

Netext WG  
Internet-Draft  
Intended status: Standards Track  
Expires: January 6, 2011

S. Krishnan, Ed.  
Ericsson  
R. Koodli  
Cisco Systems  
P. Loureiro  
NEC  
Q. Wu  
Huawei  
A. Dutta  
Telcordia  
July 5, 2010

**Localized Routing for Proxy Mobile IPv6  
draft-krishnan-netext-pmip-lr-02**

**Abstract**

Proxy Mobile IPv6 (PMIPv6) is a network based mobility management protocol that enables IP mobility for a host without requiring its participation in any mobility-related signaling. PMIPv6 requires all communications to go through the local mobility anchor. As this can be suboptimal, localized routing allows mobile nodes attached to the same or different mobile access gateways to exchange traffic by using localized forwarding or a direct tunnel between the gateways. This document proposes an initiation mechanism for localized routing.

**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 6, 2011.

**Copyright Notice**

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Initiation of Localized Routing . . . . .</a>	<a href="#">4</a>
<a href="#">2.1.</a>	<a href="#">MAG behavior . . . . .</a>	<a href="#">4</a>
<a href="#">2.2.</a>	<a href="#">LMA behavior . . . . .</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Conventions used in this document . . . . .</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Scenario A11: Two MNs attached to the same MAG and LMA . . . . .</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">Handover Considerations . . . . .</a>	<a href="#">8</a>
<a href="#">5.</a>	<a href="#">Scenario A21: Two MNs attached to different MAGs but same LMA . . . . .</a>	<a href="#">9</a>
<a href="#">5.1.</a>	<a href="#">Handover Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">6.</a>	<a href="#">Scenario A12: Two MNs attached to the same MAG with different LMAs . . . . .</a>	<a href="#">12</a>
<a href="#">6.1.</a>	<a href="#">Handover Considerations . . . . .</a>	<a href="#">13</a>
<a href="#">7.</a>	<a href="#">Scenario A22: Two MNs attached to the different MAGs with different LMAs . . . . .</a>	<a href="#">14</a>
<a href="#">8.</a>	<a href="#">IPv4 support in Localized Routing . . . . .</a>	<a href="#">15</a>
<a href="#">9.</a>	<a href="#">Message Formats . . . . .</a>	<a href="#">16</a>
<a href="#">9.1.</a>	<a href="#">Localized Routing Initiation (LRI) . . . . .</a>	<a href="#">16</a>
<a href="#">9.2.</a>	<a href="#">Localized Routing Acknowledgment (LRA) . . . . .</a>	<a href="#">17</a>
<a href="#">10.</a>	<a href="#">New Mobility Option . . . . .</a>	<a href="#">19</a>
<a href="#">10.1.</a>	<a href="#">MAG IPv6 Address . . . . .</a>	<a href="#">19</a>
<a href="#">11.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">20</a>
<a href="#">12.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">21</a>
<a href="#">13.</a>	<a href="#">Authors . . . . .</a>	<a href="#">22</a>
<a href="#">14.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">23</a>
<a href="#">15.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">24</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">25</a>



## 1. Introduction

Proxy Mobile IPv6 [[RFC5213](#)] describes the protocol operations to maintain reachability and session persistence for a Mobile Node (MN) without the explicit participation from the MN in signaling operations at the Internet Protocol (IP) layer. In order to facilitate such network-based mobility, the PMIPv6 protocol defines a Mobile Access Gateway (MAG), which acts as a proxy for the Mobile IPv6 [[RFC3775](#)] signaling, and the Local Mobility Anchor (LMA) which acts similar to a Home Agent. The LMA and the MAG establish a bidirectional tunnel for forwarding all data traffic belonging to the Mobile Nodes. In the case where both endpoints are located in the same PMIPv6 domain, this can be suboptimal and results in higher delay and congestion in the network. Moreover, it increases transport costs and traffic load at the LMA.

To overcome these issues, localized routing can be used to allow nodes attached to the same or different MAGs to directly exchange traffic by using localized forwarding or a direct tunnel between the gateways. [LR-PS] defines the problem statement for PMIPv6 localized routing. This document describes a solution for PMIPv6 localized routing. The protocol specified here assumes that each MN is attached to a MAG and that each MN's MAG has established a binding for the attached MN at its selected LMA according to [[RFC5213](#)].

