NETEXT Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2013

CJ. Bernardos, Ed.
UC3M

July 15, 2012

Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-ietf-netext-pmipv6-flowmob-04

Abstract

Proxy Mobile IPv6 allows a mobile node to connect to the same Proxy Mobile IPv6 domain through different interfaces. However, the ability of movement of selected flows from one access technology to another is missing in the basic Proxy Mobile IPv6 protocol. This document describes extensions to the Proxy Mobile IPv6 protocol that are required to support network based flow mobility over multiple physical interfaces.

The extensions required consist on the operations performed by the local mobility anchor and the mobile access gateway to manage the prefixes assigned to the different interfaces of the mobile node, as well as how the forwarding policies are handled by the network to ensure consistent flow mobility management.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

Table of Contents

1. Introduction

Proxy Mobile IPv6 (PMIPv6), specified in [RFC5213], provides network based mobility management to hosts connecting to a PMIPv6 domain. PMIPv6 introduces two new functional entities, the Local Mobility Anchor (LMA) and the Mobile Access Gateway (MAG). The MAG is the entity detecting Mobile Node's (MN) attachment and providing IP connectivity. The LMA is the entity assigning one or more Home Network Prefixes (HNPs) to the MN and is the topological anchor for all traffic belonging to the MN.

PMIPv6 allows a a mobile node to connect to the same PMIPv6 domain through different interfaces. This document specifies protocol extensions to Proxy Mobile IPv6 between the local mobility anchor and mobile access gateways to enable "flow mobility" and hence distribute specific traffic flows on different physical interfaces. It is assumed that the mobile node IP layer interface can simultaneously and/or sequentially attach to multiple MAGs, possibly over multiple media. One form to achieve this multiple attachment is described in [I-D.ietf-netext-logical-interface-support], which allows the mobile node supporting traffic flows on different physical interfaces regardless of the assigned prefixes on those physical interfaces.

In particular, this document specifies how to enable "flow mobility" in the PMIPv6 network (i.e., local mobility anchors and mobile access gateways). In order to do so, two main operations are required: i) proper prefix management by the PMIPv6 network, ii) consistent flow forwarding policies. This memo analyzes different potential use case scenarios, involving different prefix assignment requirements, and therefore different PMIPv6 network extensions to enable "flow mobility".

2. Terminology

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 [RFC2119].

The following terms used in this document are defined in the Proxy Mobile IPv6 [RFC5213]:

Local Mobility Agent (LMA).

Mobile Access Gateway (MAG).

Proxy Mobile IPv6 Domain (PMIPv6-Domain).

LMA Address (LMAA).

Proxy Care-of Address (Proxy-CoA).

Home Network Prefix (HNP).

The following terms used in this document are defined in the Multiple Care-of Addresses Registration [RFC5648] and Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support [RFC6089]:

Binding Identification Number (BID).

Flow Identifier (FID).

Traffic Selector (TS).

The following terms are defined and used in this document:

FMI (Flow Mobility Initiate). Message sent by the LMA to the MAG conveying the information required to enable flow mobility in a PMIPv6-Domain. This message is only needed when the prefixes initially assigned by the different MAGs to the mobile node are different.

FMA (Flow Mobility Acknowledgement). Message sent by the MAG in reply to an FMI message.

FMC (Flow Mobility Cache). Conceptual data structure maintained by the LMA and the MAG to support the flow mobility management operations described in this document.

3. Overview of the PMIPv6 flow mobility extensions

3.1. Use case scenarios

In contrast to a typical handover where connectivity to a physical medium is relinquished and then re-established, flow mobility assumes a mobile node can have simultaneous access to more than one network. In this specification, it is assumed that the local mobility anchor is aware of the mobile node's capabilities to have simultaneous access to both access networks and it can handle the same or a different set of prefixes on each access. How this is done is outside the scope of this specification.

