Network Working Group
Request for Comments: 5380
Obsoletes: 4140
Category: Standards Track

H. Soliman
Elevate Technologies
C. Castelluccia
INRIA
K. ElMalki
Athonet
L. Bellier
INRIA

October 2008

Hierarchical Mobile IPv6 (HMIPv6) Mobility Management

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document introduces extensions to Mobile IPv6 and IPv6 Neighbour Discovery to allow for local mobility handling. Hierarchical mobility management for Mobile IPv6 is designed to reduce the amount of signalling between the mobile node, its correspondent nodes, and its home agent. The Mobility Anchor Point (MAP) described in this document can also be used to improve the performance of Mobile IPv6 in terms of handover speed.

Table of Contents

1. Introduction

This specification introduces the concept of a hierarchical Mobile IPv6 network, utilising a new node called the Mobility Anchor Point (MAP).

Mobile IPv6 [RFC3775] allows nodes to move within the Internet topology while maintaining reachability and ongoing connections between mobile and correspondent nodes. To do this, a mobile node sends binding updates (BUs) to its home agent (HA) every time it moves.

The mobile node may send data packets via its home agent immediately after sending the binding update, but the home agent will not be able to route traffic back to the mobile node before it receives the binding update. This incurs at least half a round-trip delay before packets are again forwarded to the right place. There is an additional delay for sending data packets if the mobile node chooses to wait for a binding acknowledgement (BA). The round-trip times can be relatively long if the mobile node and its home agent are in different parts of the world.

Additional delay is also incurred if the mobile node employs route optimisation. Authenticating binding updates requires approximately 1.5 round-trip times between the mobile node and each correspondent node (for the entire return routability procedure in a best-case scenario, i.e., no packet loss). This can be done in parallel with sending binding updates to the home agent, and there are further optimisations that reduce the required 1.5 round-trips [RFC4449] [RFC4651] [RFC4866].

Nevertheless, the signalling exchanges required to update your location will always cause some disruption to active connections. Some packets will be lost. Together with link layer and IP layer connection setup delays, there may be effects to upper-layer protocols. Reducing these delays during the time-critical handover period will improve the performance of Mobile IPv6.

Moreover, in the case of wireless links, such a solution reduces the number of messages sent over the air interface to all correspondent nodes and the home agent. A local anchor point will also allow Mobile IPv6 to benefit from reduced mobility signalling with external networks.

For these reasons, a new Mobile IPv6 node, called the Mobility Anchor Point, is used and can be located at any level in a hierarchical network of routers, including the Access Router (AR). The MAP will limit the amount of Mobile IPv6 signalling outside the local domain.

The introduction of the MAP provides a solution to the issues outlined earlier, in the following way:

- The mobile node sends binding updates to the local MAP rather than the home agent (HA) (which is typically further away) and correspondent nodes (CNs).

- Only one binding update message needs to be transmitted by the mobile node (MN) before traffic from the HA and all CNs is re-routed to its new location. This is independent of the number of CNs with which the MN is communicating.

A MAP is essentially a local home agent. The aim of introducing the hierarchical mobility management model in Mobile IPv6 is to enhance the performance of Mobile IPv6 while minimising the impact on Mobile IPv6 or other IPv6 protocols. Furthermore, HMIPv6 allows mobile nodes to hide their location from correspondent nodes and home agents, while using Mobile IPv6 route optimisation. The security differences between the MAP and the home agent are described in Section 12.

It is pertinent to note that the use of the MAP does not rely on, or assume the presence of, a permanent home agent. In other words, a mobile node need not have a permanent home address or home agent in order to be HMIPv6-aware or use the features in this specification. A MAP may be used by a mobile node in a nomadic manner to achieve mobility management within a local domain. Section 6.5 describes such a scenario.

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].

In addition, the following terms are introduced:

Access Router (AR)

The AR is the mobile node’s default router. The AR aggregates the outbound traffic of mobile nodes.

Mobility Anchor Point (MAP)

A Mobility Anchor Point is a router located in a network visited by the mobile node. The MAP is used by the MN as a local HA. One or more MAPs can exist within a visited network.

Regional Care-of Address (RCoA)

An RCoA is an address allocated by the MAP to the mobile node.

HMIPv6-Aware Mobile Node

An HMIPv6-aware mobile node is a mobile node that can receive and process the MAP option received from its default router. An HMIPv6-aware mobile node must also be able to send local binding updates (binding update with the M flag set).

On-Link Care-of Address

The LCoA is the on-link CoA configured on a mobile node’s interface based on the prefix advertised by its default router. In [RFC3775], this is simply referred to as the care-of address. However, in this memo LCoA is used to distinguish it from the RCoA.

