

NetLMM Working Group
Internet-Draft
Expires: August 28, 2008

M. Liebsch
L. Le
NEC
February 25, 2008

Inter-Technology Handover for Proxy MIPv6
draft-liebsch-netlmm-intertech-proxymip6ho-01.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 28, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

The IETF is specifying Proxy Mobile IPv6 for network-based localized mobility management, which takes basic operation for registration, tunnel management and de-registration into account in a first protocol release. This document proposes an extension to Proxy Mobile IPv6 to suit MNs, which have multiple network interfaces implemented and hand over from one network interface to another network interface while remaining anchored at the same Local Mobility Anchor. This extension avoids that the LMA will prematurely forward data packets for the mobile node to the mobile node's new interface before the address configuration of this interface has completed. With the proposed extension, packet buffering at the new MAG is not necessary and packet losses during an inter-technology handover can be avoided.

Table of Contents

1.	Requirements notation	3
2.	Introduction	4
3.	Terminology and Functional Components	5
4.	Problem Statement	6
4.1.	IP address auto-configuration	6
4.2.	Delay in setting up an access link	8
4.3.	Procedural sequences on the LMA	8
4.4.	Demand for a general solution	10
5.	PMIPv6 Extensions	11
5.1.	Description and use of extensions	11
5.2.	Protocol extensions	13
5.2.1.	Impact to the binding management	13
5.2.2.	Signaling message extensions	14
6.	Security Considerations	16
7.	IANA Considerations	17
8.	Acknowledgments	18
9.	Normative References	19
Appendix A.	Proposal for a Binding Control and Indication (BCI) option	20
	Authors' Addresses	22
	Intellectual Property and Copyright Statements	23

1. Requirements notation

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](#)].

2. Introduction

The NetLMM WG is specifying Proxy Mobile IPv6 for network-based localized mobility management (NetLMM), taking basic support for registration, de-registration and handover into account in a first protocol release. This document proposes an extension to Proxy Mobile IPv6 to suit MNs, which have multiple network interfaces implemented and hand over from one network interface to another network interface while remaining anchored at the same Local Mobility Anchor. Without such extension, inter-technology handover could result in tremendous packet loss.

When an inter-technology handover occurs, the mobile node needs to perform address configuration for the new interface. For IPv6, this typically requires the mobile node to send a Router Solicitation and await a Router Advertisement. The mobile node may also need to perform duplicate address detection. During this time, the mobile node's new interface is not fully configured and not yet ready to receive data packets. If the LMA decides to forward data packets for the mobile node via the new MAG, severe packet delay and jitter as well as packet losses can occur.

This document proposes an extension to Proxy Mobile IPv6 that allows classification of a new binding as 'preliminary' and to 'activate' such binding only when the MN's new interface is fully configured and ready to receive data packets. Thanks to this extension, the LMA will not prematurely forward data packets for the MN to the new MAG before the address configuration for the MN's new interface is not completed. The result of this proposed extension is that packet buffering at the new MAG is not necessary and packet losses during an inter-technology handover can be avoided.

3. Terminology and Functional Components

- o MAG - Mobility Access Gateway. PMIPv6 functional component defined in [[I-D.ietf-netlmm-proxymip6](#)]. The MAG function is assumed to be located on the PMIPv6 domain's access routers.
- o LMA - Local Mobility Anchor. PMIPv6 functional component defined in [[I-D.ietf-netlmm-proxymip6](#)].
- o pAR - Previous Access Router. Equivalent to current access router. In a layer 3 handover situation, the access router which the mobile node is leaving.
- o nAR - New Access Router. Equivalent to handover target access router. In a layer 3 handover situation, the access router towards which the mobile node is moving.
- o pMAG - previous MAG. In a layer 3 handover situation, the MAG function located on the previous access router
- o nMAG - new MAG. In a layer 3 handover situation, the MAG function located on the new access router.
- o IF - Interface. Any network interface, which offers an MN wireless or wired access to the network infrastructure. In case an MN has multiple interfaces implemented, they are numbered (IF1, IF2, ...)
- o Inter-RAT handover. Handover between different radio access technologies.

4. Problem Statement

A problem that can occur during an inter-technology handover is that the MN's new interface or the associated access link may not be ready to receive or forward data packets immediately after the handover. This section evaluates some of these issues in more detail.

