

Mobile IPv6 Fast Handovers

Status of This Memo

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

Abstract

Mobile IPv6 enables a Mobile Node (MN) 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 (VoIP). 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	6
3.1. Addressing the Handover Latency	6
3.2. Protocol Operation	8
3.3. Protocol Operation during Network-Initiated Handover	11
4. Protocol Details	11
5. Other Considerations	15
5.1. Handover Capability Exchange	15
5.2. Determining New Care-of Address	16
5.3. Prefix Management	16
5.4. Packet Loss	17
5.5. DAD Handling	18
5.6. Fast or Erroneous Movement	19
6. Message Formats	20
6.1. New Neighborhood Discovery Messages	20
6.1.1. Router Solicitation for Proxy Advertisement (RtSolPr)	20
6.1.2. Proxy Router Advertisement (PrRtAdv)	22
6.2. Inter - Access Router Messages	25
6.2.1. Handover Initiate (HI)	25
6.2.2. Handover Acknowledge (HACK)	27
6.3. New Mobility Header Messages	28
6.3.1. Fast Binding Update (FBU)	28
6.3.2. Fast Binding Acknowledgment (FBack)	30
6.4. Unsolicited Neighbor Advertisement (UNA)	31
6.5. New Options	32
6.5.1. IP Address/Prefix Option	33
6.5.2. Link-Layer Address (LLA) Option	34
6.5.3. Mobility Header Link-Layer Address (MH-LLA) Option	35
6.5.4. Binding Authorization Data for FMIPv6 (BADF)	35
6.5.5. Neighbor Advertisement Acknowledgment (NAACK)	36
7. Related Protocol and Device Considerations	37
8. Evolution from and Compatibility with RFC 4068	38
9. Configurable Parameters	39
10. Security Considerations	39
10.1. Peer Authorization Database Entries when Using IKEv2	41
10.2. Security Policy Database Entries	42
11. IANA Considerations	42
12. Acknowledgments	43
13. References	44
13.1. Normative References	44
13.2. Informative References	45
Appendix A. Contributors	46
Appendix B. Changes since RFC 4068	46

1. Introduction

Mobile IPv6 [RFC3775] 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 link-layer procedures, 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 problems: 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 [RFC3775]. There are no special requirements for a mobile node to behave differently with respect to its standard Mobile IP operations.

This specification is applicable when a mobile node has to perform IP layer operations as a result of handovers. This specification does not address improving the link switching latency. It does not modify or optimize procedures related to signaling with the home agent of a mobile node. Indeed, while targeted for Mobile IPv6, it could be used with any mechanism that allows communication to continue despite movements. Finally, this specification does not address bulk movement of nodes using aggregate prefixes.

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 [RFC2119]. 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 in addition to those defined in [RFC3775]. The reference handover scenario is illustrated in Figure 1.




Figure 1: Reference Scenario for Handover

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 the AP's L2 address. Sometimes, AP-ID is also referred to as a Basic Service Set Identifier (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 anticipated 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 can also act as a trigger for network-initiated handover.

(AP-ID, AR-Info) tuple: Contains an access router's L2 and IP addresses, and 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, and Prefix] is called "AR-Info". See Section 5.3.

Neighborhood Discovery: The process of resolving neighborhood AP-IDs to 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.

Predictive Fast Handover: The fast handover in which an MN is able to send an FBU when it is attached to the PAR, which then establishes forwarding for its traffic (even before the MN attaches to the NAR).

Reactive Fast Handover: The fast handover in which an MN is able to send the FBU only after attaching to the NAR.

Unsolicited Neighbor Advertisement (UNA): The message in [RFC4861] with 'O' bit cleared.

Fast Neighbor Advertisement (FNA): This message from RFC 4068 [RFC4068] is deprecated. The UNA message above is the preferred message in this specification.

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.

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 the 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 process the Binding Update successfully, 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 (e.g., a "scan" in Wireless Local Area Network (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 in 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)" (see Section 6.3.2) message prior to its movement. In the event it moves without receiving an FBack, the MN can still start using NCoA after announcing its attachment through an unsolicited Neighbor Advertisement message (with the 'O' bit set to zero) [RFC4861]; NAR responds to this UNA message in case it wishes to provide a different IP address to use. In this way, NCoA configuration latency is reduced.

The information provided in the PrRtAdv message can be used even when DHCP [RFC3315] is used to configure an NCoA on the NAR's link. In this case, the protocol supports forwarding using PCoA, and the MN performs DHCP once it attaches to the NAR's link. The MN still formulates an NCoA for FBU processing; however, it MUST NOT send data packets using the NCoA in the FBU.

