

Network Working Group
Request for Comments: 4140
Category: Experimental

H. Soliman
Flarion
C. Castelluccia
INRIA
K. El Malki
Ericsson
L. Bellier
INRIA
August 2005

Hierarchical Mobile IPv6 Mobility Management (HMIPv6)

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

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	3
2. Terminology	4
3. Overview of HMIPv6	5
3.1. HMIPv6 Operation	6
4. Mobile IPv6 Extensions	8
4.1. Local Binding Update	8
5. Neighbour Discovery Extension: The MAP Option Message Format	9
6. Protocol Operation	10
6.1. Mobile Node Operation	10
6.1.1. Sending Packets to Correspondent Nodes	12
6.2. MAP Operations	12
6.3. Home Agent Operations	13
6.4. Correspondent Node Operations	13
6.5. Local Mobility Management Optimisation within a MAP Domain	13
6.6. Location Privacy	14
7. MAP Discovery	14
7.1. Dynamic MAP Discovery	14
7.1.1. Router Operation for Dynamic MAP Discovery	15
7.1.2. MAP Operation for Dynamic MAP Discovery	15
7.2. Mobile Node Operation	16
8. Updating Previous MAPs	16
9. Notes on MAP Selection by the Mobile Node	17
9.1. MAP Selection in a Distributed-MAP Environment	17
9.2. MAP Selection in a Flat Mobility Management Architecture	19
10. Detection and Recovery from MAP Failures	19
11. IANA Considerations	20
12. Security Considerations	20
12.1. Mobile Node-MAP Security	20
12.2. Mobile Node-Correspondent Node Security	22
12.3. Mobile Node-Home Agent Security	22
13. Acknowledgments	22
14. References	23
14.1. Normative References	23
14.2. Informative References	23
Appendix A: Fast Mobile IPv6 Handovers and HMIPv6	24

1. Introduction

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

Mobile IPv6 [1] allows nodes to move within the Internet topology while maintaining reachability and on-going connections between mobile and correspondent nodes. To do this a mobile node sends Binding Updates (BUs) to its Home Agent (HA) and all Correspondent Nodes (CNs) it communicates with, every time it moves. 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). In addition, one round-trip time is needed to update the Home Agent; this can be done simultaneously while updating correspondent nodes. The re-use of the home cookie (i.e., eliminating HOTI/HOT) will not reduce the number of round trip times needed to update correspondent nodes. These round trip delays will disrupt active connections every time a handoff to a new AR is performed. Eliminating this additional delay element from the time-critical handover period will significantly 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). Unlike Foreign Agents in IPv4, a MAP is not required on each subnet. 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 HA (which is typically further away) and CNs
- Only one Binding Update message needs to be transmitted by the MN before traffic from the HA and all CNs is re-routed to its new location. This is independent of the number of CNs that the MN is communicating with.

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. It also supports Fast Mobile IPv6 Handovers to help Mobile Nodes achieve seamless mobility (see

Appendix A). Furthermore, HMIPv6 allows mobile nodes to hide their location from correspondent nodes and Home Agents while using Mobile IPv6 route optimisation.

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 RFC 2119 [3].

In addition, new terms are defined below:

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 obtained by the mobile node from the visited network. An RCoA is an address on the MAP's subnet. It is auto-configured by the mobile node when receiving the MAP option.
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 (LCoA)	The LCoA is the on-link CoA configured on a mobile node's interface based on the prefix advertised by its default router. In [1], 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 operation 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 on 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 the correspondent nodes it is communicating 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 Operation

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.

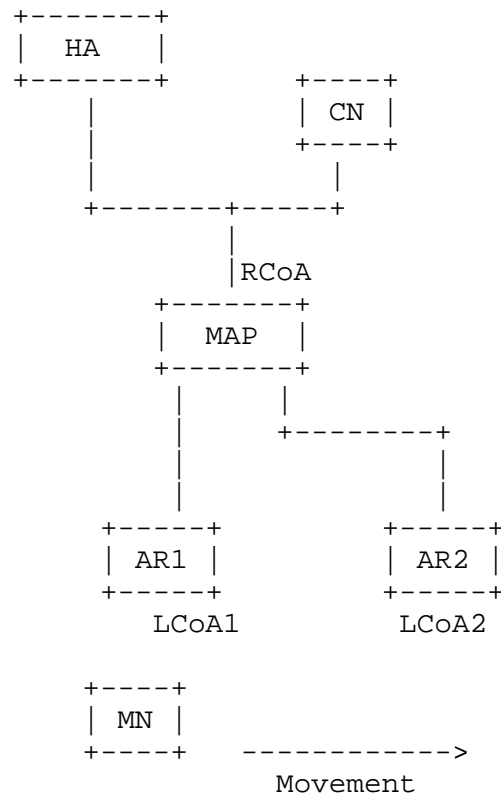


Figure 1: Hierarchical Mobile IPv6 domain

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 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 [1] 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. 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 [1]) 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 Fig 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 [1]. 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.


