

Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-ietf-netext-pmipv6-flowmob-13

Abstract

Proxy Mobile IPv6 allows a mobile node to connect to the same Proxy Mobile IPv6 domain through different interfaces. 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 described in this document consist of 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

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 August 29, 2015.

Copyright Notice

Copyright (c) 2015 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
3.2.1.	MN sharing a common set of prefixes on all MAGs	5
3.2.2.	MN with different sets of prefixes on each MAG	9
3.3.	Use of PBU/PBA signaling	11
3.4.	Use of flow-level information	12
4.	Message Formats	12
4.1.	Home Network Prefix	12
4.2.	Flow Mobility Initiate (FMI)	13
4.3.	Flow Mobility Acknowledgement (FMA)	14
5.	Conceptual Data Structures	15
5.1.	Multiple Proxy Care-of Address Registration	15
5.2.	Flow Mobility Cache	16
6.	Mobile Node considerations	17
7.	IANA Considerations	18
8.	Security Considerations	18
9.	Authors	18
10.	Acknowledgments	19
11.	References	20
11.1.	Normative References	20
11.2.	Informative References	20
	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 (HNP) to the MN and is the topological anchor for all traffic belonging to the MN.

PMIPv6 allows 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.

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

FMC (Flow Mobility Cache). Conceptual data structure 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](#)].

A combination of the two above-mentioned scenarios is also possible. 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 scenarios described before. 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](#) and [Section 3.2.2](#).

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 MUST have local policies in place to ensure that 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. The MN makes the final IP flow mobility decision, and then the LMA follows that decision and updates its forwarding state accordingly. Note that this does not prevent network initiated mobility, the network still could trigger mobility on the MN side via out-of-band mechanisms (e.g., 3GPP/ANDSF sends updated routing policies to the MN). In a given scenario and mobile node, the decision on IP flow mobility MUST be taken either by the MN or the LMA, but MUST NOT be taken by both.

[3.2.1.](#) MN sharing a common set of prefixes on all MAGs

This scenario corresponds to the first use case scenario described in [Section 3.1](#). Extensions to basic PMIPv6 [[RFC5213](#)] signaling at the time of a new attachment are needed 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 every MAG already has all the routing knowledge required to forward uplink or downlink packets, and the local mobility anchor does not need to send any kind of signaling in order to move flows across the different physical interfaces. Optionally, signaling MAY be exchanged in case the MAG needs to know about flow level information (e.g., to link flows with proper QoS paths and/or inform the mobile node).

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, default policies, etc. In this document a new Handoff Indicator (HI) value ("Attachment over a new interface sharing prefixes", value {IANA-0}) 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 an MN-LL-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 an MN-LL-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 an MN-LL-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.

