

NETEXT Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2011

CJ. Bernardos, Ed.
UC3M
March 14, 2011

Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-bernardos-netext-pmipv6-flowmob-03

Abstract

Proxy Mobile IPv6 (PMIPv6) is a network-based localized mobility management protocol that enables mobile devices to connect to a PMIPv6 domain and roam across gateways without changing their IP addresses. PMIPv6 basic specification also provides limited multi-homing support to multi-mode mobile devices. The ability of movement of selected flows from one access technology to another is missing in basic PMIPv6. This document describes enhancements to the Proxy Mobile IPv6 protocol that are required to support flow mobility over multiple physical interfaces.

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

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 September 15, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 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

1. Introduction	3
2. Terminology	3
3. Overview of the PMIPv6 flow mobility extensions	4
3.1. Use case scenarios	4
3.2. Basic Operation	5
4. Message formats	11
4.1. Flow Mobility Initiate (FMI)	11
4.2. Flow Mobility Acknowledge (FMA)	13
5. Conceptual Data Structures	14
5.1. Multiple Care-of Address Registration	14
5.2. Flow Mobility Cache	14
6. Mobile Node considerations	16
7. IANA Considerations	16
8. Security Considerations	16
9. Authors	16
10. Acknowledgments	18
11. References	18
11.1. Normative References	18
11.2. Informative References	18
Appendix A. Discussion items for IETF 80th	18
A.1. Summary of the ML discussion	19
A.2. Proposed changes for -04 version	19
Author's Address	20

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 an MN to connect to the same PMIPv6 domain through different interfaces. The "logical interface" at the IP layer may enable packet transmission and reception over different physical media. This technique can be used to achieve flow mobility, i.e., the movement of selected flows from one access technology to another. It is assumed that an IP layer interface can simultaneously and/or sequentially attach to multiple MAGs (possibly over multiple media). This document specifies protocol extensions to Proxy Mobile IPv6 between the LMA and MAGs for distributing specific traffic flows on different physical interfaces. This document assumes that a "logical interface" at the Mobile Node is capable of supporting traffic flows from different physical interfaces regardless of the assigned prefixes on those physical interfaces.

In particular, this document specifies how to manage "flow mobility" state in the PMIPv6 network (i.e. LMAs and MAGs), namely creation, refresh and cancel operation. Flow mobility is controlled by the LMA. The trigger causing the LMA to initiate a flow mobility operation is out of scope of this specification.

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 are defined and used in this document:

FMI (Flow Mobility Initiate). Message sent by the LMA to create, refresh or cancel flow mobility state in the MAG. It conveys the information required to manage the flow mobility in a PMIPv6-Domain. This message is only needed when the flow mobility operation is not triggered by the attachment of a new interface of the mobile node.

FMA (Flow Mobility Acknowledge). Message sent by the MAG in reply to an FMI message. It provides feedback about the result of a flow mobility creation, refresh or cancel operation requested in the 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

Flow mobility assumes simultaneous access to more than one network, in a contrast to a typical handover where connectivity to a physical medium is relinquished, and is re-established with another. In order to support flow mobility in a PMIPv6 network, it is required to be able to tie the different PMIPv6 mobility sessions (one per interface) to a logical interface which is hiding one or more physical interfaces. The different mobility sessions in which a mobile node may be involved can share the same set of prefixes or have different ones:

1. 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 RFC 5213.
2. 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 in RFC 5213, and the LMA needs to be able to provide the same assignment even for the simultaneous attachment (as opposed to the handover

scenario only). It is assumed for the sake of this specification that the LMA has the knowledge if the MN supports the logical interface and if to assign the same prefix(es) or different prefix(es) to both access networks. How this is done is outside of the scope of this specification.

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 above two 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.

Among the above, scenario 2 MAY need extensions to RFC 5213 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 may be necessary between the LMA and the MAG.

The scenario 1 requires flow mobility signaling whenever the LMA determines the need for 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.

