

Internet Engineering Task Force (IETF)
Request for Comments: 6463
Category: Standards Track
ISSN: 2070-1721

J. Korhonen, Ed.
Nokia Siemens Networks
S. Gundavelli
Cisco
H. Yokota
KDDI Lab
X. Cui
Huawei Technologies
February 2012

Runtime Local Mobility Anchor (LMA) Assignment Support
for Proxy Mobile IPv6

Abstract

This document describes a runtime local mobility anchor assignment functionality and corresponding mobility options for Proxy Mobile IPv6. The runtime local mobility anchor assignment takes place during a Proxy Binding Update and a Proxy Binding Acknowledgement message exchange between a mobile access gateway and a local mobility anchor. The runtime local mobility anchor assignment functionality defined in this specification can be used, for example, for load-balancing purposes.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6463>.

Copyright Notice

Copyright (c) 2012 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 and Terminology	4
2.1. Requirements	4
2.2. Terminology	4
3. Proxy Mobile IPv6 Domain Assumptions	5
4. Mobility Options	5
4.1. Redirect-Capability Mobility Option	5
4.2. Redirect Mobility Option	6
4.3. Load Information Mobility Option	7
4.4. Alternate IPv4 Care-of Address Mobility Option	9
5. Runtime LMA Assignment	9
5.1. General Operation	9
5.2. Mobile Access Gateway Operation	10
5.3. Local Mobility Anchor Operation	12
5.3.1. Co-Located rLMA and r2LMA Functions	13
5.3.2. Separate rLMA and r2LMA Functions (Proxy-MAG)	14
6. Handoff and Multi-Homing Considerations	18
7. Protocol Configuration Variables	18
8. Security Considerations	19
9. IANA Considerations	20
10. Acknowledgements	20
11. References	20
11.1. Normative References	20
11.2. Informative References	20

1. Introduction

This specification describes a runtime assignment of a local mobility anchor (LMA) for the Proxy Mobile IPv6 (PMIPv6) [RFC5213] protocol. The runtime LMA assignment takes place during a Proxy Binding Update (PBU) and a Proxy Binding Acknowledgement (PBA) message exchange between a mobile access gateway (MAG) and a LMA. The runtime LMA assignment functionality defined in this specification can be used, for example, for load-balancing purposes. MAGs and LMAs can also implement other load-balancing mechanisms that are completely transparent at the PMIPv6 protocol level and do not depend on the functionality defined in this specification.

The runtime LMA assignment functionality does not depend on the Domain Name System (DNS) or the Authentication, Authorization, and Accounting (AAA) infrastructure for the assignment of the LMA to which the mobile node (MN) is anchored. All MAGs and LMAs (either rFLMAs or r2LMAs; see Section 2.2) have to belong to the same PMIPv6 domain.

There are a number of reasons why the runtime LMA assignment is a useful addition to the PMIPv6 protocol. A few are identified below:

- o LMAs with multiple IP addresses: a cluster of LMAs or a blade architecture LMA may appear to the routing system as multiple LMAs with separate unicast IP addresses. A MAG can initially select any of the LMAs as the serving LMA using, for example, DNS- and AAA-based solutions. However, MAG's initial selection may be suboptimal from the LMA point of view and immediate runtime assignment to a "proper LMA" would be needed. The LMA could use a [RFC5142]-based approach, but that would imply unnecessary setting up of a mobility session in a "wrong LMA" with associated back-end support system interactions, additional signaling between the MAG and the LMA, and re-establishing a mobility session to the new LMA again with associated signaling.
- o Bypassing a load-balancer: a cluster of LMAs or a blade architecture LMA may have a load-balancer in front of them or integrated in one of the LMAs. The load-balancer would represent multiple LMAs during the LMA discovery phase and only its IP address would be exposed to the MAG thus hiding possible individual LMA or LMA blade IP addresses from the MAG. However, if all traffic must always go through the load-balancer, it quickly becomes a bottleneck. Therefore, a PMIPv6 protocol-level support for bypassing the load-balancer after the initial PBU/PBA exchange would greatly help scalability. Also, bypassing the load-balancer as soon as possible allows implementing load-balancers that do not maintain any MN-specific state information.

- o Independence from DNS: DNS-based load-balancing is a common practice. However, keeping MAGs up to date with LMA load status using DNS is hard, e.g., due to caching and unpredictable zone update delays [RFC6097]. Generally, LMAs constantly updating the [RFC2136] zone's master DNS server might not be feasible in a large PMIPv6 domain due to increased load on the master DNS server and additional background signaling. Furthermore, MAGs may perform (LMA) destination address selection decisions that are not in line with what the DNS administrator actually wanted [RFC3484].
- o Independence from AAA: AAA-based solutions have basically the same arguments as DNS-based solutions above. It is also typical that AAA-based solutions offload the initial LMA selection to the DNS infrastructure [RFC5779]. The AAA infrastructure does not return an IP address or a Fully Qualified domain Name (FQDN) to a single LMA; rather, it returns a FQDN representing a group of LMAs.
- o Support for IPv6 anycast addressing [RFC4291]: the current PMIPv6 specification does not specify how the PMIPv6 protocol should treat anycast addresses assigned to mobility agents. For example, a blade architecture LMA may have a unique unicast IP address for each blade and a single anycast address for all blades. A MAG could then initially send a PBU to an anycast LMA address and receive a PBA from an anycast LMA address. Once the MAG receives the unicast address of the runtime-assigned LMA blade through the initial PBU/PBA exchange, the subsequent communication continues using the unicast address.