## **2. Initiation of Localized Routing**

Since the traffic to be localized passes through both the LMA and the MAGs, it is possible, at least in some scenarios, for either of them to initiate LR. In order to eliminate ambiguity, the protocol described in this document selects the initiator of the LR based on the following rules.

### **2.1. MAG behavior**

The MAG MUST Initiate LR if both the communicating MNs are attached to it and the MNs are anchored at different LMAs. The MAG MUST NOT initiate LR in any other case.

### **2.2. LMA behavior**


The LMA MUST Initiate LR if both the communicating MNs are anchored to it. The LMA MUST NOT initiate LR in any other case.

### **3. Conventions used in this document**


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

#### 4. Scenario A11: Two MNs attached to the same MAG and LMA

In this scenario, the two Mobile Nodes involved in communication are attached to a single MAG and both are anchored at the same LMA.



The LMA initiates a localized routing session by detecting a flow between two MNs attached to the same MAG. The exact flow identification mechanism is not specified in this document, and is left open for implementations and specific deployments. An example trigger could be that an application-layer signaling entity detects the possibility of localized routing and notifies the LMA about the two flow end-points, and the LMA determines that the two end-points are attached to the same MAG. Such a trigger mechanism offers localized routing at the granularity of an individual application session, providing flexibility in usage. It is also possible that one of the mobility entities (LMA or MAG) could decide to initiate localized routing based on configured policy.



Opt1: R=0,S=0,MN1-ID,MN1-HNP,MN2-ID,MN2-HNP

Opt2: R=1,U=0,MN1-ID,MN1-HNP,MN2-ID,MN2-HNP

After detecting a possibility for localized routing, the LMA constructs a Localized Routing Initiation (LRI) message that is used to signal the intent to initiate localized routing and to convey parameters for the same. This is a Mobility Header message and it contains the MN-Identifier and the Home Network Prefix (as Mobility Header options) for each of the MNs involved. The LMA sends the LRI message to the MAG where the two MNs are attached.

The MAG starts by verifying that the two MNs are indeed attached to it. It then verifies if the EnableMAGLocalRouting flag is set to 1. If it is not, the MAG is not configured to allow localized routing and it will reject the LRI and send an LRA with status code "Localized Routing Not Allowed". It then creates Localized Routing Entries (LREs) for each direction of the communication between the two MNs. The exact form of the forwarding entries is left for the implementations to decide; however, they should contain the HNP corresponding to the destination IP address and a next-hop identifier (e.g. the layer 2 address of the next-hop). These LREs MUST override the BUL entries for the specific HNPs identified in the LRI message. Hence all traffic matching the HNPs is forwarded locally.





If a MAG is unable to make deliver packets using the LREs, it is possible that the MN is no longer attached to the MAG. Hence, the MAG SHOULD fall back to using the BUL entry, and the LMA MUST forward the received packets using its BCE.

The local forwarding is not permanent. For instance, the LMA may send a LRI message with a request to cancel an existing local forwarding service. The local forwarding also has a default lifetime, upon the expiry of which, the forwarding reverts to bidirectional tunneling. When local forwarding service ceases, the corresponding LFE entries MUST be removed.

The MAG completes the processing of the LRI message and responds with a Local Routing Acknowledgment (LRA) message. This Mobility Header message also includes the MN-ID and the HNP for each of the communicating MNs as well as an appropriate Status code indicating the outcome of LRI processing. Status code 0 indicates localized routing was successfully offered by the MAG. Any other value for Status code indicates the reason for the failure to offer localized routing service. When Status code is 0, the LMA sets a flag in the BCE corresponding to the HNPs to record that localized routing is in progress for that HNP.


The MAG may refresh the lifetime of an existing local forwarding service. For this, it sends an unsolicited LRA (U-LRA) message that contains the new lifetime value. The MAG MUST wait for the following LRI message from the LMA before it can conclude that the refresh request is granted.