The scenario 3 requires flow mobility signaling whenever the LMA determines the need for relocating flows for the new prefix(es) which are not shared across attachments.

In all the scenarios, the MAGs should be aware of the prefixes for which the MN is going to receive traffic. As a result of a flow mobility operation, these prefixes might not be limited to those delegated by the MAG upon attachment of the connected interface, and therefore in these cases, signaling is required.

The extensions described in this document support any of these aforementioned scenarios.

3.2. Basic Operation

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

When a multi-interfaced mobile node connects to a PMIPv6-domain, it performs regular attachment and as a result is able to configure an IP address (or a set of IP addresses) on the logical interface hiding the different physical interfaces. If the LMA assigns a common prefix (or set of prefixes) to the different physical interfaces attached to the domain, then all the MAGs have already all the

routing knowledge required to forward packets to the mobile node, and the LMA does not need to perform any kind of signaling in order to move flows across the different physical interfaces. Note that there should be a local policy in place that ensures that the mobile node sends outbound packets using the same physical interface from which packets belonging to the same flow are being received (the used interface might change during the lifetime of a communication). This SHOULD be enforced by the logical interface engine, and the details about how this is done are out of the scope of this document). For unidirectional outbound communications, there SHOULD be a policy at the mobile node defining which physical interface is used to send the traffic. For bidirectional outbound communications, there SHOULD be also such a policy, but its content must be consistent with the policy at the network-side (the details about how this consistency is ensured are out of the scope of this document).

In case the MAGs needs to be informed about flow mobility decisions, because of packet policing, packet enforcement, charging or similar reasons, the LMA MAY re-use the signaling defined later in this document to convey this information.




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), grouped in a unique logical interface (lif). Each physical interface is attached to a different MAG, both of them anchored and controlled by the same LMA. Since both physical interfaces are assigned the same prefix (pref1) upon attachment to the MAGs, the mobile node has one single IPv6 addresses configured on the logical interface: pref1::lif. Initially, flow X goes through MAG1 and flow Y through MAG2. The LMA, at a certain point, decides to move flow Y, so it also goes through MAG1. As show in Figure 2, no signaling between the LMA and the MAGs is needed.

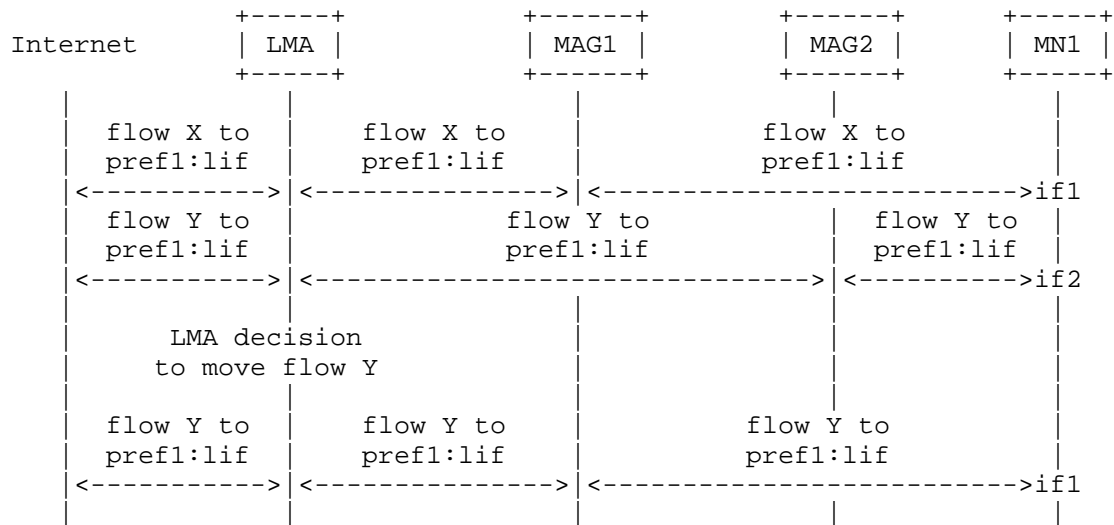


Figure 2: Flow mobility message sequence when the LMA assigns a common set of prefixes

Figure 3 shows the state of the different network entities after moving flow Y in the previous example.