If a router advertisement is used for MAP discovery, as described in this document, all ARs belonging to the MAP domain MUST advertise the MAP's IP address. The same concept (advertising the MAP's presence within its domain) should be used if other methods of MAP discovery are introduced in future.

4. Mobile IPv6 Extensions

This section outlines the extensions proposed to the binding update specified in [1].

4.1. Local Binding Update

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.




Description of extensions to the binding update:

M If set to 1 it indicates a MAP registration.

It should be noted that this is an extension to the Binding update specified in [1].

5. Neighbour Discovery Extension: The MAP Option Message Format



Fields:

Type	IPv6 Neighbor Discovery option. 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. Its default value SHOULD be set to 1 if Dynamic MAP discovery is used. The Distance MUST be set to 1 if the MAP is on the same link as the mobile node. This field need not be interpreted as the number of hops between MAP and the mobile node. The only requirement is that the meaning of the Distance field is consistently interpreted within one Domain. A Distance value of Zero MUST NOT be used.
Pref	The preference of a MAP. A 4-bit unsigned integer. A decimal value of 15 indicates the highest availability.
R	When set to 1, it indicates that the mobile node MUST form an RCoA based on the prefix in the MAP option.

Valid Lifetime The minimum value (in seconds) of both the preferred and valid lifetimes of the prefix assigned to the MAP's subnet. This value indicates the validity of the MAP's address and consequently the time for which the RCoA is valid.

Global Address One of the MAP's global addresses. The 64-bit prefix extracted from this address MUST be configured in the MAP to be used for RCoA construction by the mobile node.

Although not explicitly included in the MAP option, the prefix length of the MAP's Global IP address MUST be 64. This prefix is the one used by the mobile node to form an RCoA, by appending a 64-bit identifier to the prefix. Thus, it necessitates a static prefix length for the MAP's subnet.

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). The RCoA is formed in a stateless manner. After forming the RCoA based on the prefix received in the MAP option, 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 [1] 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 a HA) will then perform DAD (when a new binding is being created) for the mobile node's RCoA on its link and return a Binding Acknowledgement to the MN. This acknowledgement 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 MUST 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 with its HA. 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, which are connected to its 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 [1]. When communicating through the HA, the message formats in [1] can be re-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 [1], 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 a 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 inform the MAP that the mobile node has formed an RCoA (contained in the BU as a Home address). 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 [1]. 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 acts as a 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 [1].

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 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 [1].

6.4. Correspondent Node Operations

HMIPv6 is completely transparent to correspondent nodes.

6.5. Local Mobility Management Optimisation within a MAP Domain

In [1], 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.

In HMIPv6, a mobile node can use its RCoA as the source address without using a Home Address option. In other words, the RCoA can be used as a potential 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. Although such use of RCoA does not provide global mobility (i.e., communication is broken when a mobile host moves to a new MAP), 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 [1], 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 corresponding nodes and its home agent by using its RCoA in the source field of the packets that it sends. As a result, the location tracking of a mobile node by its corresponding nodes or its home agent is 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. Two different methods for MAP discovery are defined below.

Dynamic MAP Discovery is based on propagating the MAP option in Router Advertisements from the MAP to the mobile node through certain (configured) router interfaces within the routers in an operator's network. This requires manual configuration of the MAP and also that the routers receiving the MAP option allow them to propagate the option on certain interfaces. To ensure a secure communication between routers, router advertisements that are sent between routers for Dynamic MAP discovery SHOULD be authenticated (e.g., using AH, ESP, or SEND). In the case where this authentication is not possible (e.g., third party routers exist between the MAP and ARs), a network operator may prefer to manually configure all the ARs to send the MAP option, as described in this document.

Manual configuration of the MAP option information in ARs and other MAPs in the same domain is the default mechanism. It should also be possible to configure ARs and MAPs to enable dynamic mechanisms for MAP Discovery.