4.1. IP address auto-configuration

One reason of delaying availability of an MN's network interface can be that the MN needs to perform certain protocol operations to configure its new interface and these operations can take time. In order to complete the address auto-configuration on its new interface, the MN needs to send a router solicitation and awaits a router advertisement. Upon receiving a router advertisement from the new MAG, the MN can complete its address configuration and the MN's new interface is now ready to receive packets.

During the address configuration of the MN's interface, if the LMA and the new MAG have already established a tunnel and the LMA starts forwarding data packets for the MN to the new MAG, these data packets may get delayed or lost in case address configuration of the MN's new interface has not yet completed, hence it's not yet ready to receive data. However, delaying the registration signaling, which is sent from the new MAG to the LMA after the MN's new interface has attached to the network, is not possible, as the new MAG retrieves configuration data for the MN from the LMA. With host-based mobility protocols, such as Mobile IPv6 [[RFC3775](#)], MNs can easily control when a binding is updated. This is different for network-based mobility management, where hosts are not involved in IP mobility management [[RFC4831](#)].

The aforementioned problem is illustrated in Figure 1, which assumes that the HNP will be assigned under control of the LMA. Hence, the HNP option in the PBU, which is sent by the new MAG to the LMA, is set to ALL_ZERO. An MN has attached to the network with interface IF1 and receives data on this interface. When the MN's new interface IF2 comes up and is detected by the new MAG, the new MAG sends a PBU and receives a PBA from the LMA. If the LMA decides to forward data packets for the MN via the new MAG, the new MAG has to buffer these packets until address configuration of the MN's new interface has completed and the MN's new interface is ready to receive packets. While setting up IF2, the MN may not reply to address resolution signaling, as sent by the new MAG (A). If the MAG's buffer overflows or the MN cannot reply to address resolution signaling for too long, data packets for the MN are dropped and the MN can experience severe packet losses during an inter-technology handover (B).

