

Fast Handovers for Mobile IPv6

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

Mobile IPv6 enables a Mobile Node to maintain its connectivity to the Internet when moving from one Access Router to another, a process referred to as handover. During handover, there is a period during which the Mobile Node is unable to send or receive packets because of link switching delay and IP protocol operations. This "handover latency" resulting from standard Mobile IPv6 procedures, namely movement detection, new Care of Address configuration, and Binding Update, is often unacceptable to real-time traffic such as Voice over IP. Reducing the handover latency could be beneficial to non-real-time, throughput-sensitive applications as well. This document specifies a protocol to improve handover latency due to Mobile IPv6 procedures. This document does not address improving the link switching latency.

Table of Contents

1.	Introduction	3
2.	Terminology.	3
3.	Protocol Overview.	5
3.1.	Addressing the Handover Latency.	5
3.2.	Protocol Operation	7
3.3.	Protocol Operation of Network-initiated Handover	9
4.	Protocol Details	10
5.	Miscellaneous.	15
5.1.	Handover Capability Exchange	15
5.2.	Determining New Care of Address.	15
5.3.	Packet Loss.	15
5.4.	DAD Handling	16
5.5.	Fast or Erroneous Movement	16
6.	Message Formats.	17
6.1.	New Neighborhood Discovery Messages.	17
6.1.1.	Router Solicitation for Proxy Advertisement (RtSolPr)	17
6.1.2.	Proxy Router Advertisement (PrRtAdv).	20
6.2.	Inter-Access Router Messages	23
6.2.1.	Handover Initiate (HI).	23
6.2.2.	Handover Acknowledge (HACK)	25
6.3.	New Mobility Header Messages	27
6.3.1.	Fast Binding Update (FBU)	27
6.3.2.	Fast Binding Acknowledgment (FBack)	28
6.3.3.	Fast Neighbor Advertisement (FNA)	30
6.4.	New Options.	31
6.4.1.	IP Address Option	32
6.4.2.	New Router Prefix Information Option.	33
6.4.3.	Link-Layer Address (LLA) Option	34
6.4.4.	Mobility Header Link-Layer Address (MH-LLA) Option.	35
6.4.5.	Neighbor Advertisement Acknowledgment (NAACK)	35
7.	Configurable Parameters.	36
8.	Security Considerations.	37
9.	IANA Considerations.	38
10.	Acknowledgments.	39
11.	Normative References	39
12.	Contributors	39

1. Introduction

Mobile IPv6 [3] describes the protocol operations for a mobile node to maintain connectivity to the Internet during its handover from one access router to another. These operations involve movement detection, IP address configuration, and location update. The combined handover latency is often sufficient to affect real-time applications. Throughput-sensitive applications can also benefit from reducing this latency. This document describes a protocol to reduce the handover latency.

This specification addresses the following problem: how to allow a mobile node to send packets as soon as it detects a new subnet link, and how to deliver packets to a mobile node as soon as its attachment is detected by the new access router. The protocol defines IP protocol messages necessary for its operation regardless of link technology. It does this without depending on specific link-layer features while allowing link-specific customizations. By definition, this specification considers handovers that interwork with Mobile IP: once attached to its new access router, an MN engages in Mobile IP operations including Return Routability [3]. There are no special requirements for a mobile node to behave differently with respect to its standard Mobile IP operations.

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 [1]. The use of the term, "silently ignore" is not defined in RFC 2119. However, the term is used in this document and can be similarly construed.

The following terminology and abbreviations are used in this document. The reference handover scenario is illustrated in Figure 1.

Mobile Node (MN)
A Mobile IPv6 host.

Access Point (AP)
A Layer 2 device connected to an IP subnet that offers wireless connectivity to an MN. An Access Point Identifier (AP-ID) refers to the AP's L2 address. Sometimes, AP-ID is also referred to as a Base Station Subsystem ID (BSSID).

Access Router (AR)
The MN's default router.

Previous Access Router (PAR)

The MN's default router prior to its handover.

New Access Router (NAR)

The MN's default router subsequent to its handover.

Previous CoA (PCoA)

The MN's Care of Address valid on PAR's subnet.

New CoA (NCoA)

The MN's Care of Address valid on NAR's subnet.

Handover

A process of terminating existing connectivity and obtaining new IP connectivity.

Router Solicitation for Proxy Advertisement (RtSolPr)

A message from the MN to the PAR requesting information for a potential handover.

Proxy Router Advertisement (PrRtAdv)

A message from the PAR to the MN that provides information about neighboring links facilitating expedited movement detection. The message also acts as a trigger for network-initiated handover.

(AP-ID, AR-Info) tuple

Contains an access router's L2 and IP addresses, and the prefix valid on the interface to which the Access Point (identified by AP-ID) is attached. The triplet [Router's L2 address, Router's IP address, Prefix] is called "AR-Info".

Assigned Addressing

A particular type of NCoA configuration in which the NAR assigns an IPv6 address for the MN. The method by which NAR manages its address pool is not specified in this document.

Fast Binding Update (FBU)

A message from the MN instructing its PAR to redirect its traffic (toward NAR).

Fast Binding Acknowledgment (FBack)

A message from the PAR in response to an FBU.

Fast Neighbor Advertisement (FNA)

A message from the MN to the NAR to announce attachment, and to confirm the use of NCoA when the MN has not received an FBACK.

Handover Initiate (HI)

A message from the PAR to the NAR regarding an MN's handover.

Handover Acknowledge (HACK)

A message from the NAR to the PAR as a response to HI.




Figure 1: Reference Scenario for Handover

3. Protocol Overview

3.1. Addressing the Handover Latency

The ability to immediately send packets from a new subnet link depends on the "IP connectivity" latency, which in turn depends on the movement detection latency and new CoA configuration latency. Once an MN is IP-capable on the new subnet link, it can send a Binding Update to its Home Agent and one or more correspondents. Once its correspondents successfully process the Binding Update, which typically involves the Return Routability procedure, the MN can receive packets at the new CoA. So, the ability to receive packets from correspondents directly at its new CoA depends on the Binding Update latency as well as the IP connectivity latency.

The protocol enables an MN to quickly detect that it has moved to a new subnet by providing the new access point and the associated subnet prefix information when the MN is still connected to its current subnet (i.e., PAR in Figure 1). For instance, an MN may discover available access points using link-layer specific mechanisms (i.e., a "scan" in WLAN) and then request subnet information corresponding to one or more of those discovered access points. The MN may do this after performing router discovery or at any time while connected to its current router. The result of resolving an identifier associated with an access point is a [AP-ID, AR-Info] tuple, which an MN can use in readily detecting movement: when attachment to an access point with AP-ID takes place, the MN knows the corresponding new router's coordinates including its prefix, IP address, and L2 address. The "Router Solicitation for Proxy Advertisement (RtSolPr)" and "Proxy Router Advertisement (PrRtAdv)" messages (see Section 6.1) are used for aiding movement detection.