Local Binding Update

The MN sends a local binding update to the MAP in order to establish a binding between the RCoA and LCoA.

3. Overview of HMIPv6

This hierarchical Mobile IPv6 scheme introduces a new function, the MAP, and minor extensions to the mobile node operation. The correspondent node and home agent operations will not be affected.

Just like Mobile IPv6, this solution is independent of the underlying access technology, allowing mobility within or between different types of access networks.

A mobile node entering a MAP domain will receive Router Advertisements containing information about one or more local MAPs. The MN can bind its current location (on-link CoA) with an address on the MAP’s subnet (RCoA). Acting as a local HA, the MAP will receive all packets on behalf of the mobile node it is serving and will encapsulate and forward them directly to the mobile node’s current address. If the mobile node changes its current address within a local MAP domain (LCoA), it only needs to register the new address with the MAP. Hence, only the Regional CoA (RCoA) needs to be registered with correspondent nodes and the HA. The RCoA does not change as long as the MN moves within a MAP domain (see below for definition). This makes the mobile node’s mobility transparent to correspondent nodes it communicates with.

A MAP domain’s boundaries are defined by the Access Routers (ARs) advertising the MAP information to the attached mobile nodes. The detailed extensions to Mobile IPv6 and operations of the different nodes will be explained later in this document.

It should be noted that the HMIPv6 concept is simply an extension to the Mobile IPv6 protocol. An HMIPv6-aware mobile node with an implementation of Mobile IPv6 SHOULD choose to use the MAP when discovering such capability in a visited network. However, in some cases the mobile node may prefer to simply use the standard Mobile IPv6 implementation. For instance, the mobile node may be located in a visited network within its home site. In this case, the HA is located near the visited network and could be used instead of a MAP. In this scenario, the mobile node would only update the HA whenever it moves. The method to determine whether the HA is in the vicinity of the MN (e.g., same site) is outside the scope of this document.

3.1. HMIPv6 Operations

The network architecture shown in Figure 1 illustrates an example of the use of the MAP in a visited network.

In Figure 1, the MAP can help in providing seamless mobility for the mobile node as it moves from Access Router 1 (AR1) to Access Router 2 (AR2), while communicating with the correspondent node. A multi-level hierarchy is not required for a higher handover performance. Hence, it is sufficient to locate one or more MAPs (possibly covering the same domain) at any position in the operator’s network.

```
   +-------+
   |  HA   |
   +-------+       +----+
       |           | CN |
       |           +----+
       |             |
       +-------+-----+
                     |
                     |RCoA
                 +-------+
                 |  MAP  |
                 +-------+
                 |      |
                 |      +--------+
                 |               |
                 |               |
               +-----+          +-----+
               | AR1 |          | AR2 |
               +-----+          +-----+
               LCoA1             LCoA2

               +----+
               | MN |
               +----+   ------------>
                         Movement
```

Figure 1: Hierarchical Mobile IPv6 domain

Upon arrival in a visited network, the mobile node will discover the global address of the MAP. This address is stored in the Access Routers and communicated to the mobile node via Router Advertisements (RAs). A new option for RAs is defined later in this specification. This is needed to inform mobile nodes about the presence of the MAP (MAP Discovery). The discovery phase will also inform the mobile node of the distance of the MAP from the mobile node. For example, the MAP function could be implemented as shown in Figure 1, and, at the same time, also be implemented in AR1 and AR2. In this case, the mobile node can choose the first hop MAP or one further up in the hierarchy of routers. The details on how to choose a MAP are provided in Section 10.

The process of MAP Discovery continues as the mobile node moves from one subnet to the next. Every time the mobile node detects movement, it will also detect whether it is still in the same MAP domain. The Router Advertisement used to detect movement will also inform the mobile node, through Neighbour Discovery [RFC4861] and the MAP option, whether it is still in the same MAP domain. As the mobile node roams within a MAP domain, it will continue to receive the same

MAP option included in Router Advertisements from its AR. If a change in the advertised MAP’s address is received, the mobile node MUST act on the change by sending binding updates to its HA and correspondent nodes.

If the mobile node is not HMIPv6-aware, then no MAP Discovery will be performed, resulting in the mobile node using the Mobile IPv6 [RFC3775] protocol for its mobility management. On the other hand, if the mobile node is HMIPv6-aware it SHOULD choose to use its HMIPv6 implementation. If so, the mobile node will first need to register with a MAP by sending it a BU containing its home address and on-link address (LCoA). The home address used in the BU is the RCoA, which the mobile node acquires via RFC 4877 [RFC4877] Section 9 mechanisms when it first contacts a given MAP. The MAP MUST store this information in its binding cache to be able to forward packets to their final destination when received from the different correspondent nodes or HAs.