As a summary, the DNS/AAA-based approaches cannot be used to select an "appropriate" LMA at runtime. Therefore, this specification defines a solution that is applicable for LMA implementations where the IP address known to the MAG is not the best LMA of choice at runtime.

2. Requirements and Terminology

2.1. Requirements

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].

2.2. Terminology

In addition to the terminology defined in [RFC5213], the following terminology is also used:

rflMA

An LMA that receives a PBU from a MAG and decides to assign an IP mobility session with a new target LMA (r2LMA).

r2LMA

The LMA assigned to a MAG as a result of the runtime LMA assignment.

Runtime Assignment Domain

A group of LMAs that consists of at least one rflMA and one or more r2LMAs (all are part of the same PMIPv6 domain). A rflMA is allowed to assign MAGs only with r2LMAs that belong to the same runtime assignment domain. The rflMA and one or more r2LMAs may consist of multiple blades in a single network element, multiple physical network elements, or multiple LMAs distributed geographically.

3. Proxy Mobile IPv6 Domain Assumptions

The runtime LMA assignment functionality has few assumptions within the PMIPv6 domain.

Each LMA in a runtime assignment domain MUST be reachable at a unicast IP address. The rflMA and the r2LMA MUST have a prior agreement, adequate means to secure their inter-LMA communication, and an established trust relationship to perform the runtime LMA assignment.


Each LMA and MAG participating in the runtime LMA assignment is assumed to have required Security Associations (SAs) pre-established. Dynamic negotiation of the SAs using, e.g., IKEv2 [RFC5996], SHOULD be supported but is out of scope of this specification.

4. Mobility Options

In the following sections, all presented values, bit fields, and addresses are in network byte order.

4.1. Redirect-Capability Mobility Option

The Redirect-Capability mobility option has the alignment requirement of 4n. There can be zero or one Redirect-Capability mobility option in the PBU. The format of the Redirect-Capability mobility option is shown below:



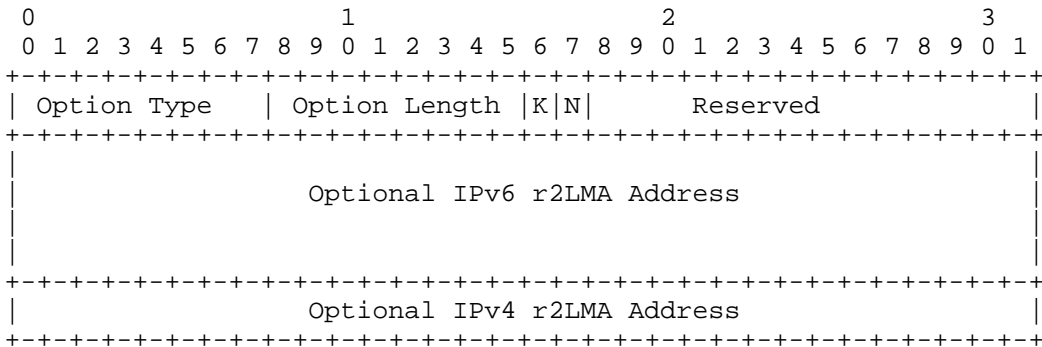
Redirect-Capability Mobility Option

- o Option Type: 8-bit identifier set to 46.
- o Option Length: 8-bit unsigned integer, representing the length of the Redirect-Capability mobility option in octets, excluding the Option Type and Length fields. The Option Length MUST be set to 2.
- o Reserved: This field is reserved for future use. This field MUST be set to zero by the sender and ignored by the receiver.

The Redirect-Capability option is used by the MAG to inform the LMA that it implements and has enabled the runtime LMA assignment functionality.

4.2. Redirect Mobility Option

The Redirect mobility option in the PBA MUST contain an unicast address of the r2LMA and the address family MUST be the same as the currently used transport between the MAG and the rLMA. There can be zero or one Redirect mobility option in the PBA. The Redirect mobility option has the alignment requirement of 4n. The format of the Redirect mobility option is shown below:



Redirect Mobility Option

- o Option Type: 8-bit identifier set to 47.

- o Option Length: 8-bit unsigned integer, representing the length of the Redirect mobility option in octets, excluding the Option Type and Length fields. If the 'K' flag is set and 'N' is unset, then the length MUST be 18. If the 'K' flag is unset and 'N' is set, then the length MUST be 6. Both the 'K' and 'N' flags cannot be set or unset simultaneously.
- o 'K' flag: This bit is set (1) if the 'Optional IPv6 r2LMA Address' is included in the mobility option. Otherwise, the bit is unset (0).
- o 'N' flag: This bit is set (1) if the 'Optional IPv4 r2LMA Address' is included in the mobility option. Otherwise, the bit is unset (0).
- o Reserved: This field is reserved for future use. MUST be set to zero by the sender and ignored by the receiver.
- o Optional IPv6 r2LMA Address: the unicast IPv6 address of the r2LMA. This value is present when the corresponding PBU was sourced from an IPv6 address.
- o Optional IPv4 r2LMA Address: the IPv4 address of the r2LMA. This value is present when the corresponding PBU was sourced from an IPv4 address (for IPv4 transport, see [RFC5844]).

