

NETEXT Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2011

CJ. Bernardos, Ed.
UC3M
October 25, 2010

Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-bernardos-netext-pmipv6-flowmob-01

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 April 28, 2011.

Copyright Notice

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

This document is subject to BCP 78 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. Prefix model	4
4. Flow mobility scenarios	5
4.1. Unique prefix per physical interface	5
4.1.1. Unique prefix per physical interface, with full flow granularity	6
4.1.2. Unique prefix per physical interface, with prefix granularity (partial handoff)	9
4.2. Shared prefix across physical interfaces (per-MN HNP set)	13
5. Message formats	15
5.1. Flow Mobility Initiate (FMI)	15
5.2. Flow Mobility Acknowledge (FMA)	17
6. Local Mobility Anchor considerations	18
6.1. Sending Flow Mobility Initiate messages	18
6.2. Receiving Flow Mobility Acknowledge messages	19
6.3. Conceptual Data Structures	20
6.4. Packet Processing considerations	22
7. Mobile Access Gateway considerations	22
7.1. Receiving Flow Mobility Initiate messages	22
7.2. Sending Flow Mobility Acknowledge messages	24
7.3. Conceptual Data Structures	24
7.4. Packet Processing considerations	26
8. Mobile Node considerations	26
9. IANA Considerations	26
10. Security Considerations	27
11. Authors	27
12. Acknowledgments	28
13. References	28
13.1. Normative References	28
13.2. Informative References	28
Author's Address	29

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 a protocol 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 and initiated 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.

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. Prefix model

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. There are multiple prefix models under which a flow mobility protocol needs to work:

1. At the time of a new attachment, the MN obtains a new prefix or a new set of prefixes. This is the default behavior with RFC 5213.
2. At the time of a new 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).
3. At the time of a new 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 needs extensions to RFC 5213 signaling at

the time of a new attachment. 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. 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 signalling is required.

The extensions described in this document support any of the three aforementioned prefix models.

4. Flow mobility scenarios

Flow mobility signaling takes place whenever the LMA decides to move a flow from one access to another. At this point, either the prefix corresponding to the flow is already valid on the target MAG, or it needs to be signaled. If the prefix is already valid, then the LMA simply relocates the flow to the target MAG; no specific signaling is required. For convenience, this scenario is called "shared prefix" scenario.

If, at the time LMA decides to perform flow handover, the prefix corresponding to the flow is not valid on the target MAG, the LMA decides to invoke flow mobility signaling specified in this document. For convenience, this scenario is called "unique prefix" scenario, since the target prefix was "unique" to begin with.

When there is signaling involved, the granularity of the flow mobility by default is at the prefix level. This means, the MAG only needs to be able to support just the signaled prefix for forwarding traffic, and is not involved in detailed flow-level inspection. The granularity of flow mobility MAY include detailed flow descriptors beyond the prefix alone, and MAG implementations MAY support flow-level inspection.

4.1. Unique prefix per physical interface

In this scenario, each physical interface of the mobile node is assigned a unique prefix (or set of prefixes). This is the default scenario supported by Proxy Mobile IPv6 as specified in RFC 5213 [RFC5213]. The LMA maintains multiple binding cache entries

(multiple mobility sessions) -- one per physical interface -- as well as routing entries -- one per assigned prefix. This scenario is shown in Figure 1.




Figure 1: Unique prefix per physical interface scenario

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 each physical interface is assigned a different prefix upon attachment (pref1 upon attachment to MAG1 and pref2 upon attachment to MAG2), the mobile node has two different IPv6 addresses configured on the logical interface: pref1::lif and pref2::lif.

4.1.1. Unique prefix per physical interface, with full flow granularity

In this particular sub-scenario, full flow mobility management granularity is supported, meaning that the LMA is able to perform per-flow routing through different MAGs (e.g., in the example above, HTTP traffic/particular flow intended to pref1::lif can be routed through MAG1, while VoIP traffic/particular flow intended to the same MN IP address -- pref1::lif -- can be routed via MAG2).

In this scenario, the LMA controls how flows are routed from the LMA to the MN, and therefore, through which interface the MN receives packets from different flows. If the LMA decides to move a particular flow from its default path (which is determined in this scenario 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 LMA can decide on which is the best MAG that should be used to forward a particular flow when the flow is initiated (e.g., based on application policy profiles) and/or during the lifetime of the flow upon receiving a trigger either based on network status or based on an event detected at the mobile node. How this decision is taken is out of scope of this specification. The FMI message contains (as explained in further detail in Section 5.1), the MN-Identifier, the Flow Identification Mobility option (specified in [I-D.ietf-mext-flow-binding]) which can convey prefix or full flow information, and the type of flow mobility operation (add flow).

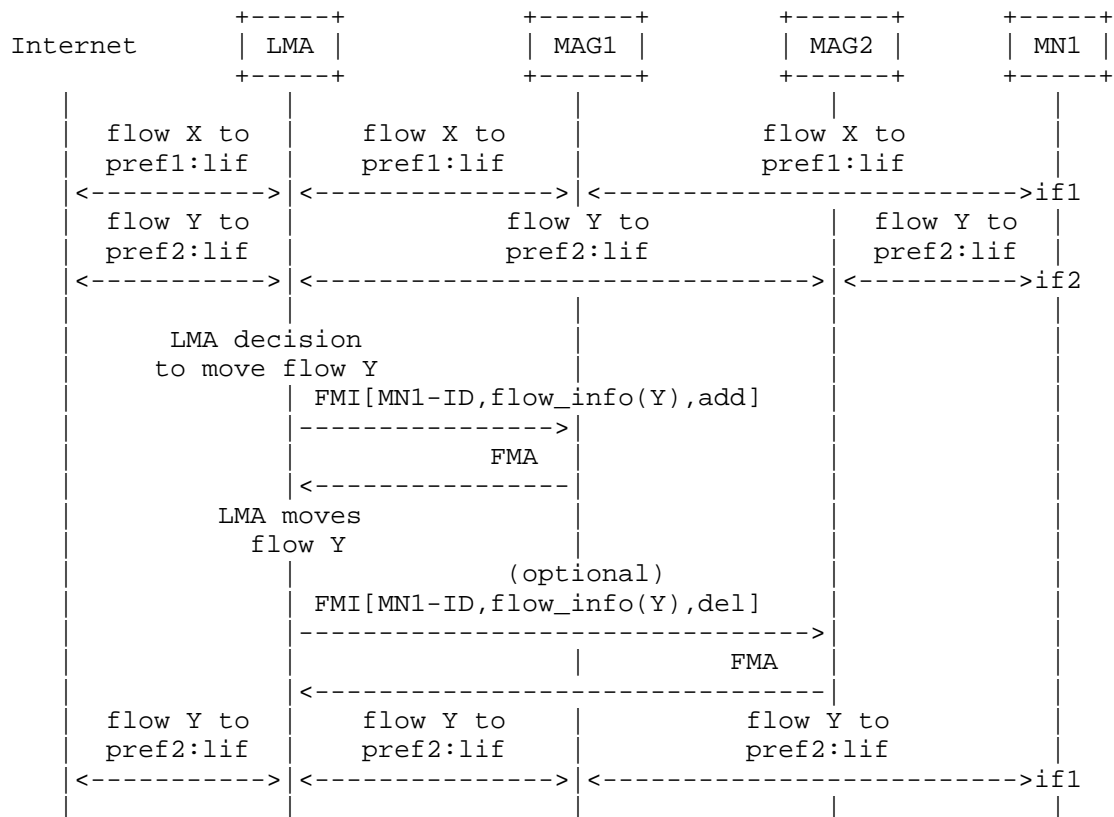


Figure 2: Flow mobility signaling for the unique prefix per physical

interface scenario, with full granularity

An example of the signaling sequence is shown in Figure 2. At the beginning MN1 has two active flows: flow X going through interface if1 and MAG1, and flow Y going through interface if2 and MAG2. At a certain moment, the LMA decides to move flow Y from interface if2 to if1. To do so, it sends a FMI message to the MAG1 where if2 (the target interface) is attached to, including the MN-ID and a Flow Mobility Option referring to flow Y. This option is defined in [I-D.ietf-mext-flow-binding] as well as the Traffic selector sub-option, which can be used to fully match a particular flow or just the MN prefix information required by the MAG to be able to route packets to the MN (the MAG by default is only aware of the prefixes delegated to the MN by the LMA upon the MN's interface attachment to the MAG, but is not aware of other MN's prefixes assigned to different interfaces). Upon reception of the message, MAG1 checks if it can forward flow Y, adds the required forwarding state so packets belonging to flow Y are delivered via the if1, and replies back to the LMA with an FMA message. Upon reception of this FMA message from MAG1, the LMA moves flow Y towards MAG1. 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.

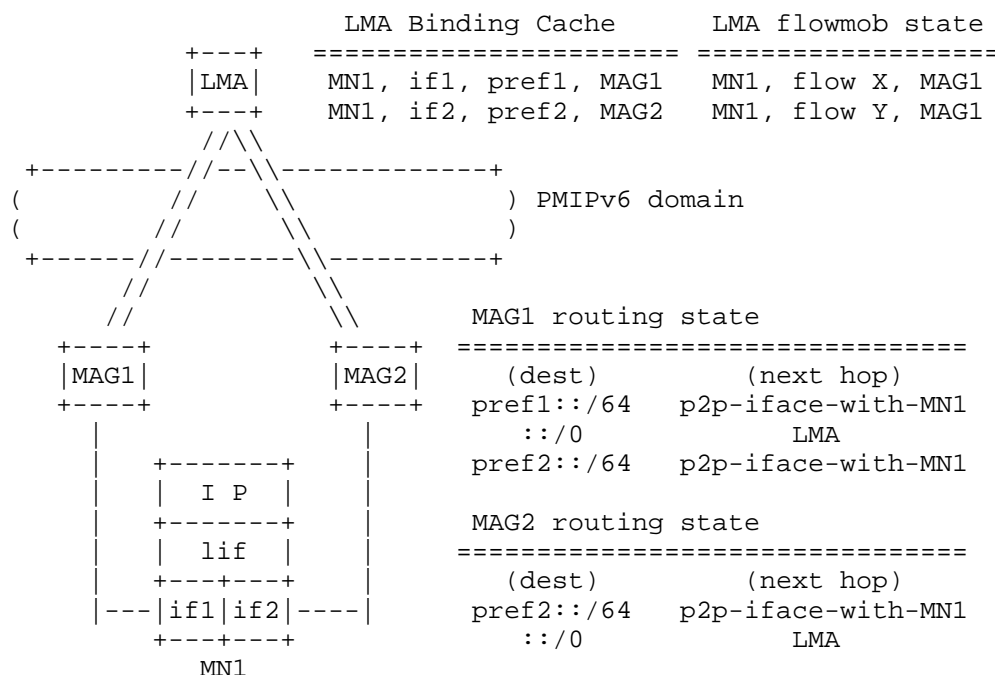


Figure 3: Unique prefix per physical interface scenario, with full flow granularity

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

4.1.2. Unique prefix per physical interface, with prefix granularity (partial handoff)

In this particular sub-scenario, only prefix mobility management granularity is supported (partial handover), meaning that it involves only transferring one (or some, but not all) of the prefixes that are allocated to an existing interface to a newly powered on interface.

This is particularly useful if the second interface is a newly connected interface, since that means there is no need for a new prefix to be allocated to the second interface. It is also useful in deployments where the operator requires only granularity of prefix instead of 5-tuples for flow mobility management, possibly due to it being less complex and having less impact to existing infrastructure. One example of operators needing only prefix granularity is 3GPP.

In such partial transfer of prefix(es) scenario (referred to as "partial handoff"), the target MAG will initiate the partial handoff

trigger using the PBU message towards the LMA. The PBU will carry the prefix(es) to be transferred in single or multiple HNP options. The partial handoff PBU message will additionally carry a new handoff indication option to indicate to the LMA that this PBU is for partial handoff of sub set of prefix(es). The LMA upon seeing this PBU from the target MAG with the partial handoff trigger embedded will update its binding cache entries. The LMA will remove the transferred prefix(es) from the binding cache associated with the connected interface and create a new binding cache for the newly powered on interface. The LMA may optionally send a message to source MAG 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. An example of the signaling sequence is shown in Figure 4.

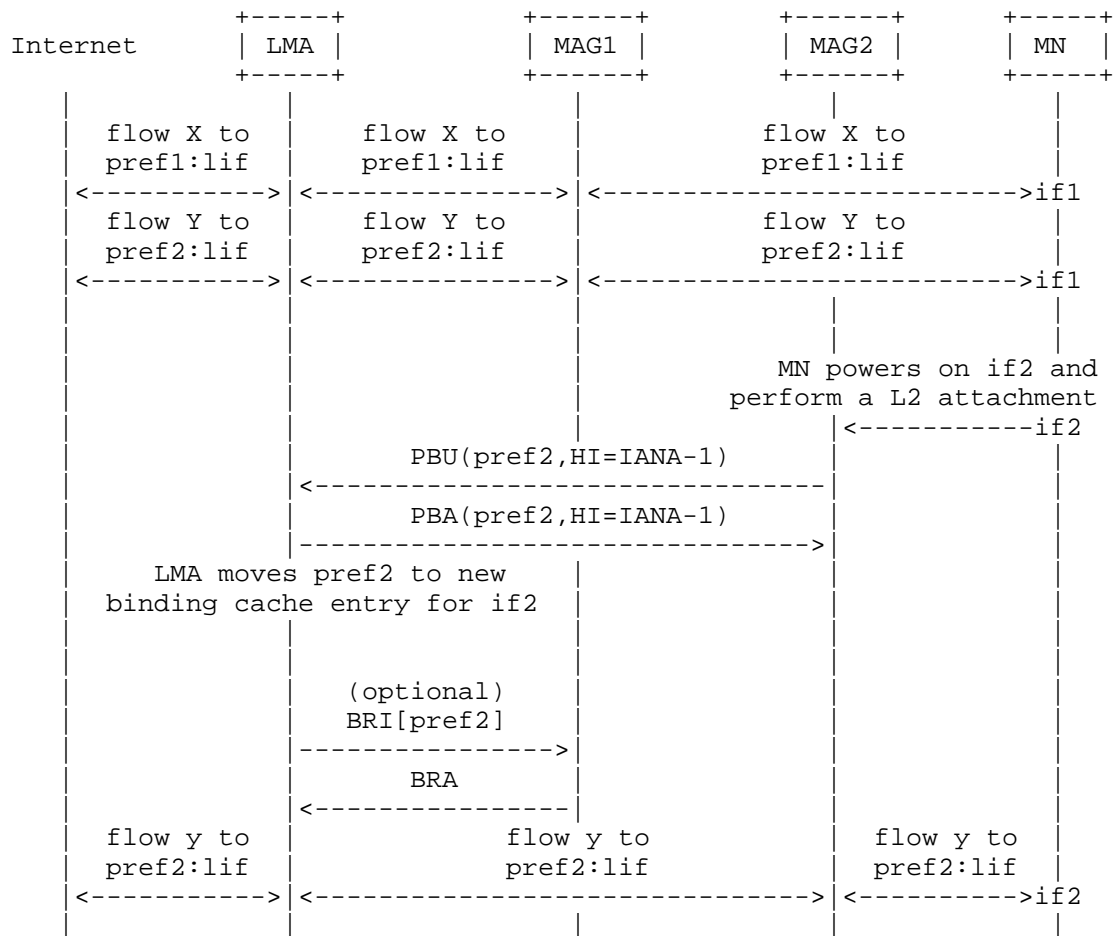


Figure 4: Message Sequence for Partial Handoff to a New Interface

In Figure 4, MAG1 is considered to proxy/advertise the two prefixes pref1 and pref2 that are assigned to if1 of the mobile node. From Figure 4 it can be seen that 2 flows (flow x, flow y) tied to prefixes pref1 and pref2 respectively are present. When the MN detects the availability of another access (e.g. WLAN access), it performs L2 attachment with MAG2 via this new access. MAG2 may receive a partial handoff trigger together with the L2 attachment. This trigger can be sent by other network entities (e.g. context transfer from MAG1 or a policy server) or sent via a layer 2 (L2) trigger from the MN during attachment. This partial handoff trigger will indicate transfer of pref2.