7.1. Dynamic MAP Discovery

The process of MAP discovery can be performed in different ways. Router advertisements are used for Dynamic MAP Discovery by introducing a new option. The access router is required to send the MAP option in its router advertisements. This option includes the distance vector from the mobile node (which may not imply the real distance in terms of the number of hops), the preference for this particular MAP, the MAP's global IP address and subnet prefix

7.1.1. Router Operation for Dynamic MAP Discovery

The ARs within a MAP domain may be configured dynamically with the information related to the MAP options. ARs may obtain this information by listening for RAs with MAP options. Each MAP in the network needs to be configured with a default preference, the right interfaces to send this option on, and the IP address to be sent. The initial value of the "Distance" field MAY be set to a default value of 1 and MUST NOT be set to zero. Routers in the MAP domain should be configured to re-send the MAP option on certain interfaces.

Upon reception of a router advertisement with the MAP option, the receiving router MUST copy the option and re-send it after incrementing the Distance field by one. If the receiving router was also a MAP, it MUST send its own option, together with the received option, in the same advertisement. If a router receives more than one MAP option for the same MAP (i.e., the same IP address in the MAP option), from two different interfaces, it MUST choose the option with the smallest distance field.

In this manner, information about one or more MAPs can be dynamically passed to a mobile node. Furthermore, by performing the discovery phase in this way, different MAP nodes are able to change their preferences dynamically based on the local policies, node overload or other load-sharing protocols being used.

7.1.2. MAP Operation for Dynamic MAP Discovery

A MAP will be configured to send its option or relay MAP options belonging to other MAPs onto certain interfaces. The choice of interfaces is done by the network administrator (i.e., manual configuration) and depends on the network topology. A default preference value of 10 may be assigned to each MAP. It should be noted that a MAP can change its preference value at any time due to various reasons (e.g., node overload or load sharing). A preference value of zero means the MAP SHOULD NOT be chosen by new mobile nodes. This value could be reached in cases of node overload or partial node failures.

The MAP option is propagated towards ARs in its domain. Each router along the path to an AR will increment the Distance field by one. If a router that is also a MAP receives advertisements from other MAPs, it MUST add its own MAP option and propagate both options to the next router or to the AR (if it has direct connectivity with the AR).

7.2. 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 revert to using the Mobile IPv6 protocol, as specified in [1].

If a multihomed 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 no MAP options are found in the router advertisement, the mobile node MUST use the Mobile IPv6 protocol, as specified in [1].

If the R flag is set, the mobile node MUST use its RCoA as the Home Address when performing the MAP registration. RCoA is then 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. Notes 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 a 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. This specification does not provide an algorithm for the distance-based MAP selection. However, such an algorithm may be introduced in future extensions utilising information about the speed of mobility from lower layers.

In a scenario where several MAPs are discovered by the mobile node in one domain, the mobile node may need some 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 uses 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 away 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 Management 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 some form of context transfer protocol between them. Alternatively, future versions of the Virtual Router Redundancy Protocol [4] or HA redundancy protocols may allow networks to recover from MAP failures.

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 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.

ARs can perform Dynamic detection of MAP failure by sending ICMP Echo request messages to the MAP regularly (e.g., every ten seconds). If no response is received, an AR may try to aggressively send echo requests to 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.

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

11. IANA Considerations

Section 4 introduces a new flag (M) to the Binding Update specified in RFC 3775.

Section 5 introduces a new IPv6 Neighbour Discovery Option called the MAP Option. IANA has assigned the Option Type value 23 for the MAP Option within the option numbering space for IPv6 Neighbour Discovery messages.

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, but 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 - Home Agent
- 3) The mobile node - correspondent node

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 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 would trust 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).

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, an LCoA and registers the binding between these two addresses with the new MAP. The MAP then verifies whether the RCoA has not been registered yet and, if so, it 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 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. However, because the RCoA is temporary and is not bound to a particular node, a mobile node does not have to initially (before the first binding update) prove that it owns its RCoA (unlike the requirement on home addresses in Mobile IPv6) when it establishes a Security Association with its MAP. A MAP only needs to ensure that a BU for a particular RCoA was issued by the same mobile node that established the Security Association for that RCoA.

The MAP does not need to have prior knowledge of the identity of the mobile node nor 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 IKE. 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 IKE [2]. 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 IKE Phase 2 negotiation. This step is identical to the use of the home address in IKE phase 2.