The Redirect option is used by the LMA to inform the MAG that the runtime LMA assignment took place and the MAG has to update its Binding Update List Entry (BULE) for the mobility session.

4.3. Load Information Mobility Option

The Load Information mobility option can be included in any PBA and is used to report priority and key load information of a LMA to a MAG (or to a 'proxy-MAG'). The Load Information mobility option has the alignment requirement of 4n. The format of the mobility option is shown below:

0								1								2								3							
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Option Type								Option Length								Priority															
																Sessions in Use															
																Maximum Sessions															
																Used Capacity															
																Maximum Capacity															


Load Information Mobility Option

- o Option Type: 8-bit identifier set to 48.
- o Option Length: 8-bit unsigned integer, representing the length of the Load Information mobility option in octets, excluding the Option Type and Length fields. The length is set to 18.
- o Priority: 16-bit unsigned integer, representing the priority of an LMA. The lower value, the higher the priority. The priority only has meaning among a group of LMAs under the same administration, for example, determined by a common LMA FQDN, a domain name, or a realm.
- o Sessions in Use: 32-bit unsigned integer, representing the number of parallel mobility sessions the LMA has in use.
- o Maximum Sessions: 32-bit unsigned integer, representing the maximum number of parallel mobility sessions the LMA is willing to accept.
- o Used Capacity: 32-bit unsigned integer, representing the used bandwidth/throughput capacity of the LMA in kilobytes per second.
- o Maximum Capacity: 32-bit unsigned integer, representing the maximum bandwidth/throughput capacity in kilobytes per second the LMA is willing to accept.

The session and capacity information can easily be used to calculate different load factors of the LMA. A MAG (or a 'proxy-MAG') MAY use the priority and load information to internally maintain priority ordering of LMAs.

4.4. Alternate IPv4 Care-of Address Mobility Option

The Alternate IPv4 Care-of Address (A4CoA) mobility option has the alignment requirement of 4n+2. The format of the mobility option is shown below:



Alternate IPv4 Care-of Address Mobility Option

- o Option Type: 8-bit identifier set to 49.
- o Option Length: 8-bit unsigned integer, representing the length of the Load Information mobility option in octets, excluding the Option Type and Length fields. The length is set to 4.
- o Alternate IPv4 Care-of Address: an IPv4 equivalent of the [RFC6275] Alternate Care-of Address option for IPv6. In the context of PMIPv6, its semantic is equivalent to the Alternate Care-of Address option for IPv6.

A MAG MAY include the Alternate IPv4 Care-of Address option in a PBU. An LMA that receives and implements the Alternate IPv4 Care-of Address option MUST echo the option as such back to the MAG in a reply PBA.

5. Runtime LMA Assignment

5.1. General Operation

During the runtime LMA assignment, the PBA is returned from the LMA Address to which the PBU was sent, i.e., from the rLMA address. After the runtime LMA assignment, all PMIPv6 communication continues directly between the MAG and the r2LMA bypassing the rLMA. The overall runtime LMA assignment flow sequence is shown in Figure 1.

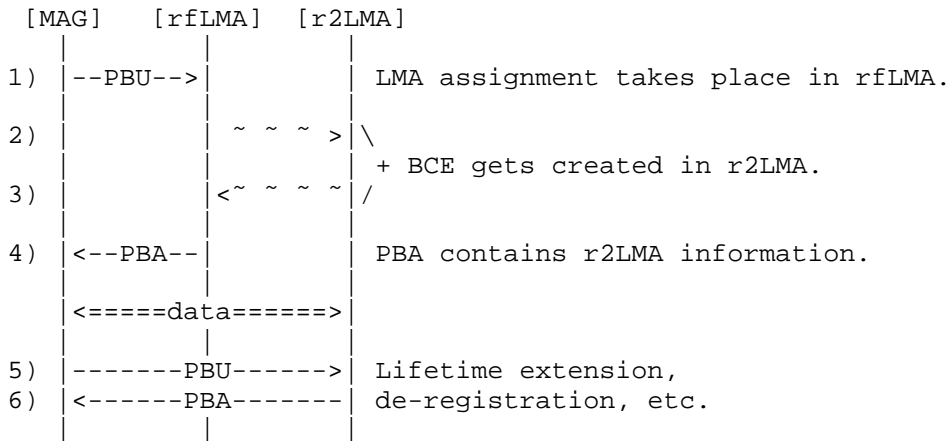


Figure 1: Runtime LMA Assignment from rFLMA to r2LMA and Setting Up a Mobility Session in the r2LMA within a Runtime Assignment Domain

The assumption in the signaling flow step 1) shown in Figure 1 is that the mobility session gets created in the r2LMA, although the rFLMA is responsible for interfacing with the MAG. There are several possible solutions for the rFLMA and the r2LMA interaction depending on, e.g., the co-location properties of the rFLMA and the r2LMA. This specification describes two:

- o Co-located rFLMA and r2LMA functions, where the 'rFLMA side of the LMA' is reachable via an anycast address or the loopback address of the LMA. See Section 5.3.1 for further details.
- o Separate rFLMA and r2LMA functions, where the rFLMA acts as a non-transparent 'proxy-MAG' to a r2LMA. See Section 5.3.2 for further details.