There are different flow mobility scenarios. In some of them the mobile node might share a common set of prefixes among all its physical interfaces, whereas in others the mobile node might have a

different subset of prefixes configured on each of the physical interfaces. The different scenarios are the following:

1. At the time of a new network attachment, the MN obtains the same prefix or the same set of prefixes as already assigned to an existing session. This is not the default behavior with basic PMIPv6 [RFC5213], and the LMA needs to be able to provide the same assignment even for the simultaneous attachment (as opposed to the handover scenario only).

2. At the time of a new network attachment, the MN obtains a new prefix or a new set of prefixes for the new session. This is the default behavior with basic PMIPv6 [RFC5213].

3. At the time of a new network attachment, the MN obtains a combination of prefix(es) in use and new prefix(es). This is a hybrid of the two above-mentioned scenarios. The local policy determines whether the new prefix is exclusive to the new attachment or it can be assigned to an existing attachment as well.

The operational description of how to enable flow mobility in each of these scenarios is provided in Section 3.2.1, Section 3.2.2 and Section 3.2.3.

The extensions described in this document support all the aforementioned scenarios.

3.2. Basic Operation

This section describes how the PMIPv6 extensions described in this document enable flow mobility support.

Both the mobile node and the local mobility anchor SHOULD have local policies in place that ensure packets are forwarded coherently for unidirectional and bidirectional communications. The details about how this consistency is ensured are out of the scope of this document.

3.2.1. MN sharing a common set of prefixes on all MAGs

This scenario corresponds to the use case scenario number 1 described in Section 3.1. This scenario needs extensions to basic PMIPv6 [RFC5213] signaling at the time of a new attachment, to ensure that the same prefix (or set of prefixes) is assigned to all the interfaces of the same mobile node that are simultaneously attached. Subsequently, no further signaling is necessary between the local mobility anchor and the mobile access gateway and flows are forwarded

according to policy rules on the local mobility anchor and the mobile node.

If the local mobility anchor assigns a common prefix (or set of prefixes) to the different physical interfaces attached to the domain, then all the MAGs already have all the routing knowledge required to forward uplink or downlink packets, and the local mobility anchor does not need to perform any kind of signaling in order to move flows across the different physical interfaces.

The local mobility anchor needs to know when to assign the same set of prefixes to all the different physical interfaces of the mobile node. This can be achieved by different means, such as policy configuration or default policies, etc. In this document a new Handoff Indicator (HI) value ("Attachment over a new interface sharing prefixes") is defined, to allow the mobile access gateway indicate to the local mobility anchor that the same set of prefixes MUST be assigned to the mobile node. The considerations of Section 5.4.1 of [RFC5213] are updated by this specification as follows:

- o If there is at least one Home Network Prefix option present in the request with a NON_ZERO prefix value, there exists a Binding Cache entry (with one all home network prefixes in the Binding Cache entry matching the prefix values of all Home Network Prefix options of the received Proxy Binding Update message), and the entry matches the mobile node identifier in the Mobile Node Identifier option of the received Proxy Binding Update message, and the value of the Handoff Indicator of the received Proxy Binding Update is equal to "Attachment over a new interface sharing prefixes".

    1. If there is a Mobile Node Link-layer Identifier Option present in the request and the Binding Cache entry matches the Access Technology Type (ATT), and MN-LL-Identifier, the request MUST be considered as a request for updating that Binding Cache entry.

    2. If there is a Mobile Node Link-layer Identifier Option present in the request and the Binding Cache entry does not match the Access Technology Type (ATT), and MN-LL-Identifier, the request MUST be considered as a request for creating a new mobility session sharing the same set of Home Network Prefixes assigned to the existing Binding Cache entry found.

    3. If there is not a Mobile Node Link-layer Identifier Option present in the request, the request MUST be considered as a request for creating a new mobility session sharing the same set of Home Network Prefixes assigned to the existing Binding

Cache entry found.