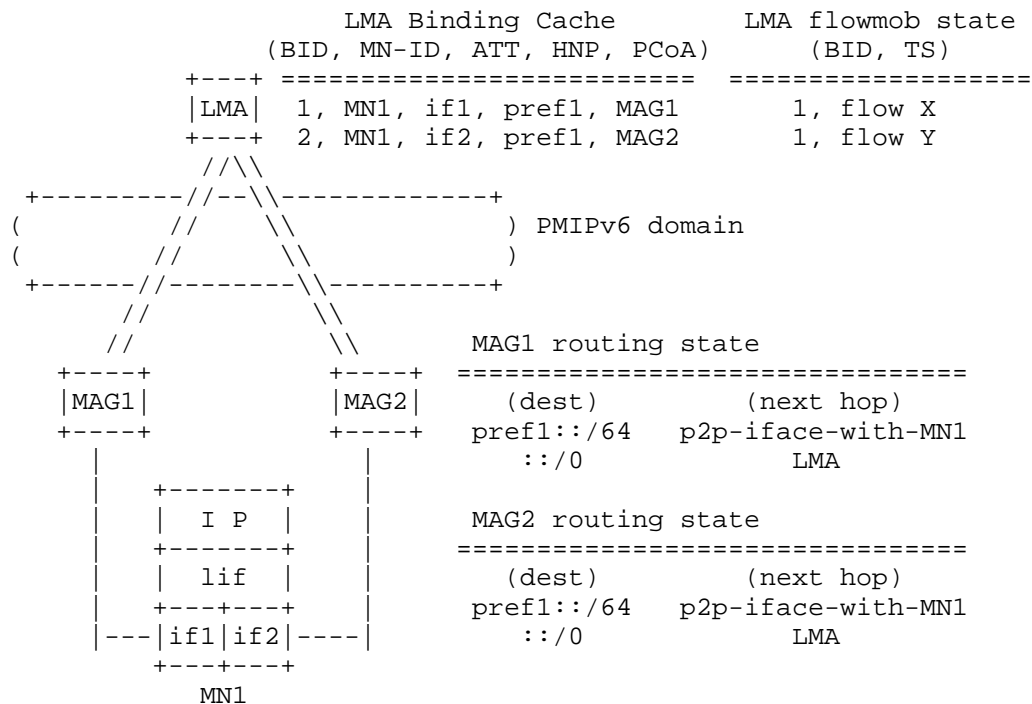


Figure 3: Data structures when the LMA assigns a common set of prefixes

A different flow mobility scenario happens when the LMA assigns different set of prefixes to physical interfaces of the same mobile node. In this case specific signaling is required between the LMA and the MAG to support this scenario. Two different possibilities are considered next.

One first possible case is the following (shown in Figure 4). 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 as a trigger by the LMA to decide perform a flow mobility decision. In this case, we consider that flows are moved with a prefix granularity, meaning that the LMA moves flows 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 LMA can move the flow by just binding pref2::/64 to MAG2. This is done by including the prefix in the PBA message, and optionally sending a message to MAG1 to remove the transferred prefix(es). This message can be a Binding Revocation Indication message [RFC5846] with the P bit set to indicate that this is revocation of PMIP prefix(es). After processing BRI, the source MAG

will send a Binding Revocation Acknowledgement (BRA) message back to LMA.

Note that this specification also supports flow mobility at a finer granularity (not just on a prefix level). This is done by including in the PBA a Flow Identification Mobility option (specified in [RFC6089]) which can convey full flow information. The MAG can also include the Flow Identification Mobility option in the PBU message that it sends to the LMA. This serves as a request for the LMA to consider the flow policy rules specified in the option.