When MAG2 receives the trigger for transfer of pref2, it will send a

PBU message to the LMA with pref2 in the HNP option and a new value IANA-1 for the HI option. This new Handoff Indicator value IANA-1 indicates that this is a partial handoff to a new interface. When the LMA receives this PBU message, it will perform the following actions. The LMA will first locate an existing binding cache entry for mobile node with pref2. If the binding cache entry is present, the LMA will remove the pref2 from this existing entry, and insert the pref2 in the newly created binding cache entry for if2.

MAG1 can optionally be informed to stop proxying for the pref2. This can be done by LMA sending a BRI notification to MAG1 to revoke pref2. MAG1 upon receiving the BRI will send a BRA back to the LMA as shown in Figure 4.

After the partial handoff is completed, the new binding cache entry created for `if2` will have `pref2` assigned and the binding cache entry for `if1` will have `pref1`. This is shown in Figure 5.

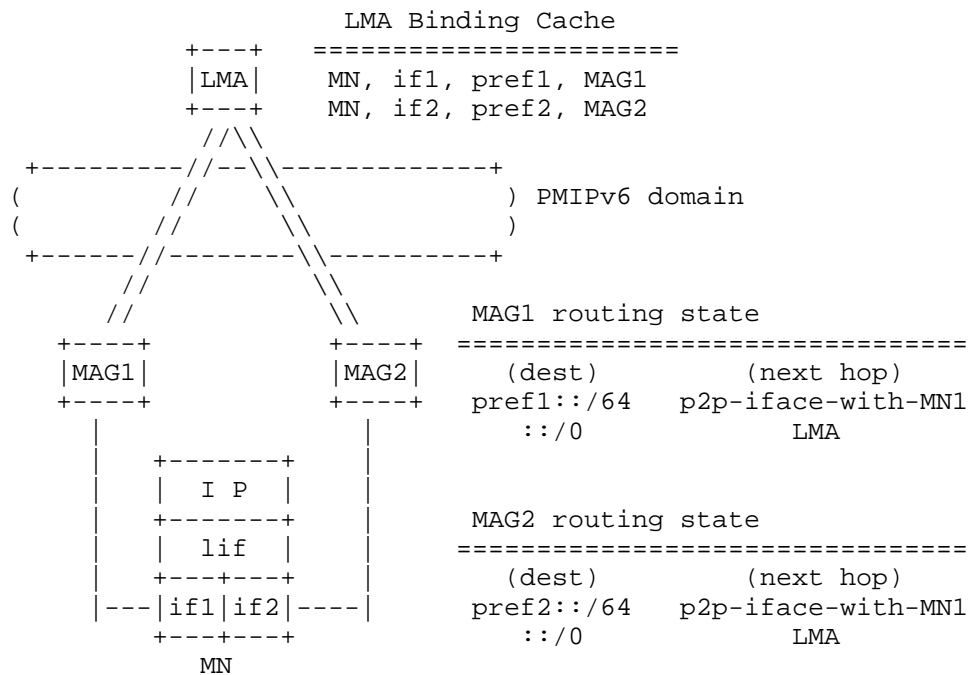


Figure 5: After Partial Handoff of pref2

4.2. Shared prefix across physical interfaces (per-MN HNP set)

In this scenario, physical interfaces of the mobile node are assigned the same prefix (or set of prefixes). The LMA maintains multiple binding cache entries (multiple mobility session) -- one per physical interface -- but they share the same HNP. How the shared prefix is routed by the LMA when there is no flow-specific state is left up to the implementation. This scenario is shown in Figure 6. Extensions to base PMIPv6 protocol as defined in RFC5213 are required to allow the LMA decide to assign the same prefix to a different interface of an MN already attached to the PMIPv6 domain.

There are different options that might be used to enable this scenario. While this is still TBD, one simple approach that can be assumed is the following. The LMA MAY know by using any mechanisms out-of-the-scope of this document that an MN has to be assigned the same prefix upon attachment of different interfaces (e.g. by consulting the MN's profile).




Figure 6: Shared prefix across physical interfaces scenario

In Figure 6, 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.

In this scenario, the LMA also decides how flows are routed from the LMA to the MN, and therefore, through which interface the MN receives packets from different flows.

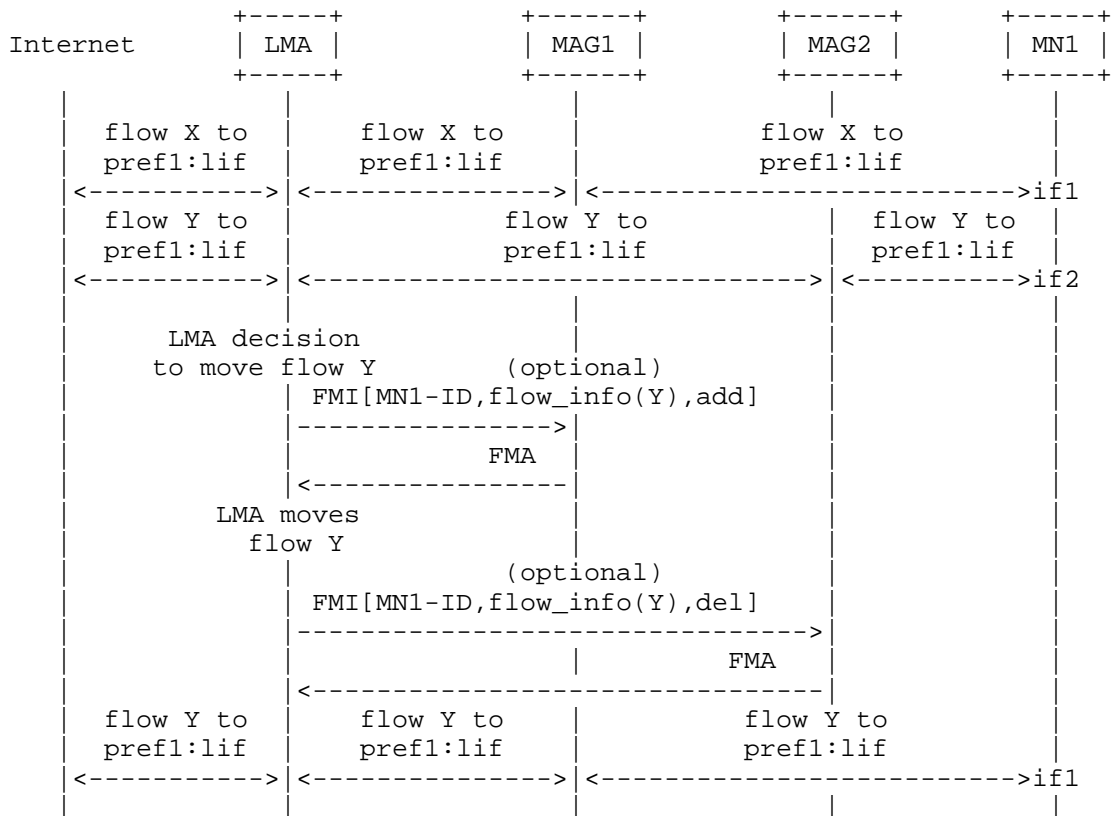


Figure 7: Flow mobility signaling for the shared prefix across physical interfaces scenario

An example of the signaling sequence is shown in Figure 7. The operation is analogous to the case of a unique prefix per physical interface. Note that in this scenario, if the target MAG does not need to perform flow-specific actions (e.g., QoS or accounting), the FMI/FMA signaling could be avoided, as no new routing state is required to forward a moved flow (since the prefix assigned to all physical interfaces is the same).

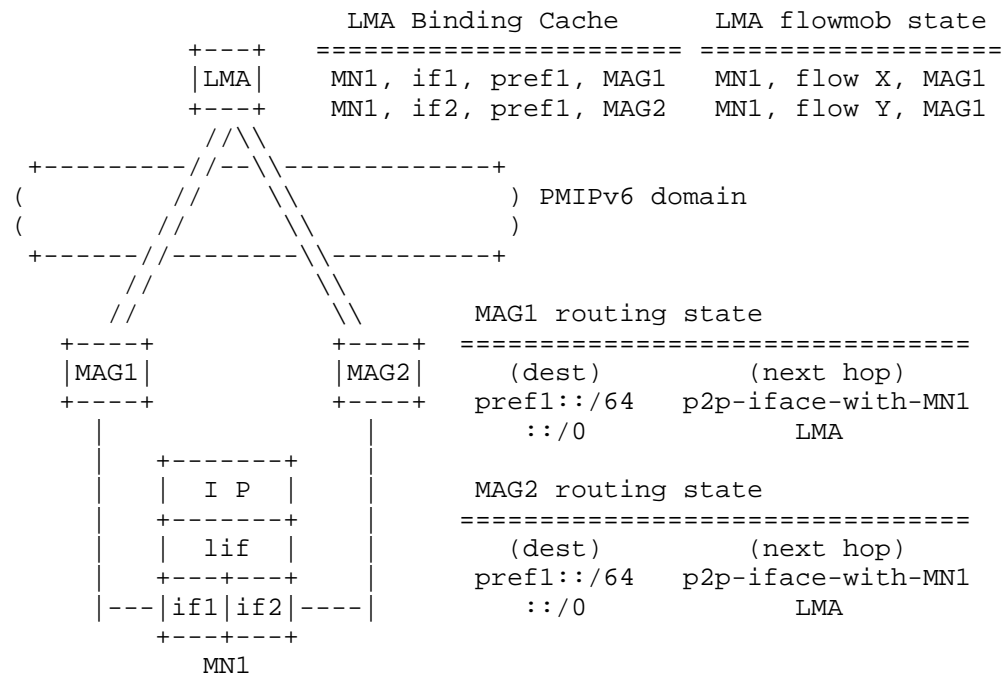


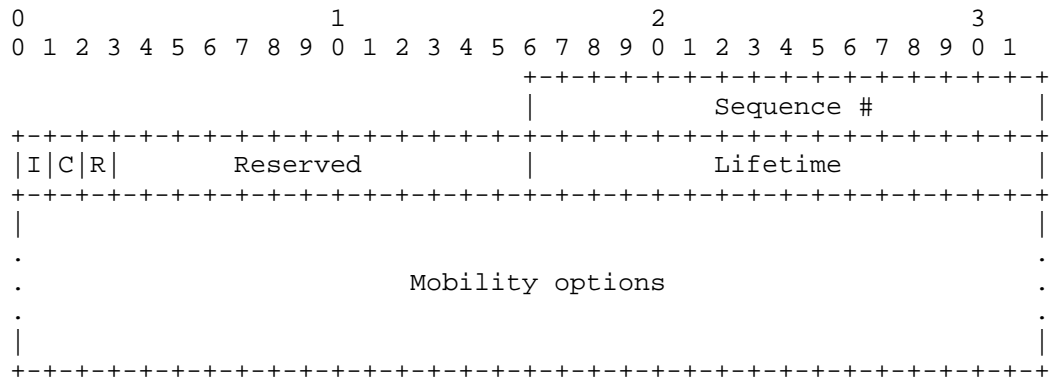
Figure 8: Shared prefix across physical interfaces scenario

Figure 8 shows the state of the different network entities after moving flow Y in the previous example.

5. Message formats

5.1. Flow Mobility Initiate (FMI)

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

**Sequence Number:**

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

'I' (initiate) flag:

Set to 1, indicates it is an FMI message.

'C' (cancel) flag:

When set to 1, indicates a request to remove state about the flow (cancel flow mobility). If set to 1, the Lifetime field MUST be set to 0.

'R' (refresh) flag:

When set to 1, indicates a request to refresh state about the flow. If the 'C' flag is set to 1, this flag should be set to 0 by the sender and ignored by the receiver.

Reserved:

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

Lifetime:

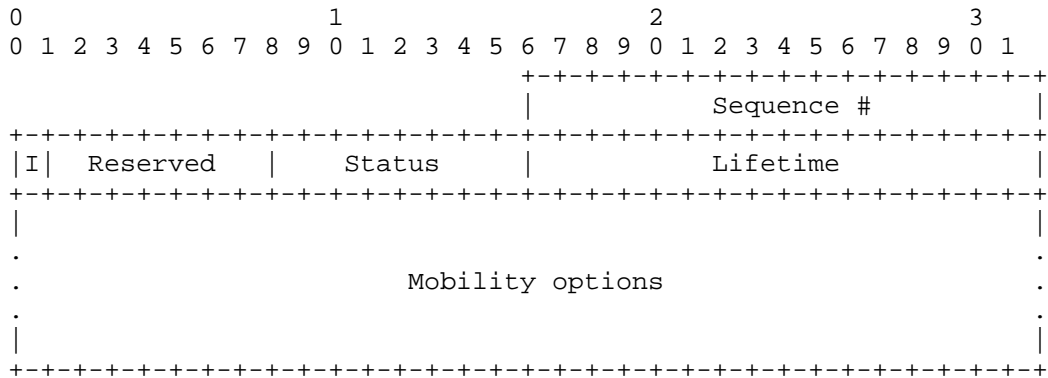
The requested time in seconds for which the LMA asks the MAG keep flow-specific state. A value of all one bits (0xffff) represents infinity.

Mobility Options:

MUST contain the MN-ID, followed by one or more Flow Identification Mobility options [I-D.ietf-mext-flow-binding].

5.2. 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 [I-D.ietf-mext-flow-binding].

6. Local Mobility Anchor considerations

This specification allows the LMA to control the distribution of specific traffic flows on different physical interfaces. This section details the LMA operations necessary to implement this specification.

6.1. Sending Flow Mobility Initiate messages

This specification allows the LMA to control and dynamically change the path used to deliver packets belonging to specific flows. This enables the LMA to have different forwarding rules for particular flows, in addition to the routes created upon regular PMIPv6 registration.

When creating a new Flow Mobility Cache entry (i.e. adding a new forwarding rule to allow flow traffic follow a different path from the one created upon regular PMIPv6 registration for the same destination prefix), the LMA includes the information needed to match the data packets to a specific flow in a Flow Identification Mobility option. This option MUST be included in an FMI message. This FMI message MUST have the cancel ('C') and refresh ('R') flags set to zero indicate that the FMI refers to a new flow. The FMA message MUST also include the MN-Identifier option of the mobile node the flow information refers to. More than one Flow Identification Mobility options MAY be included in the FMI message, but all of them MUST be subject to the same type of operation (e.g., creation of new mobility state). The LMA MUST create the corresponding Flow Mobility Cache entry upon receiving an FMA message with Status set to Success.

When canceling existing Flow Mobility state in the network (i.e. falling back to the default packet forwarding based solely on per HNP destination tunnels between LMA and MAG), the LMA sends an FMI message including the Flow Identification Mobility option(s) that refer to the flow(s) whose associated state is to be removed. This FMI message MUST have the cancel ('C') flag set to 1, and MUST include the MN-Identifier option. The LMA MAY decide to remove the corresponding Flow Mobility Cache entry at the MAG by sending this explicit signaling or by relying on the expiration of the associated timers.

The LMA MUST refresh the flow mobility state (i.e. FMC entry) at the MAG to prevent the MAG to stop forwarding specific flows upon expiration of the associated timers. When refreshing flow mobility state, the LMA sends an FMI message including the Flow Identification Mobility option(s) that refer to the flow(s) whose associated state is to be refreshed. This FMI message MUST have the refresh ('R') flag set to 1, and MUST also include the MN-Identifier option.

EDITOR'S NOTE: Retransmissions and Rate Limiting considerations TBD.

The IPv6 source address of an FMI message MUST be the LMA Address (LMAA). The destination IPv6 address of the FMI message MUST be set to the Proxy-CoA of the MAG which will create/cancel/refresh flow mobility state as indicated in the FMI message.

Security considerations stated in Section 4 of [RFC5213] "Proxy Mobile IPv6 Protocol Security" apply also for the signalling specified in this document.

EDITOR'S NOTE: Additional authentication and security requirements (if any) TBD.

6.2. Receiving Flow Mobility Acknowledge messages

Upon receiving a packet carrying a Flow Mobility Acknowledge, an LMA MUST validate the packet according to the following tests:

- o The packet meets the authentication requirements for Flow Mobility Acknowledges defined in Section XXX (TBD).
- o The IPv6 source address of the packet corresponds to the address of a MAG known by the LMA (NOTE: this is probably redundant once the security details are included).
- o The Sequence Number field matches the Sequence Number sent by the LMA to this destination address in an outstanding Flow Mobility Initiate.