Through the RtSolPr and PrRtAdv messages, the MN also formulates a prospective new CoA (NCoA) when it is still present on the PAR's link. Hence, the latency due to new prefix discovery subsequent to handover is eliminated. Furthermore, this prospective address can be used immediately after attaching to the new subnet link (i.e., NAR's link) when the MN has received a "Fast Binding Acknowledgment (FBack)" message prior to its movement. If it moves without receiving an FBack, the MN can still start using NCoA after announcing its attachment through a "Fast Neighbor Advertisement (FNA)" message. NAR responds to FNA if the tentative address is already in use thereby reducing NCoA configuration latency. Under some limited conditions in which the probability of address collision is considered insignificant, it may be possible to use NCoA immediately after attaching to the new link. Even so, all implementations MUST support and SHOULD use the mechanism specified in this document to avoid potential address conflicts.

To reduce the Binding Update latency, the protocol specifies a tunnel between the Previous CoA (PCoA) and the NCoA. An MN sends a "Fast Binding Update" message to its Previous Access Router to establish this tunnel. When feasible, the MN SHOULD send an FBU from PAR's link. Otherwise, it should be sent immediately after attachment to NAR has been detected. Subsequent sections describe the protocol mechanics. As a result, PAR begins tunneling packets arriving for PCoA to NCoA. Such a tunnel remains active until the MN completes the Binding Update with its correspondents. In the opposite direction, the MN SHOULD reverse tunnel packets to PAR until it completes the Binding Update. PAR SHOULD forward the inner packet in the tunnel to its destination (i.e., to the MN's correspondent). Such a reverse tunnel ensures that packets containing PCoA as a source IP address are not dropped due to ingress filtering. Readers

may observe that even though the MN is IP-capable on the new link, it cannot use NCoA directly with its correspondents without the correspondents first establishing a binding cache entry (for NCoA). Forwarding support for PCoA is provided through a reverse tunnel between the MN and the PAR.

Setting up a tunnel alone does not ensure that the MN receives packets as soon as it is attached to a new subnet link, unless the NAR can detect the MN's presence. A neighbor discovery operation involving a neighbor's address resolution (i.e., Neighbor Solicitation and Neighbor Advertisement) typically results in considerable delay, sometimes lasting multiple seconds. For instance, when arriving packets trigger NAR to send Neighbor Solicitation before the MN attaches, subsequent retransmissions of address resolution are separated by a default period of one second each. To circumvent this delay, an MN announces its attachment through the FNA message that allows the NAR to consider MN to be reachable. If there is no existing entry, FNA allows NAR to create one. If NAR already has an entry, FNA updates the entry while taking potential address conflicts into consideration. Through tunnel establishment for PCoA and fast advertisement, the protocol provides expedited forwarding of packets to the MN.

The protocol also provides the following important functionalities. The access routers can exchange messages to confirm that a proposed NCoA is acceptable. For instance, when an MN sends an FBU from PAR's link, FBack can be delivered after the NAR considers the NCoA acceptable for use. This is especially useful when addresses are assigned by the access router. The NAR can also rely on its trust relationship with PAR before providing forwarding support for the MN. That is, it may create a forwarding entry for the NCoA subject to "approval" from PAR which it trusts. Finally, the access routers could transfer network-resident contexts, such as access control, QoS, and header compression, in conjunction with handover. For these operations, the protocol provides "Handover Initiate (HI)" and "Handover Acknowledge (HACK)" messages. Both of these messages MUST be supported and SHOULD be used. The access routers MUST have necessary security association established by means outside the scope of this document.

3.2. Protocol Operation

The protocol begins when an MN sends an RtSolPr to its access router to resolve one or more Access Point Identifiers to subnet-specific information. In response, the access router (e.g., PAR in Figure 1) sends a PrRtAdv message containing one or more [AP-ID, AR-Info] tuples. The MN may send a RtSolPr at any convenient time, for instance as a response to some link-specific event (a "trigger") or

simply after performing router discovery. However, the expectation is that prior to sending RtSolPr, the MN will have discovered the available APs by link-specific methods. The RtSolPr and PrRtAdv messages do not establish any state at the access router; their packet formats are defined in Section 6.1.

With the information provided in the PrRtAdv message, the MN formulates a prospective NCoA and sends an FBU message when a link-specific handover event occurs. The purpose of the FBU is to authorize PAR to bind PCoA to NCoA, so that arriving packets can be tunneled to the new location of the MN. Whenever feasible, the FBU SHOULD be sent from PAR's link. For instance, an internal link-specific trigger could enable FBU transmission from the previous link. When it is not feasible, the FBU is sent from the new link. Care must be taken to ensure that the NCoA used in FBU does not conflict with an address already in use by some other node on the link. For this, FBU encapsulation within FNA MUST be implemented and SHOULD be used (see below) when the FBU is sent from NAR's link.

The format and semantics of FBU processing are specified in Section 6.3.1.

Depending on whether an FBack is received on the previous link (which clearly depends on whether the FBU was sent in the first place), there are two modes of operation.

1. The MN receives an FBack on the previous link. This means that packet tunneling is already in progress by the time the MN handovers to NAR. The MN SHOULD send FNA immediately after attaching to NAR, so that arriving and buffered packets can be forwarded to the MN right away.

Before sending an FBack to an MN, PAR can determine whether the NCoA is acceptable to the NAR through the exchange of HI and HAck messages. When assigned addressing (i.e., addresses are assigned by the router) is used, the proposed NCoA in the FBU is carried in HI, and the NAR MAY assign the proposed NCoA. Such an assigned NCoA MUST be returned in HAck, and the PAR MUST in turn provide the assigned NCoA in the FBack. If there is an assigned NCoA returned in the FBack, the MN MUST use the assigned address (and not the proposed address in the FBU) upon attaching to NAR.

2. The MN does not receive the FBack on the previous link because the MN has not sent the FBU or the MN has left the link after sending the FBU (which itself may be lost), but before receiving an FBack. Without receiving an FBack in the latter case, the MN cannot ascertain whether PAR has successfully processed the FBU. Hence, it (re)sends an FBU as soon as it attaches to NAR. To enable NAR

to forward packets immediately (when FBU has been processed) and to allow NAR to verify whether NCoA is acceptable, the MN SHOULD encapsulate the FBU in the FNA. If NAR detects that NCoA is in use when processing the FNA, for instance while creating a neighbor entry, it MUST discard the inner FBU packet and send a Router Advertisement with the "Neighbor Advertisement Acknowledge (NAACK)" option in which NAR MAY include an alternate IP address for the MN to use. This discarding avoids the rare and undesirable outcome that results from address collision. Detailed FNA processing rules are specified in Section 6.3.3.