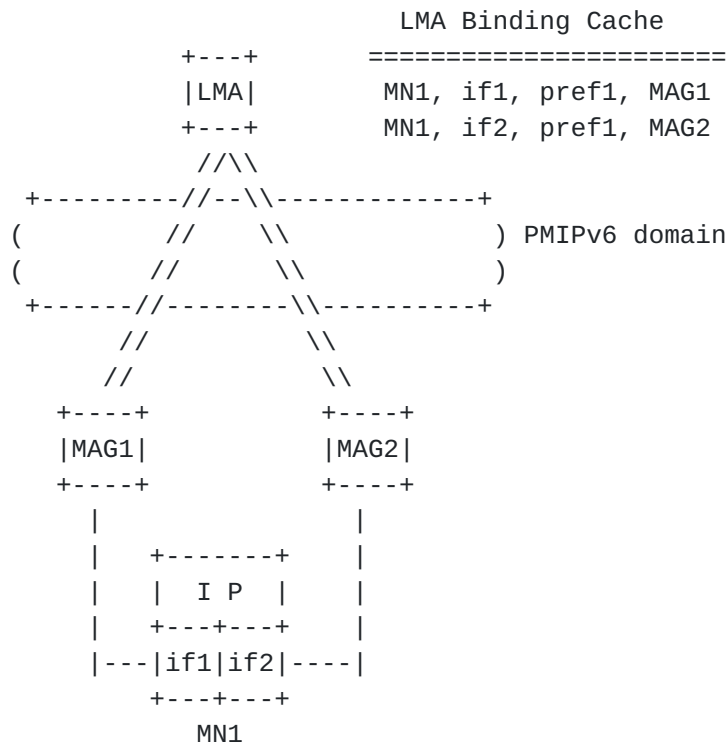


Figure 1: Shared prefix across physical interfaces scenario

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 controlled by the same local mobility anchor. Both physical interfaces are assigned the same prefix (pref1) upon attachment to the MAGs. If the IP layer at the mobile node shows one single logical interface (e.g., as described in [\[I-D.ietf-netext-logical-interface-support\]](#)), then the mobile node has one single IPv6 address configured at the IP layer: pref1::mn1. Otherwise, per interface IPv6 addresses (e.g., pref1::if1 and pref1::if2) would be configured; each address MUST be valid on every interface. We assume the first case in the following example (and in the rest of this document). Initially, flow X goes through MAG1 and flow Y through MAG2. At certain point, flow Y can be moved to also go through MAG1. Figure 2 shows the scenario in which no flow-level information needs to be exchanged, so there is no signaling between the local mobility anchor and the mobile access gateways.

Bernardos

Expires August 29, 2015

[Page 7]

Note that if different IPv6 addresses are configured at the IP layer, IP session continuity is still possible (for each of the configured IP addresses). This is achieved by the network delivering packets destined to a particular IP address of the mobile node to the right MN's physical interface where the flow is selected to be moved, and the MN also selecting the same interface when sending traffic back up link.

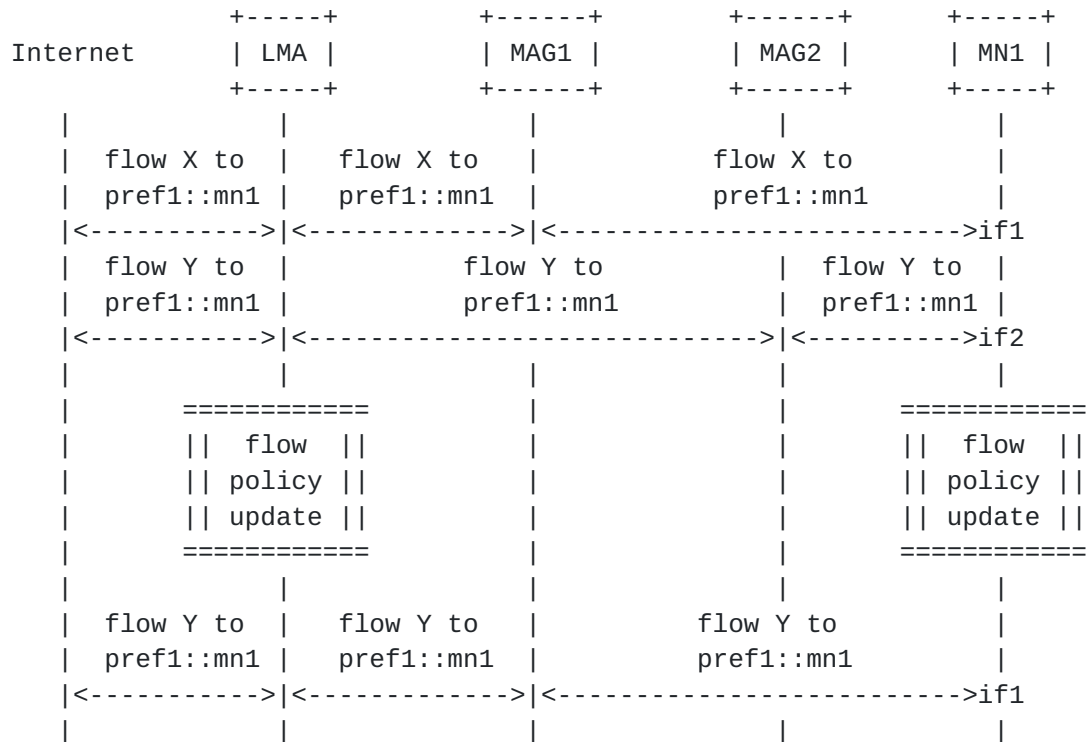


Figure 2: Flow mobility message sequence with common set of prefixes

Figure 3 shows the state of the different network entities after moving flow Y in the previous example. This document 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 MAY be used, as it does not impact on the operation of the solution in this case.



