMIPv6

   An MN in a PMIPv6 domain will decide on multicast group membership
   management completely independent of its current mobility conditions.
   It will submit MLD Report and Done messages, based on application
   triggers, using its link-local source address and multicast
   destination addresses according to [RFC3810] or [RFC2710]. These

link-local signaling messages will arrive at the currently active MAG
via one of its downstream local (wireless) links. A multicast-
unaware MAG would simply discard these MLD messages.

To facilitate multicast in a PMIPv6 domain, an MLD proxy function
[RFC4605] needs to be deployed on the MAG that selects the tunnel
interface corresponding to the MN's LMA for its upstream interface
(cf., Section 6 of [RFC5213]). Thereby, each MAG-to-LMA tunnel
interface defines an MLD proxy domain at the MAG, and it contains all
downstream links to MNs that share this specific LMA. According to
standard proxy operations, MLD Report messages will be aggregated and
then forwarded up the tunnel interface to the MN's corresponding LMA.

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

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

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

In summary, the following steps are executed on handover:

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

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

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

4. The MAG should determine whether the MN is admissible to multicast services; if it's not, then stop here.

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

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

7. The MN Membership Reports arrive at the MAG, in response either to the early query or to the query sent by the proxy instance.

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

After Re-Binding, the LMA is not required to issue a MLD General Query on the tunnel link to refresh forwarding states. Multicast state updates should be triggered by the MAG, which aggregates subscriptions of all its MNs (see the call flow in Figure 2).
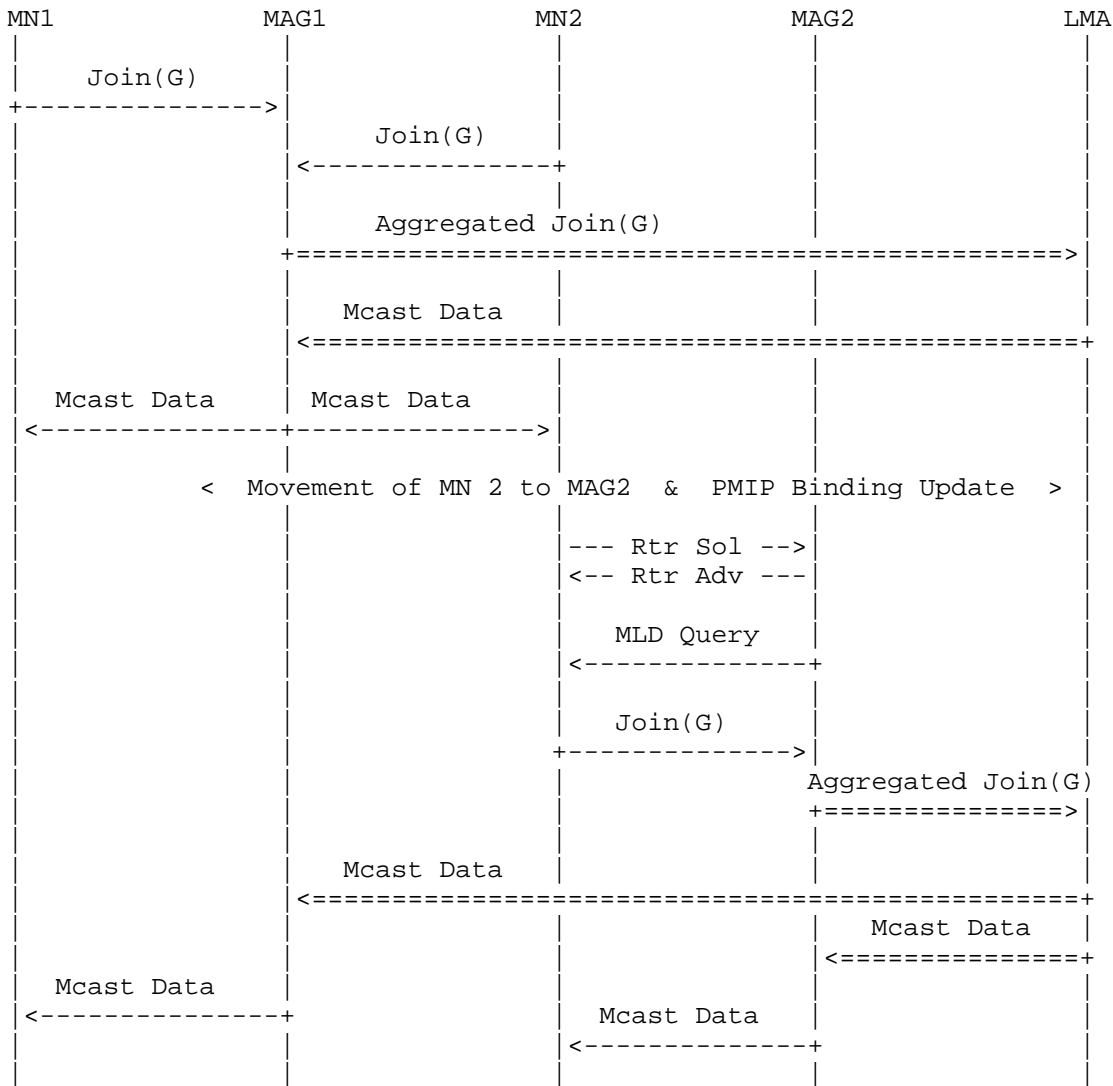
```
    MN1             MAG1            MN2            MAG2             LMA
     |               |               |               |               |
     |   Join(G)     |               |               |               |
     +-------------->|               |               |               |
     |               |   Join(G)     |               |               |
     |               |<--------------+               |               |
     |               |               |               |               |
     |               |     Aggregated Join(G)         |               |
     |               +===================================================>|
     |               |               |               |               |
     |               |   Mcast Data  |               |               |
     |               |<===================================================+
     |               |               |               |               |
     |  Mcast Data   |  Mcast Data   |               |               |
     |<--------------+-------------->|               |               |
     |               |               |               |               |
     |          <  Movement of MN 2 to MAG2  &  PMIP Binding Update  > |
     |               |               |               |               |
     |               |               |--- Rtr Sol -->|               |
     |               |               |<-- Rtr Adv ---|               |
     |               |               |               |               |
     |               |               |   MLD Query   |               |
     |               |               |<--------------+               |
     |               |               |               |               |
     |               |               |   Join(G)     |               |
     |               |               +-------------->|               |
     |               |               |               |  Aggregated Join(G)
     |               |               |               +===============>|
     |               |               |               |               |
     |               |   Mcast Data  |               |               |
     |               |<===================================================+
     |               |               |               |   Mcast Data  |
     |               |               |               |<==============+
     |  Mcast Data   |               |               |               |
     |<--------------+               |   Mcast Data  |               |
     |               |               |<--------------+               |
     |               |               |               |               |
```