In order to reduce the Binding Update latency, the protocol specifies a binding between the Previous CoA (PCoA) and NCoA. An MN sends a "Fast Binding Update" (see Section 6.3.1) message to its Previous Access Router to establish this tunnel. When feasible, the MN SHOULD send an FBU from the PAR's link. Otherwise, the MN should send the FBU immediately after detecting attachment to the NAR. An FBU message MUST contain the Binding Authorization Data for FMIPv6 (BADF) option (see Section 6.5.4) in order to ensure that only a legitimate MN that owns the PCoA is able to establish a binding. Subsequent sections describe the protocol mechanics. In any case, the result is that the 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 the PAR, again until it completes Binding Update. And, PAR MUST 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 a PCoA as a source IP address are not dropped due to ingress filtering. Even though the MN is IP-capable on the new link, it cannot use the NCoA directly with its correspondents without the correspondents first establishing a binding cache entry (for the NCoA). Forwarding support for the 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 the NAR to send Neighbor Solicitation before the MN attaches, subsequent retransmissions of address resolution are separated by a default period of one second each. In order to circumvent this delay, an MN announces its attachment immediately with an UNA message that allows the NAR to forward packets to the MN right away. 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 the

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 the PAR before providing forwarding support for the MN. That is, it may create a forwarding entry for the NCoA, subject to "approval" from the PAR, which it trusts. In addition, buffering for handover traffic at the NAR may be desirable. Even though the Neighbor Discovery protocol provides a small buffer (typically one or two packets) for packets awaiting address resolution, this buffer may be inadequate for traffic, such as VoIP, already in progress. The routers may also wish to maintain a separate buffer for servicing the handover traffic. Finally, the access routers could transfer network-resident contexts, such as access control, Quality of Service (QoS), and header compression, in conjunction with handover (although the context transfer process itself is not specified in this document). For all these operations, the protocol provides "Handover Initiate (HI)" and "Handover Acknowledge (HACK)" messages (see Section 6.2). Both of these messages SHOULD be used. The access routers MUST have the 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 message 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 an 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 an 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 to the PAR. The purpose of the FBU is to authorize the PAR to bind the PCoA to the NCoA, so that arriving packets can be tunneled to the new location of the MN. The FBU should be sent from the PAR's link whenever feasible. 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.

The format and semantics of FBU processing are specified in Section 6.3.1. The FBU message MUST contain the BADF option (see Section 6.5.4) to secure the message.

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 FBack on the previous link. This means that packet tunneling is already in progress by the time the MN handovers to the NAR. The MN SHOULD send the UNA immediately after attaching to the NAR, so that arriving as well as buffered packets can be forwarded to the MN right away. Before sending FBack to the MN, the 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 an HI message (from PAR to NAR), and NAR MAY assign the proposed NCoA. Such an assigned NCoA MUST be returned in HAck (from NAR to PAR), and PAR MUST in turn provide the assigned NCoA in FBack. If there is an assigned NCoA returned in FBack, the MN MUST use the assigned address (and not the proposed address in 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 the PAR has processed the FBU successfully. Hence, the MN (re)sends the FBU message to the PAR immediately after sending the UNA message. If the NAR chooses to supply a different IP address to use than the NCoA, it MAY send a Router Advertisement with "Neighbor Advertisement Acknowledge (NAACK)" option in which it includes an alternate IP address for the MN to use. Detailed UNA processing rules are specified in Section 6.4.

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 the "predictive" mode of operation. The scenario in which the MN sends an FBU from the NAR's link is illustrated in Figure 3. For convenience, this scenario is characterized as the "reactive" mode of operation. Note that the reactive mode also includes the case in which an FBU has been sent from the PAR's link but an FBack has not yet been received. The figure is intended to illustrate that the FBU is forwarded through the NAR, but it is processed only by the PAR.