Any Flow Mobility Acknowledge not satisfying all of these tests MUST be silently ignored.

When an LMA receives a packet carrying a valid Flow Mobility Acknowledge, the LMA MUST examine the Status field as follows:

- o If the Status field indicates that the Flow Binding Initiate was accepted (the Status field is less than 128), then the LMA MUST update the corresponding entry (or entries) in its Flow Mobility Cache, to indicate that the Flow Mobility Initiate has been acknowledged. The LMA MUST then stop retransmitting the FMI message. In addition, if the value specified in the Lifetime field in the FMA is less than the Lifetime value sent in the FMI being acknowledged, the LMA MUST subtract the difference between these two values from the remaining lifetime for the flow binding as maintained in the corresponding Flow Mobility Cache entry.

LMAs SHOULD send a new FMI well before the expiration of this period in order to extend the lifetime. This helps to avoid disruptions in communications which might otherwise be caused by network delays or clock drift.

- o If the Status Field indicates that the flow binding operation was rejected (the Status field is greater than or equal to 128), then the LMA can take steps to correct the cause of the error and retransmit the FMI (with a new Sequence Number value) subject to the rate limiting restrictions specified in Section XXX. If this is decided not to be done or it fails, then the LMA SHOULD record in its Flow Mobility Cache that future FMIs SHOULD NOT be sent to this destination. These considerations are of particular importance in case of creation/refresh of flow mobility state.
- o Additionally, for those Flow Mobility Cache entries that are newly created (not refreshed), the LMA MUST perform the actions required to ensure that the data packets matching the flow filters carried in the Traffic Selector sub-options are forwarded via the appropriate MAG. How this is done is left up to the implementation of the LMA.

6.3. Conceptual Data Structures

Each Local Mobility Anchor MUST maintain a Flow Mobility Cache. The Flow Mobility Cache MAY be implemented in any manner consistent with the external behavior described in this document. When sending a packet, the Flow Mobility Cache is searched before the Neighbor Discovery conceptual Destination Cache.

Each Flow Mobility Cache entry conceptually contains the following

fields:

- o The MN-Identifier of the mobile node for the flow this entry refers to. This field is used as primary key for searching the cache for update operations (deletion, refresh, cancel).
- o The flow filter information carried in the Traffic Selector sub-option. This information SHOULD be stored in a format that allows the LMA to forward packets matching the filter to the corresponding MAG (as indicated by the Proxy-CoA field stored in the Flow Mobility Cache entry).
- o The Proxy-CoA for which that the FMI carrying the information about this flow was sent.
- o The tunnel interface identifier (tunnel-if-id) of the bi-directional between the LMA and the MAG that MUST be used when forwarding packets belonging to the flow this entry refers to. This is internal to the local mobility anchor. The tunnel interface identifier is acquired during the tunnel creation in the standard Proxy Mobile IPv6 registration.
- o The initial value of the Lifetime field sent in the FMI.
- o The remaining lifetime of that flow binding. The lifetime value is initialized from the Lifetime field in the FMA that created or last modified this entry and is decremented until it reaches zero, at which time this entry MUST be deleted from the Flow Mobility Cache.
- o The maximum value of the Sequence Number field received in previous FMAs for this flow. The Sequence Number field is 16 bits long. Sequence Number values MUST be compared modulo 2^{16} as explained in Section 9.5.1 of [RFC3775].
- o The time at which an FMI was last sent regarding to this flow, as needed to implement the rate limiting restriction for sending FMIs.
- o The state of any retransmissions needed for FMIs referring to this flow. This state includes the time remaining until the next retransmission attempt for the FMI and the current state of the exponential back-off mechanism for retransmissions.
- o A flag specifying whether or not future FMIs should be sent to this destination.

The Flow Mobility Cache is used to determine whether a particular

packet belongs to a flow which has flow mobility state created -- and therefore needs to be processed separately from the rest of the packets -- or it can just be sent using normal packet processing rules as specified in RFC 5213.

6.4. Packet Processing considerations

The LMA MUST be able to forward packets matching the flow filters stored in the Flow Mobility Cache (carried in Traffic Selector sub-options inside the Flow Identification Mobility option carried in the FMIs) via the corresponding bi-directional tunnel.

For those packets with no matching Flow Mobility Cache, default PMIPv6 data forwarding considerations MUST be followed.

How the LMA ensures per-flow forwarding is left up to the particular implementation of the LMA.

7. Mobile Access Gateway considerations

This section details the MAG operations necessary to implement this specification.

7.1. Receiving Flow Mobility Initiate messages

This specification allows the MAG to deliver packets whose destination address could not match any destination network hosted at any interface of the MAG (i.e. a connected interface for that prefix) or sent by a mobile node with no matching binding. This enables the MAG to deliver/forward packets to/from IPv6 addresses that would not be known by a MAG not conforming to this specification.

Upon receiving a packet carrying a Flow Mobility Initiate, a MAG MUST validate the packet according to the following tests:

- o The packet meets the authentication requirements for Flow Mobility Initiates defined in Section XXX (TBD).
- o The IPv6 source address of the packet corresponds to the address of an LMA known by the MAG (NOTE: this is probably redundant once the security details are included).
- o The FMI contains one and only one MN-Identifier mobility option and one or more Flow Identification Mobility options.

Any Flow Mobility Initiate not satisfying all of these tests MUST be silently ignored.

In addition, if there is already a Flow Mobility Cache entry for that flow and the Sequence Number field stored in the entry is the same or greater than the sequence number carried in the FMI, then the MAG MUST send back an FMA with status code 127, and the last accepted sequence number in the Sequence Number field of the FMA.

When a MAG receives a packet carrying a valid Flow Mobility Initiate, the MAG MUST perform the following packet examinations:

- o If the MN-Identifier value carried into the FMI does not match any MN known to be connected to the receiving MAG, the MAG MUST send back an FMA with status code 129.
- o If the format used in any of the Traffic Selector sub-options is not supported by the receiving MAG, the MAG MUST send back an FMA with status code 131.
- o If the cancel ('C') flag is set to zero, it indicates that the FMI refers to new flow(s). The MAG SHOULD check in the Flow Mobility Cache if there is an entry referring to the flow(s) carried in the FMI. If there is already an entry for a flow with Lifetime greater than 0, then the the MAG SHOULD send back an FMA with status code 133.

The MAG MUST create the corresponding Flow Mobility state entry, and send back an FMA with status code 0 following the rules specified in Section 7.2.

- o If the refresh ('R') flag is set to 1, it indicates that the FMI refers to existing flow(s) whose state is to be refreshed. The MAG SHOULD check in the Flow Mobility Cache if there is an entry referring to the flow(s) carried in the FMI. If there is no entry, then the the MAG SHOULD send back an FMA with status code 132.

The MAG MUST update the corresponding Flow Mobility state entry or entries, and send back an FMA with status code 0 following the rules specified in Section 7.2.

- o If the cancel ('C') flag is set to 1, it indicates that the FMI refers to existing flow(s) whose state is to be removed. The MAG SHOULD check in the Flow Mobility Cache if there is an entry referring to the flow(s) carried in the FMI. If there is no entry, then the the MAG SHOULD send back an FMA with status code 132.

The MAG MUST remove the corresponding Flow Mobility state entry or entries, and send back an FMA with status code 0 following the

rules specified in Section 7.2.

7.2. Sending Flow Mobility Acknowledge messages

When constructing a packet carrying a Flow Mobility Acknowledge, the MAG MUST follow the following rules:

- o Security considerations stated in Section 4 of [RFC5213] "Proxy Mobile IPv6 Protocol Security" apply also for the signalling specified in this document.
- o EDITOR'S NOTE: Additional authentication and security requirements (if any) TBD.
- o The IPv6 source address of the packet corresponds to the egress interface of the MAG used to send the FMA to the LMA. (NOTE: this is probably redundant once the security details are included).
- o The IPv6 destination address MUST be set to the source address of the FMI being acknowledged.
- o The Sequence Number field MUST be copied from the Sequence Number given in the FMI.
- o The Lifetime field MUST be set to the remaining lifetime for the flow binding and MUST NOT be greater than the Lifetime value specified in the FMI. The MAG MAY decide to include a Lifetime value shorter than the one received in the FMI.
- o The values of the 'C' and 'R' flags MUST be copied from the values given in the FMI.
- o The Flow Identification Mobility options MUST be copied from the ones given in the FMI.

When a valid FMI is received, the MAG MUST update the Flow Mobility Cache entries accordingly as specified above. In addition, the MAG MUST perform the actions required to allow packets received from the LMA matching the flow filters stored in the Flow Mobility Cache to be delivered to the corresponding connected MN. How this forwarding is performed is up to the implementation of the MAG. Some considerations are included in Section 7.4.

7.3. Conceptual Data Structures

Each Mobile Access Gateway MUST maintain a Flow Mobility Cache. The Flow Mobility Cache MAY be implemented in any manner consistent with the external behavior described in this document. When sending a

packet, the Flow Mobility Cache is searched before the Neighbor Discovery conceptual Destination Cache.

Each Flow Mobility Cache entry conceptually contains the following fields:

- o The MN-Identifier of the mobile node for the flow this Flow Mobility Cache entry refers to. This field is used as primary key for searching the cache for update operations (deletion, refresh, cancel).
- o The flow filter information carried in the Traffic Selector sub-option. This information SHOULD be stored in a format that allows the MAG to deliver packets matching the filter to the corresponding MN (using the right interface where the MN is locally connected).
- o The Proxy-CoA from which that the FMA carrying the information about this flow was sent.
- o The tunnel interface identifier (tunnel-if-id) of the bi-directional between the LMA and the MAG that MUST be used when forwarding packets sent by the MN and belonging to the flow this entry refers to. This is internal to the MAG. The tunnel interface identifier is acquired during the tunnel creation in the standard Proxy Mobile IPv6 registration.
- o The interface identifier of the local interface of the MAG that MUST be used when delivering packets sent to the MN and matching the flow this entry refers to. This is internal to the MAG.
- o The remaining lifetime of that flow binding. The lifetime value is initialized from the Lifetime field in the FMI that created or last modified this entry and is decremented until it reaches zero, at which time this entry MUST be deleted from the Flow Mobility Cache.
- o The maximum value of the Sequence Number field received in previous FMIs for this flow. The Sequence Number field is 16 bits long. Sequence Number values MUST be compared modulo 2^{16} as explained in Section 9.5.1 of [RFC3775].

The Flow Mobility Cache is used to determine whether a particular packet belongs to a flow which has flow mobility state created -- and therefore needs to be processed separately from the rest of the packets -- or it can just be sent using normal packet processing rules as specified in RFC 5213.

7.4. Packet Processing considerations

The MAG MUST be able to forward packets matching the flow filters stored in the Flow Mobility Cache (carried in Traffic Selector sub-options inside the Flow Identification Mobility options carried in the FMIs) via the corresponding bi-directional tunnel.

For those packets with no match Flow Mobility Cache, default PMIPv6 data forwarding considerations MUST be followed.

How the MAG ensures per-flow forwarding is left up to the particular implementation of the MAG.

It might happen that the LMA initiates the movement of a flow, by sending the FMI/FMA signalling, but there is no traffic to the MN. In that case, the MN cannot learn which is the uplink interface that should be used. In order to avoid problems in this particular situation, the MAG SHOULD allow uplink traffic to pass through -- even if it does not match the flow filters stored in the FMC (at least for traffic that would be sent through that MAG in case flow mobility was not enabled).

8. 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 the logical interface at the MN "replicates" the behavior observed for downlink packets on a per-flow basis. This means that if packets belonging to flow X are received from interface if1, then the MN sends packets belonging to that flow (in the uplink) also using if1. 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.

9. IANA Considerations

TBD.

10. Security Considerations

TBD.

11. 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

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

Frank Xia

E-mail: xiayangsong@huawei.com

12. Acknowledgments

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

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5846] Muhanna, A., Khalil, M., Gundavelli, S., Chowdhury, K., and P. Yegani, "Binding Revocation for IPv6 Mobility", RFC 5846, June 2010.

13.2. Informative References

- [I-D.ietf-mext-flow-binding]
Tsirtsis, G., Soliman, H., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and NEMO Basic Support", draft-ietf-mext-flow-binding-11 (work in progress), October 2010.
- [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.

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/>

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

Y. Hong
ETRI
J. Youn
DONG-EUI Univ.
October 25, 2010

Hybrid home network prefix for multihoming in PMIPv6
draft-hong-netext-hybrid-hnp-03

Abstract

Proxy Mobile IPv6 (PMIPv6) supports multihoming where a mobile node can connect to a PMIPv6 domain through multiple interfaces for simultaneous access or inter-technology handoff. However, for an inter-technology handoff, PMIPv6 does not allow simultaneous access since all the home network prefixes associated with one interface are delivered to another interface of a mobile node. In addition, if we assume that the flow is classified by home network prefix, then the PMIPv6 cannot support flow mobility since it requires moving just some home network prefixes between interfaces. In this document, we propose a hybrid home network prefix assignment (HHNPA) scheme where both the static prefix model and the dynamic prefix model are used. By using these two different home network prefix models, it can support flow mobility that some home network prefixes are moved to another interface and some home network prefixes are remained at existing interface.

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 April 28, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to BCP 78 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. Requirements Language	4
3. Problem statement of multihoming support in PMIPv6	5
4. Hybrid home network prefix assignment	8
5. Security Considerations	13
6. IANA Considerations	14
7. References	15
7.1. Normative References	15
7.2. Informative References	15
Authors' Addresses	17

1. Introduction

Mobile IPv6 [RFC3775] and Mobile IPv4 [RFC3344] are host based IP mobility support protocols. On the other hand, Proxy Mobile IPv6 (PMIPv6) [RFC5213] is a network based IP mobility support protocol, which does not require any modifications to mobile nodes (MNs). PMIPv6 makes it possible to support mobility for IPv6 nodes without an involvement of a MN. That is, on behalf of MNs, a mobile access gateway (MAG) in the network performs the signaling for mobility management with a local mobility anchor (LMA).

PMIPv6 defines basic operations for registration, deregistration, and tunnel management. However, it does not define protocol operations for supporting seamless handover for a MN with multiple network interfaces (i.e., inter-technology handoff) [I-D.devarapalli-netext-multi-interface-support-00]. While PMIPv6 itself supports handover across different interfaces and access types, there are several issues for efficient inter-technology handoff, e.g., how to set interface identifier, how to use the same address on multiple interfaces, how to select an access technology, and how to detect and manage a handover event [I-D.krishnan-netext-intertech-ps].

PMIPv6 basically supports multihoming where a MN can connect to a PMIPv6 domain through multiple interfaces for simultaneous access and inter-technology handoff between different multiple interfaces. If a MN with multiple interfaces connects to a PMIPv6 domain over multiple access networks, the LMA will allocate a unique set of home network prefixes (HNPs) for each of the connected interfaces. However, when the MN performs an inter-technology handoff, the LMA will newly assign all the home network prefixes, which are associated with the first interface, to the second interface. Consequently, the existing mobility sessions with the second interface will be disrupted. If we assume that the flow is classified by home network prefix, then the PMIPv6 cannot support flow mobility since it requires moving just some home network prefixes between interfaces.

Fundamentally, this problem has been caused due to the fact that each of the attached interfaces must be assigned one or more unique prefixes to in PMIPv6. Therefore, to solve this problem, we propose a hybrid home network prefix assignment (HHNPA) scheme where both the static prefix model and the dynamic prefix model are employed and dynamically selected. That is, for IP session continuity after inter-technology handoff, a dynamic home network prefix is used on both interfaces.

2. Requirements Language

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [RFC2119].

3. Problem statement of multihoming support in PMIPv6

To support multiple mobility sessions for a single MN with multiple interfaces, PMIPv6 creates mobility sessions per interface and each mobility session should be managed as a separate binding cache (BC) entry. Note that although the LMA can allocate more than one home network prefix to an interface, all these prefixes are managed by one mobility session. The LMA allows a handoff between two different interfaces of a MN. In such a scenario, all the home network prefixes associated with one interface will be delivered to another interface of the MN. The decisions on when to create a new mobility session and when to update an existing mobility session are based on the handover hint included in the Proxy Binding Update (PBU) message.