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

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

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

4. Deployment Details

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

4.1. Operations of the Mobile Node

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

4.2. Operations of the Mobile Access Gateway

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

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

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

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

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

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

4.3. Operations of the Local Mobility Anchor

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

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

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

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

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

4.4. IPv4 Support

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

multicast support on the network side, an IGMP proxy function needs
to be deployed at MAGs in exactly the same way as for IPv6.
[RFC4605] defines IGMP proxy 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 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 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.

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

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

4.5. Multihoming Support

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

multicast mobility operations described in this document will
transparently preserve the association of channels to interfaces in
the following way.

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

4.6. Multicast Availability throughout the Access Network

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

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

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

4.7. A Note on Explicit Tracking

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

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

5. Message Source and Destination Address

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

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

5.1. Query

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

5.2. Report/Done

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

6. Security Considerations

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

7. Acknowledgements

   This memo follows initial requirements work presented in "Multicast
   Support Requirements for Proxy Mobile IPv6" (July 2009), and is the
   outcome of extensive previous discussions and a follow-up of several
   initial documents on the subject. The authors would like to thank
   (in alphabetical order) Jari Arkko, Luis M. Contreras, Greg Daley,
   Gorry Fairhurst, Dirk von Hugo, Liu Hui, Seil Jeon, Jouni Korhonen,
   Guang Lu, Sebastian Meiling, Akbar Rahman, Imed Romdhani, Behcet
   Sarikaya, Pierrick Seite, Stig Venaas, and Juan Carlos Zuniga for
   advice, help, and reviews of the document. Funding by the German
   Federal Ministry of Education and Research within the G-LAB
   Initiative is gratefully acknowledged.

8. References

8.1. Normative References

   [RFC2710]   Deering, S., Fenner, W., and B. Haberman, "Multicast
               Listener Discovery (MLD) for IPv6", RFC 2710,
               October 1999.

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

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

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

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

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

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

8.2. Informative References

   [AUTO-MULTICAST]
              Thaler, D., Talwar, M., Aggarwal, A., Vicisano, L., and T.
              Pusateri, "Automatic IP Multicast Without Explicit Tunnels
              (AMT)", Work in Progress, March 2010.

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

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

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

Appendix A. Initial MLD Queries on Upcoming Links

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

   Experimental tests on Linux and Cisco systems have revealed immediate
   IGMP Queries followed a link trigger event (within a fraction of 1
   ms), while MLD queries immediately followed the autoconfiguration of
   IPv6 link-local addresses at the corresponding interface.

Appendix B. State of IGMP/MLD Proxy Implementations

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

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

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

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

   Linux gproxy: IGMPv3 Proxy implementation that permits configuration
      of the upstream interface, only. Downstream interfaces are
      collected at startup without dynamic extension of this list. No
      support of multiple instances or MLD.

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

Appendix C. Comparative Evaluation of Different Approaches

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

   Neither approach establishes native multicast distribution between
   the LMA and MAG; instead, they use tunneling mechanisms. In scenario
   A, a MAG is connected to different multicast-enabled LMAs and can
   receive the same multicast stream via multiple paths depending on the
   group subscriptions of MNs and their associated LMAs. This problem,
   a.k.a. the tunnel convergence problem, may lead to redundant traffic
   at the MAGs. In contrast, scenario B configures MAGs to establish a
   tunnel to a single, dedicated multicast LMA for all attached MNs and
   relocates overhead costs to the multicast anchor. This eliminates
   redundant traffic but may result in an avalanche problem at the LMA.

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

   1. All MNs participate in distinct multicast groups.

   2. All MNs join the same multicast group.

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

Setting 1:

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

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

Setting 2:

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

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

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

Authors' Addresses

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

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

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

   EMail: mw@link-lab.net

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

   EMail: suresh.krishnan@ericsson.com