In case the mobile access gateways need to be configured to support flow mobility because of packet policing, packet enforcement, charging or similar reasons, the local mobility anchor MUST re-use the signaling defined later in this document to convey this information.

```
                                  LMA Binding Cache
                    +---+       =======================
                    |LMA|        MN1, if1, pref1, MAG1
                    +---+        MN1, if2, pref1, MAG2
                     //\\
          +---------//--\\-------------+
         (         //    \\             ) PMIPv6 domain
         (        //      \\            )
          +------//--------\\----------+
                //          \\
               //            \\
            +----+           +----+
            |MAG1|           |MAG2|
            +----+           +----+
              |                |
              |   +-------+    |
              |   |  I P  |    |
              |   +---+---+    |
              |---|if1|if2|----|
                  +---+---+
                     MN1
```

Figure 1: Shared prefix across physical interfaces scenario

Next, an example of how flow mobility works in this case is shown. In Figure 1, a mobile node (MN1) has two different physical interfaces (if1 and if2). Each physical interface is attached to a different mobile access gateway, both of them anchored and controlled by the same local mobility anchor. Since both physical interfaces are assigned the same prefix (pref1) upon attachment to the MAGs, the mobile node has one single IPv6 address configured at the IP layer: pref1::mn1. Initially, flow X goes through MAG1 and flow Y through MAG2. At certain point, flow Y can be moved to also go through MAG1. As shown in Figure 2, no signaling between the local mobility anchor and the mobile access gateways is needed.

```
                 +-----+         +------+        +------+      +-----+
   Internet      | LMA |         | MAG1 |        | MAG2 |      | MN1 |
                 +-----+         +------+        +------+      +-----+
      |             |               |               |             |
      |  flow X to  |   flow X to   |           flow X to         |
      |  pref1::mn1 |   pref1::mn1  |           pref1::mn1        |
      |<----------->|<------------->|<--------------------------->if1
      |  flow Y to  |           flow Y to           |  flow Y to  |
      |  pref1::mn1 |           pref1::mn1          |  pref1::mn1 |
      |<----------->|<----------------------------->|<----------->if2
      |             |               |               |             |
      |       ============          |               |       ============
      |       ||  flow  ||          |               |       ||  flow  ||
      |       || policy ||          |               |       || policy ||
      |       || update ||          |               |       || update ||
      |       ============          |               |       ============
      |             |               |               |             |
      |  flow Y to  |   flow Y to   |          flow Y to          |
      |  pref1::mn1 |   pref1::mn1  |          pref1::mn1         |
      |<----------->|<------------->|<--------------------------->if1
      |             |               |               |             |
```

Figure 2: Flow mobility message sequence with common set of prefixes

Figure 3 shows the state of the different network entities after moving flow Y in the previous example. This documents re-uses some of the terminology and mechanisms of the flow bindings and multiple care-of address registration specifications. Note, that in this case the BIDs shown in the figure are assigned locally by the LMA, since there is no signaling required in this scenario. In any case, alternative implementations of flow routing at the LMA could be used, as it does not impact on the operation of the solution in this case.

```
                           LMA Binding Cache        LMA flowmob state
                      (BID, MN-ID, ATT, HNP, PCoA)      (BID, TS)
                 +---+ ==========================  ===================
                 |LMA|  1, MN1, if1, pref1, MAG1       1, flow X
                 +---+  2, MN1, if2, pref1, MAG2       1, flow Y
                  //\\
       +---------//--\\-------------+
      (         //    \\             ) PMIPv6 domain
      (        //      \\            )
       +------//--------\\----------+
             //          \\
            //            \\       MAG1 routing state
         +----+           +----+  ================================
         |MAG1|           |MAG2|     (dest)         (next hop)
         +----+           +----+   pref1::/64   p2p-iface-with-MN1
           |                |         ::/0             LMA
           |                |
           |                |      MAG2 routing state
           |   +-------+    |     ================================
           |   |  I P  |    |        (dest)         (next hop)
           |   +---+---+    |      pref1::/64   p2p-iface-with-MN1
           |---|if1|if2|----|         ::/0             LMA
               +---+---+
                  MN1
```

Figure 3: Data structures with common set of prefixes

3.2.2. MN with different sets of prefixes on each MAG

A different flow mobility scenario happens when the local mobility anchor assigns different sets of prefixes to physical interfaces of the same mobile node. This covers the second and third use case scenarios described in Section 3.1. In this case, specific signaling is required between the local mobility anchor and the mobile access gateway to enable relocating flows between the different attachments, so the MAGs are aware of the prefixes for which the MN is going to receive traffic, and local routing entries are configured accordingly. Two different possibilities are considered next.