If a binding cache entry exists for a given RCoA, the MAP's IKE policy check MUST point to the SA used to install the entry. If the mobile node's credentials stored in the existing SA do not match the ones provided in the current negotiation, the MAP MUST reject the new

SA establishment request for such RCoA with an INVALID-ID-INFORMATION notification [2]. This is to prevent two different mobile nodes from registering (intentionally or not) the same RCoA. Upon receiving this notification, the mobile node SHOULD generate a new RCoA and restart the IKE negotiation. Alternatively, a MAP may decide that, if a binding cache entry already exists for a particular RCoA, no new security association should be established for such RCoA; this is independent of the mobile node credentials. This prevents the mobile node from being able to re-establish a security association for the same RCoA (i.e., to change session keys). However, this is not a major problem because the SA will typically only be used to protect signalling traffic when a MN moves, and not for the actual data traffic sent to arbitrary nodes.

Binding updates between the MAP and the mobile node MUST be protected with either AH or ESP in transport mode. When ESP is used, a non-null authentication algorithm MUST be used.

12.2. Mobile Node-Correspondent Node Security

Mobile IPv6 [1] 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 [1]. 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 [1]. The source address used in HOTI messages MUST be the mobile node's home address. 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 [1], is not impacted by this specification.

13. Acknowledgments

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 (Monash University), Francis Dupont (GET/Enst Bretagne), Gopal Dommety (Cisco), Eva Gustaffson (Ericsson), Dave Johnson (Rice University), Annika Jonsson (Ericsson), James Kempf (Docomo labs), Martti Kuparinen (Ericsson) Fergal Ladley, Gabriel Montenegro (Sun), Nick "Sharkey" Moore (Monash University) Erik Nordmark (Sun), Basavaraj Patil (Nokia), Brett Pentland (Monash University), and Alper Yegin (Samsung) for their comments on the document.

14. References

14.1. Normative References

- [1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [2] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

14.2. Informative References

- [4] Koodli, R., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.
- [5] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [6] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

Appendix A: Fast Mobile IPv6 Handovers and HMIPv6

Fast Handovers are required to ensure that the layer 3 (Mobile IP) handover delay is minimised, thus also minimising, and possibly eliminating, the period of service disruption which normally occurs when a mobile node moves between two ARs. This period of service disruption usually occurs due to the time required by the mobile node to update its HA using Binding Updates after it moves between ARs. During this time period the mobile node cannot resume or continue communications. The mechanism to achieve Fast Handovers with Mobile IPv6 is described in [5] and is briefly summarised here. This mechanism allows the anticipation of the layer 3 handover, such that data traffic can be redirected to the mobile node's new location before it moves there.

While the mobile node is connected to its previous Access Router (PAR) and is about to move to a new Access Router (NAR), the Fast Handovers in Mobile IPv6 requires in sequence:

- 1) The mobile node to obtain a new care-of address at the NAR while connected to the PAR.
- 2) New CoA to be used at NAR case: the mobile node to send a F-BU (Fast BU) to its previous anchor point (i.e., PAR) to update its binding cache with the mobile node's new CoA while still attached to PAR.
- 3) The previous anchor point (i.e., PAR) to start forwarding packets destined for the mobile node to the mobile node's new CoA at NAR (or old CoA tunnelled to NAR, if new CoA is not applicable).
- 4) Old CoA to be used at NAR case: the mobile node to send a F-BU (Fast BU) to its previous anchor point (i.e., PAR), after it has moved and attached to NAR, in order to update its binding cache with the mobile node's new CoA.

The mobile node or PAR may initiate the Fast Handover procedure by using wireless link-layer information or link-layer triggers that inform that the mobile node will soon be handed off between two wireless access points respectively attached to PAR and NAR. If the "trigger" is received at the mobile node, the mobile node will initiate the layer-3 handover process by sending a Proxy Router Solicitation message to PAR. Instead, if the "trigger" is received at PAR, then it will transmit a Proxy Router Advertisement to the appropriate mobile node, without the need for solicitations. The basic Fast Handover message exchanges are illustrated in Figure A.1.