Basic operations for multiple interfaces in PMIPv6 are as follows. First of all, the MAG decides whether to inform the LMA of the attachment of the second interface (or an inter-technology handoff), and then the MAG sends a Proxy Binding Update message. When the LMA receives the PBU message, it verifies the request message. If the PBU message includes a handoff indicator flag of 1 (i.e., initial attachment), the LMA allocates a new mobility session for second interface and thus the LMA has multiple Binding Cache entries. On the other hand, if the PBU message has a handoff indicator flag of 2 (i.e., inter-technology handoff), all the home network prefixes associated with the first interface will be associated with the second interface. Figures 1 and 2 describe the operations of two cases: initial attachment and inter-technology handoff. In particular, as shown in Figure 2, for inter-technology handoff, HNP_2 is overwritten with HNP_1 and then the existing mobility session 1 is removed from the binding cache entry at the LMA.

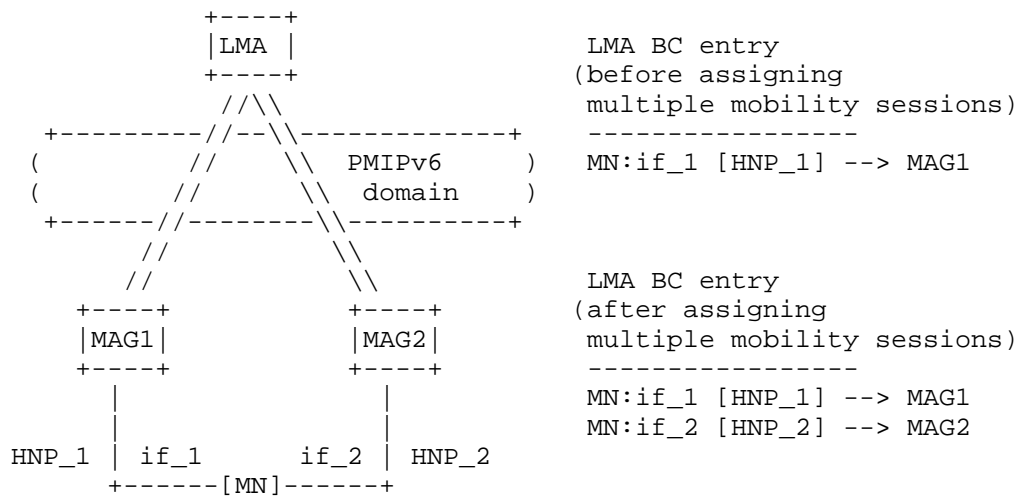


Figure 1: Multihoming in PMIPv6 : Initial attachment




Figure 2: Multihoming in PMIPv6 : Inter-technology handoff

These multihoming operations in PMIPv6 have the following problems.

First, since the LMA moves the same home network prefix assigned to the first interface to the second interface after inter-technology

handoff and updates the existing Binding Cache entry, the previous binding information of the second interface can be removed. Therefore, the existing flows through the second interface may be disrupted [I-D.devarapalli-netext-multi-interface-support-00].

In addition, since there is no way to enable flow-specific handoffs, compelled handoff of unwanted IP flows can be performed due to inter-technology handoff. The existing multihoming operations have drawbacks in terms of implementation. If we assume that the flow is classified by home network prefix, then the PMIPv6 cannot support flow mobility since it requires moving just some home network prefixes between interfaces.

As mentioned before, the MAG should set the handoff indicator flag (e.g., 1 for initial attachment or 2 for inter-technology handoff). However, the MAG does not have sufficient information on multihoming support and thus it is not an easy task to distinguish initial attachment and handoff between interfaces [I-D.krishnan-netext-intertech-ps].

Moreover, all home network prefixes are determined when the first mobility session is generated for a corresponding interface. Therefore, there is no way to add or delete a new home network prefix [I-D.jeyatharan-netext-multihoming-ps]. To address these issues, we propose a hybrid home network prefix assignment scheme in the next section.

4. Hybrid home network prefix assignment

In terms of home network prefix allocation in PMIPv6, the per-MN prefix model is mandatory whereas the shared prefix model is not supported. In the per-MN prefix model, home network prefixes are assigned to a MN and these prefixes are exclusively used by the MN; no other MNs share the home network prefix. Note that all assigned prefixes are unique and they are parts of one mobility session. If the MN simultaneously attaches to a PMIPv6 domain through multiple interfaces, each of the attached interfaces must be assigned one or more unique prefixes. Therefore, after performing an inter-technology handoff based on the per-MN prefix model, simultaneous access to the PMIPv6 domain is not allowed. If we assume that the flow is classified by home network prefix, then the PMIPv6 cannot support flow mobility since it requires moving just some home network prefixes between interfaces.

If it is able to separate prefix for the purpose of inter-technology handoff and the purpose of simultaneous access, PMIPv6 may support both interface handoff and simultaneous access at the same time and it may support flow mobility where some flow related to specific prefix only moved to another interface. Based on this idea, we propose a hybrid home network prefix assignment (HHNPA) scheme to use both the static and dynamic home network prefix models.

Figure 3 introduces the concept of the HHNPA scheme where the LMA divides home network prefixes into static home network prefixes and dynamic home network prefixes. The static home network prefix (based on the per-MN prefix model) is used for simultaneous access. On the other hand, the dynamic home network prefix is employed for only inter-technology handoff.




Figure 3: Prefix model in HHNPA

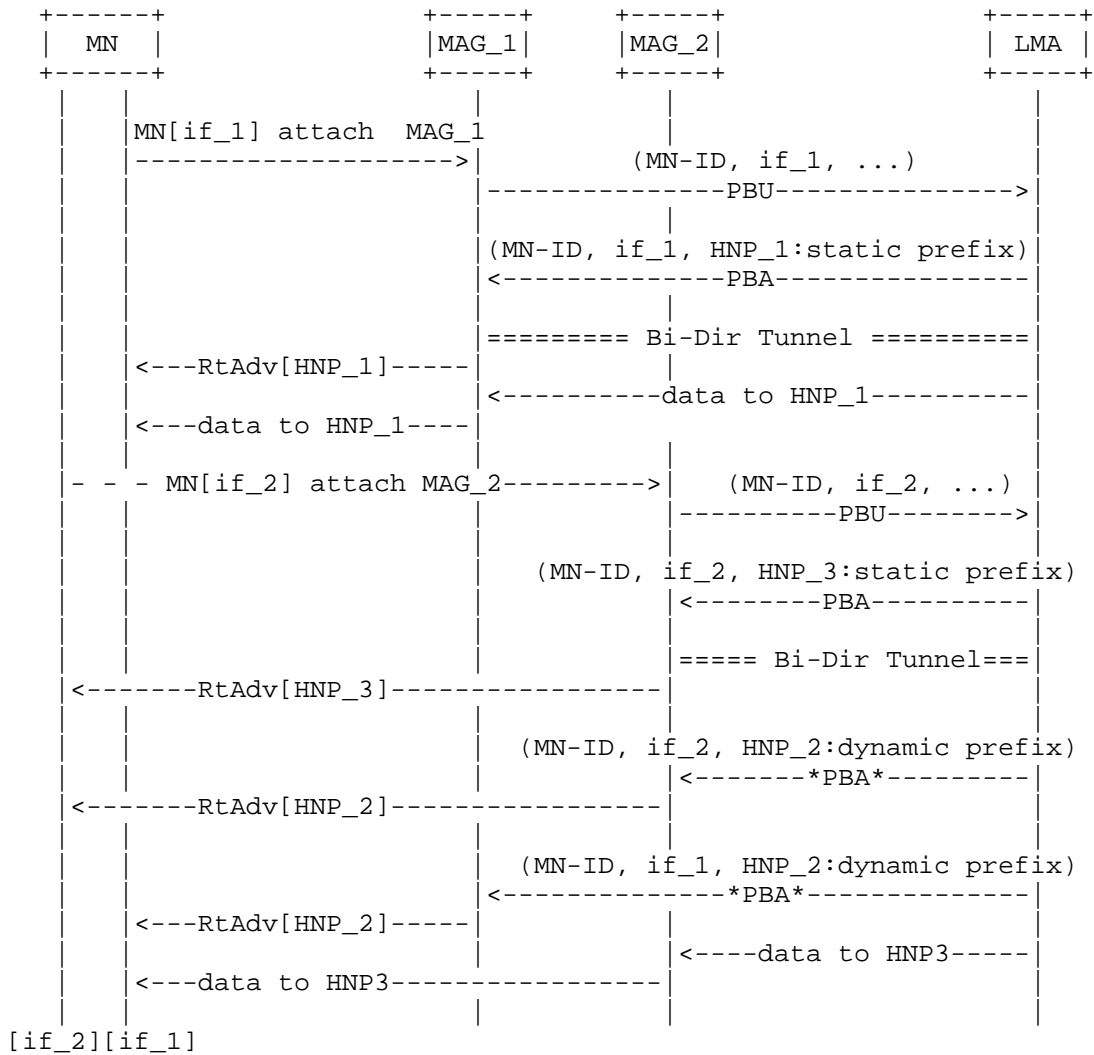


Figure 5: Message flow of HHNPA scheme

After the completion of attachment of the MN through two interfaces, the operations of inter-technology handoff of the MN are as depicted in Figure 6 and Figure 7. As shown in Figure 6, before inter-technology handoff, the dynamic prefix HNP_2 is activated in if_1. If the MN wants inter-technology handoff, the MN gives some hint of inter-technology handoff to the MAG_2. The MAG_2 which receives hint of inter-technology handoff sends a PBU message with a handoff indicator value of 2 to the LMA. The LMA which receives this PBU

message updates the Binding Cache entry filed which is related to dynamic prefix HNP_2. The updated fields of Binding Cache entry are interface and tunnel information. After the completion of updating of Binding Cache entry, the LMA sends data which is relation to HNP_2 to the MAG_2. As shown in Figure 6, after inter-technology handoff, the dynamic prefix HNP_2 is activated in if_2.




Figure 6: Usage scenario of dynamic prefix in HHNPA scheme

The HHNPA scheme has advantages of deciding a handoff indication flag. After the LMA assigns static and dynamic home network prefixes to MAGs, if MAG_2 receives handoff hint, MAG_2 checks the existence of the MN_identifier. If the MN_identifier related to the dynamic prefix HNP_2 exists in routing table, the MAG_2 acknowledges that there is a mobility session of inter-technology handoff. So, the MAG_2 can easily determine the value of handoff indication. This operation is depicted in Figure 7.

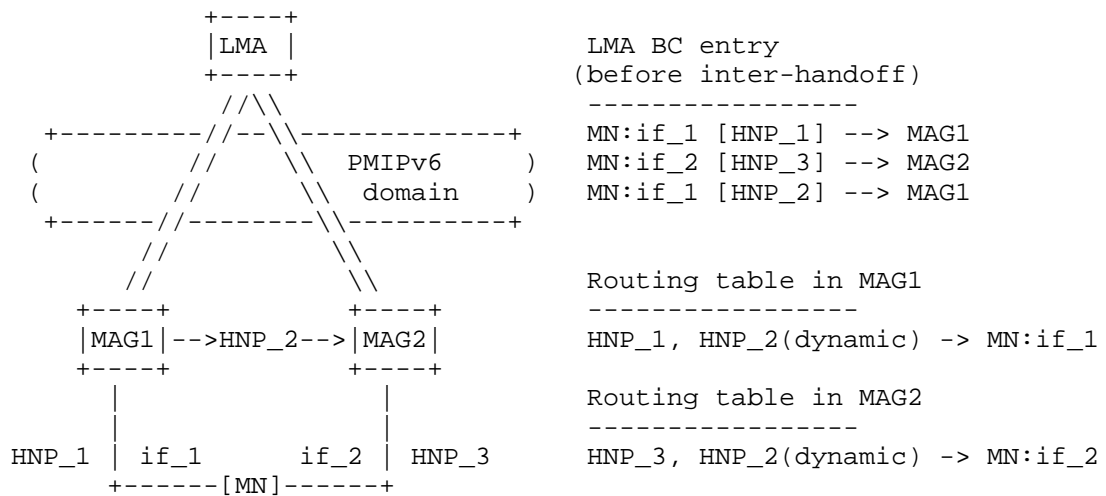


Figure 7: Setting of handoff indicator flag in HHNPA scheme

If the MN in HHNPA scheme utilizes a logical interface, it can hide the change of network interface from the host IP layer [I-D.ietf-netext-logical-interface-support-00]. If one typical application uses dynamic prefixes at the starting time of communication, the session continuity through inter-technology handoff is supported. But, if the typical application uses static prefixes at the starting time of communication, inter-technology handoff is not supported. As described in internet draft [I-D.hong-netext-scenario-logical-interface-00], if the logical interface supports the change of network interface and also the change of home network prefix (only one home network prefix is shown to IP host stack), the typical application that uses static prefix at the starting time of communication has also session continuity through inter-technology handoff.

5. Security Considerations

TBD.

6. IANA Considerations

This document has no actions for IANA.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

7.2. Informative References

- [I-D.devarapalli-netext-multi-interface-support-00]
Devarapalli, V., Kant, N., Lim, H., and C. Vogt, "Multiple Interface Support with Proxy Mobile IPv6, draft-devarapalli-netext-multi-interface-support-00", March 2009.
- [I-D.hong-netext-scenario-logical-interface-00]
Hong, Y. and J. Youn, "Scenarios of the usage of multiple home network prefixes on a logical interface, draft-hong-netext-scenario-logical-interface-00 (work in progress)", July 2010.
- [I-D.ietf-netext-logical-interface-support-00]
Melia, T. and S. Gundavelli, "Logical Interface Support for multi-mode IP Hosts, draft-ietf-netext-logical-interface-support-00 (work in progress)", July 2010.
- [I-D.jeyatharan-netext-multihoming-ps]
Jeyatharan, M. and C. Ng, "Multihoming Problem Statement in NetLMM, draft-jeyatharan-netext-multihoming-ps-01", March 2009.
- [I-D.krishnan-netext-intertech-ps]
Krishnan, S., Yokota, H., Melia, T., and C. Bernardos, "Issues with network based inter-technology handovers,

draft-krishnan-netext-intertech-ps-02", July 2009.

Authors' Addresses

Yong-Geun Hong
ETRI
161 Gajeong-Dong Yuseung-Gu
Daejeon, 305-700
Korea

Phone: +82 42 860 6557
Email: yonggeun.hong@gmail.com

Joo-Sang Youn
DONG-EUI Univ.
Busan,
Korea

Phone: +82 51 890 1993
Email: joosang.youn@gmail.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

Y. Hong
ETRI
J. Youn
DONG-EUI Univ.
October 25, 2010

Scenarios of the usage of multiple home network prefixes on a logical
interface
draft-hong-netext-scenario-logical-interface-02

Abstract

A logical interface is used to hide the existence and the change of physical interface from the IP layer of a host and it also can be used for a multiple interfaces mobile node in PMIPv6 domain. If a LMA assigns multiple home network prefixes to a multiple interfaces mobile node with a logical interface, there are various usages of multiple home network prefixes on the logical interface. As following general PMIPv6 operations described in RFC 5213, all multiple home network prefixes are shown to the IP layer. Also, the logical interface hides the existence of multiple home network prefixes and shows only one home network prefix to host IP layer. And in a LMA point of view, a LMA may acknowledge each physical interface or only logical interface of a mobile node when a multiple interfaces mobile node utilizes the logical interface. In this internet draft, we describe various scenarios of the usage of multiple home network prefixes on a logical interface.

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 April 28, 2011.

Copyright Notice

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

This document is subject to BCP 78 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. Requirements Language	5
3. Multihoming support on a logical interface in PMIPv6	6
4. The usage of multiple home network prefixes on a logical interface	8
4.1. Scenario 1 - LMA acknowledges each physical interface of a MN	8
4.2. Scenario 2 - LMA acknowledges logical interface of a MN	9
4.3. Scenario 3 - All home network prefixes are shown to IP host stack	10
4.4. Scenario 4 - Only one home network prefix is shown to IP host stack	12
5. Security Considerations	14
6. IANA Considerations	15
7. References	16
7.1. Normative References	16
7.2. Informative References	16
Authors' Addresses	17

1. Introduction

Proxy Mobile IPv6 (PMIPv6)[RFC5213] is a network based IP mobility support protocol, which does not require any modifications to mobile nodes. PMIPv6 makes it possible to support mobility for IPv6 nodes without an involvement of mobile nodes. That is, on behalf of mobile nodes, a mobile access gateway (MAG) in the network performs the signaling for mobility management with a local mobility anchor (LMA).