#### **4.1. Handover Considerations**

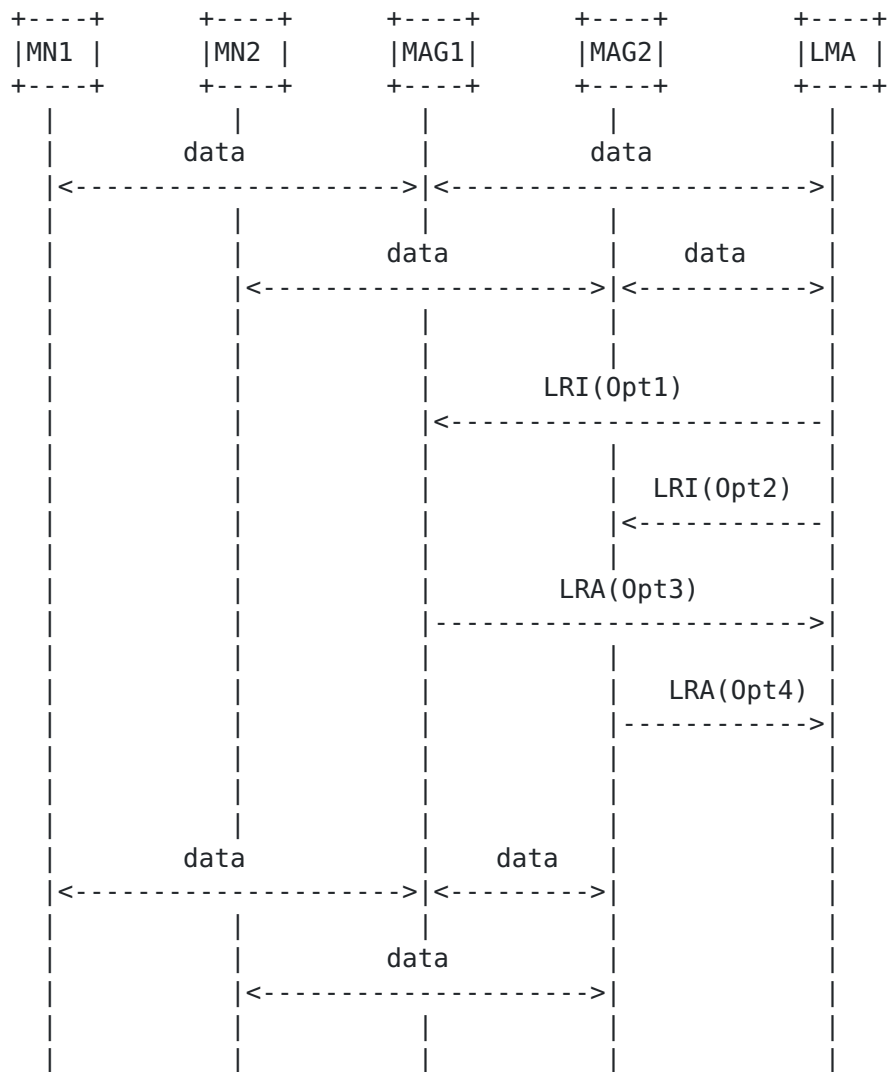
If one of the MNs, say MN1, detaches from the MAG and attaches to another MAG(say nMAG) the localized routing state needs to be re-established. When the LMA receives the PBU from nMAG for MN1, it will see that localized routing is active for for MN1. It will hence initiate LR at nMAG and update the LR state of MAG. After the handover completes, the localized routing will resemble Scenario A21.

## 5. Scenario A21: Two MNs attached to different MAGs but same LMA

The LMA may choose to support local forwarding to mobile nodes attached to two different MAGs within a single PMIPv6 domain.



As earlier, the LMA initiates LRI as a response to some trigger mechanism. In this case, however, it sends two separate LRI messages to the two MAGs. In addition to the MN-ID and the HNP options, each LRI message contains the IP Address of the counterpart MAG. When the MAG IP Address option is present, each MAG MUST create a local forwarding entry such that the packets for the MN attached to the remote MAG are sent over a tunnel associated with that remote MAG. The tunnel between the MAGs is assumed to be established by means outside the scope of this document.



Opt1: R=0,S=0,MN1-ID,MN1-HNP,MAG2-IPv6-Address

Opt2: R=0,S=0,MN2-ID,MN2-HNP,MAG1-IPv6-Address

Opt3: R=1,U=0,MN1-ID,MN1-HNP,MAG2-IPv6-Address

Opt4: R=1,U=0,MN2-ID,MN2-HNP,MAG1-IPv6-Address

As before, each MAG responds to the LRI with an LRA message. Barring the error cases, all subsequent packets are routed between the MAGs locally, without traversing the LMA.