The scenario in which an MN sends an FBU and receives an FBack on PAR's link is illustrated in Figure 2. For convenience, this scenario is characterized as "predictive" mode of operation. The scenario in which the MN sends an FBU from NAR's link is illustrated in Figure 3. For convenience, this scenario is characterized as a "reactive" mode of operation. Note that the reactive mode also includes the case in which an FBU has been sent from PAR's link but an FBack has not been received yet.

Finally, the PrRtAdv message may be sent unsolicited (i.e., without the MN first sending a RtSolPr). This mode is described in Section 3.3.

3.3. Protocol Operation of Network-initiated Handover

In some wireless technologies, the handover control may reside in the network even though the decision to undergo handover may be mutually arrived at between the MN and the network. In these networks, the PAR can send an unsolicited PrRtAdv containing the link layer address, IP address, and subnet prefixes of the NAR when the network decides that a handover is imminent. The MN MUST process this PrRtAdv to configure a new care of address on the new subnet, and MUST send an FBU to PAR prior to switching to the new link. After transmitting PrRtAdv, the PAR MUST continue to forward packets to the MN on its current link until the FBU is received. The rest of the operation is the same as that described in Section 3.2.

The unsolicited PrRtAdv also allows the network to inform the MN about geographically adjacent subnets without the MN having to explicitly request that information. This can reduce the amount of wireless traffic required for the MN to obtain a neighborhood topology map of links and subnets. Such usage of PrRtAdv is decoupled from the actual handover; see Section 6.1.2.




Figure 2: "Predictive" Fast Handover

4. Protocol Details

All descriptions refer to Figure 1.

After discovering one or more nearby access points, the MN sends **RtSolPr** to resolve access point identifiers to subnet router information. This is convenient to do after performing router discovery. However, the MN can send **RtSolPr** at any time, e.g., when one or more new access points are discovered. The MN can also send **RtSolPr** more than once during its attachment to PAR. The trigger for sending **RtSolPr** can originate from a link-specific event, such as the promise of a better signal strength from another access point coupled with fading signal quality with the current access point. Such events, often broadly referred to as "L2 triggers", are outside the scope of this document. Nevertheless, they serve as events that invoke this protocol. For instance, when a "link up" indication is obtained on the new link, protocol messages (e.g., **FNA**) can be immediately transmitted. Implementations **SHOULD** make use of such triggers whenever possible.




Figure 3: "Reactive" Fast Handover

The RtSolPr message contains one or more AP-IDs. A wildcard requests all available tuples.

As a response to RtSolPr, PAR sends a PrRtAdv message that indicates one of the following possible conditions.

1. If the PAR does not have an entry corresponding to the new access point, it MUST respond indicating that the new access point is unknown. The MN MUST stop fast handover protocol operations on the current link. The MN MAY send an FBU from its new link.
2. If the new access point is connected to the PAR's current interface (to which MN is attached), the PAR MUST respond with a Code value indicating that the new access point is connected to the current interface, but not send any prefix information. This scenario could arise, for example, when several wireless access points are bridged into a wired network. No further protocol action is necessary.
3. If the new access point is known and the PAR has information about it, then PAR MUST respond indicating that the new access point is known and supply the [AP-ID, AR-Info] tuple. If the new access point is known, but does not support fast handover, the PAR MUST indicate this with Code 3 (See Section 6.1.2).

4. If a wildcard is supplied as an identifier for the new access point, the PAR SHOULD supply neighborhood [AP-ID, AR-Info] tuples that are subject to path MTU restrictions (i.e., provide any 'n' tuples without exceeding the link MTU).

When further protocol action is necessary, some implementations MAY choose to begin buffering copies of incoming packets at the PAR. If such FIFO buffering is used, the PAR MUST continue forwarding the packets to PCoA (i.e., buffer and forward). Such buffering can be useful when the MN leaves without sending the FBU message from the PAR's link. The PAR SHOULD stop buffering after processing the FBU message. The size of the buffer is an implementation-specific consideration.

The method by which Access Routers exchange information about their neighbors, and thereby allow construction of Proxy Router Advertisements with information about neighboring subnets is outside the scope of this document.

The RtSolPr and PrRtAdv messages MUST be implemented by an MN and an access router that supports fast handovers. However, when the parameters necessary for the MN to send packets immediately upon attaching to the NAR are supplied by the link layer handover mechanism itself, use of above messages is optional on such links.

After a PrRtAdv message is processed, the MN sends an FBU at a time determined by link-specific events, and includes the proposed NCoA. The MN SHOULD send the FBU from PAR's link whenever "anticipation" of handover is feasible. When anticipation is not feasible or when it has not received an FBack, the MN sends an FBU immediately after attaching to NAR's link. This FBU SHOULD be encapsulated in an FNA message. The encapsulation allows the NAR to discard the (inner) FBU packet if an address conflict is detected as a result of (outer) FNA packet processing (see FNA processing below). In response to the FBU, the PAR establishes a binding between PCoA ("Home Address") and NCoA, and sends the FBack to the MN. Prior to establishing this binding, PAR SHOULD send an HI message to NAR, and receive HAck in response. To determine the NAR's address for the HI message, the PAR can perform the longest prefix match of NCoA (in FBU) with the prefix list of neighboring access routers. When the source IP address of the FBU is PCoA, i.e., the FBU is sent from the PAR's link, and the HI message MUST have a Code value set to 0; see Section 6.2.1. When the source IP address of the FBU is not PCoA, i.e., the FBU is sent from the NAR's link, the HI message MUST have a Code value of 1; see Section 6.2.1.

The HI message contains the PCoA, Link-Layer Address, and the NCoA of the MN. In response to processing an HI message with Code 0, the NAR

1. determines whether NCoA supplied in the HI message is a valid address for use. If it is, the NAR starts proxying [6] the address for PROXY_ND_LIFETIME during which the MN is expected to connect to the NAR. The NAR MAY use the Link-Layer Address to verify whether a corresponding IP address exists in its forwarding tables.
2. allocates NCoA for the MN when assigned addressing is used, creates a proxy neighbor cache entry, and begins defending it. The NAR MAY allocate the NCoA proposed in HI.
3. MAY create a host route entry for PCoA in case NCoA cannot be accepted or assigned. This host route entry SHOULD be implemented such that until the MN's presence is detected, either through explicit announcement by the MN or by other means, arriving packets do not invoke neighbor discovery. The NAR MAY also set up a reverse tunnel to the PAR in this case.
4. provides the status of the handover request in the Handover Acknowledge (HACK) message.

When the Code value in HI is 1, NAR MUST skip the above operations since it would have performed those operations during FNA processing. However, it SHOULD be prepared to process any other options that may be defined in the future. Sending an HI message with Code 1 allows NAR to validate the neighbor cache entry it creates for the MN during FNA processing. That is, NAR can make use of the knowledge that its trusted peer (i.e., PAR) has a trust relationship with the MN.