There are other possible implementations of the rFLMA and the r2LMA. At the end, as long as the protocol between the MAG and the rFLMA follows this specification, the co-location or inter-communication properties of the rFLMA and the r2LMA do not matter.

5.2. Mobile Access Gateway Operation

In the base PMIPv6 protocol [RFC5213], a MAG sends a PBU to an LMA; this results in creation of a Binding Cache Entry (BCE) at the LMA and the LMA sending a PBA sent back to the MAG. The MAG in turn creates a corresponding Binding Update List Entry (BULE). This specification extends the base protocol with the runtime LMA assignment functionality.

If the MAG supports the runtime LMA assignment and the functionality is also enabled (see the `EnableLMARedirectFunction` configuration variable in Section 7), then the MAG includes the Redirect-Capability mobility option in a PBU that establishes a new mobility session (i.e., Handoff Indicator Option in the PBU has the value of 1). The Redirect-Capability mobility option in the PBU is also an indication to an LMA that the MAG supports the runtime LMA assignment functionality and is prepared to be assigned with a different LMA. The runtime LMA assignment concerns always one mobility session at a time.

If the MAG receives a PBA that contains the Redirect mobility option without first including the Redirect-Capability mobility option in the corresponding PBU, then the MAG MUST ignore the option and process the PBA as described in RFC 5213.

If the MAG receives a PBA that contains the Redirect mobility option and the MAG had included the Redirect-Capability mobility option in the corresponding PBU, then the MAG MUST perform the following steps in addition to the normal [RFC5213] PBA processing:

- o The MAG updates its BULE to contain the r2LMA address included in the received Redirect mobility option.
- o If there is no SA between the MAG and the r2LMA, the MAG SHOULD initiate a dynamic creation of the SA between the MAG and the r2LMA as described in Section 4 of RFC 5213. If the dynamic SA creation fails, the MAG SHOULD log the event. The MAG MAY retry the dynamic creation of the SA, and if those also fail, the newly created BULE (and also the BUL in the r2LMA) will eventually timeout. If the failure is persistent, it can be regarded as a system-level configuration error.

The MAG is not required to send a fresh PBU to the r2LMA after a successful runtime assignment. The mobility session has already been established in the r2LMA. The MAG MUST send all user traffic to the r2LMA address. The MAG MUST send subsequent binding refresh PBUs (e.g., lifetime extension, handoff, etc.) to the r2LMA address. If there is no existing tunnel between the MAG and the r2LMA unicast address, then the MAG creates one as described in Section 6.9.1.2 of [RFC5213].

5.3. Local Mobility Anchor Operation

The text in the following sections refers to an 'LMA' when it means the combination of the rFLMA and the r2LMA, i.e., the entity where runtime LMA assignment is possible. When the text points to a specific LMA role during the runtime assignment, it uses either the 'rFLMA' or the 'r2LMA'.

If the runtime assignment functionality is enabled (see the EnableLMARedirectFunction configuration variable in Section 7) in the rFLMA but the LMA assignment is not going to take place for some reason, and the rFLMA is not willing to serve (or not capable of serving) as a normal [RFC5213] LMA for the MAG, then the rFLMA MUST reject the PBU and send back a PBA with Status Value set to 130 (Insufficient resources) error code. If the rFLMA is able to make the assignment to an r2LMA, it returns a PBA with the Redirect mobility option as defined below. Otherwise, the rFLMA MUST act as a normal [RFC5213]- or [RFC5844]-defined LMA for the MAG.

The rFLMA MUST only assign the MAG to a new r2LMA with which it knows the MAG has an SA or with which it knows the MAG can establish an SA dynamically. The rFLMA MUST NOT assign the MAG with a r2LMA that the rFLMA and the r2LMA do not have a prior agreement and an established trust relationship for the runtime LMA assignment. These SA-related knowledge issues and trust relationships are deployment specific in a PMIPv6 domain and in a runtime assignment domain, and out of scope of this specification. Possible context transfer and other coordination management between the rFLMA and the r2LMA are again deployment specific for LMAs in a runtime assignment domain. The rFLMA MUST NOT change the used transport IP address family during the runtime LMA assignment.

As a result of a successful runtime LMA assignment, the PBA MUST contain the Redirect mobility option with a valid r2LMA unicast address and the PBA Status Value indicating success.

Next, we describe two deployment and implementation models for the runtime LMA assignment. In Section 5.3.1, we describe a model where the rFLMA and r2LMA are co-located. In Section 5.3.2 we describe a model where the rFLMA acts as a non-transparent 'proxy-MAG', and where the rFLMA and the r2LMA are separate. There can be even more implementation options depending on the rFLMA and the r2LMA co-location properties, and how the inter-LMA communication is arranged.