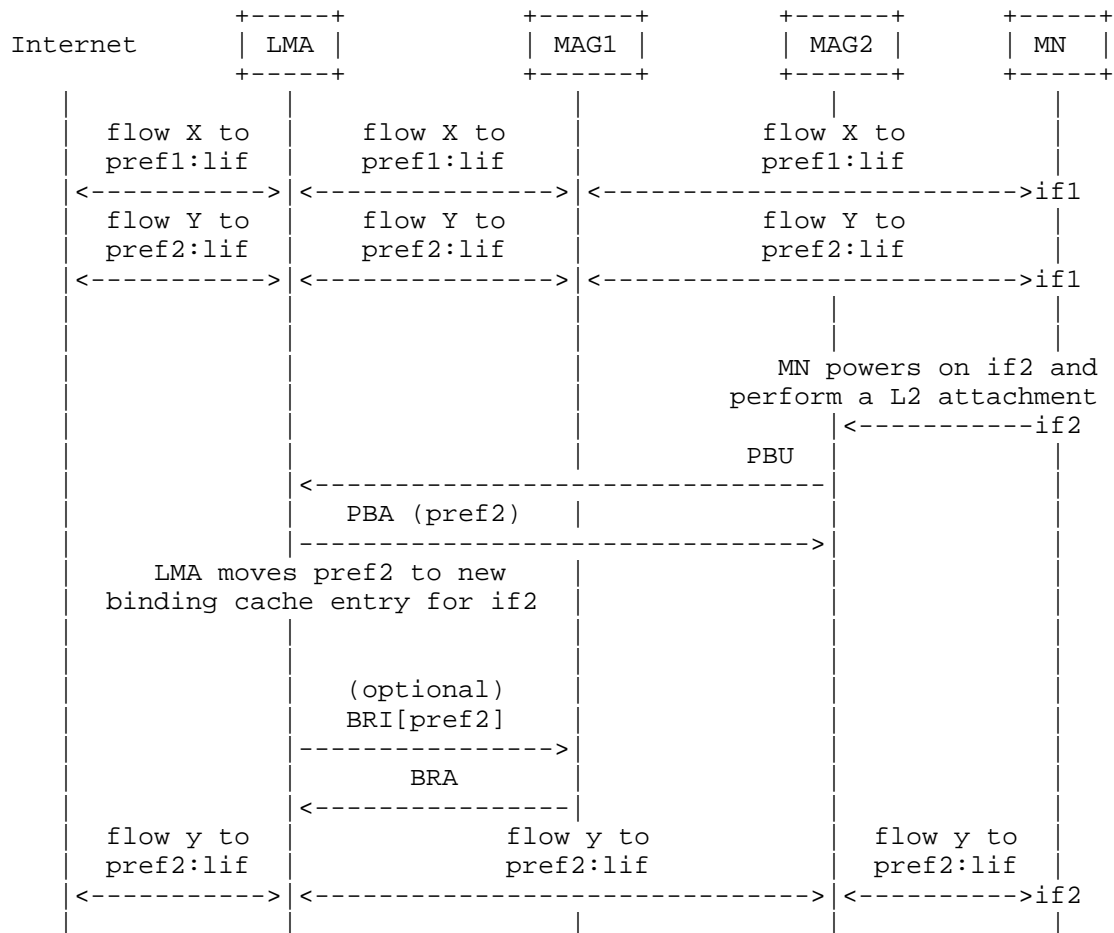


Figure 4: Flow mobility message sequence when the LMA assigns different set of prefixes per physical interface (PBU trigger)

A second possible scenario is the following. A multi-interfaced

mobile node is attached to a PMIPv6-Domain and the LMA, at a given moment, decides to move a flow. The LMA can decide to move a flow as a result of a policy change or upon receiving a trigger either based on network status or based on an event detected at the mobile node and transported via old or new MAG. How this decision is taken is out of scope of this specification. Since the LMA 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.