Due to the simultaneous usage of multiple interfaces in a mobile node and the change of network interfaces, IETF netext working group has studied the hiding access technology changes from host IP layer. Link layer implementation such as "logical interface" can hide the actually used physical interfaces and the change of physical interfaces from the IP layer. Many operating systems offer support to the usage of logical interface over multiple physical interfaces without any big efforts. In another internet draft [I-D.ietf-netext-logical-interface-support-00], operations details of the logical interface are explained and it identifies the applicability of the usage of logical interface.

PMIPv6 specifications allow mobile nodes to connect PMIPv6 domain through multiple network interfaces for simultaneous access and a LMA may allocate more than one home network prefix for a given interface of the mobile node. In basic PMIPv6 specification, the multiple home network prefixes that the LMA allocates to a mobile node are shown totally to the host IP stack of the mobile node. Also, if a mobile node utilizes a logical interface to hide the change of network interfaces, we can extend the usage of a logical interface to hide the multiplicity of home network prefixes. So we can categorize the usage of a logical interface as two cases. First, as following general PMIPv6 operations described in RFC 5213, all multiple home network prefixes are shown to host IP layer. Second, for some reasons such as the session continuity, only one home network prefix is shown to the host IP stack of the mobile node even though the LMA allocates multiple home network prefixes. In second case, the logical interface does some necessary jobs to hide the existence of multiple interfaces and also the existence of multiple home network prefixes. And the logical interface does some necessary jobs when network interfaces are changed and/or home network prefixes are changed.

Although the primary purpose of a logical interface is to hide the access technology changes from host IP layer in a mobile node, the usage of a logical interface can influence the acknowledgement of interfaces of a mobile node in a LMA point of view. If a multiple interfaces mobile node utilizes a logical interface, we can consider both cases; a LMA may acknowledge each physical interface of a

multiple interfaces mobile node; a LMA may acknowledge only a logical interface of a multiple interfaces mobile node.

2. Requirements Language

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [RFC2119].

3. Multihoming support on a logical interface in PMIPv6

In PMIPv6 specification, to provide multihoming support, it describes three key aspects as following.

- o When a mobile node connects to a PMIPv6 domain through multiple interfaces for simultaneous access, the LMA MUST allocate a mobility session for each of the attached interfaces.
- o The LMA MAY allocate more than one home network prefix for a given interface of the mobile node.
- o The LMA MUST allow for a handoff between two different interfaces of a mobile node. In such a scenario, all the home network prefixes associated with one interface (part of one mobility session) will be associated with a different interface of the mobile node.

Even though a logical interface is adopted in a multiple interface mobile node in PMIPv6 domain, the above three aspects SHOULD be applied.

For the regarding of first aspect, if a logical interface is used in a multiple interfaces mobile node, even though multiple interfaces are used, the host IP stack of the mobile node only acknowledges the logical interface. But in a LMA point of view, a LMA can acknowledge each physical interface if link-layer identifier of each physical interface is delivered to a LMA through a PBU (Proxy Binding Update) message. Or a LMA can acknowledge only logical interface if link-layer identifier of logical interface is delivered to a LMA through PBU message. Section 4.1 and section 4.2 describe these scenarios. If a LMA acknowledges each physical interface, the LMA allocates mobility sessions for each interface and each mobility session for the each physical interface is generated. But, if a LMA acknowledges only a logical interface, the LMA allocates a mobility session for the logical interface of the mobile node and one mobility session for the logical interface is generated.

For the regarding of second aspect, in basic PMIPv6 specification, if a LMA allocates multiple home network prefixes to a multiple interfaces mobile node, these multiple home network prefixes are totally shown to the host IP stack of the mobile node. This is also effective when a logical interface is used. But, due to the usage of logical interface in a multiple interfaces mobile node in PMIPv6 domain, we can extend the allocation method of multiple home network prefixes to a mobile node. As same as the role of a logical interface to hide the change of multiple interfaces, the logical interface can hide the existence and the change of multiple home

network prefixes. In section 4.3, section 4.4, we describe more detail.

For the regarding of third aspect, if a LMA acknowledges only a logical interface, this description is no longer effective because only one logical interface is shown to the LMA and the LMA assumes that the mobile node has one interface even though it has multiple physical interfaces. But, the PMPv6 specification SHOULD be extended to consider a handoff between two different physical interfaces of a mobile node with a logical interface.

4. The usage of multiple home network prefixes on a logical interface

4.1. Scenario 1 - LMA acknowledges each physical interface of a MN

The first scenario of the usage of multiple home network prefixes on a logical interface is same as described in internet draft of "Proxy Mobile IPv6 Extensions to Support Flow Mobility [I-D.bernardos-netext-pmipv6-flowmob-00]." The role of a logical interface is only to hide the existence of multiple network interfaces and the change of network interfaces for inter-technology handoff and flow mobility in a mobile node point of view. So, the logical interface is transparent to the LMA. In this scenario, link-layer identifiers of each physical interface are delivered to a LMA through proxy binding update. So, the LMA acknowledges each physical interfaces and it generates each mobility session for each physical interface.

Figure 1 shows this case. As shown in Figure 1, if the mobile node uses two physical interfaces if_1 and if_2, these two different physical interfaces are referred as different LL-ID (link-layer identifier) xxx, yyy. As a result of proxy binding update, two different mobility sessions of two physical interfaces are generated.

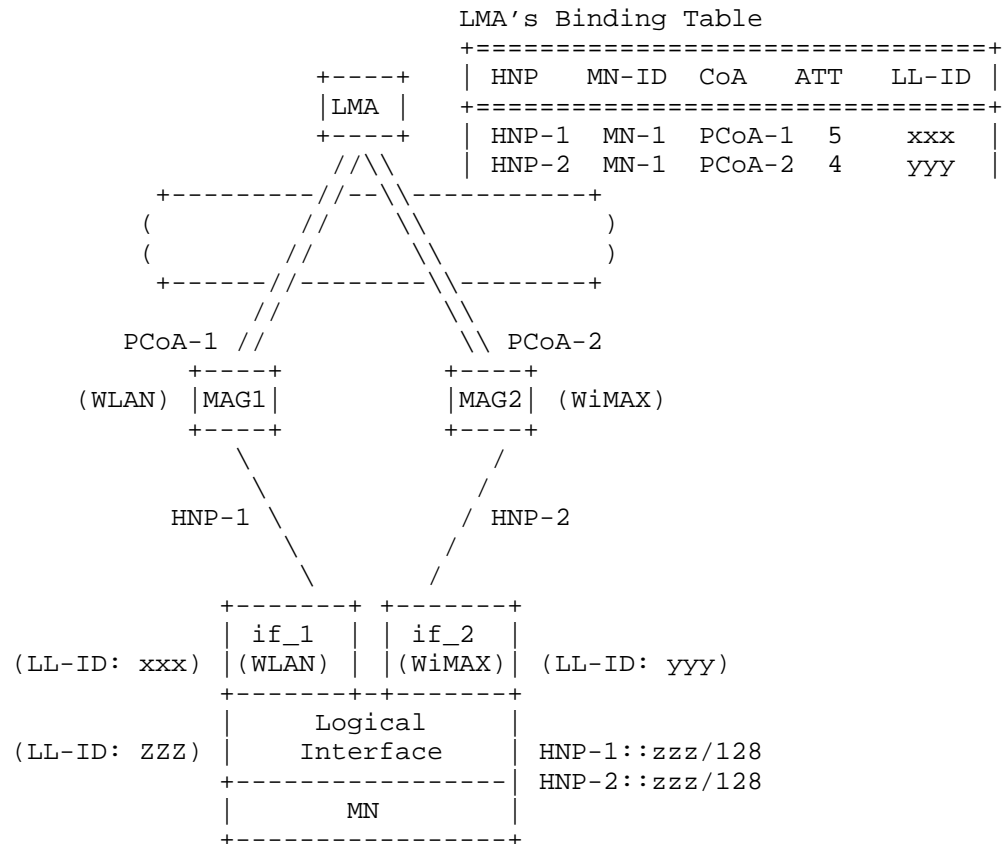


Figure 1: LMA acknowledges each physical interface of a MN

4.2. Scenario 2 - LMA acknowledges logical interface of a MN

The role of a logical interface is to hide the existence of multiple network interfaces and the change of network interfaces for inter-technology handoff and flow mobility in a mobile node point of view. Also, in a LMA point of view, the LMA acknowledges only a logical interface and it generates mobility sessions for the logical interface. In this scenario, only link-layer identifier of a logical interface is delivered to a LMA during proxy binding update even though multiple physical interfaces are used.

Figure 2 shows this case. As shown in Figure 2, the mobile node uses two physical interfaces if_1 and if_2 and these two different physical interfaces are referred as different LL-ID xxx, yyy. But, instead of using LL-IDs of each physical interface, the LL-ID of the

logical interface zzz is used in a LMA's binding table. As a result of proxy binding update, mobility session of the logical interface is generated.

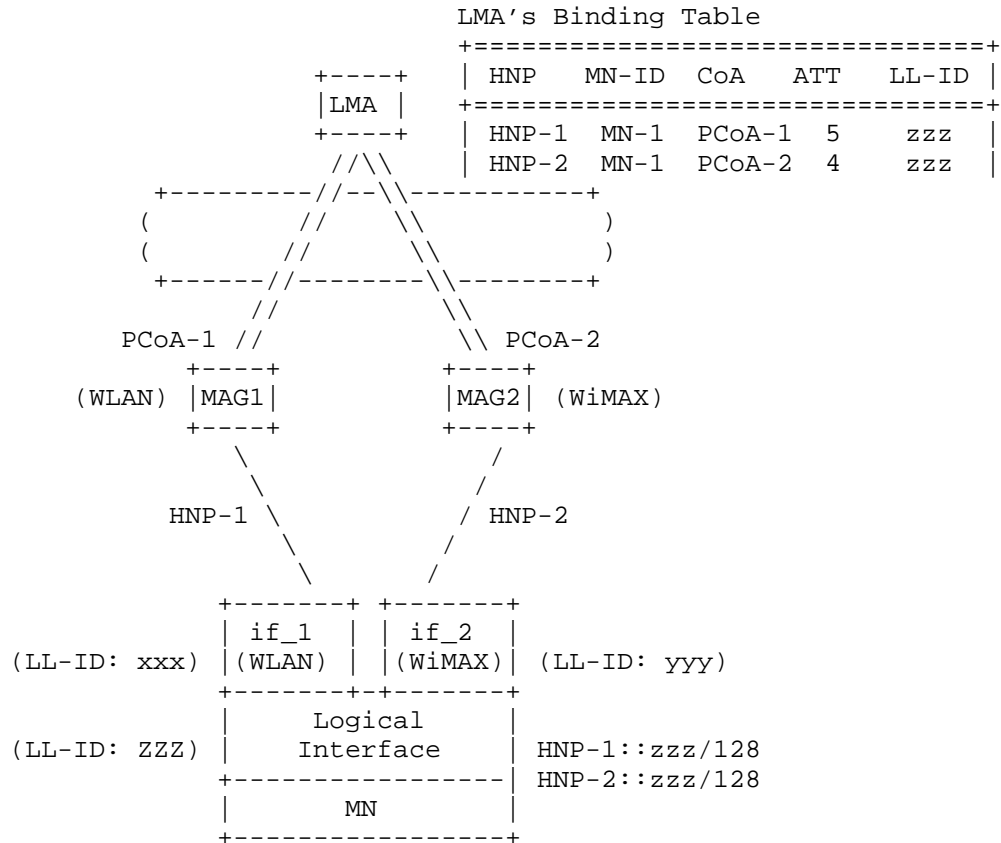


Figure 2: LMA acknowledges logical interface of a MN

4.3. Scenario 3 - All home network prefixes are shown to IP host stack

The third scenario of the usage of multiple home network prefixes on a logical interface is same as described in internet draft of "Logical Interface Support for multi-mode IP Hosts [I-D.ietf-netext-logical-interface-support-00]" and "Proxy Mobile IPv6 Extensions to Support Flow Mobility [I-D.bernardos-netext-pmipv6-flowmob-00]". The role of a logical interface is only to hide the existence of multiple network

interfaces and the change of network interfaces. The home network prefixes that a LMA allocates to a multiple interface mobile node are totally shown to host IP stack of the mobile node. The logical interface simply bypasses the home network prefixes between host IP stack and each physical interface.

Figure 3 shows this case. As shown in Figure 3, if the LMA allocates two home network prefixes HNP-1, HNP-2 to the mobile node, these two home network prefixes are shown to IP host stack of the mobile node. As a result of receiving multiple home network prefixes in IP host stack, two different IPv6 addresses HNP-1::zzz/128 and HNP-2::zzz/128 can be generated.

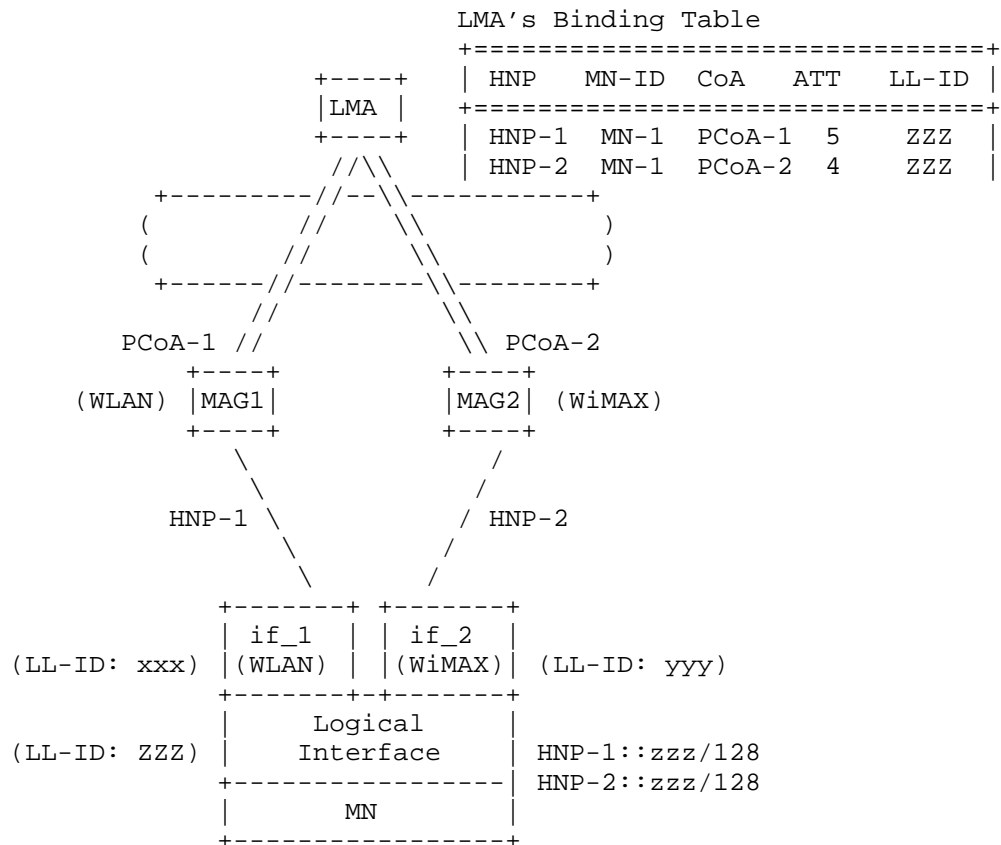


Figure 3: All HNPs are shown to IP host stack

4.4. Scenario 4 - Only one home network prefix is shown to IP host stack

In above scenario 3, basic PMIPv6 operations can be applied even though a logical interface is used. Multihoming support and inter-technology handoff support is provided without any modification of PMIPv6 operations but for flow mobility, it needs method to transfer a subset of multiple home network prefixes which are assigned to multiple interfaces of the host.

In scenario 3, one application that uses one specific home network prefix cannot utilize the multiplicity of home network prefixes. For on-going session continuity, applications cannot change their home network prefix during packet sending/receiving. Even though a logical interface supports the inter-technology handoff between different network interfaces, applications should use same home network prefix before/after handoff.