The first possibility corresponds to the use case scenario number 2 described in Section 3.1, in which a multi-interfaced mobile node obtains a different set of prefixes on each attachment. Signaling is required when a flow is to be moved from its original interface to a new one. Since the local mobility anchor cannot send a PBA message which has not been triggered in response to a received PBU message, new signaling messages are defined to cover this case. The trigger for the flow movement can be on the mobile node (e.g., by using layer-2 signaling, by explicitly start sending flow packets via a new

interface, etc.) or on the network (e.g., based on congestion and measurements performed at the network).

If the flow is being moved from its default path (which is determined by the destination prefix) to a different one, the local mobility anchor constructs a Flow Mobility Initiate (FMI) message. This message MUST be sent to the new target mobile access gateway, i.e. the one selected to be used in the forwarding of the flow. The FMI message contains (as explained in further detail in Section 4.1), the MN-Identifier, the Flow Identification Mobility option (specified in [RFC6089]) which can convey prefix or full flow information, and the type of flow mobility operation (add flow). By default, prefix information is provided. Full prefix granularity is non mandatory. Optionally, the local mobility anchor may send another FMI message, this time to remove the flow Y state at MAG2. Otherwise the flow state at MAG2 will be removed upon timer expiration. The message sequence is shown in Figure 4.

```
                 +-----+         +------+        +------+      +-----+
   Internet      | LMA |         | MAG1 |        | MAG2 |      | MN1 |
                 +-----+         +------+        +------+      +-----+
      |             |               |               |             |
      |  flow X to  |   flow X to   |           flow X to         |
      |  pref1::mn1 |   pref1::mn1  |           pref1::mn1        |
      |<----------->|<------------->|<-------------------------->if1
      |  flow Y to  |           flow Y to           |  flow Y to  |
      |  pref2::mn1 |           pref2::mn1          |  pref2::mn1 |
      |<----------->|<----------------------------->|<---------->if2
      |             |               |               |             |
      |       ============          |               |       ============
      |       ||  flow  ||          |               |       ||  flow  ||
      |       || policy ||          |               |       || policy ||
      |       || update ||          |               |       || update ||
      |       ============          |               |       ============
      |             |               |               |             |
      |             | FMI[MN1-ID,flow_info(Y),add]  |             |
      |             |-------------->|               |             |
      |             |          FMA  |               |             |
      |             |<--------------|               |             |
      |  flow Y to  |   flow Y to   |          flow Y to          |
      |  pref2::mn1 |   pref2::mn1  |          pref2::mn1         |
      |<----------->|<------------->|<-------------------------->if1
      |             |               |               |             |
      |             |           (optional)          |             |
      |             |FMI[MN1-ID,flow_info(Y),lft=0] |             |
      |             |------------------------------>|             |
      |             |               |         FMA   |             |
      |             |<------------------------------|             |
      |             |               |               |             |
```

       Figure 4: Flow mobility message sequence when the LMA assigns
     different sets of prefixes per physical interface (FMI signaling)