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 case, or a combination of scenarios, described in [Section 3.1](#). In this case, additional 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.

In this case, 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, the solution defined in this specification makes use of two mobility messages: Flow Mobility Indication and Flow Mobility Acknowledgement, which actually use the format of the Update Notifications for Proxy Mobile IPv6 defined in [\[RFC7077\]](#). The trigger for the flow movement can be on the mobile node (e.g., by using layer-2 signaling with the MAG) or on the network (e.g., based

on congestion and measurements) which then notifies the MN for the final IP flow mobility decision (as stated in [section 3.1](#)). Policy management functions (e.g., 3GPP/ANDSF) can be used for that purpose, however, how the network notifies the MN is out of the scope of this document.

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 an Flow Mobility Indication (FMI) message. This message includes a Home Network Prefix option for each of the prefixes that are requested to be provided with flow mobility support on the new MAG (note that these prefixes are not anchored by the target MAG, and therefore the MAG MUST NOT advertise them on the MAG-MN link), with the off-link bit (L) set to one. 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 MAG replies with a Flow Mobility Acknowledgement (FMA). The message sequence is shown in Figure 4.

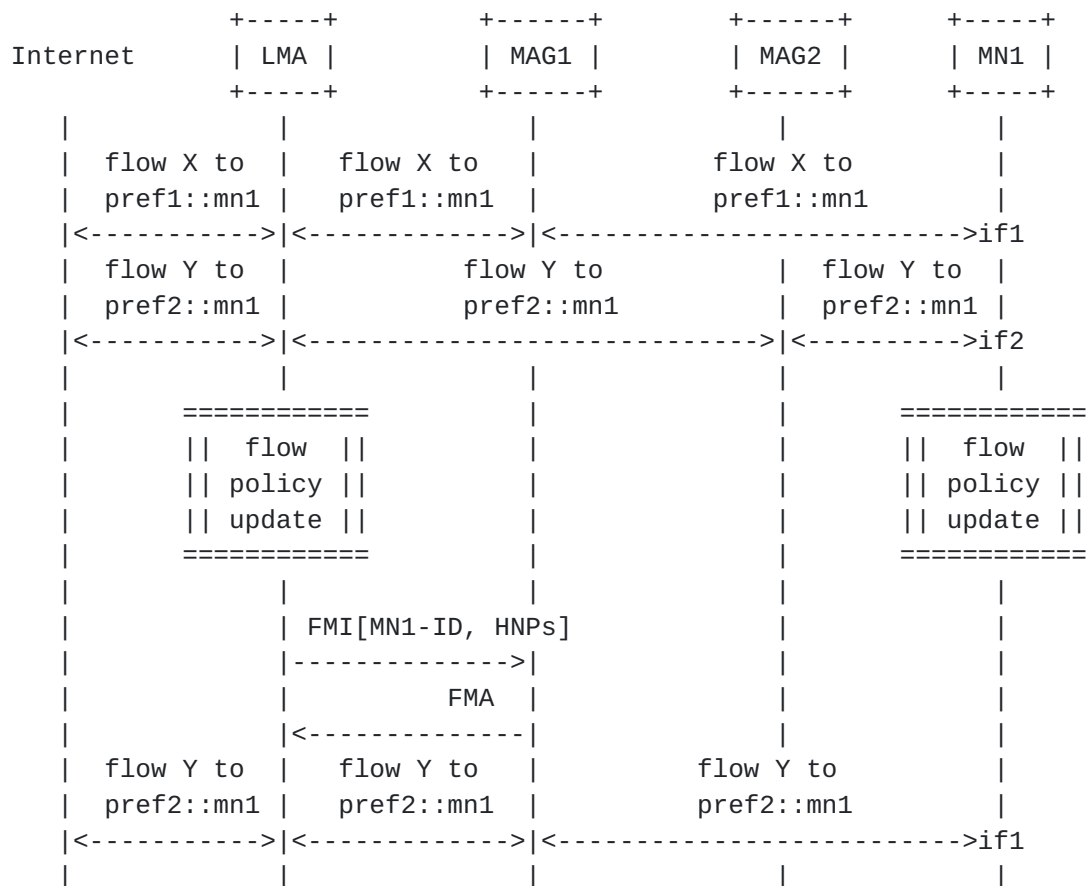


Figure 4: Flow mobility message sequence when the LMA assigns different sets of prefixes per physical interface

Bernardos

Expires August 29, 2015

[Page 10]

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.

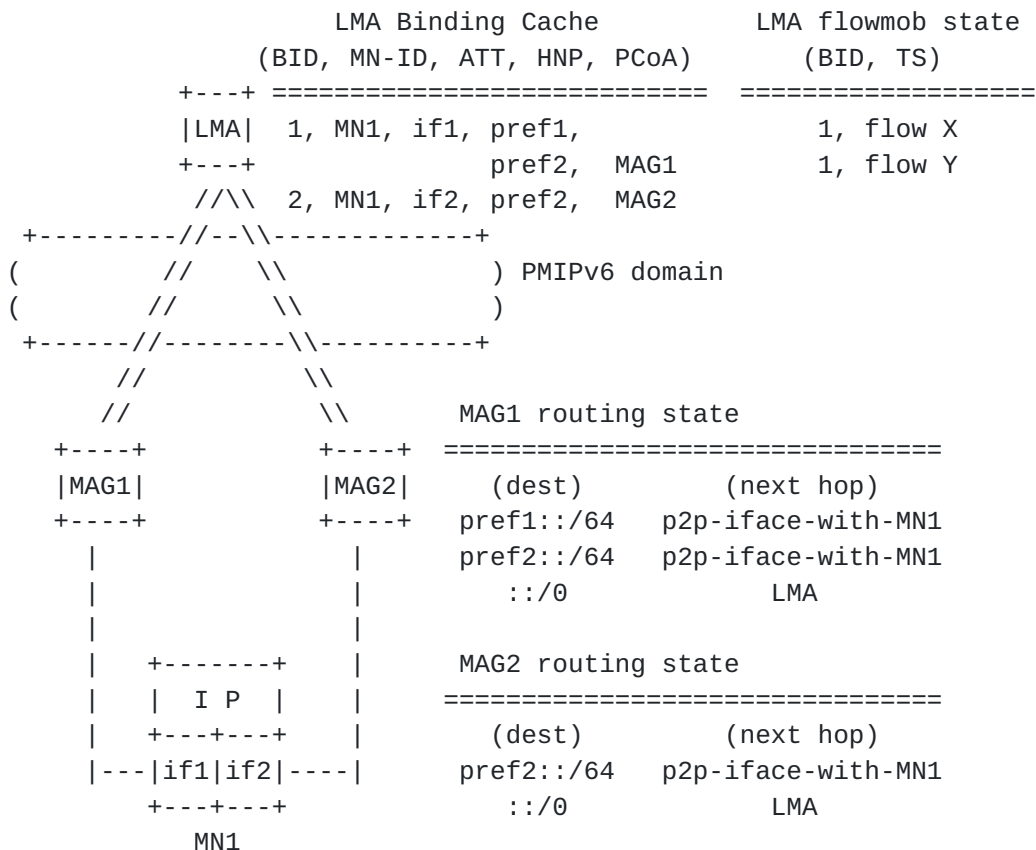


Figure 5: Data structures when the LMA assigns a different set of prefixes

3.3. Use of PBU/PBA signaling

This specification introduces the FMI/FMA signaling so the LMA can exchange with the MAG information required to enable flow mobility without waiting to a received PBU. There are however scenarios in which the trigger for flow mobility might be related to a new MN's interface attachment. In this case, the PBA sent in response to the PBU received from the new MAG can convey the same signaling that the FMI does. In this case the LMA MUST include in the PBA a Home Network Prefix option for each of the prefixes that are requested to be provided with flow mobility support on the new MAG with the off-link bit (L) set to one.