If HACK contains an assigned NCoA, the FBack MUST include it, and the MN MUST use the address provided in the FBack. The PAR MAY send the FBack to the previous link to facilitate faster reception in the event that the MN is still present. The result of the FBU and FBack processing is that PAR begins tunneling the MN's packets to NCoA. If the MN does not receive an FBack message even after retransmitting the FBU for FBU_RETRIES, it must assume that fast handover support is not available and stop the protocol operation.

When the MN establishes link connectivity with the NAR, it SHOULD send a Fast Neighbor Advertisement (FNA) message (see 6.3.3). If the MN has not received an FBack by the time the FNA is being sent, it SHOULD encapsulate the FBU in the FNA and send them together.

When the NCoA corresponding to the FNA message is acceptable, the NAR MUST

1. delete its proxy neighbor cache entry, if any is present.
2. create a neighbor cache entry and set its state to REACHABLE without overwriting an existing entry for a different layer 2 address.
3. forward any buffered packets.
4. enable the host route entry for PCoA, if any is present.

When the NCoA corresponding to the FNA message is not acceptable, the NAR MUST

1. discard the inner (FBU) packet.
2. send a Router Advertisement with the NAACK option in which it MAY include an alternate NCoA for use. This message MUST be sent to the source IP address present in the FNA using the same Layer 2 address present in the FNA.

If the MN receives a Router Advertisement with a NAACK option, it MUST use the IP address, if any, provided in the NAACK option. Otherwise, the MN should configure another NCoA. Subsequently, the MN SHOULD send an FBU using the new CoA. As a special case, the address supplied in NAACK could be PCoA itself, in which case the MN MUST NOT send any more FBUs.

Once the MN has confirmed its NCoA, it SHOULD send a Neighbor Advertisement message. This message allows MN's neighbors to update their neighbor cache entries with the MN's addresses.

Just as in Mobile IPv6, the PAR sets the 'R' bit in the Prefix Information option, and includes its 128 bit global address in the router advertisements. This allows the mobile nodes to learn the PAR's global IPv6 address. The MN reverse tunnels its packets to the same global address of PAR. The tunnel end-point addresses must be configured accordingly. When PAR receives a reverse tunneled packet, it must verify if a secure binding exists for the MN identified by PCoA in the tunneled packet, before forwarding the packet.

5. Miscellaneous

5.1. Handover Capability Exchange

The MN expects a PrRtAdv in response to its RtSolPr message. If the MN does not receive a PrRtAdv message even after RTSOLPR_RETRIES, it must assume that PAR does not support the fast handover protocol and stop sending RtSolPr messages.

Even if an MN's current access router is capable of fast handover, the new access router to which the MN attaches may be incapable of fast handover. This is indicated to the MN during "runtime", through the PrRtAdv message with a Code value of 3 (see Section 6.1.2).

5.2. Determining New Care of Address

Typically, the MN formulates its prospective NCoA using the information provided in a PrRtAdv message and sends the FBU. The PAR MUST use the NCoA present in the FBU in its HI message. The NAR MUST verify if the NCoA present in HI is already in use. In any case, NAR MUST respond to HI using a HAcK, in which it may include another NCoA to use, especially when assigned address configuration is used. If there is a CoA present in HAcK, the PAR MUST include it in the FBack message.

If a PrRtAdv message carries an NCoA, the MN MUST use it as its prospective NCoA.

5.3. Packet Loss

Handover involves link switching, which may not be exactly coordinated with fast handover signaling. Furthermore, the arrival pattern of packets is dependent on many factors, including application characteristics, network queuing behaviors, etc. Hence, packets may arrive at the NAR before the MN is able to establish its link there. These packets will be lost unless they are buffered by the NAR. Similarly, if the MN attaches to the NAR and then sends an FBU message, packets arriving at the PAR will be lost unless they are buffered. This protocol provides an option to indicate a request for buffering at the NAR in the HI message. When the PAR requests this feature (for the MN), it SHOULD also provide its own support for buffering.

5.4. DAD Handling

Duplicate Address Detection (DAD) was defined in [7] to avoid address duplication on links when stateless address auto-configuration is used. The use of DAD to verify the uniqueness of an IPv6 address configured through stateless auto-configuration adds delays to a handover.

The probability of an interface identifier duplication on the same subnet is very low, however it cannot be ignored. In this document, certain precautions are proposed to minimize the effects of a duplicate address occurrence.

In some cases, the NAR may already have the knowledge required to assess whether the MN's address is a duplicate before the MN moves to the new subnet. For example, the NAR can have a list of all nodes on its subnet, perhaps for access control, and by searching this list, it can confirm whether the MN's address is a duplicate. The result of this search is sent back to the PAR in the HAck message. If such knowledge is not available at the NAR, it may indicate this by not confirming the NCoA in the HAck message. The NAR may also indicate this in the NAACK option in response to the FNA message. In such cases, the MN would have to follow the address configuration procedure according to [6] after attaching to the NAR.

5.5. Fast or Erroneous Movement

Although this specification is for fast handover, the protocol is limited in terms of how fast an MN can move. Ping-Pong is a special case of fast movement, where an MN moves between the same two access points rapidly. Another instance of the same problem is erroneous movement, i.e., the MN receives information prior to a handover that it is moving to a new access point, but it is either moved to a different one or it aborts movement altogether. All of the above behaviors are usually the result of link layer idiosyncrasies and thus are often resolved at the link layer itself.

IP layer mobility, however, introduces its own limits. IP layer handovers should occur at a rate suitable for the MN to update the binding of, at least, its HA and preferably that of every CN with which it is in communication. An MN that moves faster than necessary for this signaling to complete, which may be a few seconds, may start losing packets. The signaling cost over the air interface and in the network may increase significantly, especially in the case of rapid movement between several access routers. To avoid the signaling overhead, the following measures are suggested.

An MN returning to the PAR before updating the necessary bindings when present on the NAR MUST send a Fast Binding Update with the Home Address equal to the MN's PCoA and a lifetime of zero to the PAR. The MN should have a security association with the PAR since it performed a fast handover to the NAR. The PAR, upon receiving this Fast Binding Update, will check its set of outgoing (temporary fast handover) tunnels. If it finds a match, it SHOULD tear down that tunnel (i.e., stop forwarding packets for this MN and start delivering packets directly to the node instead). The MN SHOULD NOT attempt to use any of the fast handover mechanisms described in this specification and SHOULD revert back to standard Mobile IPv6.

Temporary tunnels for the purpose of fast handovers should use short lifetimes (a small number of seconds or less). The lifetime of such tunnels should be enough to allow an MN to update all its active bindings. The default lifetime of the tunnel should be the same as the lifetime value in the FBU message.