The state in the network after moving a flow, for the case the LMA assigns a different set of prefixes is shown in Figure 5.

```
                           LMA Binding Cache          LMA flowmob state
                      (BID, MN-ID, ATT, HNP, PCoA)       (BID, TS)
                 +---+ ============================  ===================
                 |LMA|  1, MN1, if1, pref1,               1, flow X
                 +---+             pref2,  MAG1           1, flow Y
                  //\\  2, MN1, if2, pref2,  MAG2
       +---------//--\\-------------+
      (         //    \\             ) PMIPv6 domain
      (        //      \\            )
       +------//--------\\----------+
             //          \\
            //            \\       MAG1 routing state
          +----+           +----+  ================================
          |MAG1|           |MAG2|     (dest)         (next hop)
          +----+           +----+   pref1::/64   p2p-iface-with-MN1
            |                 |      pref2::/64   p2p-iface-with-MN1
            |                 |         ::/0             LMA
            |                 |
            |   +-------+    |      MAG2 routing state
            |   |  I P  |    |     ================================
            |   +---+---+    |        (dest)         (next hop)
            |---|if1|if2|----|      pref2::/64   p2p-iface-with-MN1
                +---+---+              ::/0             LMA
                   MN1
```

Figure 5: Data structures when the LMA assigns a different set of prefixes

The second possibility corresponds to the use case scenario number 3 described in Section 3.1, in which upon new physical interface attachment, the MN obtains a combination of prefix(es) in use and new prefix(es). Here, the mobile node is already attached to the PMIPv6-Domain via MAG1. At a certain moment, the mobile node attaches a new interface (if2) to MAG2. MAG2 sends a PBU which is then used by the LMA to enable flow mobility. In this case, we consider that flows are moved with a prefix granularity, meaning that flows are moved by moving prefixes among the different MAGs the mobile node is attached to. In this example, flow Y is bound to pref2::/64 and therefore the flow can be moved by just binding pref2::/64 to MAG2. This is done by including the prefix in the PBA message. The scenario is shown in Figure 6.

Optionally, a Binding Revocation Indication message [RFC5846] with the P bit set MAY be sent to MAG1 to indicate that this is a revocation of PMIP prefix(es). After processing BRI, the source MAG MUST send a Binding Revocation Acknowledgement (BRA) message back to the LMA.

```
                 +-----+           +-------+       +-------+     +-----+
   Internet      | LMA |           | MAG1 |        | MAG2 |      | MN  |
                 +-----+           +-------+       +-------+     +-----+
      |             |                 |               |             |
      |  flow X to  |    flow X to    |           flow X to         |
      |  pref1::mn1 |    pref1::mn1   |           pref1::mn1        |
      |<----------->|<--------------->|<-------------------------->if1
      |  flow Y to  |    flow Y to    |           flow Y to         |
      |  pref2::mn1 |    pref2::mn1   |           pref2::mn1        |
      |<----------->|<--------------->|<-------------------------->if1
      |             |                 |               |             |
      |             |                 |               |             |
      |             |                 |            MN powers on if2 and
      |             |                 |           performs L2 attachment
      |             |                 |               |<-----------if2
      |             |                 |          PBU  |             |
      |             |<--------------------------------|             |
      |             |   PBA (pref2)   |               |             |
      |             |-------------------------------->|             |
      |     LMA moves pref2 to new    |               |             |
      |  binding cache entry for if2  |               |             |
      |             |                 |               |             |
      |  flow y to  |             flow y to           |  flow y to  |
      |  pref2::mn1 |             pref2::mn1          |  pref2::mn1 |
      |<----------->|<------------------------------->|<---------->if2
      |             |                 |               |             |
      |             |   (optional)    |               |             |
      |             |   BRI[pref2]    |               |             |
      |             |---------------->|               |             |
      |             |       BRA       |               |             |
      |             |<----------------|               |             |
      |             |                 |               |             |
```

      Figure 6: Flow mobility message sequence with different set of
              prefixes per physical interface (PBU signaling)

In case flow mobility is needed with a finer granularity (e.g., flow level instead of full prefix), a Flow Identification Mobility option (specified in [RFC6089]) that can convey full flow information MUST be included in the PBA. The MAG MAY also include the Flow Identification Mobility option in the PBU message that it sends to the LMA. This serves as a request from MAG to LMA to consider the flow policy rules specified in the option. In this case, no prefix is removed from any MAG because the movement is performed at a flow level.