The protocol does not require any synchronization between the MAGs before local forwarding begins. Each MAG begins its local forwarding independent of the other.



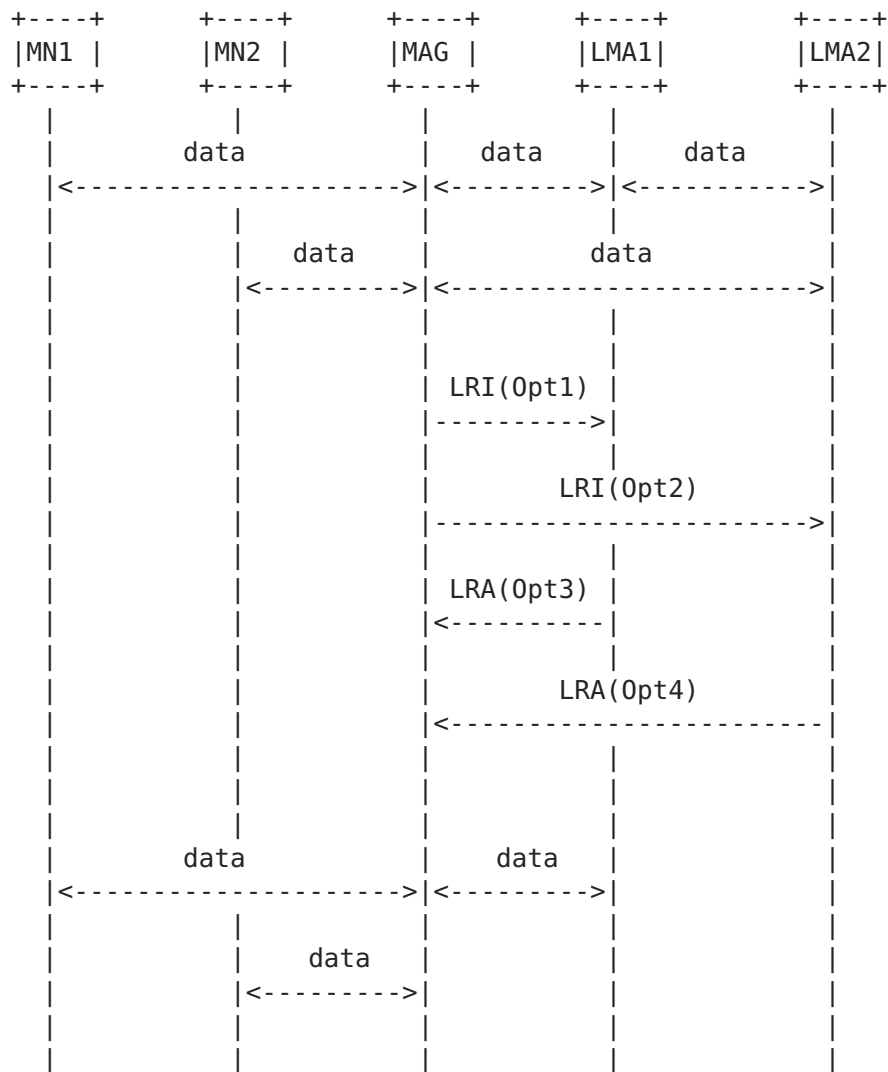
### **5.1. Handover Considerations**

If one of the MNs, say MN1, detaches from its current MAG (in this case MAG1) and attaches to another MAG (say nMAG1) the localized routing state needs to be re-established. When the LMA receives the PBU from nMAG1 for MN1, it will see that localized routing is active for MN1. It will hence initiate LR at nMAG1 and update the LR state of MAG2 to use nMAG1 instead of MAG1.

The MAG sends an LRI message containing the MN-ID, HNP and the counterpart LMA address to each LMA. Each LMA makes decision to support local forwarding independently, based on, among others, policy configuration for the counterpart LMA. Each LMA MUST respond to the LRI message with an LRA message. Only after it receives both the LRA messages each with Status value set to zero (success) from the two different LMAs, the MAG MUST conclude that it can provide local forwarding support for the two Mobile Nodes.







