Netext Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2010

D. Oulai
S. Krishnan
Ericsson

October 26, 2009

Optimized Local Routing for PMIPv6
draft-oulai-netext-opt-local-routing-01.txt

Status of this Memo

This Internet-Draft will expire on April 29, 2010.

Copyright Notice

Abstract

Base Proxy Mobile IPv6 requires all communications to go through the local mobility anchor. As this can be suboptimal, local routing has been defined to allow mobile nodes attached to the same or different mobile access gateways to exchange traffic by using local forwarding or a direct tunnel between the gateways. This document proposes an initiation method and fast handover mechanisms for local routing. The solutions aim at reducing handover delay and packet loss.

Table of Contents

1. Introduction

Base Proxy Mobile IPv6 requires all communications to go through the LMA [RFC5213]. 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 this situation, local routing has been defined to allow nodes attached to the same or different mobile access gateways to exchange traffic by using local forwarding or a direct tunnel between the gateways. [I-D.ietf-netext-pmip6-lr-ps] details the local routing problem statement.

```
                           +----+
              +------------|LMA |----------+
              |            +----+          |
              |              |             |
              |              |             |
              |              |             |
              |              |             |
           +----+          +----+        +----+
           |MAG1|==========|MAG2|        |MAG3|
           +----+          +----+        +----+
              |                            |
              ==============================
           +----+          +----+
           |MN1 |          |MN2 | -------->
           +----+          +----+
```

Figure 1: Local Routing Scenario

Figure 1 defines a scenario where local routing could be used. MN1 and MN2 are communicating and are attached to the same PMIPv6 domain. MN1 is attached to MAG1 and MN2 is initially attached to MAG2. The LMA triggers the local routing by exchanging messages with the MAGs. A tunnel is created between MAG1 and MAG2.

Regarding handover, suppose MN2 moves from MAG2 to MAG3. Based on PMIPv6 specification, MN2 does an attach on MAG3. MAG3 performs authentication then sends a PBU to the LMA. If there is no local routing, the traffic is switched towards MAG3 at this point. However, local routing imply that traffic from MN1 do not reach LMA but is routed directly towards MAG2. Therefore, the LMA has to send a message to MAG1 in order to switch traffic. This results in worst handover performance than base PMIPv6.

This document proposes an initiation method and two fast handover mechanisms for PMIPv6 local routing. The solutions aim at reducing handover delay and packet loss. They reused some PFMIPv6 concepts [I-D.ietf-mipshop-pfmipv6] .

2. 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].

3. Terminology

All the general mobility-related terms used in this document are to be interpreted as defined in the Mobile IPv6 [RFC3775] and Proxy Mobile IPv6 [RFC5213] specifications as well as in [I-D.ietf-netext-pmip6-lr-ps].

Local Routing Indication (LRI): message sent to the MAG by the LMA or a peer MAG in a local routing session. This message triggers local routing.

Local Routing Acknowledgement (LRA): message sent by the MAG to the LMA or a peer MAG in a local routing session. This message acknowledge a LRI message.

4. Operations

The operations described here are based on the scenario presented in Figure 1.

4.1. Local Routing Initialization

When the LMA notices that Local Routing should be initiated between two MNs, it sends one LRI message to each MAG where the MNs are attached. The LRI messages contains information to identify both MNs, the target MAG which will be the other endpoint of the tunnel and some data needed to establish a security association between the MAGs. Note that we can have pre-established security associations between the MAGs of the same PMIPv6 domain. The LRI MAY also includes GRE Keys and a Tunnel mode option to indicate which type of tunneling SHOULD be used between the MAGs. The Tunnel mode is determined based on the information the LMA have on the MAGs capabilities. Note that the tunnel mode MAY also be preconfigured between MAGs.

If the MAG accepts to apply local routing, it sends a LRA message to the LMA to acknowledge the Local Routing. At this point, the LMA MUST not modify the binding of the MN as other communications MAY still go through the LMA. The LMA MUST record that there is an ongoing Local Routing session and also the involved MAGs and MNs. How the charging is done when local routing is in place is out of scope of this document.