But, if a logical interface supports the change of network interface and also the change of home network prefix, applications can utilize the multiplicity of home network prefixes. As a same manner of the handling of network interface, only one home network prefix is shown to host IP stack of the mobile node. And the logical interface manages the relation between multiple home network prefixes and the home network prefix that is shown to host IP stack and dynamically update them. In this case, only one logical interface and one home network prefix is shown to host IP stack of the mobile node.

Figure 4 shows this case. As shown in Figure 4, if the LMA allocates two home network prefixes HNP-1, HNP-2 to the mobile node, only one home network prefix is shown to IP host stack of the mobile node. As a result of receiving single home network prefix in IP host stack, one IPv6 addresses HNP-1::zzz/128 can be generated.

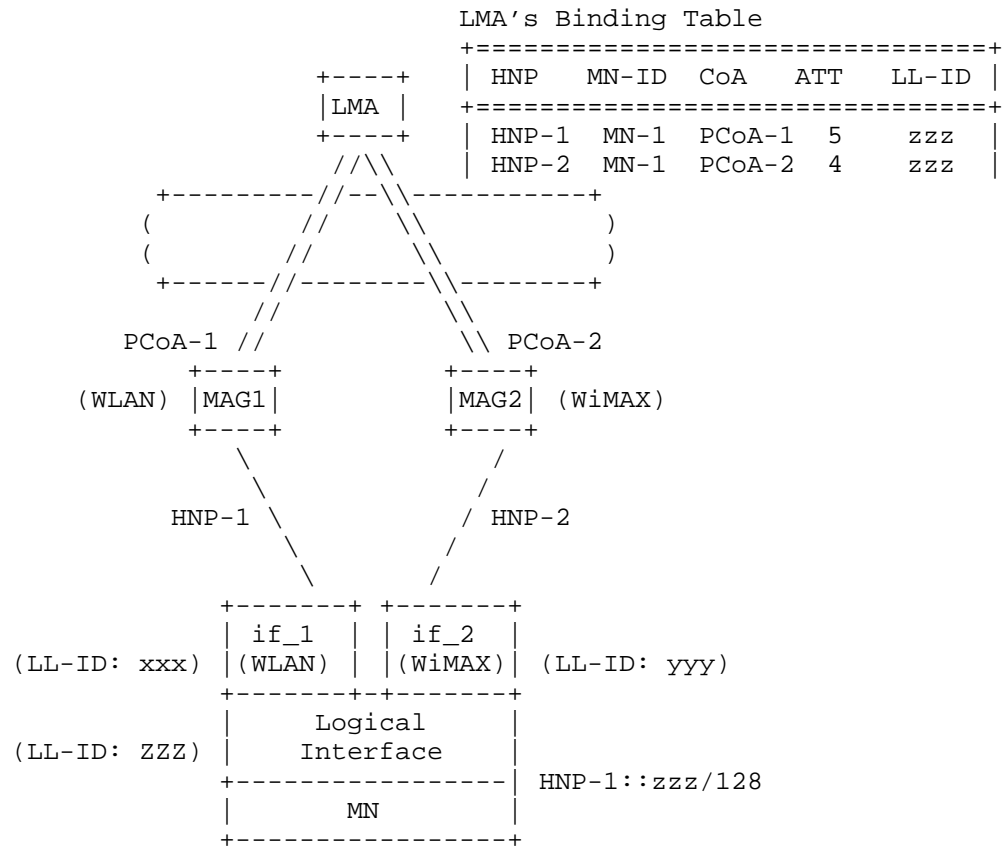


Figure 4: Only one HNP is shown to IP host stack

5. Security Considerations

TBD.

6. IANA Considerations

This document has no actions for IANA.

7. References

7.1. 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.

7.2. Informative References

- [I-D.bernardos-netext-pmipv6-flowmob-00]
Bernardos, CJ., Jeyatharan, M., Koodli, R., Melia, T., and F. Xia, "Proxy Mobile IPv6 Extensions to Support Flow Mobility, draft-melia-bernardos-netext-pmipv6-flowmob-00 (work in progress)", July 2010.
- [I-D.ietf-netext-logical-interface-support-00]
Melia, T. and S. Gundavelli, "Logical Interface Support for multi-mode IP Hosts, draft-ietf-netext-logical-interface-support-00 (work in progress)", September 2010.

Authors' Addresses

Yong-Geun Hong
ETRI
161 Gajeong-Dong Yuseung-Gu
Daejeon, 305-700
Korea

Phone: +82 42 860 6557
Email: yonggeun.hong@gmail.com

Joo-Sang Youn
DONG-EUI Univ.
Busan,
Korea

Phone: +82 51 890 1993
Email: joosang.youn@gmail.com

NETEXT WG
Internet-Draft
Intended status: Informational
Expires: April 27, 2011

T. Melia, Ed.
Alcatel-Lucent
S. Gundavelli, Ed.
Cisco
October 24, 2010

Logical Interface Support for multi-mode IP Hosts
draft-ietf-netext-logical-interface-support-01.txt

Abstract

A Logical Interface is a software semantic internal to the host operating system. This semantic is available in all popular operating systems and is used in various protocol implementations. The Logical Interface support is desirable on the mobile node operating in a Proxy Mobile IPv6 domain, for leveraging various network-based mobility management features such as inter-technology handoffs, multihoming and flow mobility support. This document explains the operational details of Logical Interface construct and the specifics on how the link-layer implementations hide the physical interfaces from the IP stack and from the network nodes. Furthermore, this document identifies the applicability of this approach to various link-layer technologies and analyzes the issues around it when used in context with various mobility management features.

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 April 27, 2011.

Copyright Notice

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

This document is subject to BCP 78 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	4
2. Requirements Language	5
3. Terminology	6
4. Hiding link layer technologies - Approaches and Applicability	7
4.1. Link-layer Abstraction - Approaches	7
4.2. Applicability Statement	8
4.2.1. Link layer support	8
4.2.2. Logical Interface	9
5. Logical Interface Operation	10
5.1. Logical Interface Link Layer Configuration	11
5.2. Bring up a new physical interface	12
5.3. Link Scoped Traffic	13
5.3.1. Unicast Traffic	13
5.3.2. Multicast Traffic	13
5.4. Global Scoped Traffic	14
5.5. Logical Interface Conceptual Data Structures	14
5.6. MTU considerations	15
6. Logical Interface Use-cases in Proxy Mobile IPv6	16
6.1. Multihoming Support	16
6.2. Inter-Technology Handoff Support	17
6.3. Flow Mobility Support	19
7. IANA Considerations	20
8. Security Considerations	21
9. Authors	22
10. Acknowledgements	22

11. Appendix	22
12. References	23
12.1. Normative References	23
12.2. Informative References	23
Authors' Addresses	23

1. Introduction

Proxy Mobile IPv6 [RFC5213] is a network-based mobility protocol. Some of the key goals of the protocol include support for multihoming, inter-technology handoffs and flow mobility support. The PMIPv6 extensions chartered in the NETEXT WG) allow the mobile node to attach to the network using multiple interfaces (simultaneously or sequentially), or to perform handoff between different interfaces of the mobile node. However, for supporting these features, the mobile node is required to be activated with specific software configuration that allows the mobile node to either perform inter-technology handoffs between different interfaces, attach to the network using multiple interfaces (sequentially or simultaneously), or perform flow movement from one access technology to another. This document analyzes from the mobile node's perspective a specific approach that allows the mobile node to leverage these mobility features. Specifically, it explores the use of the Logical Interface support, a semantic available on most operating systems.

A Logical Interface is a construct internal to the operating system. It is an approach where the link-layer implementations hide the physical interfaces from the IP stack and from the network nodes. This semantic is widely available in all popular operating systems. Many applications such as Mobile IP client [RFC3775], IPsec VPN client [RFC4301] and L2TP client [RFC3931] all rely on this semantic for their protocol implementation and the same semantic can also be useful in this context. Specifically, the mobile node can use the logical interface configuration for leveraging various network-based mobility management features provided by the Proxy Mobile IPv6 domain [RFC5213]. The rest of the document provides the operational details of the Logical Interface on the mobile node and the inter-working between a mobile node using logical interface and network elements in the Proxy Mobile IPv6 domain. It also analyzes the issues involved with this approach and characterizes the contexts in which such use is appropriate.

2. Requirements Language

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [RFC2119].

3. Terminology

This document uses the following terms:

PIF Physical Interface: a network device providing IP connectivity (e.g. an Ethernet card, a WLAN card, an LTE interface).

LIF Logical Interface: a virtual interface hiding to the IP stack the heterogeneous wired/wireless network devices.

VLL-ID Virtual Link Layer ID: a virtual MAC address configured on the LIF. It can be randomly generated or configured based on the MAC of one of the PIF.

UE (see Figure 2).




Figure 2: 3GPP LTE/EPC data plane architecture (GTP option)

- o Logical interface: this refers to solutions (see Figure 3) that logically group/bond several physical interfaces so they appear to the upper layers (i.e. IP) as one single interface (where application sockets bind). Depending on the OS support, it might be possible to use more than one physical interface at a time -- so the node is simultaneously attached to different media -- or just to provide a fail-over mode. Controlling the way the different media is used (simultaneous, sequential attachment, etc) is not trivial and requires additional intelligence and/or configuration at the logical interface device driver. An example of this type of solution is the Logical interface, which is defined in this document, or the bonding driver (a Linux implementation).

4.2. Applicability Statement

We now focus on the applicability of the above solutions against the following requirements:

- o multi technology support
- o sequential vs. simultaneous access

4.2.1. Link layer support

Link layer mobility support applies to cases when the same link layer technology is used and mobility can be fully handled at these layers. One example is the case where several 802.11 APs are deployed in the

same subnet and all of them share higher layer resources such as DHCP server, IP gateway, etc. In this case the APs can autonomously (or with the help of a central box) communicate and control the STA association changes from one AP to another, without the STA being aware of the movement. This type of scenario is applicable to cases when the different points of attachment (i.e. APs) belong to the same network domain, e.g. Enterprise, hotspots from same operator, etc.

This type of solution does not typically allow for simultaneous attachment to different access networks, and therefore can only be considered for inter-access technology handovers, but not for flow mobility. Existing RFC 5213 handover hint mechanisms could benefit from link layer information (e.g. triggers) to detect and identify MN handovers.

Link layer support is not applicable when two different access technologies are involved (e.g. 802.11 WLAN and 802.16 WiMAX) and the same is true when the same access technology expands over multiple network domains. This solution does not impose any change at the IP layer since changes in the access technology occur at layer two.

4.2.2. Logical Interface

The use of a logical interface allows the mobile node to provide a single interface view to the layers above IP (thus not changing the IP layer itself). Upper layers can bind to this interface, which hides inner inter-access technology handovers or data flow transfers among different physical interfaces.

This type of solution may support simultaneous attachment, in addition to sequential attachment. It requires additional support at the node and the network in order to benefit from simultaneous attachment. For example special mechanisms are required to enable addressing a particular interface from the network (e.g. for flow mobility). In particular extensions to PMIPv6 are required in order to enable the network (i.e., the MAG and LMA) to deal with physical interfaces, instead to IP interfaces as current RFC5213 does. RFC5213 assumes that each physical interface capable of attaching to a MAG is an IP interface, while the logical interface solution groups several physical interfaces under the same IP logical interface.

It is therefore clear that the Logical Interface approach satisfies the multi technology and the sequential vs: simultaneous access support.

5. Logical Interface Operation

On most operating systems, a network interface is associated with a physical device that provides the capability for transmitting and receiving network packets. In some cases a network interface can also be implemented as a logical interface which does not feature any packet transmission or reception capabilities, but relies on other network interfaces for such capabilities. A logical interface can be realized by that means. General overview of a logical interface is shown in Figure 3.

The logical interface allows heterogeneous attachment while leaving the change in the media transparent to the IP stack. Simultaneous and sequential network attachment procedures are possible enabling inter-technology and flow mobility scenarios. Through link awareness the logical interface can keep consistent neighbor caches and move flows across access networks transparently to the upper layers.

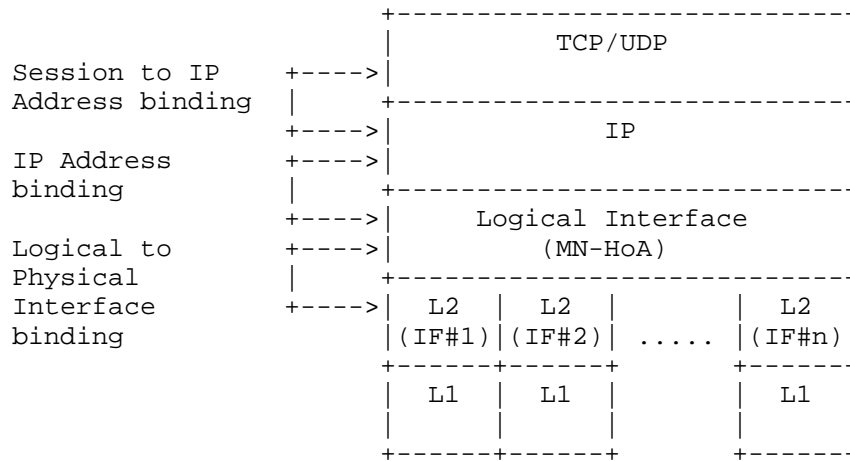


Figure 3: General overview of logical interface

From the perspective of the IP stack and the applications, a Logical interface is just another interface. A host does not see any difference between a Logical and a physical interface. All interfaces are represented as software objects to which IP address configuration is bound. However, the Logical interface has some special properties which are essential for enabling inter-technology handover and flow-mobility features. Following are those properties:

- o P1: Logical interface has a relation to a set of physical interfaces (sub-interfaces) on the host. Sub-interfaces can be attached/detached to the Logical Interface at any time (i.e. upon

L2 hints).

- o P2: The Logical Interface may or may not use the same link layer identifier as the physical interfaces (i.e. some technologies might not allow changing the link layer ID) .
- o P3: The Logical Interface has the path awareness of an IPv4/IPv6 link through a sub-interface.
- o P4: The logical Interface may manage heterogeneous links. As such, different MTUs may be announced on different links. The Logical Interface should be able to configure a common value (e.g., the minimum value observed by any link)".
- o P5: Send/Receive operations of a Logical interface are managed dynamically and are tied to the sub-interfaces (i.e. dynamic mapping not be visible to the applications).
- o P6: The Logical interface should transmit uplink packets on the same physical interface on which the downlink packet was received for the particular prefix/flow.

5.1. Logical Interface Link Layer Configuration

The logical interface has a virtual link-layer identifier (VLL-ID) that is not associated with any physical interfaces (see P2). This VLL-ID can be a representative of those of the physical interfaces or can be independently assigned. The usage of the VLL-ID is as follows:

- o Used for the neighbor discovery operation
- o Used to configure the IP address for this logical interface when the SLAAC is applied
- o Stored in the BCE at the Local Mobility Anchor via the Link-layer Identifier Option of the PBU
- o Used as the source link-layer address for sending packets from this logical interface

In order to support the above usage, all the physical interfaces SHOULD be able to send packets with the VLL-ID as the source link-layer address and SHOULD be able to receive packets with the VLL-ID as the destination link-layer address (the promiscuous mode).

If some of the connected wireless links do not allow sending packets with an arbitrary link-layer address, then the link-layer of the

corresponding PIF MUST be used instead. When receiving packets, whose destination LL-ID is that of the PIF, that LL-ID MUST be replaced with the VLL-ID before it appears to the IP layer.

5.2. Bring up a new physical interface

When a new PIF is enabled the following applies (see P1):

Bring up a new PIF: When a physical interface is enabled, a link-local address is formed by configuring the well-known link-local prefix FE80::/64 to the interface identifier. This address is a unicast address and has link-only scope. The MN can use this address to reach its neighbors.

Sending Neighbor Solicitation: When the MN wants to send a unicast packet but does not know the neighbor's link-layer address, it will perform address resolution by sending Neighbor Solicitation message through all of the enabled PIFs which are hidden by the LIF.

Receiving Neighbor Solicitation and Sending Neighbor Advertisement: When the LIF receives a Neighbor Solicitation message from a PIF, it will send a Neighbor Advertisement response message via the same PIF. The LIF may also send unsolicited Neighbor Advertisement message via all enabled PIFs in order to propagate new information quickly.

Sending Router Solicitation and receiving Router Advertisement: The LIF sends Router Solicitation messages through all enabled PIFs. The Router Advertisement messages are returned to the LIF through the PIFs that the Router Solicitation messages are sent from. The source link-layer address used in Neighbor Solicitation, Neighbor Advertisement and Router Solicitation is the link-layer identifier of the LIF.