Opt1: R=0,S=0,MN1-ID,MN1-HNP

Opt2: R=0,S=0,MN2-ID,MN2-HNP

Opt3: R=1,U=0,MN1-ID,MN1-HNP

Opt4: R=1,U=0,MN2-ID,MN2-HNP

### 6.1. Handover Considerations

If one of the MNs, say MN1, detaches from its current MAG (in this case MAG1) and attaches to another MAG (say nMAG1) the localized routing state needs to be re-established. After the handover completes, the localized routing will resemble Scenario A22.



## **7. Scenario A22: Two MNs attached to the different MAGs with different LMAs**

This scenario will not be covered in this document since PMIPv6 does not define any form of inter-LMA communications. When a supported scenario, such as Scenario A12, morphs into Scenario A22 the node that initiated the localized routing session SHOULD tear it down in order to prevent lasting packet loss. This can result in transient packet loss when routing switches between the localized path into the normal path through the LMAs. In applications that are loss sensitive, this can lead to observable service disruptions. In deployments where Scenario A22 is possible, it is recommended that localized routing not be initiated when packet-loss-sensitive applications are in use.

## **8. IPv4 support in Localized Routing**

PMIPv6 MNs can use an IPv4 HoA as described in [[I-D.ietf-netlmm-pmip6-ipv4-support](#)]. In order to support the setup and maintenance of localized routes for these IPv4 HoAs in PMIPv6, MAGs must add the IPv4 HoAs into their LREs. The MAGs MUST also support encapsulation of IPv4 packets as described in [[I-D.ietf-netlmm-pmip6-ipv4-support](#)]. The localized routing protocol messages MUST include a IPv4 HoA option in their signaling messages in order to support IPv4 addresses for localized routing.

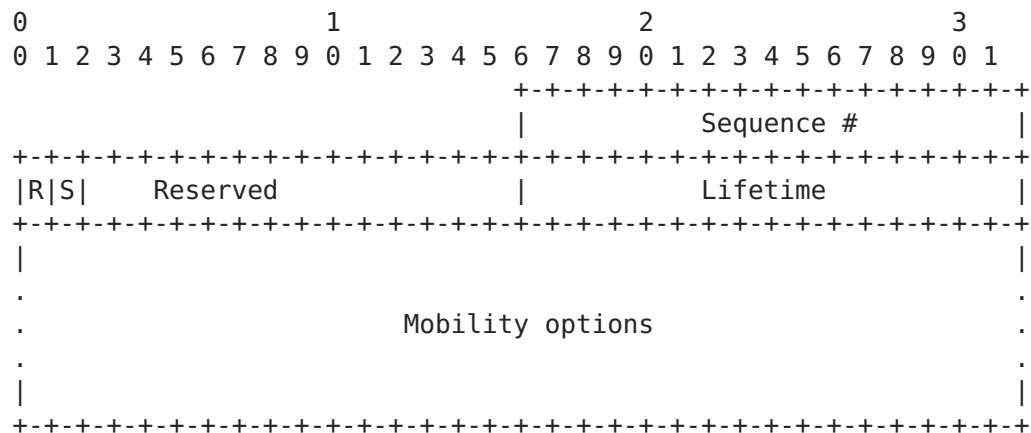
If the transport network between the PMIPv6 entities involved in localized routing is IPv4-only, the LRI and LRA messages MUST be encapsulated similar to the PBU/PBA messages as specified in [[I-D.ietf-netlmm-pmip6-ipv4-support](#)]. The encapsulation mode used SHOULD be identical to the one used to transport PBU and PBA messages.

## **9. Message Formats**

All the Localized routing messages use a new mobility header type (TBA1). The LRI message requests creation or deletion of localized routing state and the LRA message acknowledges the creation or deletion of such localized routing state.

### **9.1. Localized Routing Initiation (LRI)**

The LMA sends an LRI message to a MAG to request local forwarding for a pair of MNs. The MAG may also send this message to request the two LMAs for offering local forwarding as described in [Section 6](#) .



Sequence Number: A monotonically increasing integer. Set by a sending node in a request message, and used to match a reply to the request.

'R' flag: Set to 0, indicates it is an LRI message.

'S' flag: When set to 1, indicates a request to stop local routing.