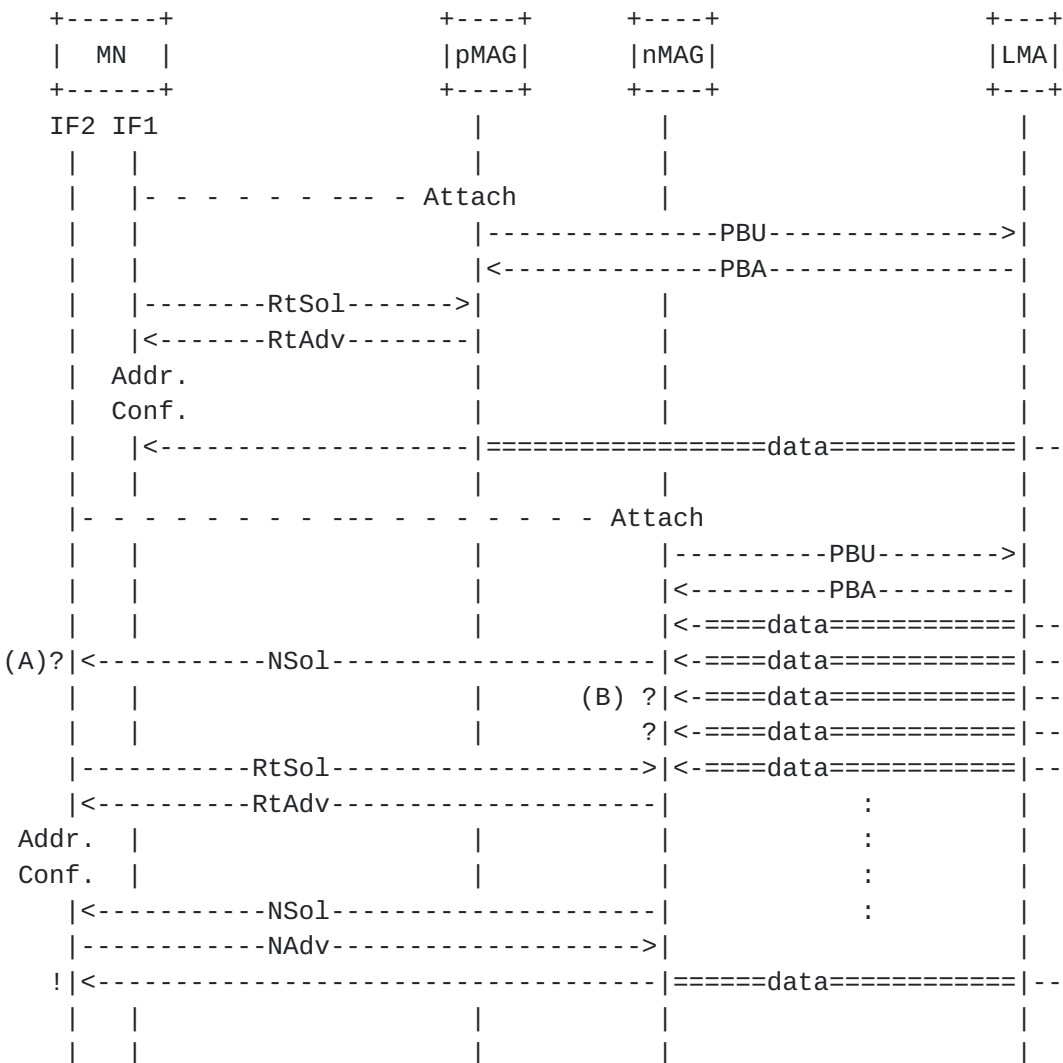


Figure 1: Issue with inter-RAT mobility.

Duplicate Address Detection (DAD) tests affect the delay, during which an MN's interface does not reply to Neighbor Solicitations, considerably. During such delay, MAGs must buffer all packets for the addressed MN until Neighbor Discovery has been performed completely to avoid packet loss. Normally, routers consider a neighbor queue size of 3 packets. Using this default setting on MAGs can easily result in packet loss in case completion of the Neighbor Discovery procedure is delayed.

Even if the MN does not perform DAD tests (`DAD_TRANSMITS = 0`), some aspects of MN address configuration according to [RFC4862] can cause MNs missing the initial Neighbor Solicitation being sent from its new MAG. First of all, this can happen due to processing delay on MNs and in case the Neighbor Solicitation follows immediately after the

Router Advertisement, which carries the HNP assigned to the MN's IF2. Relevant for successful Neighbor Discovery is that the MN's IF2 has joined the auto-configured IPv6 address' Solicited Node Multicast Address. According to [\[RFC4862\]](#), MNs delay joining a Solicited Node Multicast Address according to a randomly chosen interval between 0 and MAX_RTR_SOLICITATION_DELAY seconds. By default, the parameter MAX_RTR_SOLICITATION_DELAY is set to 1 second. If MNs follow this procedure, the MN's IF2 will very likely miss the initial Neighbor Solicitation.

In case the MN's IF2 does not respond to an initial Neighbor Solicitation, the new MAG follows the procedure in [\[RFC4861\]](#) and waits RETRANS_TIMER milliseconds before it attempts to address a further Neighbor Solicitation to the Solicited Node Multicast Address of IF2. By default, RETRANS_TIMER is set to 1000 ms, which requires the new MAG to buffer all packets destined to the MN's IF2 during at least 1 second to avoid packet loss.

[4.2.](#) Delay in setting up an access link

Another risk for a delay in forwarding data packets from a new MAG to the MN's IF2 can be some latency in setting up a particular access technology's radio bearer or access specific security associations after the new MAG received the MN's HNP from the LMA via the PBA signaling message.

In case an access technology needs the MN's IP address or HNP to set up a radio bearer between an MN's IF2 and the network infrastructure, the responsible network component might have to wait until the nMAG has received the associated information from the LMA in the Proxy Binding Acknowledgment. Delay in forwarding packets from the nMAG to the MN's IF2 depends now on the latency in setting up the radio bearer.

A similar problem can occur in case the set up of a required security association between the MN's IF2 and the network takes time and such set up can be performed only after the MN's IP address or HNP is available on the nMAG.

[4.3.](#) Procedural sequences on the LMA

In case of an intra-technology handover it's beneficial for an MN to receive data packets after a handover as soon as possible after the MN's IF attached to the infrastructure and to the associated new MAG. Delay in address configuration is not an issue, as the IF has an IPv6 address configured and has already joined the associated Solicited Node Multicast Address. Hence, the MN's IF is prepared to reply to initial Neighbor Solicitations being sent by the new MAG.