The mobile node will always need to know the original sender of any received packets to determine if route optimisation is required. This information will be available to the mobile node because the MAP does not modify the contents of the original packet. Normal processing of the received packets (as described in [RFC3775]) will give the mobile node the necessary information.

To use the network bandwidth in a more efficient manner, a mobile node may decide to register with more than one MAP simultaneously and to use each MAP address for a specific group of correspondent nodes. For example, in Figure 1, if the correspondent node happens to exist on the same link as the mobile node, it would be more efficient to use the first hop MAP (in this case assume it is AR1) for communication between them. This will avoid sending all packets via the "highest" MAP in the hierarchy and thus will result in more efficient usage of network bandwidth. The mobile node can also use its current on-link address (LCoA) as a CoA, as specified in [RFC3775]. Note that the mobile node MUST NOT present an RCoA from a MAP’s subnet as an LCoA in a binding update sent to another MAP. The LCoA included in the binding update MUST be the mobile node’s address, derived from the prefix advertised on its link.

4. Mobile IPv6 Extension - Local Binding Update

This section outlines the extensions proposed to the binding update specified in [RFC3775].

A new flag is added: the M flag, which indicates MAP registration. When a mobile node registers with the MAP, the M and A flags MUST be set to distinguish this registration from a BU being sent to the HA or a correspondent node.

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |            Sequence #         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |A|H|L|K|M|      Reserved       |            Lifetime           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Local Binding Update

M

If set to 1, it indicates a MAP registration.

5. Neighbour Discovery Extension: The MAP Option

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |     Type      |    Length     |  Dist |  Pref |R|  Reserved   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                      Valid Lifetime                           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   +                                                               +
   |                                                               |
   +                  Global IP Address for MAP                    +
   |                                                               |
   +                                                               +
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: The MAP option

Type

IPv6 Neighbour Discovery option. Its value is 23.

Length

8-bit unsigned integer. The length of the option and MUST be set to 3.

Dist

A 4-bit unsigned integer identifying the distance between MAP and the receiver of the advertisement, measure in the number of hops and starting from 1 if the MAP is on the same link as the mobile node. A distance value of zero MUST NOT be used.

Pref

The preference field, used as an indicator of operator preference. A 4-bit unsigned integer. A decimal value of 15 indicates the highest preference. When comparing two potential MAPs, the mobile node SHOULD inspect this field as a tie-breaker for MAPs that have equal Dist values.

R

When set to 1, it indicates that the mobile node is allocated the RCoA by the MAP based on Section 9 of [RFC4877].

Valid Lifetime

The minimum value (in seconds) of both the valid lifetime of the prefix used to form the MAP’s address and that used to form the RCoA. This value indicates the validity of the MAP’s address and the RCoA.

Global Address

One of the MAP’s global addresses.

6. Protocol Operation

This section describes the HMIPv6 protocol. In HMIPv6, the mobile node has two addresses, an RCoA on the MAP’s link and an on-link CoA (LCoA). This RCoA is formed in a stateless manner by combining the mobile node’s interface identifier and the subnet prefix received in the MAP option.

As illustrated in this section, this protocol requires updating the mobile nodes’ implementation only. The HA and correspondent node are unchanged. The MAP performs the function of a "local" HA that binds the mobile node’s RCoA to an LCoA.

6.1. Mobile Node Operation

When a mobile node moves into a new MAP domain (i.e., its MAP changes), it needs to configure two CoAs: an RCoA on the MAP’s link and an on-link CoA (LCoA). After employing [RFC4877] to acquire an RCoA, the mobile node sends a local BU to the MAP with the A and M flags set. The local BU is a BU defined in [RFC3775] and includes the mobile node’s RCoA in the Home Address option. No alternate-CoA option is needed in this message. The LCoA is used as the source address of the BU. This BU will bind the mobile node’s RCoA (similar to a home address) to its LCoA. The MAP (acting as an HA) will then return a binding acknowledgement to the MN. This acknowledgement either identifies the binding as successful or contains the appropriate fault code. No new error codes need to be supported by the mobile node for this operation. The mobile node MUST silently ignore binding acknowledgements that do not contain a routing header type 2, which includes the mobile node’s RCoA.

Following a successful registration with the MAP, a bi-directional tunnel between the mobile node and the MAP is established. All packets sent by the mobile node are tunnelled to the MAP. The outer header contains the mobile node’s LCoA in the source address field, and the MAP’s address in the destination address field. The inner header contains the mobile node’s RCoA in the source address field, and the peer’s address in the destination address field. Similarly, all packets addressed to the mobile node’s RCoA are intercepted by the MAP and tunnelled to the mobile node’s LCoA.