Reserved: This field is unused. MUST be set zero.

Lifetime: The requested time in seconds for which the sender wishes to have local forwarding. A value of 0xffff (all ones) indicates an infinite lifetime.

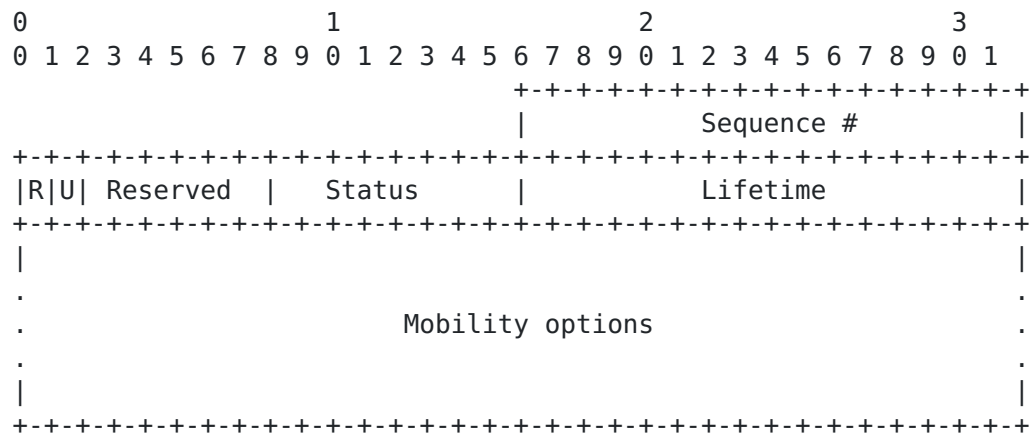
Mobility Options: MUST contain the MN-ID, followed by one or more HNPs for each of the MNs. For instance, for Mobile Nodes MN-1 and MN-2 with identifiers MN1-ID, MN2-ID and Home Network Prefixes MN1-HNP and MN2-HNP, the following tuple in the following order MUST be present: [MN1-ID, MN1-HNP], [MN2-ID, MN2-HNP]. The MN-ID and HNP options are the same as in [\[RFC5213\]](#). MAY contain the remote MAG IPv6 address option, which is identical to the HNP option except for Prefix Length equal to 128 bits.

The LRI message SHOULD be re-transmitted if a corresponding LRA message is not received within LRA\_WAIT\_TIME time units, up to a maximum of LRI\_RETRIES, each separated by LRA\_WAIT\_TIME time units.

## 9.2. Localized Routing Acknowledgment (LRA)

A MAG sends an LRA message to the LMA as a response to the LRI message. An LMA may also send this message to a MAG as a response to the LRI message as described in [Section 6](#).





Sequence Number: is copied from the sequence number field of the LRI message being responded to.

'R' flag: Set to 1, indicates it is an LRA message.

'U' flag: When set to 1, the LRA message is sent unsolicited. The Lifetime field indicates a new requested value. The MAG MUST wait for the regular LRI message to confirm that the request is acceptable to the LMA.

Reserved: This field is unused. MUST be set zero.

Status:

0: Success

128: Localized Routing Not Allowed

129: MN not attached

Lifetime: The time in seconds for which the local forwarding is supported. Typically copied from the corresponding field in the LRI message.

Mobility Options: When Status code is 0, MUST contain the [MN-ID, HNP] tuples in the same order as in the LRI message. When Status code is 1, MUST contain only those [MN-ID, HNP] tuples for which local forwarding is supported. The MN-ID and HNP options are the same as in [\[RFC5213\]](#).