5.3.1. Co-Located rLMA and r2LMA Functions

In this solution approach, the rLMA and the r2LMA are part of the same 'co-located LMA', and may even be using the same physical network interface. The rLMA is reachable via an anycast or a loopback address of the LMA. Each r2LMA is reachable via its unicast address. Figure 2 illustrates example signaling flows for the solution.

The MAG-LMA SA is between the MAG and the rLMA (i.e., the anycast or the loopback address of the LMA). How this SA has been set up is out of scope of this specification, but a manual SA configuration is one possibility.

The rLMA becomes active when the runtime LMA assignment functionality is enabled (see the `EnableLMARedirectFunction` configuration variable in Section 7). When the rLMA receives a PBU destined to it, and the PBU contains the `Redirect-Capability` mobility option, then the 'co-located LMA' MUST create a mobility session in a r2LMA role using the procedures described in [RFC5213]. If there is no existing tunnel between the MAG and the r2LMA unicast address, then the r2LMA creates one as described in Section 5.3 of [RFC5213]. The r2LMA used for accepting and anchoring the mobility session MUST also have the runtime LMA assignment functionality enabled (see the `EnableLMARedirectAcceptFunction` configuration variable in Section 7).

If the mobility session creation succeeded, then the 'co-located LMA' in the rLMA role sends a PBA to the MAG. The PBA is sourced using the rLMA (anycast or loopback) address. The PBA MUST contain the r2LMA unicast address (IPv6 or IPv4) in the `Redirect` mobility option.

If the PBU is received on the r2LMA unicast address, then the PBU is processed as described in RFC 5213 and the response PBA MUST NOT contain the `Redirect` mobility option.

If the PBU is received on the rLMA address and there is no `Redirect-Capability` mobility option in the PBU, then the 'co-located LMA' MAY choose to be a LMA for the MAG (assuming the rLMA address is not an anycast address). Otherwise, the rLMA MUST reject the PBU and send back a PBA in a rLMA role with `Status Value` set to 130 (`Insufficient resources`) error code (as mentioned in Section 5.3).