According to the procedure described in [[I-D.ietf-netlmm-proxymip6](#)], an LMA updates its BCE and associated forwarding information after it received a PBU from a MAG and before it sends a PBA. In case of an inter-technology handover, such procedure can be disadvantageous because data packets, which arrive at the new MAG, may initiate the Neighbor Discovery Procedure before the new MAG can actually send the MN's HNP for IF2 in the previously solicited Router Advertisement. This will result in ignoring the initial Neighbor Solicitation and may cause an overflow in the nMAG's neighbor queue. In this context, we characterize such a procedure on the LMA as 'Forward-before-Signaling' (FbS), which is illustrated in Figure 2.

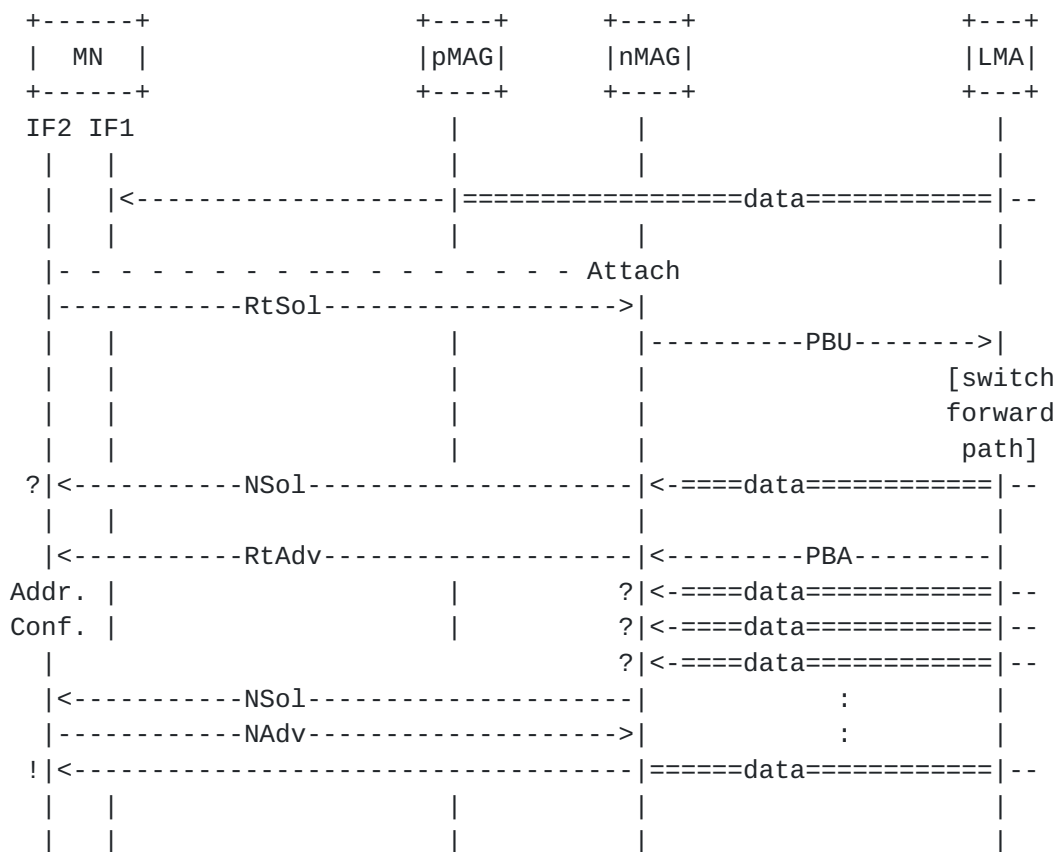


Figure 2: Issue with Forward-before-Signaling for inter-RAT mobility.

In case of an inter-technology handover, it could be beneficial if an LMA proceeds according to 'Forward-after-Signaling' (FaS). Such procedure considers replying to a PBU before the LMA switches the associated BCE's forwarding information. This could at least increase the probability, that the MN's IF2 receives HNP information before any request for address resolution.

4.4. Demand for a general solution

To reduce the risk of packet loss, some settings on an MN could be chosen appropriately to speed up the process of network interface configuration. This could include omitting DAD tests (DAD_TRANSMITS = 0) and setting the parameter MAX_RTR_SOLICITATION_DELAY to 0, which allows joining the configured Solicited Node Multicast Address immediately. The latter optimization is acceptable in particular for point-to-point links and when the network assigns individual prefixes to MNs.