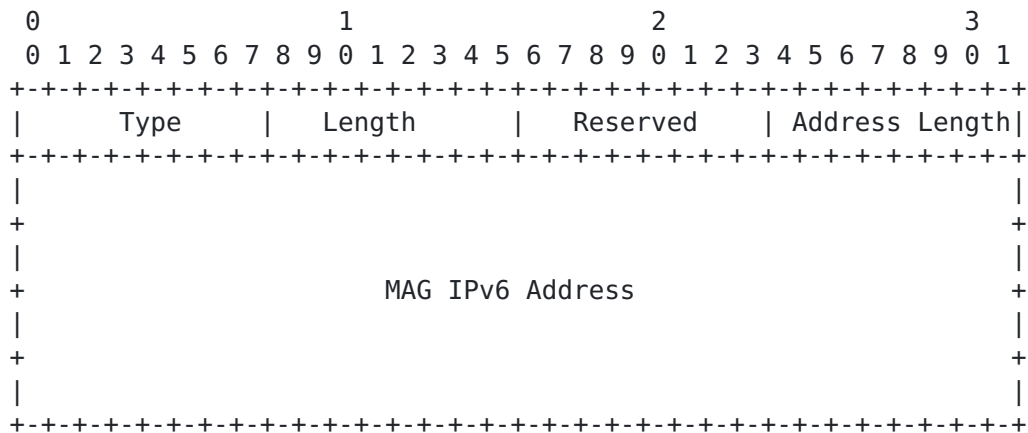




## 10. New Mobility Option

### 10.1. MAG IPv6 Address

The MAG IPv6 address mobility option contains the IPv6 address of a MAG involved in the localized routing. The MAG IPv6 address option has an alignment requirement of  $8n+4$ .

Type  
TBA3

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the type and length fields. This field **MUST** be set to 18.

Reserved (R)

This 8-bit field is unused for now. The value **MUST** be initialized to 0 by the sender and **MUST** be ignored by the receiver.

## Address Length

This field **MUST** be set to 128.

MAG IPv6 Address

A 16 byte field containing the MAG's IPv6 Address.

## **11. Security Considerations**

The protocol specified in this document uses the same security association between the LMA and the MAG to protect the LRI and LRA messages. No new security risks are identified. Support for integrity protection using IPsec is required, but support for confidentiality is not necessary.

## **12. IANA Considerations**

The Localized Routing Initiation, described in [Section 9.1](#) and the Local Routing Acknowledgment, described in [Section 9.2](#) require a single Mobility Header Type (TBA1) from the Mobility Header Types registry at <http://www.iana.org/assignments/mobility-parameters>

The MAG IPv6 Address and the LMA IPv6 Address require a Mobility Option Type each (TBA2 and TBA3) from the Mobility Options registry at <http://www.iana.org/assignments/mobility-parameters>

### **13. Authors**

This draft merges ideas from five different drafts addressing the PMIP localized routing problem. The authors of these drafts are listed below (in alphabetical order)

Kuntal Chowdhury <kchowdhury@starentnetworks.com>

Ashutosh Dutta <adutta@research.telcordia.com>

Rajeev Koodli <rkoodli@starentnetworks.com>

Suresh Krishnan <suresh.krishnan@ericsson.com>

Marco Liebsch <marco.liebsch@nw.neclab.eu>

Paulo Loureiro <paulo.loureiro@nw.neclab.eu>

Desire Oulai <desire.oulai@videotron.com>

Behcet Sarikaya <sarikaya@ieee.org>

Qin Wu <sunseawq@huawei.com>

Hidetoshi Yokota <yokota@kddilabs.jp>

#### **14. Acknowledgments**

The authors would like to thank Sri Gundavelli, Julien Abeille, Tom Taylor, Kent Leung, Mohana Jeyatharan, Jouni Korhonen, Glen Zorn, Ahmad Muhanna and Zoltan Turanyi for their comments and suggestions.

## **15. Normative References**

- [I-D.ietf-netext-pmip6-lr-ps]  
Liebsch, M., Jeong, S., and W. Wu, "PMIPv6 Localized Routing Problem Statement", [draft-ietf-netext-pmip6-lr-ps-00](#) (work in progress), September 2009.
- [I-D.ietf-netlmm-pmip6-ipv4-support]  
Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", [draft-ietf-netlmm-pmip6-ipv4-support-18](#) (work in progress), February 2010.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), August 2008.

Authors' Addresses

Suresh Krishnan (editor)  
Ericsson  
8400 Blvd Decarie  
Town of Mount Royal, Quebec  
Canada

Email: suresh.krishnan@ericsson.com

Rajeev Koodli  
Cisco Systems

Email: rkoodli@cisco.com

Paulo Loureiro  
NEC

Email: paulo.loureiro@nw.neclab.eu

Qin Wu  
Huawei

Email: Sunseawq@huawei.com

Ashutosh Dutta  
Telcordia

Email: adutta@research.telcordia.com