3.4. Use of flow-level information

This specification does not mandate flow-level information to be exchanged between the LMA and the MAG to provide flow mobility support. It only requires the LMA to keep flow-level state ([Section 5.2](#)). However, there are scenarios in which the MAG might need to know which flow(s) is/are coming within a prefix that has been moved, to link it/them to proper QoS path(s) and optionally inform the MN about it. This section describes the extensions used to include flow-level information in the signaling defined between the LMA and the MAG.

This specification re-uses some of the mobility extensions and message formats defined in [\[RFC5648\]](#) and [\[RFC6089\]](#), namely the Flow Identification Mobility Option and the Flow Mobility Sub-Options.

In case the LMA wants to convey flow-level information to the MAG, it MUST include in the FMI (or the PBA) a Flow Identification Mobility Option for all the flows that the MAG needs to be aware with flow granularity. Each Flow Identification Option MUST include a Traffic Selector Sub-Option including such flow-level information.

To remove a flow binding state at the MAG, the LMA simply sends a FMI (or PBA if it is in response to a PBU) message that includes flow identification options for all the flows that need to be refreshed, modified, or added, and simply omits those that need to be removed.

Note that even if a common set of prefixes is used, providing the MAG with flow-level information requires signaling to be exchanged in this case between the LMA and the MAG. This is done sending a FMI message (or a PBA if it is sent in response to a PBU).

4. Message Formats

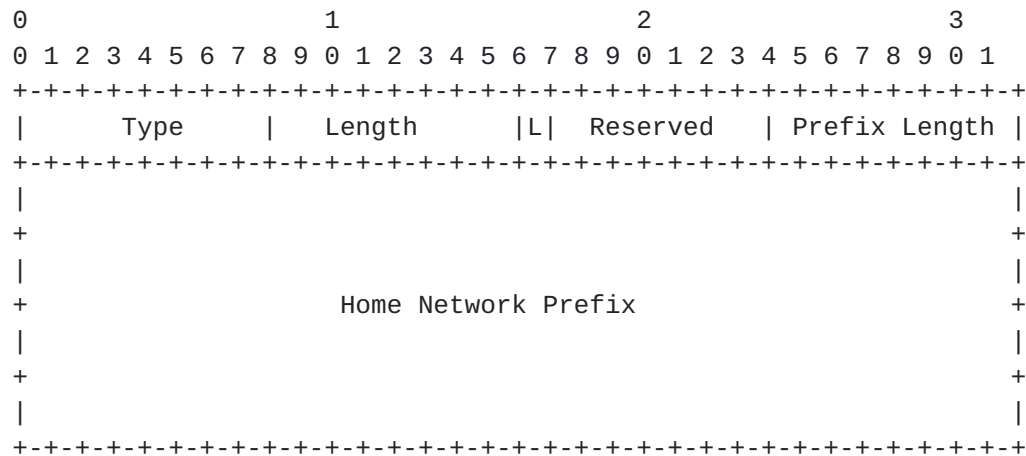
This section defines modifications to the Proxy Mobile IPv6 [\[RFC5213\]](#) protocol messages.

This specification requires implementation of UPN [\[RFC7077\]](#) and UPA [\[RFC7077\]](#) messages with the specific Notification Reason and Status Code values as defined by this document. This document does not require implementation of any other aspects of [\[RFC7077\]](#).

4.1. Home Network Prefix

A new flag (L) is included in the Home Network Prefix option to indicate to the Mobile Access Gateway whether the conveyed prefix has to be hosted on-link or not on the point-to-point interface with the mobile node. A prefix is hosted off-link for the flow mobility

purposes defined in this document. The rest of the Home Network Prefix option format remains the same as defined in [\[RFC5213\]](#).