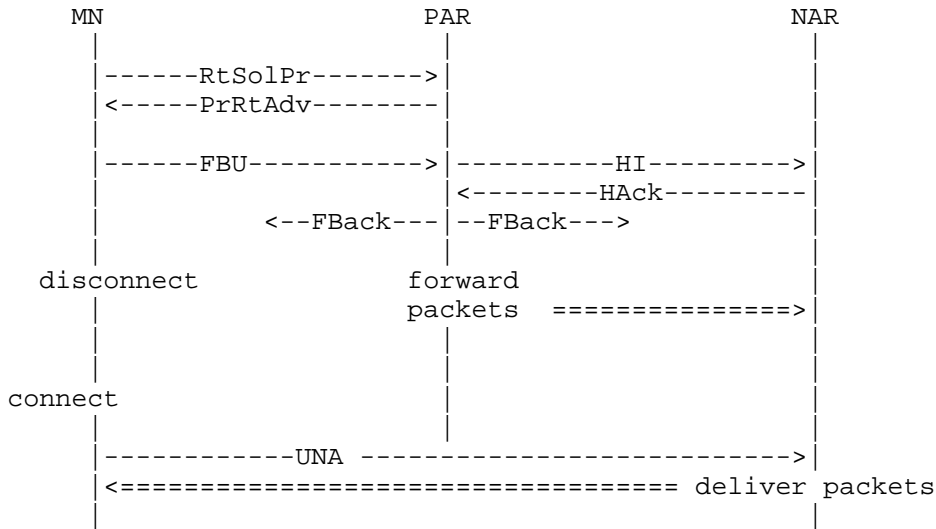


Figure 2: Predictive Fast Handover

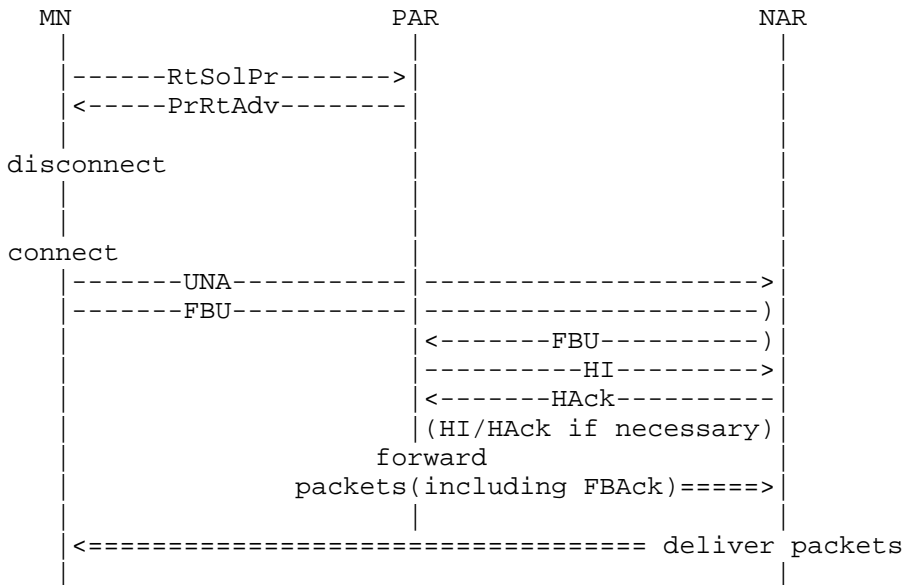


Figure 3: Reactive Fast Handover

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

3.3. Protocol Operation during 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 such networks, the PAR can send an unsolicited PrRtAdv containing the link-layer address, IP address, and subnet prefix 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 the 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.

4. Protocol Details

All descriptions refer to Figure 1.

After discovering one or more nearby access points, the MN sends RtSolPr to the PAR in order to resolve access point identifiers to subnet router information. A convenient time to do this is 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., UNA) can be transmitted immediately. Implementations SHOULD make use of such triggers whenever available.

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

As a response to RtSolPr, the 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 the 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 First in First Out (FIFO) buffering is used, the PAR MUST continue forwarding the packets to the PCoA (i.e., buffer and forward). While the protocol does not forbid such an implementation support, care must be taken to ensure that the PAR continues forwarding packets to the PCoA (i.e., uses a buffer and forward approach). The PAR SHOULD stop buffering once it begins forwarding packets to the NCoA.

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 the 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 the 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. In response to the FBU, the PAR establishes a binding between the PCoA ("Home Address") and the NCoA, and sends the FBack to the MN. Prior to establishing this binding, the PAR SHOULD send an HI message to the NAR, and receive HAcK in response. In order 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 the PCoA, i.e., the FBU is sent from the PAR's link, 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 the NCoA supplied in the HI message is unique before beginning to defend it. It sends a Duplicate Address Detection (DAD) probe [RFC4862] for NCoA to verify uniqueness. However, in deployments where the probability of address collisions is considered extremely low (and hence not an issue), the parameter DupAddrDetectTransmits (see [RFC4862]) is set to zero on the NAR, allowing it to avoid performing DAD on the NCoA. The NAR similarly sets DupAddrDetectTransmits to zero in other deployments where DAD is not a concern. Once the NCoA is determined to be unique, the NAR starts proxying [RFC4861] the address for PROXY_ND_LIFETIME during which the MN is expected to connect to the NAR. In case there is already an NCoA present in its data structure (for instance, it has already processed an HI message earlier), the NAR MAY verify if the LLA is the same as its own or that of the MN itself. If so, the NAR MAY allow the use of the NCoA.