Figure A.1: Fast Mobile IPv6 Handover Protocol

The mobile node obtains a new care-of address while connected to PAR by means of router advertisements containing information from the NAR (Proxy Router Advertisement, which may be sent due to a Proxy Router Solicitation). The PAR will validate the mobile node's new CoA by sending a Handover Initiate (HI) message to the NAR. The new CoA sent in the HI message is formed by appending the mobile node's current interface identifier to the NAR's prefix. Based on the response generated in the Handover Acknowledge (HAcK) message, the PAR will either generate a tunnel to the mobile node's new CoA (if the address was valid) or generate a tunnel to the NAR's address (if the address was already in use on the new subnet). If the address was already in use on the new subnet, it is assumed that there will be no time to perform another attempt to configure the mobile node with a CoA on the new link. Therefore, the NAR will generate a host route for the mobile node using its old CoA. Note that message 1a may precede message 2b or occur at the same time.

In [5], the ARs act as local Home Agents, which hold binding caches for the mobile nodes and receive Binding Updates. This makes these ARs function like the MAP specified in this document. Also, it is quite possible that the ARs are not directly connected, but communicate through an aggregation router. Therefore, such an aggregation router is also an ideal position for the MAP functionality. These are two ways of integrating the HMIPv6 and Fast Handover mechanisms. The first involves placing MAPs in place of the ARs, which is a natural step. The second scenario involves placing the MAP in an aggregation router "above" the ARs. In this case, [5] specifies forwarding of packets between PAR and NAR. This could be inefficient in terms of delay and bandwidth efficiency because packets will traverse the MAP-PAR link twice and packets will arrive out of order at the mobile node. Using the MAP in the aggregation

router would improve the efficiency of Fast Handovers, which could make use of the MAP to redirect traffic, thus saving delay and bandwidth between the aggregation router and the PAR.




Figure A.2: Fast Mobile IPv6 Handover Protocol using HMIPv6

In Figure A.2, the HI/HAcK messages now occur between the MAP and NAR in order to check the validity of the newly requested care-of address and to establish a temporary tunnel should the new care-of address not be valid. Therefore, the same functionality of the Fast Handover procedure is kept, but the anchor point is moved from the PAR to the MAP.

As in the previous Fast Handover procedure, in the network-determined case the layer-2 "triggers" at the PAR will cause the PAR to send a Proxy Router Advertisement to the mobile node with the MAP option. In the mobile-determined case, this is preceded by a Proxy Router Solicitation from the mobile node. The same layer-2 trigger at PAR in the network-determined case could be used to independently initiate Context Transfer (e.g., QoS) between PAR and NAR. In the mobile-determined case, the trigger at PAR could be replaced by the reception of a Proxy Router Solicitation or F-BU. Context Transfer is being worked on in the IETF Seamoby WG.

The combination of Fast Handover and HMIPv6 allows the anticipation of the layer 3 handoff, such that data traffic can be efficiently redirected to the mobile node's new location before it moves there. However, it is not easy to determine the correct time to start forwarding traffic from the MAP to the mobile node's new location, which has an impact on how smooth the handoff will be. The same issues arise in [5] with respect to when to start forwarding between PAR and NAR. Packet loss will occur if this is performed too late or too early with respect to the time in which the mobile node detaches from PAR and attaches to NAR. Such packet loss is likely to occur if the MAP updates its binding cache upon receiving the anticipated F-BU, because it is not known exactly when the mobile node will perform or complete the layer-2 handover to NAR, relative to when the mobile node transmits the F-BU. Also, some measure is needed to support the case in which the mobile node's layer-2 handover unexpectedly fails (after Fast Handover has been initiated) or when the mobile node moves quickly back-and-forth between ARs (ping-pong). Simultaneous bindings [6] provide a solution to these issues. In [6], a new Simultaneous Bindings Flag is added to the Fast Binding Update (F-BU) message and a new Simultaneous Bindings suboption is defined for the Fast Binding Acknowledgement (F-Back) message. Using this enhanced mechanism, upon layer-3 handover, traffic for the mobile node will be sent from the MAP to both PAR and NAR for a certain period, thus isolating the mobile node from layer-2 effects such as handover timing, ping-pong, or handover failure and providing the mobile node with uninterrupted layer-3 connectivity.

Authors' Addresses

Hesham Soliman
Flarion Technologies

EMail: h.soliman@flarion.com

Claude Castelluccia
INRIA Rhone-Alpes
655 avenue de l'Europe
38330 Montbonnot Saint-Martin
France

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

Karim El Malki
Ericsson AB
LM Ericssons Vag. 8
126 25 Stockholm
Sweden

EMail: karim@elmalki.homeip.net

Ludovic Bellier
INRIA Rhone-Alpes
655 avenue de l'Europe
38330 Montbonnot Saint-Martin
France

EMail: ludovic.bellier@inria.fr

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