Off-link Home Network Prefix Flag (L):

The Off-link Home Network Prefix Flag is set to indicate to the Mobile Access Gateway that the home network prefix conveyed in the option is not to be hosted on-link, but has to be considered for flow mobility purposes and therefore added to the Mobile Access Gateway routing table. If the flag is set to 0, the Mobile Access Gateway assumes that the home network prefix has to be hosted on-link.

4.2. Flow Mobility Initiate (FMI)

The FMI message (Figure 6) used in this specification is the Update Notification (UPN) message specified in [RFC7077]. The message format, transport and security consideration are as specified in [RFC7077]. The format of the message is specified in [Section 4.1 of \[RFC7077\]](#) and is shown below for completeness. This specification does not modify the UPN message, however, it defines new Notification Reason values for use in this specification.

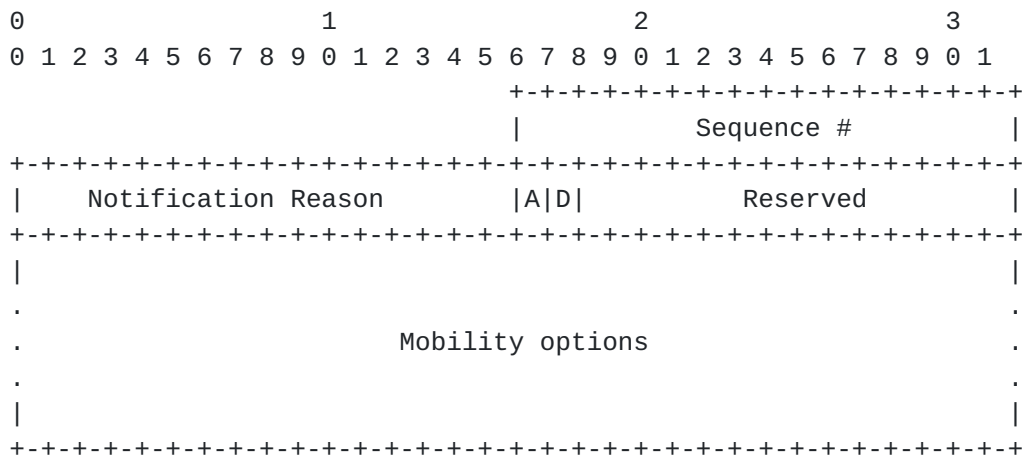


Figure 6: Flow Mobility Initiate message format

This document defines the following new Notification Reason value for use in Update Notification message:

Notification Reason:

{IANA-1} - FLOW-MOBILITY. Request to add/refresh the prefix(es) conveyed in the Home Network Prefix options included in the message to the set of prefixes for which flow mobility is provided.

The Mobility Options field of an FMI MUST contain the MN-ID, followed by one or more Home Network Prefixes options. Prefixes for which flow mobility was provided that are not present in the message MUST be removed from the set of flow mobility enabled prefixes.

4.3. Flow Mobility Acknowledgement (FMA)

The FMA message (Figure 7) used in this specification is the Update Notification Ack (UPA) message specified in [Section 4.2 of \[RFC7077\]](#). The message format, transport and security consideration are as specified in [\[RFC7077\]](#). The format of the message is specified in [Section 4.2 of \[RFC7077\]](#) and is shown below here for completeness. This specification does not modify the UPA message, however, it defines new Status Code values for use in this specification.