Also following an FaS procedure on the LMA in case of an inter-technology handover can reduce the risk of packet loss. Furthermore, increasing the neighbor queue size on MAGs could help to reduce packet loss, but does not solve the jitter problem. Preferable is a general solution based on an extension to the Proxy MIPv6 base protocol.

This document proposes an extension to Proxy Mobile IPv6 that provides the new MAG with a signaling mechanism to notify the LMA when the MN's new interface is fully configured and ready to receive data packets. Thanks to this signaling mechanism, the LMA will not prematurely forward data packets for the MN to the new MAG before the address configuration of the MN's new interface is not complete. The result of this proposed extension is that excessive packet buffering at the new MAG is not necessary and packet losses during an inter-technology handover can be avoided.

5. PMIPv6 Extensions

This section describes an extension to [[I-D.ietf-netlmm-proxymip6](#)] to solve the problem of using a handover target technology before it is ready to receive and send IP packets. The key extension is to allow classification of an MN's binding cache entry on the LMA as 'preliminary' or 'active'. Classification can be performed either locally on the LMA or on a network component, which is not colocated with the LMA. An example about how classification can be controlled and signaled is described in [Section 5.1](#), whereas [Section 5.2](#) describes associated extensions to existing signaling and states.

5.1. Description and use of extensions

In an inter-technology handover scenario, an MN activates a network interface, which is different from the currently used one. The new interface should serve as interface through which the MN sends and receives packets. After a handover, the MN shuts down the previously used interface. The extension to the base PMIPv6 protocol as described in this document forwards downlink packets to an MN's new interface only when the new interface is ready to receive and send IP packets.

Figure 3 illustrates the components of a PMIPv6 network, which comprises the MN's LMA as well as its pMAG and the nMAG. The pMAG is supposed to support an access network, which conforms to the technology of Interface 1 (IF1) of the MN, whereas the nMAG supports an access network, which conforms to the technology of IF2.

The MN attaches to the PMIPv6 network with IF1 according to the procedure in [[I-D.ietf-netlmm-proxymip6](#)]. The MN starts receiving data packets on IF1. When the MN activates IF2 to prepare an inter-technology handover, the nMAG receives an attach indication and sends the PBU to the LMA to retrieve configuration information for the MN (e.g. HNP). The LMA is able to identify a technology change for the MN's binding, for example by means of processing the Access Technology Type option, which comes along with the PBU. Also the MN Interface Identifier option helps the LMA also to identify an interface change. Alternative mechanisms to learn about the nature of a handover are possible, such as the explicit notification from the nMAG or from any other network component. This requires the LMA to maintain information about the MN's currently used technology into the MN's binding cache entry, as it is considered in [[I-D.ietf-netlmm-proxymip6](#)].

If the LMA detects a technology change for the same MN, it enters the nMAG as forwarding information into the binding cache, but classifies this binding as 'preliminary' (A). The status 'preliminary' causes

the LMA to keep using the previous binding as forwarding information until the preliminary binding will be activated. The LMA acknowledges the PBU by means of sending a PBA to the nMAG. The PBA carries the MN's HNP as well as a flag, which indicates the status 'preliminary' to the nMAG (B). The nMAG learns through the PBA, that an explicit activation of the binding is required. If explicit activation of the binding is expected from the nMAG, the nMAG should find out when the MN's new interface has been set up and configured completely. For example, the nMAG could listen to Neighbor Solicitations from the MN's IF2, which indicate that the MN has configured a global address at IF2 and performs duplicate address detection. Details about how the nMAG finds out that IF2 is ready for being used is out of scope of this version of the document.

As soon as IF2 has been set up completely (C), the nMAG sends a further PBU to the LMA to 'activate' the previously as 'preliminary' classified binding. The PBU should indicate explicit activation of a binding. The activation causes the LMA to change forwarding information to refer to the nMAG and adjust the MN's binding cache entry accordingly (D). From that moment, downlink packets will be forwarded towards the MN's IF2 (E).

Note: As the status 'preliminary' needs to be explicitly changed to 'active' before the LMA uses the new forwarding information, this example makes use of a further PBU/PBA handshake between the nMAG and the LMA. Alternative mechanisms are possible to activate the new state on the LMA.