2. allocates the 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 the PCoA (on the interface to which the MN is attaching to) in case the 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 SHOULD 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 to the PAR.

When the Code value in HI is 1, the NAR MUST skip the above operations. Sending an HI message with Code 1 allows the NAR to validate the neighbor cache entry it creates for the MN during UNA processing. That is, the NAR can make use of the knowledge that its trusted peer (i.e., the 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 as well 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 the 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.

As soon as the MN establishes link connectivity with the NAR, it:

1. sends an UNA message (see Section 6.4). If the MN has not received an FBack by the time UNA is being sent, it SHOULD send an FBU message following the UNA message.
2. joins the all-nodes multicast group and the solicited-node multicast group corresponding to the NCoA.
3. starts a DAD probe for NCoA, see [RFC4862].

When a NAR receives an UNA message, it:

1. deletes its proxy neighbor cache entry, if it exists, updates the state to STALE [RFC4861], and forwards arriving and buffered packets.

2. updates an entry in INCOMPLETE state [RFC4861], if it exists, to STALE and forwards arriving and buffered packets. This would be the case if NAR had previously sent a Neighbor Solicitation that went unanswered perhaps because the MN had not yet attached to the link.

The buffer for handover traffic should be linked to this UNA processing. The exact mechanism is implementation dependent.

The NAR may choose to provide a different IP address other than the NCoA. This is possible if it is proxying the NCoA. In such a case, it:

1. MAY send a Router Advertisement with the NAACK option in which it includes an alternate IP address for use. This message MUST be sent to the source IP address present in UNA using the same Layer 2 address present in UNA.

If the MN receives an IP address in the NAACK option, it MUST use it and send an FBU using the new CoA. As a special case, the address supplied in NAACK could be the PCoA itself, in which case the MN MUST NOT send any more FBUs. The Status codes for the NAACK option are specified in Section 6.5.5.

Once the MN has confirmed its NCoA (either through DAD or when provided for by the NAR), it SHOULD send a Neighbor Advertisement message with the 'O' bit set, to the all-nodes multicast address. This message allows MN's neighbors to update their neighbor cache entries.

For data forwarding, the PAR tunnels packets using its global IP address valid on the interface to which the MN was attached. The MN reverse tunnels its packets to the same global address of PAR. The tunnel end-point addresses must be configured accordingly. When the PAR receives a reverse tunneled packet, it must verify if a secure binding exists for the MN identified by the PCoA in the tunneled packet, before forwarding the packet.

5. Other Considerations

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 the PAR does not support the fast handover protocol and stop sending any more RtSolPr messages.

Even if an MN's current access router is capable of providing fast handover support, 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, the 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. However, the MN itself does not have to wait on PAR's link for this exchange to take place. It can handover any time after sending the FBU message; sometimes it may be forced to handover without sending the FBU. In any case, it can still confirm using NCoA from NAR's link by sending the UNA message.

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

When DHCP is used, the protocol supports forwarding for PCoA only. In this case, the MN MUST perform DHCP operations once it attaches to the NAR even though it formulates an NCoA for transmitting the FBU. This is indicated in the PrRtAdv message with Code = 5.

5.3. Prefix Management

As defined in Section 2, the Prefix part of "AR-Info" is the prefix valid on the interface to which the AP is attached. This document does not specify how this Prefix is managed, its length and assignment policies. The protocol operation specified in this document works regardless of these considerations. Often, but not necessarily always, this Prefix may be the aggregate prefix (such as /48) valid on the interface. In some deployments, each MN may have its own per-mobile prefix (such as a /64) used for generating the NCoA. Some point-to-point links may use such a deployment.

When per-mobile prefix assignment is used, the "AR-Info" advertised in PrRtAdv still includes the (aggregate) prefix valid on the interface to which the target AP is attached, unless the access routers communicate with each other (using HI and HAck messages) to

manage the per-mobile prefix. The MN still formulates an NCoA using the aggregate prefix. However, an alternate NCoA based on the per-mobile prefix is returned by NAR in the HAcK message. This alternate NCoA is provided to the MN in either the FBack message or in the NAAcK option.

5.4. 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 until the FBU is processed will be lost unless they are buffered. This protocol provides an option to indicate 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.