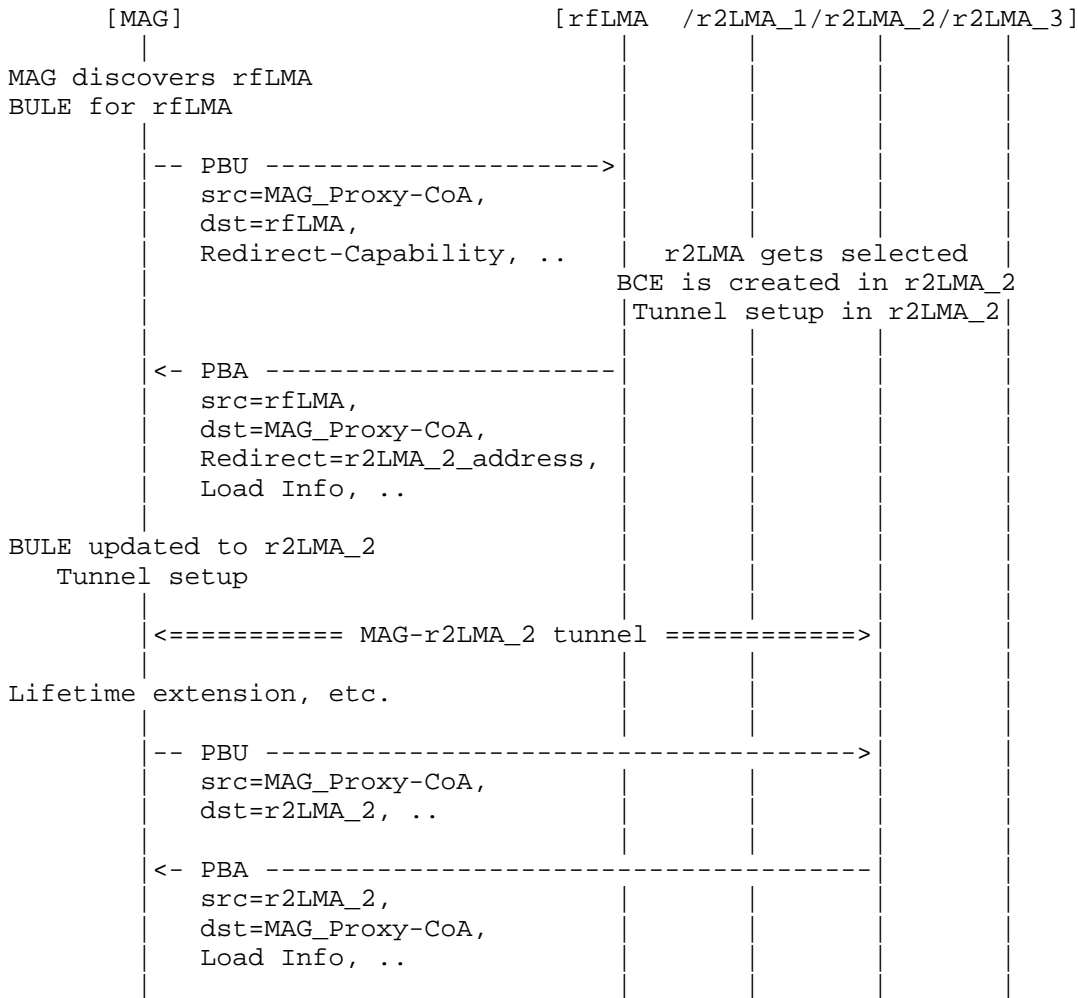


Figure 2: Co-Located rflMA and r2LMA Example

5.3.2. Separate rflMA and r2LMA Functions (Proxy-MAG)

In this solution approach, the rflMA and the r2LMA are two isolated functions, and may even be physically separate networking nodes. The r2LMA can be any [RFC5213]- or [RFC5844]-compliant LMA that doesn't have any knowledge of this specification when IPv6 transport is used. In case of IPv4 transport, the [RFC5844]-compliant LMA MUST also implement the Alternate IPv4 Care-of Address option (see Section 4.4). Figure 3 illustrates example signaling flows for the solution.

The rLMA is actually a non-transparent 'proxy-MAG' that shows up as an LMA implementing this specification towards the MAG, and as a base [RFC5213]-compliant MAG to the r2LMA. (See [RFC2616] for a generic definition of a non-transparent proxy; although it's for HTTP, the idea also applies here.) This type of operation is also referred to as 'chaining' in other contexts. The protocol between the 'proxy-MAG' and the r2LMA is the base [RFC5213] PMIPv6 protocol.

The MAG-LMA SA is between the MAG and the rLMA, and [RFC5213] SA considerations apply fully. The MAG has no knowledge of the 'proxy-MAG'-r2LMA SA. [RFC5213] considerations regarding the SA between the 'proxy-MAG' and the r2LMA apply fully. It is also possible that 'proxy-MAG'-r2LMA security is arranged using other means than IPsec, for example, using layer-2 VPNs.

When the rLMA receives a PBU, and the PBU contains the Redirect-Capability mobility option, then the rLMA in a 'proxy-MAG' role:

- o Processes the PBU using the procedures described in RFC 5213 except that no mobility session gets created. Instead, the rLMA creates a proxy state based on the received PBU.
- o Assigns a r2LMA to the MAG.
- o Creates a new PBU', which includes all non-security related mobility options from the original PBU and an Alternate Care-of Address (ACoA) option containing the Proxy Care-of Address of the original MAG. If the original PBU already included an ACoA option, then the content of the ACoA option in the PBU' MUST be the same as in the original PBU.

Note, in case of IPv4 transport [RFC5844], the Alternate IPv4 Care-of Address (A4CoA) option MUST be used and contain the IPv4 Proxy Care-of Address of the original MAG.

- o Sends the new PBU' sourced from its 'proxy-MAG' IPv6 or IPv4 Proxy Care-of Address and destined to the r2LMA address using the procedures described in RFC 5213 (or RFC 5844 in case of IPv4 transport).

The r2LMA processes the received PBU' using the procedures described in RFC 5213 or RFC 5844. In case of IPv4 transport, the r2LMA uses the IPv4 Proxy Care-of Address from the Alternate IPv4 Care-of Address option for the tunnel setup and the creation of the BCE. The reply PBA' MUST be destined to the source address of the received PBU', i.e., the Care-of Address the 'proxy-MAG'.

Once the rLMA in a 'proxy-MAG' role receives a reply PBA' from the r2LMA and the mobility session creation succeeded in the r2LMA, the rLMA sends a PBA to the original MAG. The PBA is sourced from the rLMA address and destined to the MAG (IPv6 or IPv4) Proxy Care-of Address. The PBA MUST contain the r2LMA (IPv6 or IPv4) unicast address in the Redirect mobility option. Other non-security-related mobility options (including the Load Information option) are copied from the PBA' to the PBA as such.

If one of these errors occurs:

- o the PBA' Status Value indicates that the mobility session creation failed in the r2LMA. For example, the Status Value in the PBA' is set to 130 (Insufficient resources), or
- o there was no PBA' response from the r2LMA, or
- o the PBA' did not include the Alternate IPv4 Care-of Address option although it was included in the corresponding PBU' (when using IPv4 transport),

then the rLMA SHOULD assign the MAG to a new r2LMA and rerun the procedure for sending the PBU' described earlier for the new r2LMA. The number and order of r2LMA reassignment attempts is controlled by the local policy and the amount of known r2LMAs in the rLMA. When the rLMA in a 'proxy-MAG' role concludes the mobility session creation failed with r2LMA(s), the rLMA MUST set the Status Value in the PBA as received from the latest contacted PBA' Status Value or to 130 (Insufficient resources) in case of no responses from rLMA(s), and send the reply PBA to the MAG. The PBA is sourced from the rLMA address and destined to the MAG Proxy Care-of Address. Other possible non-security-related mobility options (including the Load Information option) are copied from the PBA' to the PBA as such.

Once the rLMA has sent the reply PBA to the MAG, it can remove the proxy state. Subsequent traffic between the MAG and the r2LMA will bypass the rLMA (assuming the mobility session creation succeeded in the r2LMA).