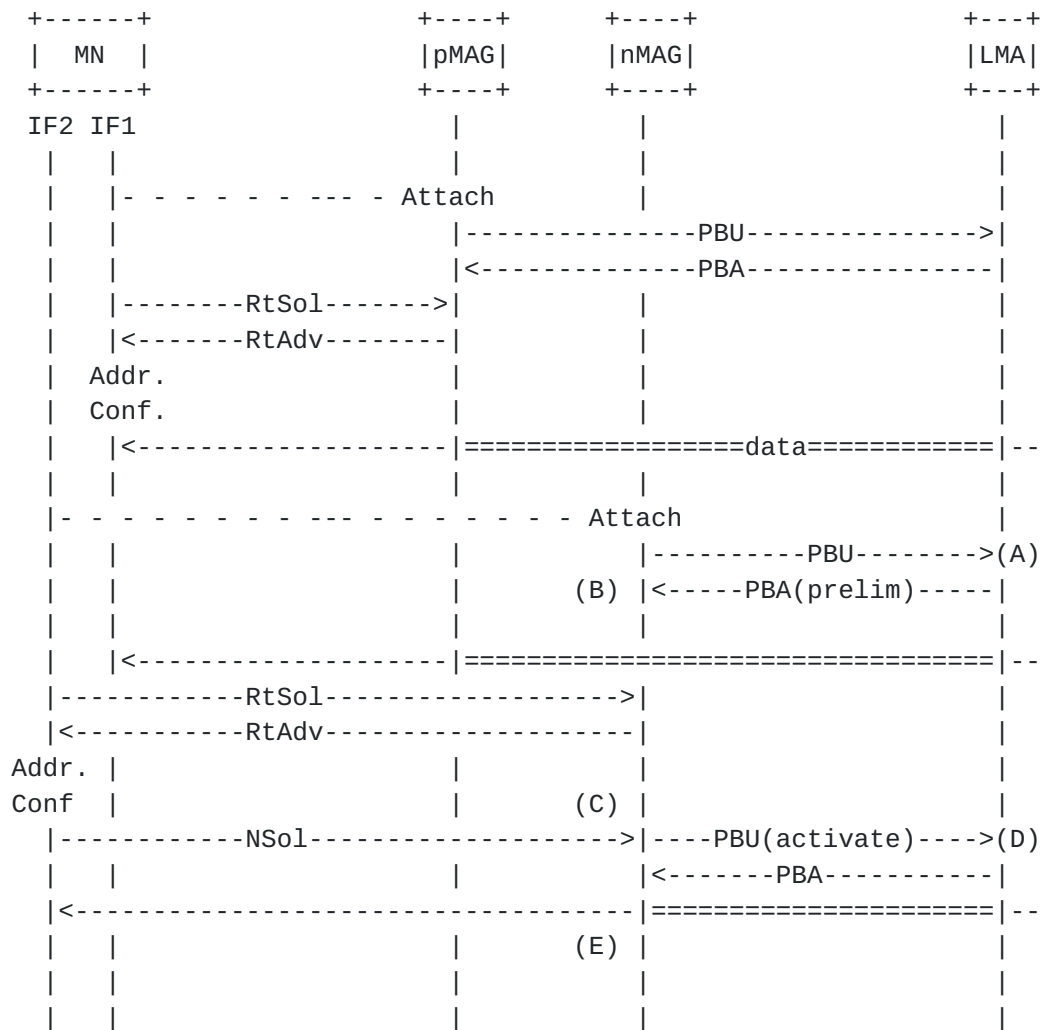


Figure 3: Solution for inter-RAT mobility.

5.2. Protocol extensions

5.2.1. Impact to the binding management

The extension proposed in [Section 5.1](#) requires maintenance of two forwarding entries on the LMA during an inter-technology handover procedure. One of these entries is marked as 'active', whereas the other one is 'preliminary'. Extensions to an MN's entry in the LMA's binding cache should comprise the interface identifier of the associated tunnel towards a MAG, as well as the associated access technology. The interface identifier and access technology type info can be taken from the PBU received at the LMA and linked to each forwarding entry accordingly.

MAGs should maintain the status of an MN's binding in their binding update list. This is in particular important in case a new MAG needs to explicitly activate a binding after the associated MN's new interface has proven to be ready for IP traffic.

5.2.2. Signaling message extensions

5.2.2.1. Proxy Binding Update message

The PBU message should carry a binding status flag to allow indication and control of a binding status. In the PBU, the flag can in particular be used from a MAG to signal activation of a binding after the MN's new interface or the associated access link has been set up completely. If a MAG is not sure about the status of a binding, the MAG should assume that sending the PBU will result in an active binding status and learn about the real status from the PBA sent by the LMA.