The effect of erroneous movement is typically limited to the loss of packets since routing can change and the PAR may forward packets toward another router before the MN actually connects to that router. If the MN discovers itself on an unanticipated access router, a Fast Binding Update to the PAR SHOULD be sent. Since Fast Binding Updates are authenticated, they supercede the existing binding and packets MUST be redirected to the newly confirmed location of the MN.

6. Message Formats

All the ICMPv6 messages have a common Type specified in [4]. The messages are distinguished based on the Subtype field (see below). The values for the Subtypes are specified in Section 9. For all the ICMPv6 messages, the checksum is defined in [2].

6.1. New Neighborhood Discovery Messages

6.1.1. Router Solicitation for Proxy Advertisement (RtSolPr)

Mobile Nodes send Router Solicitation for Proxy Advertisement in order to prompt routers for Proxy Router Advertisements. All the Link-Layer Address options have the format defined in 6.4.3.




Figure 4: Router Solicitation for Proxy (RtSolPr) Message

IP Fields:

Source Address

An IP address assigned to the sending interface.

Destination Address

The address of the Access Router or the all routers multicast address.

Hop Limit 255. See RFC 2461.

Authentication Header

If a Security Association for the IP Authentication Header exists between the sender and the destination address, then the sender SHOULD include this header. See RFC 2402 [5].

ICMP Fields:

Type The Experimental Mobility Protocol Type. See [4].

Code 0

Checksum The ICMPv6 checksum.

Subtype 2

Reserved MUST be set to zero by the sender and ignored by the receiver.

Identifier MUST be set by the sender so that replies can be matched to this Solicitation.

Valid Options:

Source Link-Layer Address

When known, the Link-Layer Address of the sender SHOULD be included using the Link-Layer Address option. See the LLA option format below.

New Access Point Link-Layer Address

The Link-Layer Address or identification of the access point for which the MN requests routing advertisement information. It MUST be included in all RtSolPr messages. More than one such address or identifier can be present. This field can also be a wildcard address with all bits set to zero.

Future versions of this protocol may define new option types. Receivers MUST silently ignore any options that they do not recognize and continue processing the rest of the message.

Including the source LLA option allows the receiver to record the sender's L2 address so that neighbor discovery can be avoided when the receiver needs to send packets back to the sender (of the RtSolPr message).

When a wildcard is used for a New Access Point LLA, no other New Access Point LLA options must be present.

A Proxy Router Advertisement (PrRtAdv) message should be received by the MN in response to a RtSolPr. If such a message is not received in a timely manner (no less than twice the typical round trip time (RTT) over the access link or 100 milliseconds if RTT is not known), it SHOULD resend the RtSolPr message. Subsequent retransmissions can be up to RTSOLPR_RETRIES, but MUST use an exponential backoff in which the timeout period (i.e., 2xRTT or 100 milliseconds) is doubled prior to each instance of retransmission. If Proxy Router Advertisement is not received by the time the MN disconnects from the PAR, the MN SHOULD send an FBU immediately after configuring a new CoA.

When RtSolPr messages are sent more than once, they MUST be rate limited with MAX_RTSOLPR_RATE per second. During each use of a RtSolPr, exponential backoff is used for retransmissions.

6.1.2. Proxy Router Advertisement (PrRtAdv)

Access routers send Proxy Router Advertisement messages gratuitously if the handover is network-initiated or as a response to a RtSolPr message from an MN, providing the Link-Layer Address, IP address, and subnet prefixes of neighboring routers. All the Link-Layer Address options have the format defined in Section 6.4.3.

IP Fields:

Source Address

MUST be the Link-Local Address assigned to the interface from which this message is sent.

Destination Address

The Source Address of an invoking Router Solicitation for a Proxy Advertisement or the address of the node the Access Router is instructing to handover.




Figure 5: Proxy Router Advertisement (PrRtAdv) Message

Hop Limit 255. See RFC 2461 [6].

Authentication Header

If a Security Association for the IP Authentication Header exists between the sender and the destination address, the sender SHOULD include this header. See RFC 2402 [5].

ICMP Fields:

Type The Experimental Mobility Protocol Type. See RFC 4065 [4].

Code 0, 1, 2, 3 or 4. See below.

Checksum	The ICMPv6 checksum.
Subtype	3
Reserved	MUST be set to zero by the sender and ignored by the receiver.
Identifier	Copied from the Router Solicitation for Proxy Advertisement or set to Zero if unsolicited.

Valid Options in the following order:

Source Link-Layer Address

When known, the Link-Layer Address of the sender SHOULD be included using the Link-Layer Address option. See the LLA option format below.

New Access Point Link-Layer Address

The Link-Layer Address or identification of the access point is copied from the RtSolPr message. This option MUST be present.

New Router's Link-Layer Address

The Link-Layer Address of the Access Router for which this message is proxied. This option MUST be included when Code is 0 or 1.

New Router's IP Address

The IP address of NAR. This option MUST be included when Code is 0 or 1.

New Router Prefix Information Option.

Specifies the prefix of the Access Router for which the message is proxied and is used for address auto-configuration. This option MUST be included when Code is 0 or 1. However, when this prefix is the same as that used in the New Router's IP Address option (above), the Prefix Information option need not be present.

New CoA Option

MAY be present when a PrRtAdv is sent unsolicited. PAR MAY compute a new CoA using NAR's prefix information and the MN's L2 address, or by any other means.

Future versions of this protocol may define new option types. Receivers MUST silently ignore any options they do not recognize and continue processing the message.

Currently, Code values 0, 1, 2, 3 and 4 are defined.

A Proxy Router Advertisement with Code 0 means that the MN should use the [AP-ID, AR-Info] tuple (present in the options above) for movement detection and NCoA formulation. In this case, the Option-Code field in the New Access Point LLA option is 1, reflecting the LLA of the access point for which the rest of the options are related. Multiple tuples may be present.

A Proxy Router Advertisement with Code 1 means that the message is sent unsolicited. If a New CoA option is present following the New Router Prefix Information option, the MN SHOULD use the supplied NCoA and send the FBU immediately or else stand to lose service. This message acts as a network-initiated handover trigger; see Section 3.3. The Option-Code field in the New Access Point LLA option (see below) in this case is 1 reflecting the LLA of the access point for which the rest of the options are related.

A Proxy Router Advertisement with Code 2 means that no new router information is present. Each New Access Point LLA option contains an Option-Code value (described below) that indicates a specific outcome.

- When the Option-Code field in the New Access Point LLA option is 5, handover to that access point does not require a change of CoA. No other options are required in this case.
- When the Option-Code field in the New Access Point LLA option is 6, the PAR is not aware of the Prefix Information requested. The MN SHOULD attempt to send an FBU as soon as it regains connectivity with the NAR. No other options are required in this case.
- When the Option-Code field in the New Access Point LLA option is 7, it means that the NAR does not support fast handover. The MN MUST stop fast handover protocol operations. No other options are required in this case.