Note also that this initiation mechanism can be extended to work in scenarios where the MAGs are attached to different LMAs. In this case, extra signalling will be required between the LMAs before they initiate the local routing process. How the LMAs discover each others is out of scope of this document.

```
+----+      +----+      +----+      +----+        +----+        +----+
|MN1 |      |MN2 |      |MAG1|      |MAG2|        |MAG3|        |LMA |
+----+      +----+      +----+      +----+        +----+        +----+
  |           |           |           |             |             |
  |        data           |           |  data       |             |
  |<--------------------->|<------------------------------------->|
  |           |           |           |             |             |
  |           |         data          |            data           |
  |           |<--------------------->|<------------------------->|
  |           |           |           |             |        LR decision
  |           |           | LRI(MN1,MN2,MAG2,Security,Tun_mode,GRE_Keys)
  |           |           |<--------------------------------------|
  |           |           |           |             |             |
  |           |           | LRI(MN2,MN1,MAG1,Security,Tun_mode,GRE_Keys)
  |           |           |           |<--------------------------|
  |           |           |           |             |             |
  |           |           |           |      LRA(MN1,MN2,MAG2)    |
  |           |           |-------------------------------------->|
  |           |           |           |             |             |
  |           |           |           |      LRA(MN2,MN1,MAG1)    |
  |           |           |           |-------------------------->|
  |           |           |           |             |             |
  |        data           |    data   |             |             |
  |<--------------------->|<--------->|             |             |
  |           |           |           |             |             |
  |           |         data          |             |             |
  |           |<--------------------->|             |             |
  |           |           |           |             |             |
  |           |           |           |             |             |
```

Figure 2: Local Routing Initialization

4.2. Handover

The handover mechanisms proposed here are proactive. They are based on the fact that the moving MN is able to send information to the MAG it is still attached to about the target access point or MAG. Those information are used by the old MAG to identify the target MAG. Note that a MN which does not support this specification performs base PMIPv6 handover and the LMA has to restart the Local Routing session later.

4.2.1. Non optimized handover

In this mechanism, before moving to MAG3, MN2 sends a Handover Initiate message to MAG2. This message MAY be L2 or L3 specific and

MUST provide MAG2 with useful information to resolve MAG3's address. For example, in 3GPP LTE networks, the Mobility Management Entity (MME) may generate this message towards the old Serving Gateway with the result that no modification is needed on the MN side. The details of this message is out of scope for this document.

If the target MAG is different from MAG2, MAG2 sends a refresh PBU message to the LMA containing a new flag indicating a possible movement. When receiving this PBU, the LMA MUST not update the BCE but rather abort the Local Routing sessions by sending a deregistration LRI message to MAG1. Therefore, MAG1 switches back the traffic towards the LMA. At this time, the traffic is asymmetric as the local routing from MAG2 to MAG1 is unaffected. When the LMA receives the PBU indicating the new attach point of MN2, it can start over a new Local Routing session between MN1 and MN2. Figure 3 shows this mechanism. Even if we refer to this method as non-optimized, it allows to reduce packet loss and handover delay. Another option could be for the LMA to send to MAG1 a LRI to directly set up the tunnel towards MAG3.