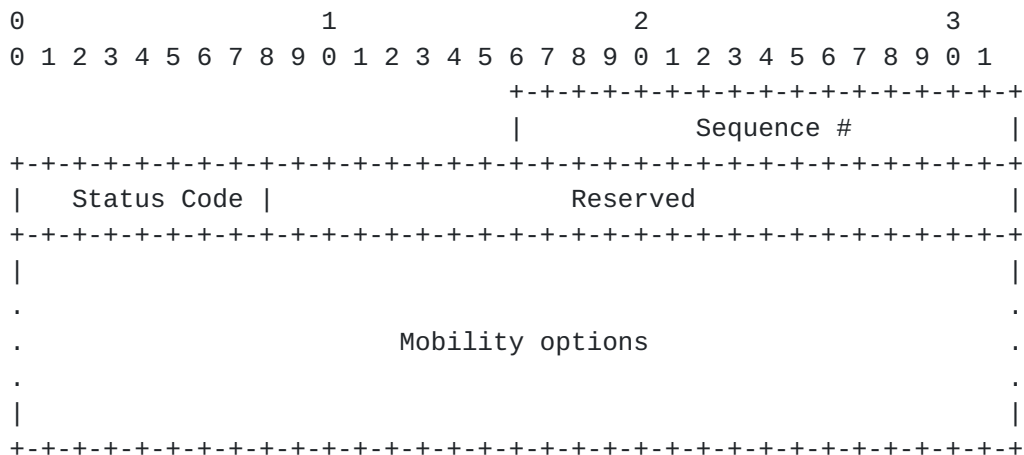


Figure 7: Flow Mobility Acknowledgement message format

This document defines the following new status code values for use in Update Notification Acknowledgement message.

Status Code:

0: Success.

{IANA-2}: Reason unspecified.

{IANA-3}: MN not attached.

When Status code is 0, the Mobility Options field of an FMA MUST contain the MN-ID, followed by one or more Home Network Prefixes options.

5. Conceptual Data Structures

This section summarizes the extensions to Proxy Mobile IPv6 that are necessary to manage flow mobility.

5.1. Multiple Proxy Care-of Address Registration

The binding cache structure of the local mobility anchor is extended to allow multiple proxy care-of address (Proxy-CoA) registrations, 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 a mobile node is equal to the number of the MN's interfaces attached to any MAGs.

This specification re-uses the extensions defined in [[RFC5648](#)] to manage multiple registrations, but in the context of Proxy Mobile

IPv6. The binding cache is therefore extended to include more than one proxy care-of address and to associate each of them with a binding identifier (BID). Note that the BID is a local identifier, assigned and used by the local mobility anchor to identify which entry of the flow mobility cache is used to decide how to route a given flow.

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 8: Extended Binding Cache

Figure 8 shows an example of extended binding cache, containing two binding cache entries (BCEs) of a mobile node MN1 attached to the network using two different access technologies. Both of the two attachments share the same prefix (HNP1) and are bound to two different Proxy-CoAs (two MAGs).

5.2. Flow Mobility Cache

Each local mobility anchor MUST maintain a flow mobility cache (FMC) as shown in Figure 9. The flow mobility cache is a conceptual list of entries that is separate from the binding cache. This conceptual list contains an entry for each of the registered flows. This specification re-uses the format of the flow binding list defined in [\[RFC6089\]](#). Each 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 9: 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 entry of the flow mobility cache points to a specific binding cache entry identifier (BID). When a flow is moved, the local mobility anchor 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 de-registered.

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 mobile access gateway, 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 new value for the Handoff Indicator {IANA-0} and a new flag (L) in the Home Network Prefix option. A new Notification Reason value for use in Update Notification message ({IANA-1}) and new Status Values for use in Update Notification Acknowledgement message ({IANA-2} to {IANA-4}) are also defined.

8. Security Considerations

The protocol signaling extensions defined in this document share the same security concerns of Proxy Mobile IPv6 [[RFC5213](#)] and do not pose any additional security threats to those already identified in [[RFC5213](#)] and [[RFC7077](#)].

The mobile access gateway and the local mobility anchor MUST use the IPsec security mechanism mandated by Proxy Mobile IPv6 [[RFC5213](#)] to secure the signaling described in this document.

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, Juan-Carlos Zuniga and Dirk von Hugo 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 under grant agreement n. FP7-317941 (iJOIN project).

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.
- [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.
- [RFC7077] Krishnan, S., Gundavelli, S., Liebsch, M., Yokota, H., and J. Korhonen, "Update Notifications for Proxy Mobile IPv6", [RFC 7077](#), November 2013.

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-10](#) (work in progress), September 2014.

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/>