As an alternative to a binding status flag in the PBU and the PBA and for discussion, we propose the use of a new Binding Control and Indication (BCI) option. The format and use of such an option is described in [Appendix A](#).

Figure 4 illustrates the Proxy Binding Update message with a binding status flag.

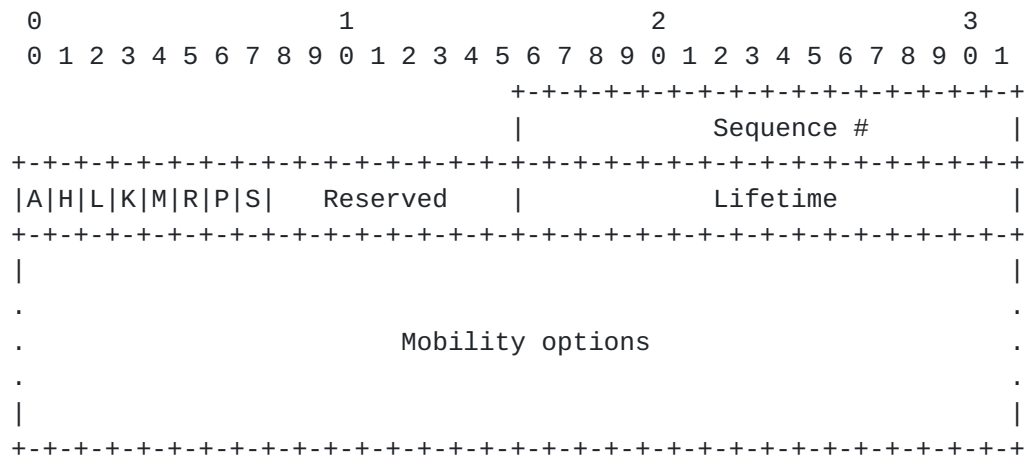


Figure 4

Binding Status Flag (S): A new flag (S) is included in the Proxy Binding Update message to indicate and control the status of a binding. A value of 1 indicates the status PRELIMINARY, whereas a value of 0 allows a MAG to ACTIVATE a binding. In case a MAG does

not know about the status of a binding when a Proxy Binding Update is sent, it MUST set the S-flag to a value of 0.

The other fields of the Proxy Binding Update message are used according to [[I-D.ietf-netlmm-proxymip6](#)].

5.2.2.2. Proxy Binding Acknowledgment message

Also the PBA message should carry the binding status flag to allow indication of a binding status. In case the LMA classifies a binding as 'preliminary', it indicates the status to the MAG, which has sent the PBU, in the PBA. After a MAG explicitly activated a binding, which has been classified as 'preliminary' beforehand, the LMA should indicate the result of processing the PBU in the PBA by means of setting the binding status flag appropriately.

Figure 5 illustrates the Proxy Binding Acknowledgment message with a binding status flag.

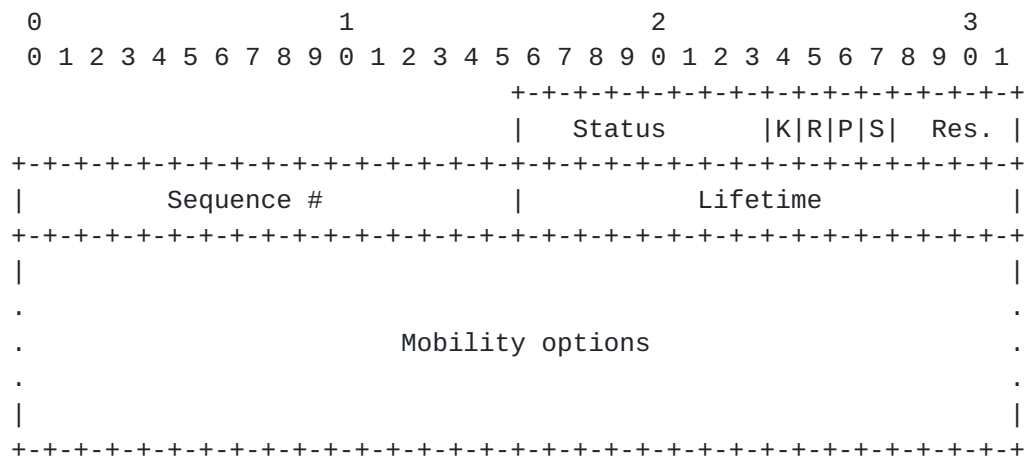


Figure 5

Binding Status Flag (S): A new flag is included in the Proxy Binding Acknowledgment message to indicate the status of a binding. A value of 1 indicates the status PRELIMINARY, whereas a status ACTIVE is indicated with value 0.