Note also that this method can be applied without any specific fast mobility mechanisms. The difference will be that instead of resolving MAG3's address, MAG2 detects MN2's detach and send a deregistration PBU as recommended by [RFC5213].

```
 +----+     +----+      +----+      +----+      +----+      +----+
 |MN1 |     |MN2 |      |MAG1|      |MAG2|      |MAG3|      |LMA |
 +----+     +----+      +----+      +----+      +----+      +----+
   |          |           |           |           |           |
   |        data          |    data   |           |           |
   |<-------------------->|<--------->|           |           |
   |          |        data           |           |           |
   |          |<--------------------->|           |           |
   |    HO decision       |           |           |           |
   |        HO_Init (MAG3 or Access Point)        |           |
   |          |---------------------->|           |           |
   |          |           |           |          PBU          |
   |          |           |           |---------------------->|
   |          |           |           |           |           |
   |          |           |   Dereg LRI(MN1,MN2,MAG2)         |
   |          |           |<----------------------------------|
   |          |           |    LRA(MN1,MN2,MAG2)              |
   |          |           |---------------------------------->|
   |        data          |           |  data     |           |
   |--------------------->|---------------------------------->|
   |          |         data          |      data             |
   |          |<----------------------|<----------------------|
   |          |            data       |           |           |
   |          |---------------------->|           |           |
   |         data         |    data   |           |           |
   |<---------------------|<----------|           |           |
   |          |           |           |           |           |
   |          |           | Attach    |           |           |
   |          |<--------------------------------->|           |
   |          |           |           |           |    PBU    |
   |          |           |           |           |---------->|
   |          |           |           |           |           |
   |          |           | LRI(MN1,MN2,MAG3,Security,Tun_mode,GRE_Keys)
   |          |           |<----------------------------------|
   |          |           |          LRA(MN1,MN2,MAG3)        |
   |          |           |---------------------------------->|
   |          |           | LRI(MN2,MN1,MAG1,Security,Tun_mode,GRE_Keys)
   |          |           |           |           |<----------|
   |          |           |           |          LRA(MN1,MN2,MAG1)
   |          |           |           |           |---------->|
   |        data          |         data          |           |
   |<-------------------->|<--------------------->|           |
   |          |           |     data  |           |           |
   |          |<--------------------------------->|           |
   |          |           |           |           |           |
```

Figure 3: Non optimized handover

4.2.2. Optimized handover

We assume here that Local Routing sessions MUST be maintained after handover. Before moving to MAG3, MN2 sends a Handover Initiate message to MAG2. MAG2 resolves MAG3 address. Then it sends a LRI message to inform MAG1 that MN2 will be attached to MAG3. It also sends a LRI message to inform MAG3 to start a Local Routing session with MAG1. To perform this, there need to be SA between each MAG and all its neighbour which support Local Routing and where a MN may move. This SA is used to exchange LRI/LRA messages between MAG2 and MAG3. After this, the Local Routing session is established between MAG1 and MAG3. MN2 can resume communications right after attachment with MAG3 if MAG3 advertises MN2's prefix before completing the PBU/PBA exchange or packet may be buffered until the PBA is received by MAG3. Figure 4 shows this mechanism. If this mechanism is used, MAG3 MUST include an indication in the PBU to let the LMA know that the LR session is now between MAG1 and MAG3.

In this method, MAG2 suggest a tunnelling mode and GRE Keys to MAG1 and MAG3. MAG2 SHOULD suggest the same tunnelling mode and GRE keys he was using unless he has specific information allowing a more accurate choice. In any case, if MAG1 and MAG3 rejects the suggestion of MAG2, they MUST directly negotiate the tunnelling mode and the GRE keys.

Another option could be for the MAG2 to send a LRI only to MAG3 in order to establish temporary Local Routing session between MAG2 and MAG3. Therefore, packets follow the path MAG1-MAG2-MAG3 during handover and the normal Local Routing path is restored after handover. However, this imply an additional tunnelling step and requires the LMA to establish a direct LR session between MAG1 and MAG3.

In case the MAGs are attached to different LMAs, extra inter-LMA signalling is required to updated the LR states in each LMA.

```
    +----+     +----+      +----+      +----+     +----+      +----+
    |MN1 |     |MN2 |      |MAG1|      |MAG2|     |MAG3|      |LMA |
    +----+     +----+      +----+      +----+     +----+      +----+
      |          |           |           |          |           |
      |        data          |    data   |          |           |
      |<-------------------->|<--------->|          |           |
      |          |           |           |          |           |
      |          |         data          |          |           |
      |          |<--------------------->|          |           |
      |     HO decision      |           |          |           |
      |          HO_Init (MAG3 or Access Point)     |           |
      |          |---------------------->|          |           |
      |          | LRI(MN1,MN2,MAG3,Security,Tun_mode,GRE_Keys) |
      |          |           |<----------|          |           |
      |          |         LRI(MN1,MN2,MAG1,Security,Tun_mode,GRE_Keys)
      |          |           |           |--------->|           |
      |          |          LRA(MN1,MN2,MAG3)       |           |
      |          |           |---------->|          |           |
      |          |           |         LRA(MN1,MN2,MAG1)        |
      |          |           |           |<---------|           |
      |          |           |   Attach  |          |           |
      |          |<-------------------------------->|           |
      |        data          |         data         |           |
      |<-------------------->|<-------------------->|           |
      |          |           |           |          |           |
      |          |           |     data  |          |           |
      |          |<-------------------------------->|           |
      |          |           |           |          |    PBU    |
      |          |           |           |          |---------->|
      |          |           |           |          |    PBA    |
      |          |           |           |          |<----------|
      |          |           |           |          |           |
```