A Proxy Router Advertisement with Code 3 means that new router information is only present for a subset of access points requested. The Option-Code field values (defined above including a value of 1) distinguish different outcomes for individual access points.

A Proxy Router Advertisement with Code 4 means that the subnet information regarding neighboring access points is sent unsolicited, but the message is not a handover trigger, unlike when the message is sent with Code 1. Multiple tuples may be present.

When a wildcard AP identifier is supplied in the RtSolPr message, the PrRtAdv message should include any 'n' [Access Point Identifier, Link-Layer Address option, Prefix Information Option] tuples corresponding to the PAR's neighborhood.

6.2. Inter-Access Router Messages

6.2.1. Handover Initiate (HI)

The Handover Initiate (HI) is an ICMPv6 message sent by an Access Router (typically PAR) to another Access Router (typically NAR) to initiate the process of a MN's handover.




Figure 6: Handover Initiate (HI) Message

IP Fields:

Source Address

The IP address of the PAR.

Destination Address

The IP address of the NAR.

Hop Limit

255. See RFC 2461 [6].

Authentication Header

The authentication header MUST be used when this message is sent. See RFC 2402 [5].

ICMP Fields:

Type	The Experimental Mobility Protocol Type. See RFC 4065 [4].
Code	0 or 1. See below
Checksum	The ICMPv6 checksum.
Subtype	4
S flag	Assigned address configuration flag. When set, this message requests a new CoA to be returned by the destination. May be set when Code = 0. MUST be 0 when Code = 1.
U flag	Buffer flag. When set, the destination SHOULD buffer any packets moving toward the node indicated in the options of this message. Used when Code = 0, SHOULD be set to 0 when Code = 1.
Reserved	MUST be set to zero by the sender and ignored by the receiver.
Identifier	MUST be set by the sender so replies can be matched to this message.

Valid Options:

Link-Layer Address of MN

The Link-Layer Address of the MN that is undergoing handover to the destination (i.e., NAR). This option MUST be included so that the destination can recognize the MN.

Previous Care of Address

The IP address used by the MN while attached to the originating router. This option SHOULD be included so that a host route can be established if necessary.

New Care of Address

The IP address the MN wishes to use when connected to the destination. When the 'S' bit is set, the NAR MAY assign this address.

The PAR uses a Code value of 0 when it processes an FBU with PCoA as a source IP address. The PAR uses a Code value of 1 when it processes an FBU whose source IP address is not PCoA.

If a Handover Acknowledge (HACK) message is not received as a response in a short time period (no less than twice the typical RTT between source and destination, or 100 milliseconds if RTT is not known), the Handover Initiate SHOULD be resent. Subsequent retransmissions can be up to HI_RETRIES, but MUST use exponential backoff in which the timeout period (i.e., 2xRTT or 100 milliseconds) is doubled during each instance of retransmission.

6.2.2. Handover Acknowledge (HACK)

The Handover Acknowledgment message is a new ICMPv6 message that MUST be sent (typically by NAR to PAR) as a reply to the Handover Initiate message.




Figure 7: Handover Acknowledge (HACK) Message

IP Fields:

Source Address

Copied from the destination address of the Handover Initiate Message to which this message is a response.

Destination Address

Copied from the source address of the Handover Initiate Message to which this message is a response.

Hop Limit 255. See RFC 2461 [6].

Authentication Header

The authentication header MUST be used when this message is sent. See RFC 2402 [5].

ICMP Fields:

Type	The Experimental Mobility Protocol Type. See RFC 4065 [4].
Code	0: Handover Accepted, NCoA valid 1: Handover Accepted, NCoA not valid 2: Handover Accepted, NCoA in use 3: Handover Accepted, NCoA assigned (used in Assigned addressing) 4: Handover Accepted, NCoA not assigned (used in Assigned addressing) 128: Handover Not Accepted, reason unspecified 129: Administratively prohibited 130: Insufficient resources
Checksum	The ICMPv6 checksum.
Subtype	5
Reserved	MUST be set to zero by the sender and ignored by the receiver.
Identifier	Copied from the corresponding field in the Handover Initiate message to which this message is a response.

Valid Options:

New Care of Address

If the S flag in the Handover Initiate message is set, this option MUST be used to provide NCoA the MN should use when connected to this router. This option MAY be included, even when the 'S' bit is not set, e.g., Code 2 above.

Upon receiving an HI message, the NAR MUST respond with a Handover Acknowledge message. If the 'S' flag is set in the HI message, the NAR SHOULD include the New Care of Address option and a Code 3.

The NAR MAY provide support for PCoA (instead of accepting or assigning NCoA), establish a host route entry for PCoA, and set up a tunnel to the PAR to forward MN's packets sent with PCoA as a source IP address. This host route entry SHOULD be used to forward packets once the NAR detects that the particular MN is attached to its link.

When responding to an HI message containing a Code value 1, the Code values 1, 2, and 4 in the HAck message are not relevant.

Finally, the new access router can always refuse handover, in which case it should indicate the reason in one of the available Code values.

6.3. New Mobility Header Messages

Mobile IPv6 uses a new IPv6 header type called Mobility Header [3]. The Fast Binding Update, Fast Binding Acknowledgment, and Fast Neighbor Advertisement messages use the Mobility Header.

6.3.1. Fast Binding Update (FBU)

The Fast Binding Update message is identical to the Mobile IPv6 Binding Update (BU) message. However, the processing rules are slightly different.




Figure 8: Fast Binding Update (FBU) Message

IP fields:

Source Address

The PCoA or NCoA

Destination Address

The IP address of the Previous Access Router

A flag

MUST be set to one to request that PAR send a Fast Binding Acknowledgment message.

H flag

MUST be set to one. See RFC 3775 [3].

L flag

See RFC 3775 [3].

K flag See RFC 3775 [3].

Reserved This field is unused. MUST be set zero.

Sequence Number See RFC 3775 [3].

Lifetime See RFC 3775 [3].

Mobility Options MUST contain an alternate CoA option set to the NCoA when an FBU is sent from PAR's link.

The MN sends an FBU message any time after receiving a PrRtAdv message. If the MN moves prior to receiving a PrRtAdv message, it SHOULD send an FBU to the PAR after configuring NCoA on the NAR according to Neighbor Discovery and IPv6 Address Configuration protocols.

The source IP address is PCoA when the FBU is sent from PAR's link, and the source IP address is NCoA when sent from NAR's link. When the FBU is sent from NAR's link, it SHOULD be encapsulated within an FNA.

The FBU MUST also include the Home Address Option, and the Home Address is PCoA. An FBU message MUST be protected so that PAR is able to determine that the FBU message is sent by a genuine MN.

6.3.2. Fast Binding Acknowledgment (FBack)

The Fast Binding Acknowledgment message is sent by the PAR to acknowledge receipt of a Fast Binding Update message in which the 'A' bit is set. The Fast Binding Acknowledgment message SHOULD NOT be sent to the MN before the PAR receives a HAcK message from the NAR. The Fast Binding Acknowledgment MAY also be sent to the MN on the old link.