The other fields of the Proxy Binding Acknowledgment message are used according to [[I-D.ietf-netlmm-proxymip6](#)].

6. Security Considerations

Signaling between MAGs and LMAs as well as information carried by PBU and PBA messages is protected and authenticated according to the mechanisms described in [[I-D.ietf-netlmm-proxymip6](#)].

In case MAGs or LMAs make use of a further protocol interface to an external component, the associated protocol must be protected and information must be authenticated. Such protocol interface may for example be used by an LMA or a MAG to learn about the nature of a handover or about the status of an MN's interface.

7. IANA Considerations

In case a BCI option as proposed in [Appendix A](#) will be used for binding control and status indication, there is a need to assign a type to identify the BCI option.

8. Acknowledgments

May thanks to Jun Awano for his valuable comments and the associated fruitful discussion about binding control and status indications.

9. Normative References

- [I-D.ietf-netlmm-proxymip6]
Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K.,
and B. Patil, "Proxy Mobile IPv6",
[draft-ietf-netlmm-proxymip6-11](#) (work in progress),
February 2008.
- [I-D.muhamma-netlmm-pmipv6-support-transient-coa]
Muhamma, A., Krishnan, S., Leung, K., and B. Patil, "Proxy
MIPv6 support for transient registrations",
[draft-muhamma-netlmm-pmipv6-support-transient-coa-00](#) (work
in progress), February 2008.
- [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.
- [RFC4831] Kempf, J., "Goals for Network-Based Localized Mobility
Management (NETLMM)", [RFC 4831](#), April 2007.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
"Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#),
September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
Address Autoconfiguration", [RFC 4862](#), September 2007.

Appendix A. Proposal for a Binding Control and Indication (BCI) option

[Section 5.2.2](#) proposes adding a Binding Status flag to the Proxy Binding Update message. In the context of preliminary bindings and possible future extensions, which may require binding control and binding status indication, we propose a new Binding Control and Indication (BCI) option. The use of such an option allows to omit the Binding Status flags in the PBU and the PBA messages.

The BCI option is meant to complement the Handoff Indicator (HI) option, which is specified in [[I-D.ietf-netlmm-proxymip6](#)], and allows clear separation between handoff type indications and binding status indications. Furthermore, the BCI option can be used to control the binding according to a pre-defined procedure. One example is the controlled activation of a preliminary binding from a MAG. Further application scenarios are possible, such as the control of bindings to coordinate uplink traffic, as proposed in [[I-D.muhamma-netlmm-pmipv6-support-transient-coa](#)].




Figure 6

Type: Identifies the BCI option. To be assigned by IANA.

Length: 8-bit unsigned integer indicating the length of the option in octets, excluding the Type and the Length fields.

Binding Control and Indication (BCI) Mask: Mask to control and indicate the status of a binding. For this draft, the following values are proposed:

[0x0000]: Active Binding

[0x0001]: Preliminary Binding

Other bits of this mask must set to 0 and be ignored until their meaning is specified.

Binding ID (BID): In case an LMA distinguishes BCEs and possibly multiple bindings per MN by means of a binding identifier, this field can be used to refer to a particular binding by means of the BID field. In case binding identifiers are not used beyond local management on an LMA, this field must be set to 0.

Note 1: Using a BCI Mask allows simultaneous setting of multiple bits for control and indication in this option.

Note 2: In case the binding identifiers are not used, alternatively this field and the preceding Reserved field can be omitted. The Length field of this option must be adjusted accordingly.

Authors' Addresses

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

Phone: +49 6221 4342146
Email: liebsch@nw.necclab.eu

Long Le
NEC Laboratories Europe
NEC Europe Ltd.
Kurfuersten-Anlage 36
69115 Heidelberg,
Germany

Phone: +49 6221 4342224
Email: le@nw.necclab.eu

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