Figure 4: Optimized handover

4.3. IPv4 considerations

In case of IPv4 Home Addresses used over IPv6 transport, the proposed mechanisms works well and the chosen tunnelling mode MUST be IPv4 over IPv6. The GRE Keys MUST also be present to differentiate private IPv4 HoAs.

If IPv4 transport is involved, the tunnelling mode suggested by MAG2 MAY not be accurate as NATs may be involved. In this case, unless there are preconfigured tunnels, the MAGs MUST directly negotiate the tunnelling mode.

5. Messages formats

This protocol requires two new messages, Local Routing Indication (LRI) and Local Routing Acknowledgement (LRA) messages. They use MH type (IANA-TBD). Figure 5 shows the MH. If the Local Routing type is set to 1, it is a Local Routing Indication message (Section 5.1). If it is set to 2, it is a Local Routing Acknowledgement message Section 5.2).

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Payload Proto |  Header Len   |   MH Type     |   Reserved    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Checksum            |    LR Type    |               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+               +
|                                                               |
.                  Local Routing Message Data                   .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Local Routing Message

Local Routing Type

8-bit unsigned integer. It defines the type of Local Routing message. Assigned values are:

0 Reserved
1 Local Routing Indication Message
2 Local Routing Acknowledgement Message

Local Routing Message Data

It follows the format defined for the specific Local Routing message (Section 5.1 and Section 5.2).

5.1. Local Routing Indication message

The LR type is set to 1. The format of the corresponding Local Routing Message Data is depicted in Figure 6. There is a 2n alignment for this message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |  LR  Type = 1 |A|  Reserved   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Sequence #          |         Lifetime              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                        Mobility options                       .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: Local Routing Indication Message

Acknowledge bit (A)

Set to 1 to request a LRA message from the target MAG.

Reserved

Unused bits. Must be set to 0.

Sequence #

A 16-bit unsigned integer used by the LMA or a MAG to match a LRI message with a LRA message.

Lifetime

A 16-bit unsigned integer used to define the lifetime of the Local Routing session.

Mobility options

Variable length field. The Mobility header MUST be an integer multiple of 8 octets long. There MUST be at least options present to provide the source MN's address, the destination MN's address and the target MAG address. This field can include options for security and tunneling mode indication.

5.2. Local Routing Acknowledge message

The LRA is sent by the target MAG in response to a LRI message with the A bit set. The LR type is set to 2. The format of the corresponding Local Routing Message Data is depicted in Figure 7. There is a 2n alignment for this message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |  LR  Type = 2 |    Reserved   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Sequence #          |         Lifetime              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                         Mobility options                      .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: Local Routing Acknowledgement Message

Reserved

Unused bits. Must be set to 0.

Sequence #

A 16-bit unsigned integer used by the requesting entity to match a LRI message with a LRA message. MUST be copied from the sequence # receive in the LRI message.

Lifetime

A 16-bit unsigned integer. MUST be copied from the lifetime receive in the LRI message.

Mobility options

Variable length field. The Mobility header MUST be an integer multiple of 8 octets long. The options present in the LRI message MUST be copied in the LRA message.

6. Security Considerations

Message between the LMA and the MAGs MUST be exchanged using the security associations existing for PBU/PBA messages. If inter-MAG signaling messages are used, there MUST be a SA between the MAGs as a malicious node can trigger a Local Routing session update.

7. IANA Considerations

This specification requires type to be assigned for LRI and LRA message.

8. Normative References

[I-D.ietf-mipshop-pfmipv6]
Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", draft-ietf-mipshop-pfmipv6-09 (work in progress), September 2009.

[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.

[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

Desire Oulai
Ericsson
8400 Blvd Decarie
Town of Mount Royal, Quebec
Canada

Email: desire.oulai@ericsson.com

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

Email: Suresh.Krishnan@ericsson.com