This specification allows a mobile node to use more than one RCoA if it received more than one MAP option. In this case, the mobile node MAY perform the binding update procedure for each RCoA. In addition, the mobile node MUST NOT use one RCoA (e.g., RCoA1) derived from a MAP’s prefix (e.g., MAP1) as a care-of address in its binding update to another MAP (e.g., MAP2). This would force packets to be encapsulated several times (twice in this example) on their path to the mobile node. This form of multi-level hierarchy will reduce the protocol’s efficiency and performance.

After registering with the MAP, the mobile node MUST register its new RCoA with its HA by sending a BU that specifies the binding (RCoA, home address), as in Mobile IPv6. The mobile node’s home address is used in the Home Address option and the RCoA is used as the care-of address in the source address field. The mobile node may also send a similar BU (i.e., that specifies the binding between the home address and the RCoA) to its current correspondent nodes.

The mobile node SHOULD wait for the binding acknowledgement from the MAP before registering the RCoA with other nodes (e.g., CN or HA, if available). It should be noted that when binding the RCoA with the HA and correspondent nodes, the binding lifetime MUST NOT be larger than the mobile node’s binding lifetime with the MAP, which is received in the binding acknowledgement.

In order to speed up the handover between MAPs and reduce packet loss, a mobile node SHOULD send a local BU to its previous MAP, specifying its new LCoA. Packets in transit that reach the previous MAP are then forwarded to the new LCoA.

The MAP will receive packets addressed to the mobile node’s RCoA (from the HA or correspondent nodes). Packets will be tunnelled from the MAP to the mobile node’s LCoA. The mobile node will de-capsulate the packets and process them in the normal manner.

When the mobile node moves within the same MAP domain, it should only register its new LCoA with its MAP. In this case, the RCoA remains unchanged.

Note that a mobile node may send a BU containing its LCoA (instead of its RCoA) to correspondent nodes. If these nodes are connected to the same link, packets will then be routed directly, without going through the MAP.

6.1.1. Sending Packets to Correspondent Nodes

The mobile node can communicate with a correspondent node through the HA, or in a route-optimised manner, as described in [RFC3775]. When communicating through the HA, the message formats in [RFC3775] are used.

If the mobile node communicates directly with the correspondent node (i.e., the CN has a binding cache entry for the mobile node), the mobile node MUST use the same care-of address used to create a binding cache entry in the correspondent node (RCoA) as a source address. According to [RFC3775], the mobile node MUST also include a Home Address option in outgoing packets. The Home Address option MUST contain the mobile node’s home address.

6.2. MAP Operations

The MAP acts like an HA; it intercepts all packets addressed to registered mobile nodes and tunnels them to the corresponding LCoA, which is stored in its binding cache.

A MAP has no knowledge of the mobile node’s home address. The mobile node will send a local BU to the MAP with the M and A flags set. The aim of this BU is to bind the RCoA (contained in the BU as a home address) to the mobile node’s LCoA. If successful, the MAP MUST return a binding acknowledgement to the mobile node, indicating a successful registration. This is identical to the HA operation in [RFC3775]. No new error codes are introduced for HMIPv6. The binding acknowledgement MUST include a routing header type 2 that contains the mobile node’s RCoA.

The MAP MUST be able to accept packets tunnelled from the mobile node, with the mobile node being the tunnel entry point and the MAP being the tunnel exit point.

The MAP employs [RFC4877] Section 9 procedures for the allocation of RCoA, and subsequently acts as an HA for the RCoA. Packets addressed to the RCoA are intercepted by the MAP, using proxy Neighbour Advertisement, and then encapsulated and routed to the mobile node’s LCoA. This operation is identical to that of the HA described in [RFC3775].

A MAP MAY be configured with the list of valid on-link prefixes that mobile nodes can use to derive LCoAs. This is useful for network operators that need to stop mobile nodes from continuing to use the MAP after moving to a different administrative domain. If a mobile

node sent a binding update containing an LCoA that is not in the MAP’s "valid on-link prefixes" list, the MAP could reject the binding update using existing error code 129 (administratively prohibited).

6.3. Home Agent Operations

The support of HMIPv6 is completely transparent to the HA’s operation. Packets addressed to a mobile node’s home address will be forwarded by the HA to its RCoA, as described in [RFC3775].

6.4. Correspondent Node Operations

HMIPv6 is completely transparent to correspondent nodes.

6.5. Local Mobility Management Optimisation within a MAP Domain

In [RFC3775], it is stated that for short-term communication, particularly communication that may easily be retried upon failure, the mobile node MAY choose to directly use one of its care-of addresses as the source of the packet, thus not requiring the use of a Home Address option in the packet. Such use of the CoA will reduce the overhead of sending each packet due to the absence of additional options. In addition, it will provide an optimal route between the mobile node and correspondent node.