Figure 9: Fast Binding Acknowledgment (FBack) Message

IP fields:

Source Address

The IP address of the Previous Access Router.

Destination Address

The NCoA

Status

8-bit unsigned integer indicating the disposition of the Fast Binding Update. Values of the Status field that are less than 128 indicate that the Binding Update was accepted by the receiving node. The following such Status values are currently defined:

0 Fast Binding Update accepted

1 Fast Binding Update accepted but NCoA is invalid. Use NCoA supplied in "alternate" CoA

Values of the Status field that are greater than or equal to 128 indicate that the Binding Update was rejected by the receiving node. The following such Status values are currently defined:

128 Reason unspecified

129 Administratively prohibited

130 Insufficient resources

131 Incorrect interface identifier length

'K' flag

See RFC 3775 [3].

Reserved

An unused field. MUST be set to zero.

Sequence Number

Copied from the FBU message for use by the MN in matching this acknowledgment with an outstanding FBU.

Lifetime

The granted lifetime in seconds for which the sender of this message will retain a binding for traffic redirection.

Mobility Options

MUST contain an "alternate" CoA if Status is 1.

6.3.3. Fast Neighbor Advertisement (FNA)

A MN sends a Fast Neighbor Advertisement to announce itself to the NAR. When the Mobility Header Type is FNA, the Payload Proto field may be set to IPv6 to assist FBU encapsulation.




Figure 10: Fast Neighbor Advertisement (FNA) Message

IP fields:**Source Address**

NCoA

Destination Address

NAR's IP Address

Mobility Options

MUST contain the Mobility Header Link-Layer Address of the MN in the MH-LLA option format. See Section 6.4.4.

The MN sends a Fast Neighbor Advertisement to the NAR, as soon as it regains connectivity on the new link. Arriving or buffered packets can be immediately forwarded. If NAR is proxying NCoA, it creates a neighbor cache entry in REACHABLE state. If there is no entry, it creates one and sets it to REACHABLE. If there is an entry in the INCOMPLETE state without a Link-Layer Address, it sets it to

REACHABLE. During the process of creating a neighbor cache entry, NAR can also detect if NCoA is in use, thus avoiding address collisions. Since the FBU is encapsulated within the FNA when sent from NAR's link, NAR drops the FBU if it detects a collision.

The combination of NCoA (present in source IP address) and the Link-Layer Address (present as a Mobility Option) SHOULD be used to distinguish the MN from other nodes.

6.4. New Options

All the options are of the form shown in Figure 11.

The Type values are defined from the Neighbor Discovery options space. The Length field is in units of 8 octets, except for the Mobility Header Link-Layer Address option, whose Length field is in units of octets in accordance with Section 6.2 in [3]. Option-Code provides additional information for each of the options (See individual options below).




Figure 11: Option Format

6.4.1. IP Address Option

This option is sent in the Proxy Router Advertisement, the Handover Initiate, and Handover Acknowledge messages.




Figure 12: IPv6 Address Option

Type	17
Length	The size of this option in 8 octets including the Type, Option-Code, and Length fields.
Option-Code	1 Old Care-of Address 2 New Care-of Address 3 NAR's IP address
Prefix Length	The Length of the IPv6 Address Prefix.
Reserved	MUST be set to zero by the sender and MUST be ignored by the receiver.
IPv6 Address	The IP address for the unit defined by the Type field.

6.4.2. New Router Prefix Information Option

This option is sent in the PrRtAdv message to provide the prefix information valid on the NAR.




Figure 13: New Router Prefix Information Option

Type	18
Length	The size of this option in 8 octets including the Type, Option-Code, and Length fields.
Option-Code	0
Prefix Length	8-bit unsigned integer. The number of leading bits in the Prefix that are valid. The value ranges from 0 to 128.
Reserved	MUST be set to zero by the sender and MUST be ignored by the receiver.
Prefix	An IP address or a prefix of an IP address. The Prefix Length field contains the number of valid leading bits in the prefix. The bits in the prefix after the prefix length are reserved and MUST be initialized to zero by the sender and ignored by the receiver.

6.4.3. Link-Layer Address (LLA) Option




Figure 14: Link-Layer Address Option

Type	19
Length	The size of this option in 8 octets including the Type, Option-Code, and Length fields.
Option-Code	0 wildcard requesting resolution for all nearby access points 1 Link-Layer Address of the New Access Point 2 Link-Layer Address of the MN 3 Link-Layer Address of the NAR (i.e., Proxied Originator) 4 Link-Layer Address of the source of the RtSolPr or PrRtAdv message 5 The access point identified by the LLA belongs to the current interface of the router 6 No prefix information available for the access point identified by the LLA 7 No fast handovers support available for the access point identified by the LLA
LLA	The variable length Link-Layer Address.

Depending on the size of the individual LLA option, appropriate padding **MUST** be used to ensure that the entire option size is a multiple of 8 octets.

The New Access Point Link-Layer Address contains the Link-Layer Address of the access point for which handover is about to be attempted. This is used in the Router Solicitation for the Proxy Advertisement message.

The MN Link-Layer Address option contains the Link-Layer Address of an MN. It is used in the Handover Initiate message.

The NAR (i.e., Proxied Originator) Link-Layer Address option contains the Link-Layer Address of the Access Router to which the Proxy Router Solicitation message refers.

6.4.4. Mobility Header Link-Layer Address (MH-LLA) Option

This option is identical to the LLA option, but is carried in the Mobility Header messages (i.e., FNA). In the future, other Mobility Header messages may also make use of this option. For instance, including this option in FBU allows PAR to obtain the MN's LLA readily. The format of the option when the LLA is 6 bytes is shown in Figure 15. When the LLA size is different, the option MUST be aligned appropriately. See Section 6.2 in [3].




Figure 15: Mobility Header Link-Layer Address Option

Type	7
Length	The size of this option in octets not including the Type, Length, and Option-Code fields.
Option-Code	2 Link-Layer Address of the MN
LLA	The variable length Link-Layer Address.

6.4.5. Neighbor Advertisement Acknowledgment (NAACK)




Figure 16: Neighbor Advertisement Acknowledgment Option

Type	20
------	----

Length	8-bit unsigned integer. Length of the option, in 8 octets. The length is 1 when NCoA is not supplied. The length is 3 when NCoA is supplied (immediately following the Reserved field).
Option-Code	0
Status	8-bit unsigned integer indicating the disposition of the Fast Neighbor Advertisement message. The following Status values are currently defined: 1 The New CoA is invalid. 2 The New CoA is invalid; use the supplied CoA. The New CoA MUST be present following the Reserved field. 128 Link Layer Address unrecognized.
Reserved	MUST be set to zero by the sender and MUST be ignored by the receiver.