It should be noted that since all the MAGs appear to the MN with the same IPv6 link-local and link-layer addresses (and that only the MAG shares the physical links with the MN) the ND caches of the LIF at the MN do not need complex extensions nor internal state kept at the LIF to be able to send traffic via multiple PIFs associated to the same LIF. The LIF engine would be able to generate the whole L2 frame and deliver it to the right PIF(s). No change in the L2 frame is needed at the LIF.

5.3. Link Scoped Traffic

The following section analyzes both unicast and multicast traffic handled by the LIF (see P3 and P5).

5.3.1. Unicast Traffic

Link-local unicast traffic generated by the LIF is sent through all PIFs associated to the LIF. As an example, Neighbor Advertisements messages generated by the LIF are sent through all PIFs. From the viewpoint of ND, this does not suppose any problem, as all the PIFs are logically grouped under the same LIF.

When receiving, the LIF receives all the traffic received via any of the PIFs associated to the LIF, and processes it normally. For example, Neighbor Solicitations are received and processed by the LIF without any modification (adding/updating the ND cache). Since in PMIPv6 only point-to-point interfaces are supported between the MN and the MAG, and all the MAGs show the same IPv6 link-local and link-layer addresses, this mode of operation of the LIF does not add any issue from the point of view of ND.

5.3.2. Multicast Traffic

Link-local multicast traffic generated by the LIF is sent through all PIFs associated to the LIF. As an example, Router Solicitation messages generated by the LIF are sent through all PIFs. This might cause multiple messages being received in response, though that would not cause any issue. When receiving, traffic from all PIFs is delivered to the LIF, which processes it normally. Examples of this traffic could be Router Advertisements or Neighbor Solicitations. As a result of the reception of certain types of link-local multicast traffic, the LIF might need to generate and send a (unicast) response. In this case, there are two possible approaches that can be followed:

- o Proceed as specified in Section 5.3.1 and send unicast responses via all PIFs associated to the LIF
- o Keep state at the LIF so replies can be sent via the proper interface only.




Figure 4: Link scoped traffic management by the LIF and its relation to Neighbor Discovery

5.4. Global Scoped Traffic

The following section analyzes both unicast and multicast traffic handled by the LIF (see P3 and P5).

For global-scoped traffic, the same assumptions taken in [RFC5213] for unicast traffic apply. Beyond these assumptions, the MULTIMOB WG is looking at ways to enhance the handling of multicast traffic in a PMIPv6 domain. The use of Logical Interface in the mobile node does not affect any of the aforementioned scenarios.

5.5. Logical Interface Conceptual Data Structures

The LIF has populates its neighbor cache according to standard operations. It should be noted that given the specificity of the PMIPv6 protocol there is only one entry being all the MAGs configured with the same link local address. The LIF has one default route in its routing table pointing to the link local address of the MAG (it should be noted that the same as before applies). The prefix list contains all the prefixes received during the attachment phase. The destination cache may contain multiple entries but the next hops is the same for all entries pointing the link local address of the MAG.

The LIF should maintain the following data structures as depicted in

figure Figure 5

LIF TABLE		FLOW table	
PIF_ID	FLOW RoutingPolicies Home Network Prefix Link Layer Address Status	FLOW ID	PIF_ID
PIF_ID	FLOW RoutingPolicies Home Network Prefix Link Layer Address Status	FLOW_ID	PIF_ID
....		

Figure 5

The LIF table maintains the mapping between the LIF and each PIF associated to the LIF (see P3). For each PIF entry the table should store the associated Routing Policies, the Home Network Prefix received during the SLAAC procedure, the configured Link Layer Address (as described above) and the Status of the PIF (active, not active).

The FLOW table allows a LIF to properly route each IP flow to the right interface. It assumed that the LIF can identify flows traversing its PIFs and can map accordingly to any of the PIF. For locally generated traffic the LIF performs interface selection. For traffic of an existing flow received from the network on a different PIF than the one locally stored, the LIF should interpret as an explicit flow mobility trigger and update the PIF_ID parameter in the corresponding table (see P6).

5.6. MTU considerations

The LIF SHOULD be configured with the maximum MTU value that is supported by all interfaces (see P4).

6. Logical Interface Use-cases in Proxy Mobile IPv6

This section explains how the Logical interface support on the mobile node can be used for enabling some of the Proxy Mobile IPv6 protocol features.

6.1. Multihoming Support

A mobile node with multiple interfaces can attach simultaneously to the Proxy Mobile IPv6 domain. Each of the attachment links are assigned a unique set of IPv6 prefixes. If the host is configured to use Logical interface over the physical interface through which it is attached, following are the related considerations.

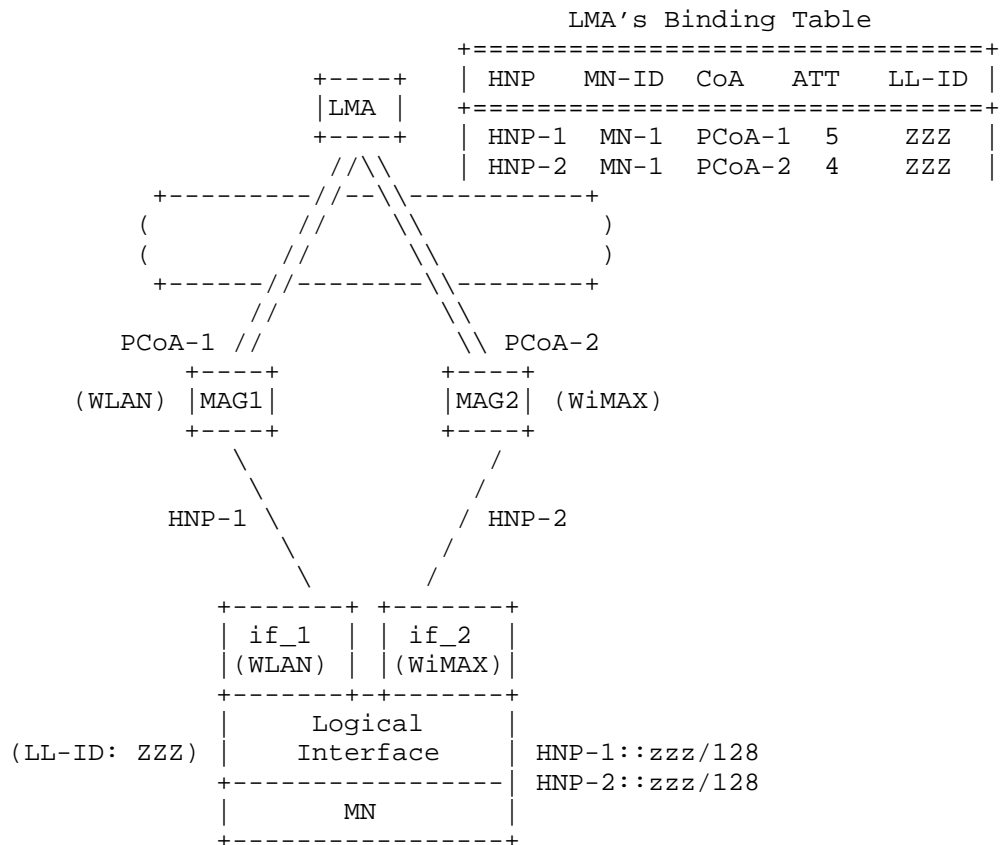


Figure 6: Multihoming Support

- o The mobile node detects the advertised prefixes from the MAG1 and MAG2 as the on link prefixes on the link to which the Logical interface is attached.
- o The mobile node can generate address configuration using stateless auto configuration mode from any of those prefixes.
- o The applications can be bound to any of the addresses bound to the Logical interface and that is determined based on the source address selection rules.
- o The host has path awareness for the hosted prefixes based on the received Router Advertisement messages. Any packets with source address generated using HNP_1 will be routed through the interface if_1 and for packets using source address from HNP_2 will be routed through the interface if_2.

6.2. Inter-Technology Handoff Support

The Proxy Mobile IPv6 protocol enables a mobile node with multiple network interfaces to move between access technologies, but still retaining the same address configuration on its attached interface. The protocol enables a mobile node to achieve address continuity during handoffs. If the host is configured to use Logical interface over the physical interface through which it is attached, following are the related considerations.




Figure 7: Inter-Technology Handoff Support

- o When the mobile node performs an handoff between if_1 and if_2, the change will not be visible to the applications of the mobile node. It will continue to receive Router Advertisements from the network, but from a different sub-interface path.
- o The protocol signaling between the network elements will ensure the local mobility anchor will switch the forwarding for the advertised prefix set from MAG1 to MAG2.
- o The MAG2 will host the prefix on the attached link and will include the home network prefixes in the Router Advertisements that it sends on the link.

6.3. Flow Mobility Support

For supporting flow mobility support, there is a need to support vertical handoff scenarios such as transferring a subset of prefix(es) (hence the flows associated to it/them) from one interface to another. The mobile node can support this scenario by using the Logical interface support. This scenario is similar to the Inter-technology handoff scenario defined in Section 6.2, only a subset of the prefixes are moved between interfaces.

Additionally, IP flow mobility in general initiates when the LMA decides to move a particular flow from its default path to a different one. The LMA can decide on which is the best MAG that should be used to forward a particular flow when the flow is initiated e.g. based on application policy profiles) and/or during the lifetime of the flow upon receiving a network-based or a mobile-based trigger.

As an example of mobile-based triggers, the LMA could receive input (e.g. by means of a layer 2.5 function via L3 signaling [RFC5677]) from the MN detecting changes in the mobile wireless environment (e.g. weak radio signal, new network detected, etc.). Upon receiving these triggers, the LMA can initiate the flow mobility procedures. For instance, when the mobile node only supports single-radio operation (i.e. one radio transmitting at a time), only sequential (i.e. not simultaneous) attachment to different MAGs over different media is possible. In this case layer 2.5 signaling can be used to perform the inter-access technology handover and communicate to the LMA the desired target access technology, MN-ID, Flow-ID and prefix.

7. IANA Considerations

This specification does not require any IANA Actions.

8. Security Considerations

This specification explains the operational details of Logical interface on an IP host. The Logical Interface implementation on the host is not visible to the network and does not require any special security considerations.

9. Authors

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

Carlos Jesus Bernardos Cano

cjbc@it.uc3m.es

Antonio De la Oliva

aoliva@it.uc3m.es

Yong-Geun Hong

yonggeun.hong@gmail.com

Kent Leung

kleung@cisco.com

Tran Minh Trung

trungtm2909@gmail.com

Hidetoshi Yokota

yokota@kddilabs.jp

Juan Carlos Zuniga

JuanCarlos.Zuniga@InterDigital.com

10. Acknowledgements

The authors would like to acknowledge prior discussions on this topic in NETLMM and NETEXT working groups. The authors would also like to thank Joo-Sang Youn, Pierrick Seite, Rajeev Koodli, Basavaraj Patil, Julien Laganier for all the discussions on this topic.

11. Appendix

TBD

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

12.2. Informative References

- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC5677] Melia, T., Bajko, G., Das, S., Golmie, N., and JC. Zuniga, "IEEE 802.21 Mobility Services Framework Design (MSFD)", RFC 5677, December 2009.

Authors' Addresses

Telemaco Melia (editor)
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

Email: telemaco.melia@alcatel-lucent.com

Sri Gundavelli (editor)
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: sgundave@cisco.com

Netext Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2011

Rajeev Koodli
Cisco Systems
October 18, 2010

Multi-access Indicator for Flow Mobility
draft-koodli-netext-multiaccess-indicator-00.txt

Abstract

When a Mobile Node attaches to the mobile network using multiple access networks, it is important for the Mobile Network Gateway to know whether the Mobile Node is capable of simultaneous multi-access, so that the former can distribute the traffic flows using the most appropriate interface. This document defines a new EAP attribute which can be used for such an indication to the Mobile Network Gateway.

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 April 21, 2011.

Copyright Notice

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

This document is subject to BCP 78 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. Protocol Overview and Extensions	3
3. IANA Considerations	5
4. Security Considerations	5
5. Acknowledgement	5
6. Informative References	5
Author's Address	6

1. Introduction

With multi-access, a Mobile Node (MN) may be simultaneously attached to a mobile network. For instance, in the 3GPP architecture, a MN may be attached to the same Mobile Network Gateway (called PGW) via 4G cellular LTE technology as well as the Wireless LAN (WiFi) technology. Such simultaneous access provides opportunity to distribute traffic based on the most appropriate access for the type of traffic in question (and potentially based on the Time Of the Day). In order to accomplish this flow distribution (or flow mobility), the MNG needs to know that the MN's attachment is for multi-access (and not handover) purposes and that the MN has necessary host abstractions to support prefix sharing across access interfaces. This document defines an attribute to be used during the EAP-AKA authentication process so that the 3GPP AAA server understands the MN's capabilities. Subsequently, the 3GPP AAA server provides the MN's capabilities in the Diameter message to the PGW, which can then make the policy decision to perform flow mobility accordingly.

2. Protocol Overview and Extensions

In the 3GPP architecture [3gpp.4g-2], two types of "non-3GPP" accesses are supported. In the trusted access model, the access network is considered trustworthy by the 3GPP network operator. An example of a trusted network is another cellular network such as CDMA or WiMax, but may also include broadband wireline network with WiFi access (such as a residential access network operated by the 3GPP cellular service provider). In the untrusted access model, the 3GPP service provider does not possess a trust relationship with the non-3GPP access provider. An example includes WiFi hot spot access.

In the trusted access model, the MN communicates with an Access Network Gateway (ANG). In the untrusted access model, the MN communicates with a node called ePDG that serves as a secure data termination point. In both the trusted and untrusted access models, the MN performs EAP-AKA or EAP-AKA' authentication. In the trusted access model, the EAP-AKA messages are transported from the MN to the ANG over a protocol specific to the access network. In the untrusted access model, the EAP-AKA messages are transported from the MN to the ePDG over IKEv2. Both the ANG and the ePDG communicate with the (3GPP) AAA server using Diameter. This is shown in Figure 1, which we explain further below.

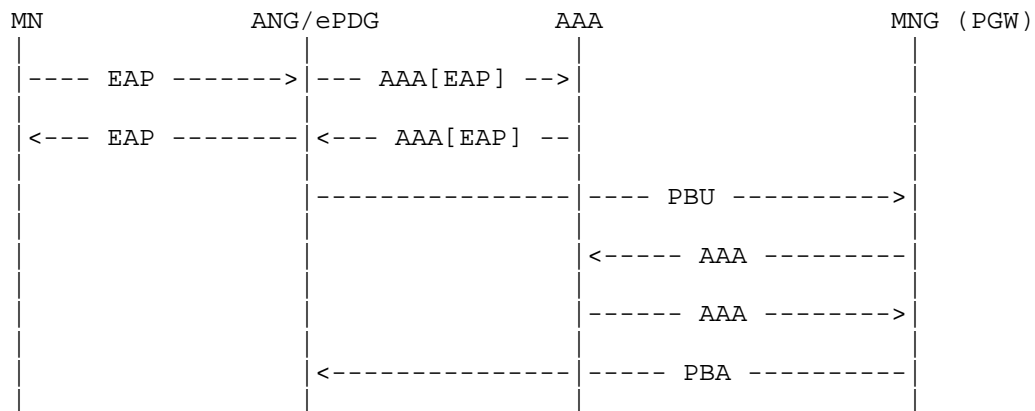


Figure 1: Authentication and Registration

- o The MN attaches to the non-3GPP access network
- o The MN performs EAP-AKA or EAP-AKA' authentication. The EAP messages are sent over IKEv2 when the MN is connected using untrusted access.
- o As a part of the EAP-AKA procedure, the MN responds with EAP-Response/AKA-Challenge message. In this message, the MN includes a new attribute AT_MA_IND which indicates that the MN's attachment is for multi-access purposes, and that the MN supports the necessary abstraction (Logical Interface) for flow mobility.
- o The ANG/ePDG forward the EAP message to the AAA server using the Diameter protocol message Authorization Request (AAR) message specified in [RFC4005].
- o The 3GPP AAA server verifies through subscription records (at HSS) that the the MN is allowed to use flow mobility. Subsequently, the AAA server provides the result in a new EAP attribute AT_MA_STATUS in the EAP-Request/AKA-Notification message (which is used for indicating the IP Mobility Selection mode). The EAP message sent using the Diameter Authorization Answer (AAA) message to the ANG/ePDG, which forwards the EAP message to the MN.
- o The ANG/ePDG sends the PMIP6 PBU message

- o The PGW contacts the AAA server to update the PGW's address for the MN's connection
- o The AAA server provides the MN's multi-access indication and Logical Interface capability in a new MN_MULTIACCESS (0x0000000000000003) flag in the MIP6-Feature-Vector AVP [RFC5447].
- o The PGW now understands that the MN is capable of flow mobility. It provides a prefix in the PBA accordingly. For instance, it may provide a new prefix as well as one or more of the already-assigned prefixes in the PBA.
- o The ANG/ePDG MUST be able to provide forwarding support for the prefixes provided in the PBA, regardless of the type of attachment indicated in the PBU message.