HMIPv6-aware mobile nodes can use their RCoA as the source address without using a Home Address option. In other words, the RCoA can be used as a source address for upper layers. Using this feature, the mobile node will be seen by the correspondent node as a fixed node while moving within a MAP domain.

This usage of the RCoA does not have the cost of Mobile IPv6 (i.e., no bindings or Home Address options are sent over the Internet), but still provides local mobility management to the mobile nodes with near-optimal routing. Although such use of RCoA does not provide global mobility (i.e., communication is broken when a mobile node changes its RCoA), it would be useful for several applications (e.g., web browsing). The validity of the RCoA as a source address used by applications will depend on the size of a MAP domain and the speed of the mobile node. Furthermore, because the support for BU processing in correspondent nodes is not mandated in [RFC3775], this mechanism can provide a way of obtaining route optimisation without sending BUs to the correspondent nodes.

Enabling this mechanism can be done by presenting the RCoA as a temporary home address for the mobile node. This may require an implementation to augment its source address selection algorithm with the knowledge of the RCoA in order to use it for the appropriate applications.

6.6. Location Privacy

In HMIPv6, a mobile node hides its LCoA from its correspondent nodes and its home agent by using its RCoA in the source field of the packets that it sends. As a result, address-based location tracking of a mobile node by its correspondent nodes or its home agent is more difficult because they only know its RCoA and not its LCoA.

7. MAP Discovery

This section describes how a mobile node obtains the MAP address and subnet prefix, and how ARs in a domain discover MAPs.

This specification requires network administrators to manually configure the MAP option information in ARs; future mechanisms may be defined to allow MAPs to be discovered dynamically.

7.1. Mobile Node Operation

When an HMIPv6-aware mobile node receives a Router Advertisement, it should search for the MAP option. One or more options may be found for different MAP IP addresses. A mobile node SHOULD register with the MAP having the highest preference value. A MAP with a preference value of zero SHOULD NOT be used for new local BUs (i.e., the mobile node can refresh existing bindings but cannot create new ones). However, a mobile node MAY choose to register with one MAP over another, depending on the value received in the distance field, provided that the preference value is above zero.

A MAP option containing a valid lifetime value of zero means that this MAP MUST NOT be selected by the MN. A valid lifetime of zero indicates a MAP failure. When this option is received, a mobile node MUST choose another MAP and create new bindings. Any existing bindings with this MAP can be assumed to be lost. If no other MAP is available, the mobile node MUST NOT attempt to use HMIPv6.

If a multi-homed mobile node has access to several ARs simultaneously (on different interfaces), it SHOULD use an LCoA on the link defined by the AR that advertises its current MAP.

A mobile node MUST store the received option(s) in order to choose at least one MAP to register with. Storing the options is essential, as they will be compared to other options received later for the purpose of the movement detection algorithm.

If the R flag is set, the mobile node MUST place its RCoA in place of the home address in the binding update message. This causes the RCoA to be bound to the LCoA in the MAP’s binding cache.

A mobile node MAY choose to register with more than one MAP simultaneously, or use both the RCoA and its LCoA as care-of addresses simultaneously with different correspondent nodes.

8. Updating Previous MAPs

When a mobile node moves into a new MAP domain, the mobile node may send a BU to the previous MAP requesting it to forward packets addressed to the mobile node’s new CoA. An administrator MAY restrict the MAP from forwarding packets to LCoAs outside the MAP’s domain. However, it is RECOMMENDED that MAPs be allowed to forward packets to LCoAs associated with some of the ARs in neighbouring MAP domains, provided that they are located within the same administrative domain.

For instance, a MAP could be configured to forward packets to LCoAs associated with ARs that are geographically adjacent to ARs on the boundary of its domain. This will allow for a smooth inter-MAP handover as it allows the mobile node to continue to receive packets while updating the new MAP, its HA and, potentially, correspondent nodes.

9. Note on MAP Selection by the Mobile Node

HMIPv6 provides a flexible mechanism for local mobility management within a visited network. As explained earlier, a MAP can exist anywhere in the operator’s network (including the AR). Several MAPs can be located within the same domain independently of each other. In addition, overlapping MAP domains are also allowed and recommended. Both static and dynamic hierarchies are supported.

When the mobile node receives a Router Advertisement including a MAP option, it should perform actions according to the following movement detection mechanisms. In a hierarchical Mobile IP network, such as the one described in this document, the mobile node should be:

- "Eager" to perform new bindings.

- "Lazy" in releasing existing bindings.