3.2.3. MN with combination of prefix(es) in use and new prefix(es) on each MAG

This scenario is a hybrid of the ones described in Section 3.2.1 and Section 3.2.2. It requires flow mobility signaling to enable relocating flows for the new prefix(es) which are not shared across attachments.

4. Message formats

This section defines extensions to the Proxy Mobile IPv6 [RFC5213] protocol messages.

4.1. Flow Mobility Initiate (FMI)

The LMA sends an FMI message to a MAG to enable flow mobility. It is a Mobility Header 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
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |          Sequence #           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|I|        Reserved             |           Lifetime            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                        Mobility options                       .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Sequence Number:

A monotonically increasing integer. Set by the LMA sending then initiate message, and used to match a reply in the Acknowledgement.

'I' (initiate) flag:

Set to 1, indicates it is an FMI message.

Reserved:

This field is unused. MUST be set to zero by the sender.

Lifetime:

The requested time in seconds for which the LMA asks the MAG keep flow-specific state. A value of all one bits (0xffff) represents infinity. If set to 0, it indicates a request to remove state about the flow (cancel flow mobility)

Mobility Options:

MUST contain the MN-ID, followed by one or more Flow Identification Mobility options [RFC6089].

4.2. Flow Mobility Acknowledgement (FMA)

The MAG sends an FMI message to the LMA as a response to the FMI message. It is a Mobility Header 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
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |           Sequence #          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|I|  Reserved   |    Status     |           Lifetime            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                        Mobility options                       .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Sequence Number:

A monotonically increasing integer. Copied from the value set by the sending LMA in the FMI message being acknowledged by this FMA message.

'I' flag:

Set to 0, indicates it is an FMA message.

Reserved:

This field is unused. MUST be set to zero by the sender.

Status (values to be assigned by IANA):

??: Success.

??: Reason unspecified.

??: MN not attached.

??: Sequence number out of window.

??: Traffic Selector format unsupported.

??: No existing Flow Mobility Cache entry.

Lifetime:

The requested time in seconds for which the MAG keeps flow-specific state. A value of all one bits (0xffff) represents infinity.

Mobility Options:

When Status code is 0, MUST contain the MN-ID, followed by one or more Flow Identification Mobility options [RFC6089].

5. Conceptual Data Structures

5.1. Multiple Care-of Address Registration

The local mobility anchor is extended to allow a mobile node to register multiple proxy care of address (Proxy-CoA) and support the mobile node use the same address (prefix) beyond a single interface and mobile access gateway. The LMA maintains multiple binding cache entries for an MN. The number of binding cache entries for an MN is equal to the number of the MN's interfaces attaching to any MAGs.

BID-PRI	BID	MN-ID	ATT	HNP(s)	Proxy-CoA
20	1	MN1	WiFi	HNP1,HNP2	IP1 (MAG1)
30	2	MN1	3GPP	HNP1,HNP3	IP2 (MAG2)

Figure 7: Extended Binding Cache

Figure 7 shows two Binding Cache Entries of the MN1 when it is attached to the network using two different access technologies. Both of the two attachments share HNP1 and are bounded to two different Proxy-CoAs.

5.2. Flow Mobility Cache

Each LMA MUST maintain a flow mobility cache (FMC) as shown in Figure 8. This table MUST contain an entry for each flow sent from the MN. A flow binding entry includes the following fields:

- Flow Identifier Priority (FID-PRI).
- Flow Identifier (FID).
- Traffic Selector (TS).
- Binding Identifier (BID).
- Action.
- Active/Inactive.

FID-PRI	FID	TS	BIDs	Action	A/I
10	2	TCP	1	Forward	Active
20	4	UDP	1,2	Forward	Inactive

Figure 8: Flow Mobility Cache

The BID field contains the identifier of the binding cache entry which packets matching the flow information described in the TS field will be forwarded to. When a flow is decided to be moved, the affected BID(s) of the table are updated.