If the rLMA receives a PBU with no Redirect-Capability mobility option in the PBU, then the PBU is processed as described in Section 5.3, i.e., the rLMA may or may not act as an [RFC5213] or [RFC5844] LMA to the MAG.

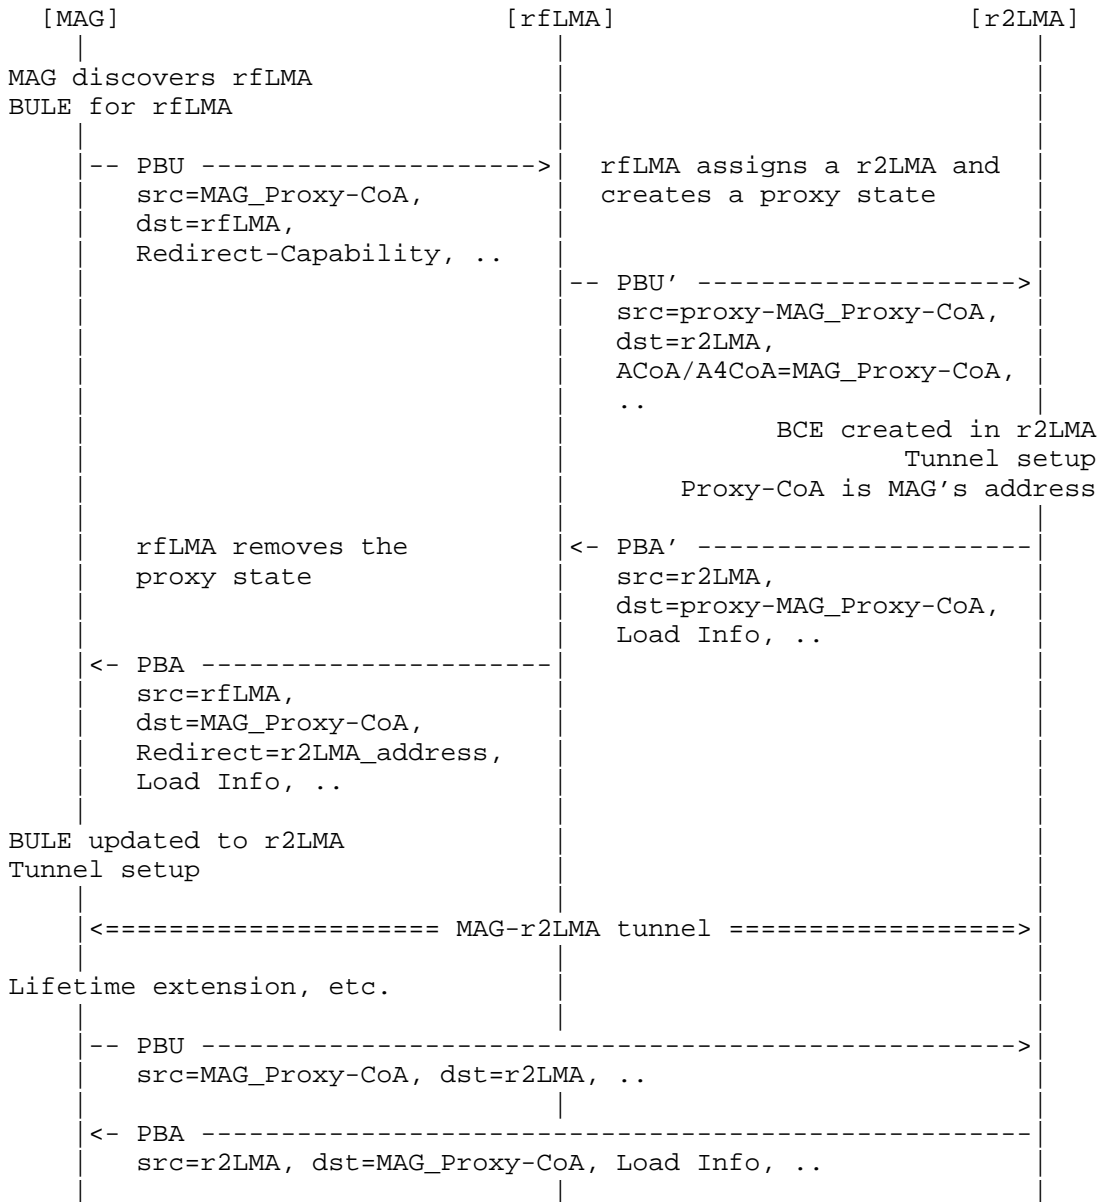


Figure 3: Separate rFLMA and r2LMA ('proxy-MAG') Example

6. Handoff and Multi-Homing Considerations

A MN can be multi-homed, i.e., have network connectivity over multiple interfaces connected to one or more accesses. If PMIPv6-based handovers between multiple interfaces or accesses are desired, then a single LMA should have a control over all possible multi-homed mobility sessions the MN has. Once the MN has established one mobility session with one LMA, the subsequent mobility sessions of the same MN would be anchored to the LMA that was initially assigned. If each mobility session over a different interface (and possibly a MAG) has no requirements for PMIPv6-based handovers between accesses or interfaces, then the rest of the considerations in this section do not apply.