The above means that the mobile node should register with any "new" MAP advertised by the AR (Eager). The method by which the mobile node determines whether the MAP is a "new" MAP is described in Section 9.1. The mobile node should not release existing bindings until it no longer receives the MAP option (or receives it with a lifetime of zero) or the lifetime of its existing binding expires (Lazy). This Eager-Lazy approach, described above, will assist in providing a fallback mechanism in case of the failure of one of the MAP routers, as it will reduce the time it takes for a mobile node to inform its correspondent nodes and HA about its new care-of address.

9.1. MAP Selection in Distributed MAP Environment

The mobile node needs to consider several factors to optimally select one or more MAPs, where several MAPs are available in the same domain.

There are no benefits foreseen in selecting more than one MAP and forcing packets to be sent from the higher MAP down through a hierarchy of MAPs. This approach may add forwarding delays and eliminate the robustness of IP routing between the highest MAP and the mobile node; therefore, it is prohibited by this specification. Allowing more than one MAP ("above" the AR) within a network should not imply that the mobile node forces packets to be routed down the hierarchy of MAPs. However, placing more than one MAP "above" the AR can be used for redundancy and as an optimisation for the different mobility scenarios experienced by mobile nodes. The MAPs are used independently of each other by the MN (e.g., each MAP is used for communication to a certain set of CNs).

In terms of the distance-based selection in a network with several MAPs, a mobile node may choose to register with the furthest MAP to avoid frequent re-registrations. This is particularly important for fast mobile nodes that will perform frequent handoffs. In this scenario, the choice of a more distant MAP would reduce the probability of having to change a MAP and informing all correspondent nodes and the HA.

In a scenario where several MAPs are discovered by the mobile node in one domain, the mobile node may need sophisticated algorithms to be able to select the appropriate MAP. These algorithms would have the mobile node speed as an input (for distance-based selection) combined with the preference field in the MAP option. However, this specification proposes that the mobile node use the following algorithm as a default, where other optimised algorithms are not available. The following algorithm is simply based on selecting the MAP that is most distant, provided that its preference value did not reach a value of zero. The mobile node operation is shown below:

1. Receive and parse all MAP options.

2. Arrange MAPs in a descending order, starting with the furthest MAP (i.e., MAP option having largest Dist field).

3. Select first MAP in list.

4. If either the preference value or the valid lifetime fields are set to zero, select the following MAP in the list.

5. Repeat step (4) while new MAP options still exist, until a MAP is found with a non-zero preference value and a non-zero valid lifetime.

Implementing the steps above would result in mobile nodes selecting, by default, the most distant or furthest available MAP. This will continue until the preference value reduces to zero. Following this, mobile nodes will start selecting another MAP.

9.2. MAP Selection in a Flat Mobility Architecture

Network operators may choose a flat architecture in some cases where a Mobile IPv6 handover may be considered a rare event. In these scenarios, operators may choose to include the MAP function in ARs only. The inclusion of the MAP function in ARs can still be useful to reduce the time required to update all correspondent nodes and the HA. In this scenario, a mobile node may choose a MAP (in the AR) as an anchor point when performing a handoff. This kind of dynamic hierarchy (or anchoring) is only recommended for cases where inter-AR movement is not frequent.

10. Detection and Recovery from MAP Failures

This specification introduces a MAP that can be seen as a local home agent in a visited network. A MAP, like a home agent, is a single point of failure. If a MAP fails, its binding cache content will be lost, resulting in loss of communication between mobile and correspondent nodes. This situation may be avoided by using more than one MAP on the same link and by utilising a form of context transfer protocol between them. However, MAP redundancy is outside the scope of this document.

In cases where such protocols are not supported, the mobile node would need to detect MAP failures. The mobile node can detect this situation when it receives a Router Advertisement containing a MAP option with a lifetime of zero. The mobile node should then start the MAP Discovery process and attempt to register with another MAP. After it has selected and registered with another MAP, it will also need to inform correspondent nodes and the home agent if its RCoA has changed. Note that in the presence of a protocol that transfers binding cache entries between MAPs for redundancy purposes, a new MAP may be able to provide the same RCoA to the mobile node (e.g., if both MAPs advertise the same prefix in the MAP option). This would save the mobile node from updating correspondent nodes and the home agent.

Access Routers can be triggered to advertise a MAP option with a lifetime of zero (indicating MAP failure) in different ways:

- By manual intervention.

- In a dynamic manner.

One way of performing dynamic detection of MAP failure can be done by probing the MAP regularly (e.g., every 10 seconds). If no response is received, an AR MAY try to aggressively probe the MAP for a short period of time (e.g., once every 5 seconds for 15 seconds); if no reply is received, a MAP option may be sent with a valid lifetime value of zero. The exact mechanisms for probing MAPs is outside the scope of this document. The above text simply shows one example of detecting failures.