3. IANA Considerations

This document defines a new flag for the MIP6-Feature-Vector AVP in RFC 5447, which may need IANA assignment.

4. Security Considerations

This documents defines a new EAP attribute to extend the capability of EAP-AKA protocol as specified in Section 8.2 of RFC 4187 [RFC4187]. This attribute is passed from the MN to the AAA server. The document does not specify any new messages or options to the EAP-AKA protocol.

5. Acknowledgement

Thanks to Aeneas-Dodd Noble for flow mobility discussions.

6. Informative References

- [3gpp.4g-2] "Architecture Enhancements for non-3GPP accesses, 3GPP TS 23.402 8.7.0, December 2009.", .
- [RFC4005] Calhoun, P., Zorn, G., Spence, D., and D. Mitton, "Diameter Network Access Server Application", RFC 4005,

August 2005.

[RFC4187] Arkko, J. and H. Haverinen, "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)", RFC 4187, January 2006.

[RFC5447] Korhonen, J., Bournelle, J., Tschofenig, H., Perkins, C., and K. Chowdhury, "Diameter Mobile IPv6: Support for Network Access Server to Diameter Server Interaction", RFC 5447, February 2009.

Author's Address

Rajeev Koodli
Cisco Systems
USA

Email: rkoodli@cisco.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 22, 2011

S. Gundavelli
Cisco
M. Liebsch
NEC
October 19, 2010

PMIPv6 inter-working with WiFi access authentication
draft-liebsch-netext-pmip6-authiwk-00.txt

Abstract

Proxy Mobile IPv6, the IETF's protocol for network-based mobility management, requires a completed and successful authentication of the mobile node before it is registered at the mobility anchor. This document describes inter-working between access authentication mechanisms, such as IEEE 802.1X, and the Proxy Mobile IPv6 protocol to enable trusted WiFi access to a network-based mobility management domain. Furthermore, the use of authentication method specific identifiers for unique identification of mobile nodes during setup and maintenance of their mobility session is described, following recommendations of related standards organizations, such as 3GPP and the WiMAX Forum.

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 April 22, 2011.

Copyright Notice

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

This document is subject to BCP 78 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. Conventions and Terminology	5
3. Functional Objectives	6
4. Inter-working with IEEE 802.1X EAP	9
4.1. General use with authentication against a RADIUS Server	9
5. Security Considerations	11
6. IANA Considerations	12
7. Normative References	13
Authors' Addresses	14

1. Introduction

Proxy Mobile IPv6 (PMIPv6) [RFC5213] represents the IETF's protocol for network-based mobility management and is being deployed in various standards, such as the 3rd Generation Partnership Project (3GPP), to complement host mobility. According to the PMIPv6 standard, mobile nodes (MN) do not require a secure interface to the mobility anchor (LMA), as there is no direct signaling for mobility management between the MN and the LMA, but the Mobility Access Gateway (MAG) sets up and maintains a mobility binding on the LMA on behalf of the host by means of a Proxy Binding Update (PBU). [RFC5213] requires a successful authentication of the MN before the MAG sends a PBU to the LMA to set up a mobility binding for the MN. Furthermore, it assumes the MAG to be informed about a mobile node identifier (MN-Identifier), which unambiguously identifies the MN during the mobility session. Such MN-Identifier can be a static identifier or a temporary identifier, which may be derived from a static identifier.

This document intends to provide guidelines for PMIPv6 to inter-work with access authentication protocols which have been designed for IEEE 802-type of link technologies. Initial versions of this document focus on IEEE 802.1X and its recommendation to use the Extensible Authentication Protocol (EAP) [RFC3748]. Based on the procedure for general inter-working, more specific use cases are documented for discussion and reference. These use cases include the use of the Wireless LAN technology according to the IEEE 802.11 standard to provide trusted access to 3GPP's packet core network. So far, WLAN has been considered as untrusted access being even provided by third parties and MNs connect through WLAN to the mobile operator network through an established secure tunnel. Stepping towards WLAN trusted access avoids the overhead of an established IPsec tunnel with a packet data gateway in the operator's core network, but requires inter-working between WLAN access authentication and the operator's authentication and identification mechanisms. In the context of trusted WLAN access and network-based mobility management, WLAN security is being used to protect traffic on the wireless link whereas the trust relationship between a MAG and the LMA is used to convey traffic through the operator's core network.

The first version of this document discusses inter-working between IEEE 802.1X EAP and PMIPv6 as well as some specific use cases for trusted WLAN access in 3GPP's evolved packet core, which are based on recommended authentication schemes, such as EAP-AKA [RFC5448]. Further use cases with different EAP authentication schemes as well as inter-working between PMIPv6 and web authentication will be added to future versions of this document. Prior to describing details of PMIPv6 inter-working with various access authentication schemes in

Section 4, Section 3 describes functional objectives to enable trusted WLAN access to mobile operator networks and efficient inter-working between WiFi access authentication and operators' mobility management as well as policy and AAA infrastructure.

2. Conventions and 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].

This document uses the terminology of [RFC5213]. The following additional terms are used in the context of this draft:

- o AAA -- Authentication, Authorization and Accounting
- o EAP -- Extensible Authentication Protocol
- o PCC -- Policy and Charging Control
- o PMK -- Pairwise Master Key

3. Functional Objectives

Major motivation and objective to document inter-working between WiFi access authentication and PMIPv6 is to describe complete system operation, message sequences and identification schemes for network-based mobility management using PMIPv6 including IEEE 802.11-based access as proven and widely accepted radio technology and associated authentication mechanisms. Inter-action between access authentication and mobility management allows the specification of missing components in [RFC5213], mainly referring to MAG operation being triggered by successful MN authentication and MN identification.

The relevance of WiFi radio access is proven by various standards' initiative in specifying inter-working with IEEE 802.11-based technology. One example is the 3GPP's interest in supporting traffic offload to WLAN networks. Another example is the WiMax Forum's Network Architecture, which consider a WiFi-WiMAX inter-working function to enable access to the WiMAX network through WiFi radio access and to support handover between WiFi and WiMAX radio access.

The PMIPv6 standard [RFC5213] assumes a completed and successful access authentication of MNs (or their subscriber) before the MAG registers the MN at an LMA by means of a PBU. One objective of this document is to analyze relevant access authentication schemes and to document the operation of PMIPv6 in dependency of these authentication mechanisms. The EAP procedure as IEEE 802.X recommendation is being considered most relevant at this time. Web-authentication is a further popular access authentication scheme, which can be analyzed and inter-working with PMIPv6 can be specified, even though manual subscriber inter-action during access authentication conflicts with automatic and seamless operation, e.g. during dual radio handover from 3GPP access to WiFi access.

A further objective is to analyze the details of preferred authentication schemes, taking 3GPP and WiMAX Forum recommendations into account, and to document the use of common identifiers for access authentication and PMIPv6-based mobility management. Such identifier-specific inter-working must take further requirements, such as unique identification of a MN during the mobility session, into account. Some identifiers, which are generated during access authentication, are unique for an MN, but are not stable and valid beyond a certain radio access point. In such case, the MAG must use a different identifier or resolve such temporary identifier into a unique identifier which is valid beyond a single access point and MAG.

A further goal is to analyze inter-working between access

authentication schemes and PMIPv6 during handover, which may also imply a change in the radio access technology. Treatment of authentication methods, keys and identifiers and associated inter-working with PMIPv6 operation is documented.

Figure 1 depicts a high-level view of a WiFi network being integrated into a mobile operator network as trusted access. Instead of using a Security and Mobility Gateway, such as the 3GPP's Packet Data Gateway (PDG), which terminates an IPsec tunnel with the UE, the system relies on concatenated protected links between the UE and the WiFi access network, as well as between the WiFi access network and the LMA. The illustrated setup assumes a MAG function to be co-located with the WiFi Access Point or a WiFi Controller (Ctrlr). Inter-working between WiFi access authentication, PMIPv6 operation and the operator network's AAA and PCC (Policy and Charging Control) infrastructure is achieved by means of associated interfaces with the LMA. Future extensions may consider a direct policy configuration interface with the WiFi access network controller. This version of the inter-working document does not assume a direct policy control interface between the WiFi access network and the operator's PCC system. If needed, the PMIPv6 protocol interface may be proposed to convey associated information. Policy configuration in the WiFi access network is considered out of scope of this documentation.

Figure 1: Integration of the WiFi radio technology to provide trusted access to mobile operator networks

4. Inter-working with IEEE 802.1X EAP

4.1. General use with authentication against a RADIUS Server

IEEE 802.1X recommends EAP for access authentication, which can make use of an Authentication Server using for example the RADIUS protocol between the Authenticator and the Authentication Server. [RFC3579] specifies RADIUS extensions to convey EAP attributes between an Authenticator and the RADIUS server. Figure 2 depicts general inter-working between PMIPv6 and IEEE 802.1X using EAP.

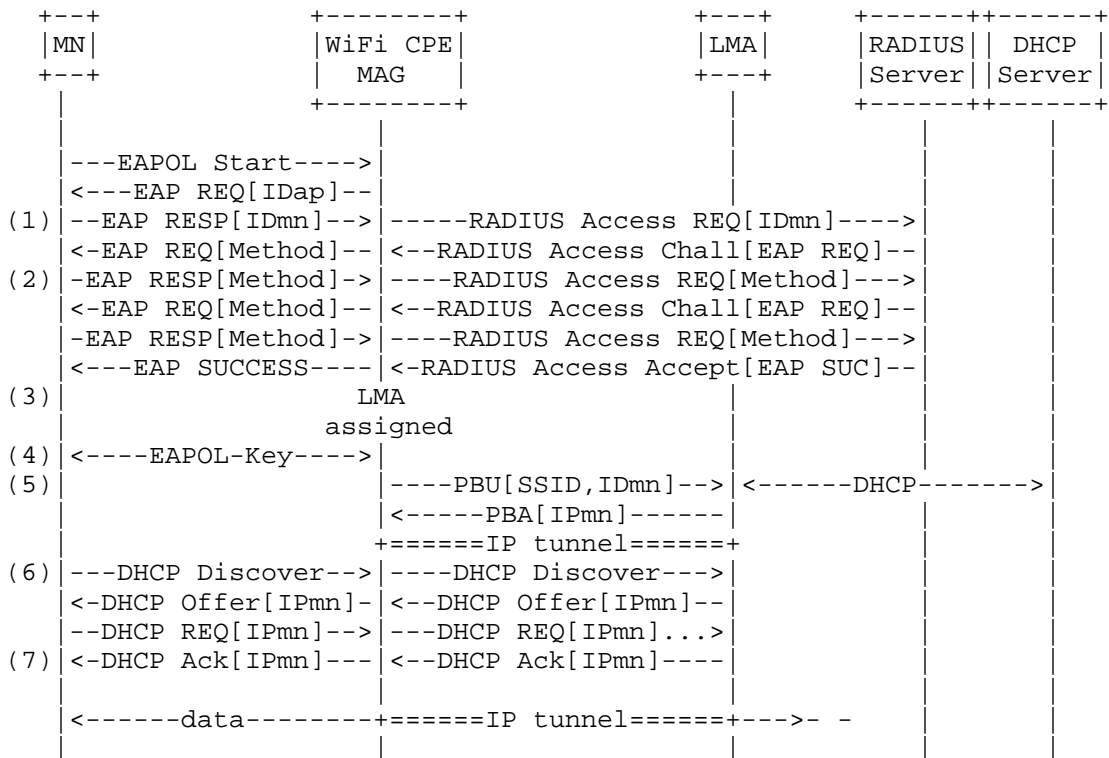


Figure 2: PMIPv6 inter-working with WPA2-802.1X access authentication against a RADIUS server

After the MN has associated with a WiFi Access Point, the EAPOL procedure starts (1). EAP attributes are mapped by the WiFi AP/Ctrlr between EAPOL on the wireless link and RADIUS operation on the link towards the RADIUS server. The RADIUS server selects one or multiple

authentication methods, which are performed with the MN in a challenge-response procedure (2). As a result of a successful EAP procedure, the RADIUS server may assign an LMA to the MN and signal the LMA identifier or the LMA IP address to the MAG function in the WiFi access network (3). The MN and the WiFi Access Point can now negotiate the Session Key to protect the wireless access (4). At that time, the MAG can take the EAP success as trigger to initiate the PBU registration of the MN with the LMA (5). The keys and identifiers being used and generated differ between the EAP and authentication method. In general, the MAG should not use the generated Session Key or security association identifier, as scope is limited to the MN's association with the Access Point. More suitable is an identifier being negotiated during the authentication procedure with the RADIUS server, e.g. based on the Pairwise Master Key (PMK) or any identifier which derives from the PMK without including single Access Point specific information, such as the AP's MAC address. One example, which will be described in more detail in future versions of this document, is the use of the International Mobile Subscriber Identity (IMSI) to derive a NAI at the Authentication Server. This IMSI-based NAI is then used as MN-Identifier in the PBU. Such approach is being proposed in 3GPP for trusted access to the mobile operator network through non-3GPP type radio access networks [3GPP-TS23.402] [3GPP-TS33.402].

As a result of the MN's registration, the LMA performs DHCP with a DHCP server to retrieve a valid IP address for the MN (IPmn). The assigned IP address is then signaled to the MAG in the PBA. The MN learns about this IP address from the DHCP procedure (6). After successful completion of the DHCP procedure (7), the MN can use the protected wireless link to communicate with the network infrastructure.

5. Security Considerations

This document analyzes and documents inter-working between WiFi access authentication and PMIPv6 mobility management to enable trusted access to a mobile operator network which uses network-based mobility management. The document refers to standard operation of PMIPv6 [RFC5213] as well as well accepted WiFi authentication mechanisms, such as EAP using a RADIUS server as authentication server, without introducing new messages or message sequences. Solely the inter-working of access authentication and PMIPv6 is described by means of message sequence charts. Furthermore, the use of identifiers, which are built during access authentication, for MN identification in the PMIPv6-based mobility management protocol is described. Hence, the documented inter-working should not introduce any new security threats.

6. IANA Considerations

This document is based on standardized protocols for WiFi access authentication and network-based mobility management. No additional protocol messages and options are specified so far in this document.

7. Normative References

- [3GPP-TS23.402]
"3GPP TS 23.402; 3rd Generation Partnership Project;
Technical Specification Group Services and System Aspects;
Architecture enhancements for non-3GPP accesses (Release
10)", <<http://www.3gpp.org>>.
- [3GPP-TS33.402]
"3GPP TS 33.402; 3rd Generation Partnership Project;
Technical Specification Group Services and System Aspects;
3GPP System Architecture Evolution (SAE); Security aspects
of non-3GPP accesses (Release 9)", <<http://www.3gpp.org>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication
Dial In User Service) Support For Extensible
Authentication Protocol (EAP)", RFC 3579, September 2003.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H.
Levkowetz, "Extensible Authentication Protocol (EAP)",
RFC 3748, June 2004.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K.,
and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5448] Arkko, J., Lehtovirta, V., and P. Eronen, "Improved
Extensible Authentication Protocol Method for 3rd
Generation Authentication and Key Agreement (EAP-AKA')",
RFC 5448, May 2009.

Authors' Addresses

Sri Gundavelli
Cisco
170 West Tasman Drive
San Jose, CA 95134,
USA

Email: sgundave@cisco.com

Marco Liebsch
NEC Laboratories Europe
Kurfuersten-Anlage 36
D-69115 Heidelberg,
Germany

Email: liebsch@neclab.eu