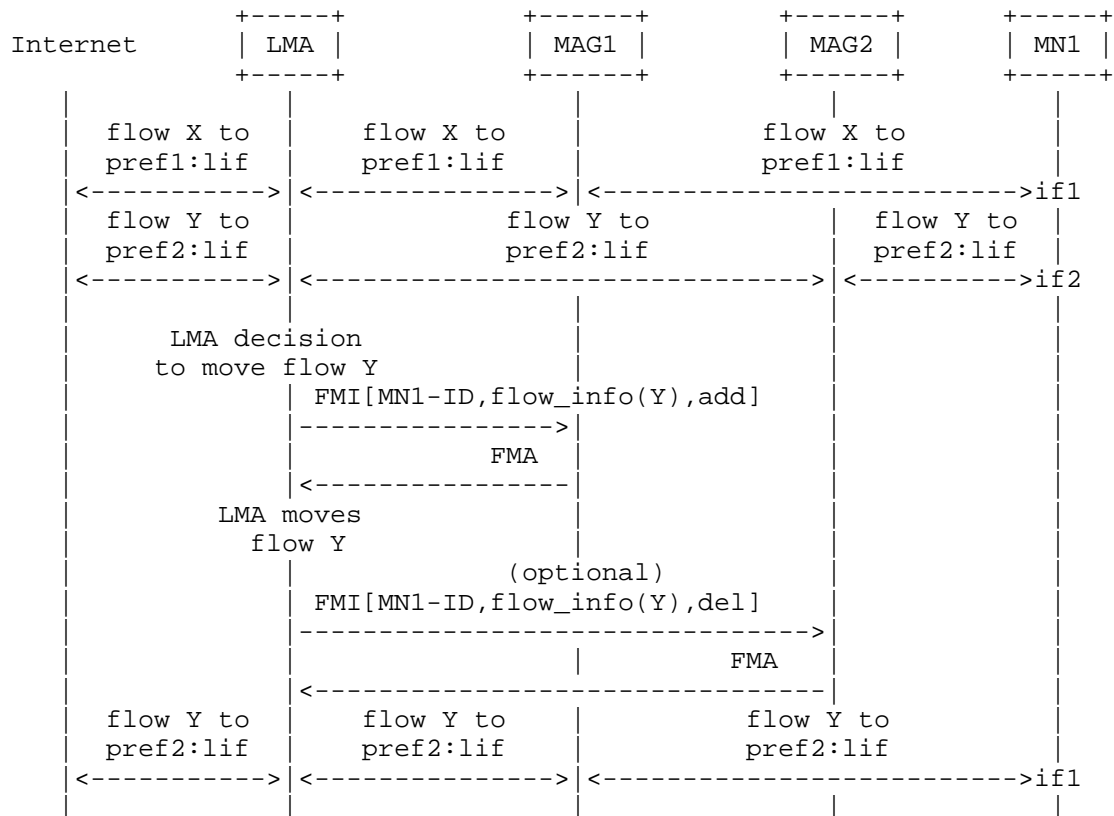


Figure 5: Flow mobility message sequence when the LMA assigns different set of prefixes per physical interface (FMI trigger)

If the LMA decides to move a particular flow from its default path (which is determined by the destination prefix) to a different one, it constructs a Flow Mobility Initiate (FMI) message. This message is sent to the new target MAG, i.e. the one selected to be the 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). Optionally, the LMA 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 5.

The state in the network after moving a flow, for the case the LMA assigns a different set of prefixes is shown in Figure 6.