This specification does not mandate a particular recovery mechanism. However, any mechanism between the MAP and an AR SHOULD be secure to allow for message authentication, integrity protection, and protection against replay attacks.

Note that the above suggestion for detecting MAP failure may not detect MAP failures that might take place between probes, i.e.,if a MAP reboots between probes.

11. Tunelling Impacts on MTU

This specification requires the mobile node to tunnel outgoing traffic to the MAP. Similarly, the MAP tunnels inbound packets to the mobile node. If the mobile node has a home agent elsewhere on the Internet, this will result in double encapsulations of inbound and outbound packets. This may have impacts on the mobile node’s path MTU. Hence, mobile nodes MUST consider the encapsulation of traffic between the node and the MAP when calculating the available MTU for upper layers.

12. Security Considerations

This specification introduces a new concept to Mobile IPv6, namely, a Mobility Anchor Point that acts as a local home agent. It is crucial that the security relationship between the mobile node and the MAP is strong; it MUST involve mutual authentication, integrity protection, and protection against replay attacks. Confidentiality may be needed for payload traffic, such as when the mobile node is unwilling to reveal any traffic to the access network beyond what is needed for the mobile node to attach to the network and communicate with a MAP. Confidentiality is not required for binding updates to the MAP. The absence of any of these protections may lead to malicious mobile nodes impersonating other legitimate ones or impersonating a MAP. Any of these attacks will undoubtedly cause undesirable impacts to the mobile node’s communication with all correspondent nodes having knowledge of the mobile node’s RCoA.

Three different relationships (related to securing binding updates) need to be considered:

1. The mobile node - MAP

2. The mobile node - correspondent node

3. The mobile node - home agent

12.1. Mobile Node - MAP Security

In order to allow a mobile node to use the MAP’s forwarding service, initial authorisation (specifically for the service, not for the RCoA) MAY be needed. Authorising a mobile node to use the MAP service can be done based on the identity of the mobile node exchanged during the security association (SA) negotiation process. The authorisation may be granted based on the mobile node’s identity or based on the identity of a Certificate Authority (CA) that the MAP trusts. For instance, if the mobile node presents a certificate signed by a trusted entity (e.g., a CA that belongs to the same administrative domain, or another trusted roaming partner), it would be sufficient for the MAP to authorise the use of its service. Note that this level of authorisation is independent of authorising the use of a particular RCoA. Similarly, the mobile node trusts the MAP if it presents a certificate signed by the same CA or by another CA that the mobile node is configured to trust (e.g., a roaming partner). It is likely that some deployments would be satisfied with the use of self-signed certificates for either the mobile node or the MAP or both. This guarantees that the mobile node and the MAP are authenticated for address allocation and future binding updates without the need for identity authentication. Hence, the use of trusted third-party certificates is not required by this specification.

It is important to note that in this specification, authentication and authorisation are effectively the same thing. All the MAP needs in order to allocate the mobile node an RCoA is to authenticate the mobile node and verify that it belongs to a trusted group (based on its certificate).

IKEv2 MUST be supported by the mobile node and the MAP. IKEv2 allows the use of Extensible Authentication Protocol (EAP) as a mechanism to bootstrap the security association between the communicating peers.

Hence, EAP can be used with IKEv2 to leverage the Authentication, Authorization, and Accounting (AAA) infrastructure to bootstrap the SA between the mobile node and the MAP. Such a mechanism is useful in scenarios where an administrator wishes to avoid the configuration and management of certificates on mobile nodes. A MAP MAY support the use of EAP over IKEv2.

If EAP is used with IKEv2, the EAP method runs between the mobile node and a AAA server. Following a successful authentication, the resulting keying material can be used to bootstrap IKEv2 between the MAP and the mobile node. The specification of which EAP methods should be used or how keys are transported between the MAP and the AAA server is outside the scope of this document.

HMIPv6 uses an additional registration between the mobile node and its current MAP. As explained in this document, when a mobile node moves into a new domain (i.e., served by a new MAP), it obtains an RCoA and an LCoA and registers the binding between these two addresses with the new MAP. The MAP then verifies the BU and creates a binding cache entry with the RCoA and LCoA. Whenever the mobile node gets a new LCoA, it needs to send a new BU that specifies the binding between its RCoA and its new LCoA. This BU needs to be authenticated; otherwise, any host could send a BU for the mobile node’s RCoA and hijack the mobile node’s packets.

The MAP does not need to have prior knowledge of the identity of the mobile node or its home address. As a result, the SA between the mobile node and the MAP can be established using any key establishment protocols such as IKEv2. A return routability test is not necessary.