One possible solution already supported by this specification is applying the runtime LMA assignment only for the very first initial attach a multi-homed MN does towards a PMIPv6 domain. After the initial attach, the assigned r2LMA address has been stored in the policy profile. For the subsequent mobility sessions of the multi-homed MN, the same assigned r2LMA address would be used and there is no need to contact the rLMA. Ensuring the discovery of the same r2LMA each time relies on the MN having an identity that can always point to the same policy profile, independent of the access that is used.

MAGs have a control over selectively enabling and disabling the runtime assignment of the LMA. If the multi-homed MN is attached to a PMIPv6 domain via multiple MAGs, the assigned r2LMA address should be stored in the remote policy store and downloaded as a part of the policy profile download to a MAG. Alternatively, MAGs can share policy profile information using other means. In both cases, the actual implementation of the policy profile information sharing is specific to a PMIPv6 deployment and out of scope of this specification.

7. Protocol Configuration Variables

This specification defines two configuration variables that control the runtime LMA assignment functionality within a PMIPv6 domain.

EnableLMARedirectFunction

This configuration variable is available in both a MAG and in a rLMA. When set to TRUE (i.e., enabled), the PMIPv6 node enables the runtime LMA assignment functionality. The default value is FALSE (i.e., disabled).

EnableLMARedirectAcceptFunction

This configuration variable is available in a r2LMA. When set to TRUE (i.e., enabled), the r2LMA is able to accept runtime LMA assignment mobility sessions from a r1LMA. The default value is FALSE (i.e., disabled).

Note that the MAG and LMA configuration variables from Sections 9.1 and 9.2 of [RFC5213] do not apply for an LMA when it is in an r1LMA role.

8. Security Considerations

The security considerations of PMIPv6 signaling described in RFC 5213 apply to this document. An incorrectly configured LMA may cause unwanted runtime LMA assignment attempts to non-existing LMAs or to other LMAs that do not have and will not have an SA with the MAG. Consequently, the MAG will experience failed binding updates or unsuccessful creation of mobility sessions. An incorrectly configured LMA may also cause biased load distribution within a PMIPv6 domain. This document also assumes that the LMAs that participate in runtime LMA assignment have adequate prior agreement and trust relationships between each other.

If the SAs between MAGs and LMAs are manually keyed (as may be needed by the scenario described in Section 5), then the anti-replay service of ESP-protected PMIPv6 traffic cannot typically be provided. This is, however, deployment specific to a PMIPv6 domain.

If a PMIPv6 domain deployment with a runtime LMA assignment requires that a r1LMA has to modify a PBU/PBA in any way, e.g., by changing the source and destination IP address or any other field of the encapsulating IP packet, then the security mechanism (such as possible authentication options) used to protect the PBU/PBA MUST NOT cover the outer IP packet on those parts that might get modified. Alternatively, the r1LMA can do all required security processing on the PBU/PBA, and the communication between the r1LMA and the r2LMA would be unprotected at the PMIPv6 protocol level. In this case, the runtime assignment domain MUST implement an adequate level of security using other means, such as layer-2 VPNs.

9. IANA Considerations

New mobility options for use with PMIPv6 are defined in the [RFC6275] "Mobility Options" registry. The mobility options are defined in Section 4:

Redirect-Capability Mobility Option	46
Redirect Mobility Option	47
Load Information Mobility Option	48
Alternate IPv4 Care-of Address	49

10. Acknowledgements

The author would like to thank Basavaraj Patil, Domagoj Premec, Ahmad Muhanna, Vijay Devarapalli, Rajeev Koodli, Yungui Wang, Pete McCann, and Qin Wu for their discussion of this document. A special thanks to Qian Li for her detailed feedback on the protocol details.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

11.2. Informative References

- [RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

- [RFC5142] Haley, B., Devarapalli, V., Deng, H., and J. Kempf, "Mobility Header Home Agent Switch Message", RFC 5142, January 2008.
- [RFC5779] Korhonen, J., Bournelle, J., Chowdhury, K., Muhanna, A., and U. Meyer, "Diameter Proxy Mobile IPv6: Mobile Access Gateway and Local Mobility Anchor Interaction with Diameter Server", RFC 5779, February 2010.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6097] Korhonen, J. and V. Devarapalli, "Local Mobility Anchor (LMA) Discovery for Proxy Mobile IPv6", RFC 6097, February 2011.

Authors' Addresses

Jouni Korhonen (editor)
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
Finland

E-Mail: jouni.nospam@gmail.com

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

E-Mail: sgundave@cisco.com

Hidetoshi Yokota
KDDI Lab
2-1-15 Ohara, Fujimino
Saitama 356-8502
Japan

E-Mail: yokota@kddilabs.jp

Xiangsong Cui
Huawei Technologies
Huawei Building, No. 156 Beiqing Rd.
Z-park, Shi-Chuang-Ke-Ji-Shi-Fan-Yuan
Hai-Dian District, Beijing 100095
P.R. China

E-Mail: Xiangsong.Cui@huawei.com