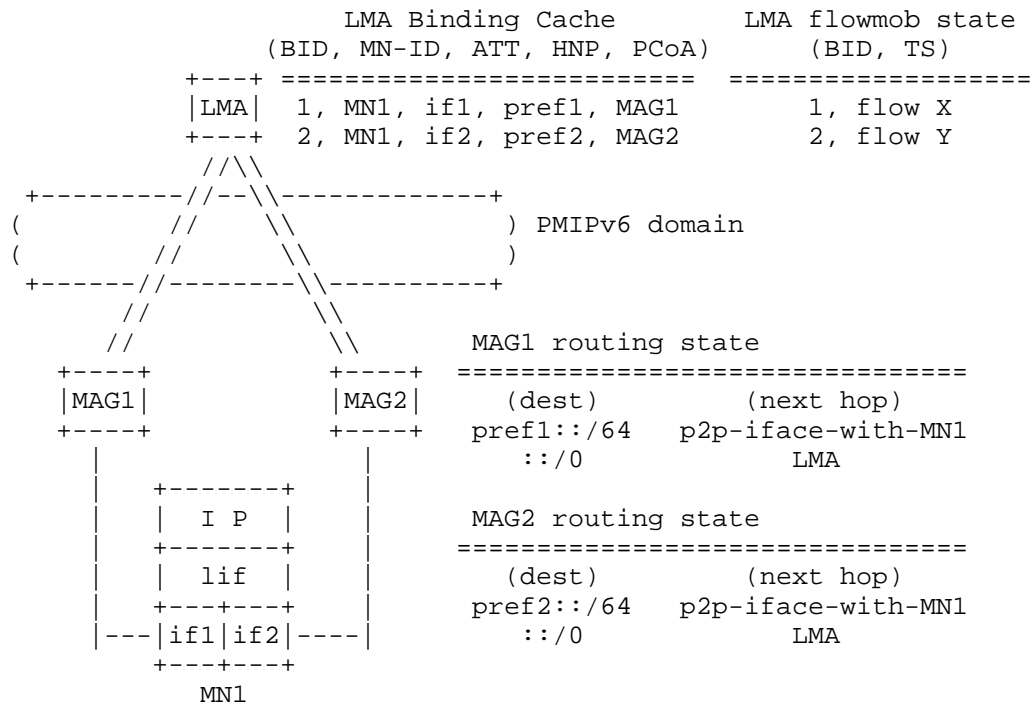


Figure 6: Data structures when the LMA assigns a different set of prefixes

4. Message formats

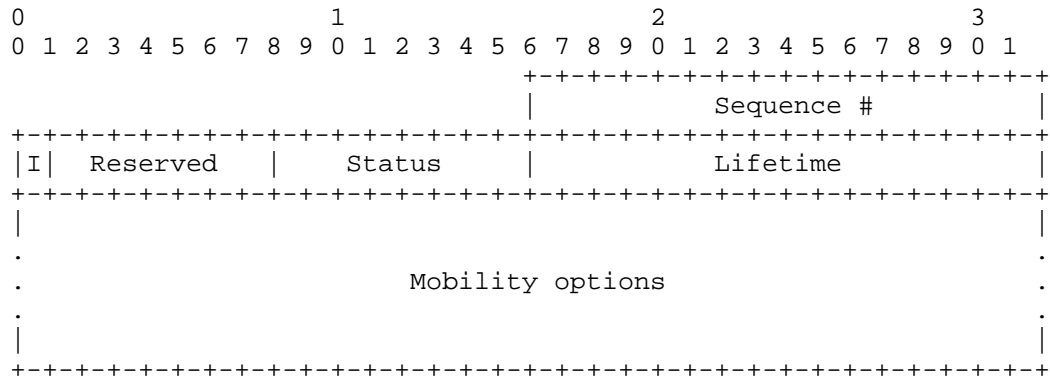
4.1. Flow Mobility Initiate (FMI)

The LMA sends an FMI message to a MAG to inform about a particular flow movement (LMA initiated). It is a Mobility Header message.

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

4.2. Flow Mobility Acknowledge (FMA)

The MAG sends an FMI message to the LMA as a response to the FMI message. It is a Mobility Header message.



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:

0: Success.

128: Reason unspecified.

129: MN not attached.

130: Sequence number out of window.

131: Traffic Selector format unsupported.

132: No existing Flow Mobility Cache entry.

133: Already 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 LMA is extended to allow a mobile node to register multiple proxy care of address (Proxy-CoA). The LMA maintains multiple binding cache entries for a MN. The number of binding cache entries of a MN is equal to the number of the MN's interfaces attaching to the MAG.