The MAP needs to set the SA for the RCoA (not the LCoA). This can be performed with IKEv2 [RFC4306]. The mobile node uses its LCoA as the source address, but specifies that the RCoA should be used in the SA.

This is achieved by using the RCoA as the identity in the IKE CHILD_SA negotiation. This step is identical to the use of the home address in IKE CHILD_SA when negotiating with the home agent.

The IPsec Peer Authorization Database (PAD) entries and configuration payloads described in [RFC4877] for allocating dynamic home addresses SHOULD be used by the MAP to allocate the RCoA for mobile nodes. Binding updates between the MAP and the mobile node MUST be protected with either Authentication Header (AH) or Encapsulating Security Payload (ESP) in transport mode. When ESP is used, a non-null authentication algorithm MUST be used.

The Security Policy Database (SPD) entries in both the home agent and the mobile node are identical to those set up for the home agent and mobile node, respectively, as outlined in [RFC4877].

12.2. Mobile Node - Correspondent Node Security

Mobile IPv6 [RFC3775] defines a return routability procedure that allows mobile and correspondent nodes to authenticate binding updates and acknowledgements. This specification does not impact the return routability test defined in [RFC3775]. However, it is important to note that mobile node implementers need to be careful when selecting the source address of the HoTI and CoTI messages, defined in [RFC3775]. The source address used in HoTI messages SHOULD be the mobile node’s home address unless the mobile node wishes to use the RCoA for route optimisation. The packet containing the HoTI message is encapsulated twice. The inner encapsulating header contains the RCoA in the source address field and the home agent’s address in the destination address field. The outer encapsulating header contains the mobile node’s LCoA in the source address field and the MAP’s address in the destination field.

12.3. Mobile Node - Home Agent Security

The security relationship between the mobile node and its home agent, as discussed in [RFC3775], is not impacted by this specification.

The relationship between the MAP and the mobile node is not impacted by the presence of a home agent.

13. IANA Considerations

Both the MAP option and M flag were allocated for RFC 4140 and will continue to be used by this specification.

14. Acknowledgements

The authors would like to thank Conny Larsson (Ericsson) and Mattias Pettersson (Ericsson) for their valuable input to this document. The authors would also like to thank the members of the French RNRT MobiSecV6 project (BULL, France Telecom, and INRIA) for testing the first implementation and for their valuable feedback. The INRIA HMIPv6 project is partially funded by the French government.

In addition, the authors would like to thank the following members of the working group, in alphabetical order: Samita Chakrabarti (Sun), Gregory Daley, Gopal Dommety (Cisco), Francis Dupont (GET/Enst Bretagne), Eva Gustaffson (Ericsson), Dave Johnson (Rice University), Annika Jonsson (Ericsson), James Kempf (Docomo labs), Martti

Kuparinen (Ericsson), Fergal Ladley, Gabriel Montenegro (Microsoft), Nick "Sharkey" Moore, Vidya Narayanan (Qualcomm), Erik Nordmark (Sun), Basavaraj Patil (Nokia), Brett Pentland (NEC), Thomas Schmidt, and Alper Yegin (Samsung) for their comments on the document.

15. References

15.1. 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.

[RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[RFC4877] Devarapalli, V. and F. Dupont, "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture", RFC 4877, April 2007.

15.2. Informative References

[RFC4449] Perkins, C., "Securing Mobile IPv6 Route Optimization Using a Static Shared Key", RFC 4449, June 2006.

[RFC4651] Vogt, C. and J. Arkko, "A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization", RFC 4651, February 2007.

[RFC4866] Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6", RFC 4866, May 2007.

Appendix A. Changes from RFC 4140

- Dynamic MAP Discovery was removed.

- Updated the security section to use IKEv2 instead of IKEv1.

- The document clarified that HMIPv6 can be used without the need for a home agent.

- Several editorials throughout the document.

- IKEv2 only is now used to allocate the RCoA.

RFC 4140 was implemented and interop tested by at least two different organisations. A test suite including test cases for RFC 4140 was also developed by Ericsson and run against both implementations. No major issues were found. The scalability of Dynamic MAP Discovery, defined in RFC 4140, was seen as inappropriate for large-scale deployments and prone to loops. It was removed from this specification.

At this time, there is no publicly known deployment of this specification.

Authors’ Addresses

Hesham Soliman
Elevate Technologies

EMail: hesham@elevatemobile.com

Claude Castelluccia
INRIA

Phone: +33 4 76 61 52 15
EMail: claude.castelluccia@inria.fr

Karim ElMalki
Athonet

EMail: karim@elmalki.homeip.net

Ludovic Bellier
INRIA

EMail: ludovic.bellier@inria.fr

Full Copyright Statement

Intellectual Property