The NAR responds to the FNA with the NAACK option to notify the MN to use a different NCoA if there is address collision. If the NCoA is invalid, the Router Advertisement MUST use the NCoA as the destination address but use the L2 address present in the FNA. The MN SHOULD use the NCoA if it is supplied with the NAACK option. If the NAACK indicates that the Link-Layer Address is unrecognized, the MN MUST NOT use the NCoA or PCoA and SHOULD start the process of acquiring an NCoA at the NAR immediately.

New option types may be defined in the future.

7. Configurable Parameters

Parameter Name	Default Value	Definition
-----	-----	-----
RTSOLPR_RETRIES	3	Section 6.1.1
MAX_RTSOLPR_RATE	3	Section 6.1.1
FBUS_RETRIES	3	Section 4
PROXY_ND_LIFETIME	1.5 seconds	Section 6.2.2
HI_RETRIES	3	Section 6.2.1

8. Security Considerations

The following security vulnerabilities are identified, and suggested solutions are mentioned.

1. Insecure FBU: In this case, packets meant for one address could be stolen, or redirected to some unsuspecting node. This concern is the same as that in an MN and Home Agent relationship.

Hence, the PAR MUST ensure that the FBU packet arrived from a node that legitimately owns the PCoA. The access router and its hosts may use any available mechanism to establish a security association that MUST be used to secure FBU. The current version of this protocol does not specify how this security association is established. However, future work may specify this security association establishment.

If an access router can ensure that the source IP address in an arriving packet could only have originated from the node whose Link-Layer Address is in the router's neighbor cache, then a bogus node cannot use a victim's IP address for malicious redirection of traffic. Such an operation is recommended at least on neighbor discovery messages including the RtSolPr message.

2. Secure FBU, malicious or inadvertent redirection: In this case, the FBU is secured, but the target of binding happens to be an unsuspecting node due to inadvertent operation or malicious intent. This vulnerability can lead to an MN with a genuine security association with its access router redirecting traffic to an incorrect address.

However, the target of malicious traffic redirection is limited to an interface on an access router with which the PAR has a security association. The PAR MUST verify that the NCoA to which PCoA is being bound actually belongs to NAR's prefix. To do this, HI and HAcK message exchanges are to be used. When NAR accepts NCoA in HI (with Code = 0), it proxies NCoA so that any arriving packets are not sent on the link until the MN attaches and announces itself through FNA. Therefore, any inadvertent or malicious redirection to a host is avoided. It is still possible to jam NAR's buffer with redirected traffic. However, since NAR's handover state corresponding to NCoA has a finite (and short) lifetime corresponding to a small multiple of anticipated handover latency, the extent of this vulnerability is arguably small.

3. Sending an FBU from NAR's link: A malicious node may send an FBU from NAR's link providing an unsuspecting node's address as NCoA. Since the FBU is encapsulated in the FNA, NAR should detect the

collision with an address in use when processing the FNA, and then drop the FBU. When NAR is unable to detect address collisions, there is a vulnerability that redirection can affect an unsuspecting node.

9. IANA Considerations

This document defines four new experimental ICMPv6 messages that use the Experimental Mobility Protocol ICMPv6 format [4]. These four new Subtype value assignments out of the Experimental Mobility Protocol Subtype Registry [4] have been assigned as follows:

Subtype	Description	Reference
-----	-----	-----
2	RtSolPr	Section 6.1.1
3	PrRtAdv	Section 6.1.2
4	HI	Section 6.2.1
5	HAck	Section 6.2.2

This document defines four new Neighbor Discovery [6] options that have received Type assignments from IANA.

Option-Type	Description	Reference
-----	-----	-----
17	IP Address Option	Section 6.4.1
18	New Router Prefix Information Option	Section 6.4.2
19	Link-Layer Address Option	Section 6.4.3
20	Neighbor Advertisement Acknowledgment Option	Section 6.4.5

This document defines three new Mobility Header messages that have received type allocations from the Mobility Header Types registry at <http://www.iana.org/assignments/mobility-parameters>:

1. Fast Binding Update, described in Section 6.3.1
2. Fast Binding Acknowledgment, described in Section 6.3.2, and
3. Fast Neighbor Advertisement, described in Section 6.3.3.

This document defines a new Mobility Option which has received type assignments from the Mobility Options Type registry at <http://www.iana.org/assignments/mobility-parameters>:

1. Mobility Header Link-Layer Address option, described in Section 6.4.4.

10. Acknowledgments

The editor would like to thank all those who have provided feedback on this specification, but can only mention a few here: Martin Andre, Vijay Devarapalli, Youn-Hee Han, Emil Ivov, Suvidh Mathur, Koshiro Mitsuya, Gabriel Montenegro, Takeshi Ogawa, Sun Peng, YC Peng, Domagoj Premec, and Jonathan Wood. The editor would like to acknowledge a contribution from James Kempf to improve this specification. The editor would also like to thank the [mipshop] working group chair Gabriel Montenegro and the erstwhile [mobile ip] working group chairs Basavaraj Patil and Phil Roberts for providing much support for this work.

11. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.
- [3] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [4] Kempf, J., "Instructions for Seamoby and Experimental Mobility Protocol IANA Allocations", RFC 4065, July 2005.
- [5] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [6] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [7] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

12. Contributors

This document originated in the fast handover design team effort. The members of this design team in alphabetical order were: Gopal Dommety, Karim El-Malki, Mohammed Khalil, Charles Perkins, Hesham Soliman, George Tsirtsis, and Alper Yegin.

The design team member's contact information:

Gopal Dommety
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134

Phone: +1 408 525 1404
Email: gdommety@cisco.com

Karim El Malki
Ericsson Radio Systems AB
LM Ericssons Vag. 8
126 25 Stockholm
SWEDEN

Phone: +46 8 7195803
Fax: +46 8 7190170
Email: Karim.El-Malki@era.ericsson.se

Mohamed Khalil
Nortel Networks

Email: mkhalil@nortelnetworks.com

Charles E. Perkins
Communications Systems Lab
Nokia Research Center
313 Fairchild Drive
Mountain View, California 94043
USA

Phone: +1-650 625-2986
Fax: +1 650 625-2502
Email: charliep@iprg.nokia.com

Hesham Soliman
Flarion Technologies

Email: H.Soliman@flarion.com

George Tsirtsis
Flarion Technologies

EMail: G.Tsirtsis@flarion.com

Alper E. Yegin
Samsung Advanced Institute of Technology
75 West Plumeria Drive
San Jose, CA 95134
USA

Phone: +1 408 544 5656
EMail: alper.yegin@samsung.com

Author's Address

Rajeev Koodli, Editor
Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94043 USA

Phone: +1 650 625 2359
Fax: +1 650 625 2502
EMail: Rajeev.Koodli@nokia.com

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.