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 attaches 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 contains entry for each flow sent from the MN. A flow binding entry includes the following fields:

- o Flow Identifier - Priority (FID-PRI)-
- o Flow Identifier (FID).
- o Traffic Selector (TS).
- o Binding Identifier (BID).
- o Action.
- o 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 BIDs field contains the identifier of the binding cache entry that all of the packets matching the flow information described in the TS field will be forwarded to. When the flow mobility occurs, the BIDs will be updated with new binding cache entry identifier.

Similar to flow binding described in [RFC6089], each flow binding entry points to a specific binding cache entry identifier (BID). When the LMA decides to move a flow, it 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 keep a similar structure to the FMC (details TBD).

6. Mobile Node considerations

This specification assumes the MN implements the logical interface model. The "logical interface" at the IP layer hides the use of different physical media from the IP stack, enabling the MN to send and receive packets over different interfaces. This document assumes the MN behaves as stated in the applicability statement document [I-D.ietf-netext-logical-interface-support]. In particular, it is assumed that -- for the case of bidirectional traffic -- the logical interface at the MN "replicates" the behavior observed for downlink packets on a per-flow basis. This means that the MN sends UL Flow X on the same interface which received the DL Flow X. It also means that if the LMA moves flow X during its lifetime, the MN will follow that change, upon the reception of packets of flow X via a different interface.

This specification only supports flow mobility between different physical interfaces belonging to the same logical interface. If an MN has several logical interfaces, flow mobility across different logical interfaces is not supported.

7. IANA Considerations

TBD.

8. Security Considerations

TBD.

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 discussions on this topic.

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.
- [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-01 (work in progress), October 2010.

Appendix A. Discussion items for IETF 80th

This appendix tries to serve as basis for the discussion in the IETF 80th on flow mobility. It includes a summary of the major issues/comments raised on the NETEXT mailing list, as well as a proposed plan for a future revision of the document.

A.1. Summary of the ML discussion

Here we list (in no particular order) some of the issues raised on the NETEXT mailing list:

- o Lack of realistic scenario for applicability: no use-case/client for LMA-initiated mobility, no real-life scenario where the LMA would receive flow mobility policies.
- o Consistency of policy rules between the MN and LMA does not ensure that the LMA knows what decision the MN took because the LMA does not necessarily know the context in which the MN is.
- o Discrepancies on the solution approach: dynamic attachment/detachment of interfaces from sessions (new prefixes cannot be added to sessions) vs dynamic prefix management. It's being argued that the draft changes the basics of RFC5213 session management.
- o Discrepancies on the solution approach: requirement on the existence of L2 triggers to aid in the dynamic attachment/detachment of interfaces from sessions for flow mobility purposes.
- o How does the LMA know channel condition of each radio, applications requirements of apps running in the UE?.
- o Source of triggers for flow mobility: MAG, LMA or both?

A.2. Proposed changes for -04 version

Based on the discussion on the ML list, a possible way to modify this document in -04 version is the following. We define two different approaches, based on the L2 signaling support:

1. L2 signaling based. When an MN attaches to a new MAG, it can use extended L2 signaling (e.g., HI=FM) to indicate that the attachment is for flow mobility. In this case, same prefix is assigned to the new interface (which is added to the existing mobility session). Alternatively, a new prefix can also be added to the session (this is up to the policy configured). Now new signaling is required between MAG and LMA, just a new HI value, the extended L2 signaling in place and updating the state machines of MAG and LMA to support this new behavior.
2. IP based. If no extended L2 signaling is available (i.e., no HI=FM), MAGs create new sessions upon new MN interface attachment. The LMA manages the prefixes of the session (decides

to assign the same of an already attached interface or a new one) as well as the movement of them (with a prefix/flow granularity). The trigger for the movement of a flow is out of scope (MAG triggers are considered). This is basically the operation described in the current version of the draft.

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/>