Similar to flow binding described in [RFC6089], each flow binding entry points to a specific binding cache entry identifier (BID). When a flow is moved, the LMA simply updates the pointer of the flow binding entry with the BID of the interface to which the flow will be moved. The traffic selector (TS) in flow binding table is defined as in [RFC6088]. TS is used to classify the packets of flows basing on specific parameters such as service type, source and destination address, etc. The packets matching with the same TS will be applied the same forwarding policy. FID-PRI is the order of precedence to take action on the traffic. Action may be forward or drop. If a binding entry becomes 'Inactive' it does not affect data traffic. An entry becomes 'Inactive' only if all of the BIDs are deregistered.

The Mobile Access Gateway MAY also maintain a similar data structure. In case no full flow mobility state is required at the MAG, the

Binding Update List (BUL) data structure is enough and no extra conceptual data entries are needed. In case full per-flow state is required at the MAG, it SHOULD also maintain a Flow Mobility Cache structure.

6. Mobile Node considerations

This specification assumes that the mobile node IP layer interface can simultaneously and/or sequentially attach to multiple MAGs, possibly over multiple media. The mobile node MUST be able to enforce uplink policies to select the right outgoing interface. One form to achieve this multiple attachment is described in [I-D.ietf-netext-logical-interface-support], which allows the mobile node supporting traffic flows on different physical interfaces regardless of the assigned prefixes on those physical interfaces.

7. IANA Considerations

This specification defines two new mobility header types (Flow Mobility Initiate and Flow Mobility Acknowledgement) and a new value for the Handoff Indicator.

8. Security Considerations

The protocol signaling extensions defined in this document share the same security concerns of Proxy Mobile IPv6 [RFC5213]. The new Flow Mobility Initiate and Flow Mobility Acknowledgement messages exchanged between the mobile access gateway and the local mobility anchor MUST be protected using IPsec using the established security association between them.

9. Authors

This document reflects contributions from the following authors (in alphabetical order).

Kuntal Chowdhury

E-mail: Kchowdhu@cisco.com

Vijay Devarapalli

E-mail: vijay@wichorus.com

Sri Gundavelli

E-mail: sgundave@cisco.com

Youn-Hee Han

E-mail: yhhan@kut.ac.kr

Yong-Geun Hong

E-mail: yonggeun.hong@gmail.com

Mohana Dahamayanthi Jeyatharan

E-mail: mohana.jeyatharan@sg.panasonic.com

Rajeev Koodli

E-mail: rkoodli@cisco.com

Kent Leung

E-mail: kleung@cisco.com

Telemaco Melia

E-mail: Telemaco.Melia@alcatel-lucent.com

Bruno Mongazon-Cazavet

E-mail: Bruno.Mongazon-Cazavet@alcatel-lucent.com

Chan-Wah Ng

E-mail: chanwah.ng@sg.panasonic.com

Behcet Sarikaya

E-mail: sarikaya@ieee.org

Tran Minh Trung

E-mail: trungtm2909@gmail.com

Frank Xia

E-mail: xiayangsong@huawei.com

10. Acknowledgments

The authors would like to thank Juan-Carlos Zuniga, Pierrick Seite, Julien Laganier for all the useful discussions on this topic.

The authors would also like to thank Marco Liebsch and Juan-Carlos Zuniga for their reviews of this document.

The work of Carlos J. Bernardos has also been partially supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project) and by the Ministry of Science and Innovation of Spain under the QUARTET project (TIN2009-13992-C02-01).

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

[RFC5648] Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", RFC 5648, October 2009.

[RFC5846] Muhanna, A., Khalil, M., Gundavelli, S., Chowdhury, K., and P. Yegani, "Binding Revocation for IPv6 Mobility", RFC 5846, June 2010.

[RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, January 2011.

[RFC6089] Tsirtsis, G., Soliman, H., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support", RFC 6089, January 2011.

11.2. Informative References

[I-D.ietf-netext-logical-interface-support]
Melia, T. and S. Gundavelli, "Logical Interface Support for multi-mode IP Hosts", draft-ietf-netext-logical-interface-support-05 (work in

progress), April 2012.

Author's Address

Carlos J. Bernardos (editor)
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/