Whereas buffering can enable a smooth handover, the buffer size and the rate at which buffered packets are eventually forwarded are important considerations when providing buffering support. There are a number of aspects to consider:

- o Some applications transmit less data over a given period of data than others, and this implies different buffering requirements. For instance, Voice over IP typically needs smaller buffers compared to high-resolution streaming video, as the latter has larger packet sizes and higher arrival rates.
- o When the mobile node appears on the new link, having the buffering router send a large number of packets in quick succession may overtax the resources of the router, the mobile node itself, or the path between these two.

In particular, transmitting a large amount of buffered packets in succession can congest the path between the buffering router and the mobile node. Furthermore, nodes (such as a base station) on the path between the buffering router and the mobile node may drop such packets. If a base station buffers too many such packets, they may contribute to additional jitter for packets arriving behind them, which is undesirable for real-time communication.

- o Since routers are not involved in end-to-end communication, they have no knowledge of transport conditions.

- o The wireless connectivity of the mobile node may vary over time. It may achieve a smaller or higher bandwidth on the new link, signal strength may be weak at the time it just enters the area of this access point, and so on.

As a result, it is difficult to design an algorithm that would transmit buffered packets at appropriate spacing under all scenarios. The purpose of fast handovers is to avoid packet loss. Yet, draining buffered packets too fast can, by itself, cause loss of the packets, as well as blocking or loss of following packets meant for the mobile node.

This specification does not restrict implementations from providing specialized buffering support for any specific situation. However, attention must be paid to the rate at which buffered packets are forwarded to the MN once attachment is complete. Routers implementing this specification MUST implement at least the default algorithm, which is based on the original arrival rates of the buffered packets. A maximum of 5 packets MAY be sent one after another, but all subsequent packets SHOULD use a sending rate that is determined by metering the rate at which packets have entered the buffer, potentially using smoothing techniques such as recent activity over a sliding time window and weighted averages [RFC3290].

It should be noted, however, that this default algorithm is crude and may not be suitable for all situations. Future revisions of this specification may provide additional algorithms, once enough experience of the various conditions in deployed networks is attained.

5.5. DAD Handling

Duplicate Address Detection (DAD) was defined in [RFC4862] 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. Hence, the protocol specified in this document SHOULD only be used in deployments where the probability of such address collisions is extremely low or it is not a concern (because of the address management procedure deployed). The protocol requires the NAR to send a DAD probe before it starts defending the NCoA. However, this DAD delay can be turned off by setting DupAddrDetectTransmits to zero on the NAR [RFC4862].

This document specifies messages that can be used to provide duplicate-free addresses, but the document does not specify how to create or manage such duplicate-free addresses. In some cases, the NAR may already have the knowledge required to assess whether or not the MN's address is a duplicate before the MN moves to the new subnet. For example, in some deployments, the NAR may maintain a pool of duplicate-free addresses in a list for handover purposes. In such cases, the NAR can provide this disposition in the HAcK message (see Section 6.2.2) or in the NAcK option (see Section 6.5.5).

5.6. Fast or Erroneous Movement

Although this specification is for fast handover, the protocol is limited in terms of how fast an MN can move. A special case of fast movement is ping-pong, 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 either moves 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 Home Agent 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 of the order of 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 terminate that tunnel; i.e., start delivering packets directly to the node instead. In order for the PAR to process such an FBU, the lifetime of the security association has to be at least that of the tunnel itself.

Temporary tunnels for the purposes of fast handovers should use short lifetimes (of the order of at most a few tens 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, it SHOULD send a new Fast Binding Update to the PAR. This FBU supersedes the existing binding at the PAR, and the packets will be redirected to the newly confirmed location of the MN.

6. Message Formats

All the ICMPv6 messages have a common Type specified in [RFC4443]. The messages are distinguished based on the Subtype field (see below). For all the ICMPv6 messages, the checksum is defined in [RFC4443].

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 Section 6.5.2.




Figure 4: Router Solicitation for Proxy Advertisement (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.

ICMP Fields:

Type: 154

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 (LLA) 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. See the LLA option below.

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 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 an 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 an 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 an 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 6.4.3.




Figure 5: Proxy Router Advertisement (PrRtAdv) Message

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 Proxy Advertisement or the address of the node the access router is instructing to handover.

Hop Limit: 255. See RFC 2461.

ICMP Fields:

Type: 154

Code: 0, 1, 2, 3, 4, or 5. 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 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 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 for. This option MUST be included when the Code is 0 or 1.

New Router's IP Address: The IP address of the NAR. This option MUST be included when the Code is 0 or 1.

New Router Prefix Information Option: Specifies the prefix of the access router the message is proxied for and is used for address auto-configuration. This option MUST be included when the Code is 0 or 1. However, when this prefix is the same as what is 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 PrRtAdv is sent unsolicited. The PAR MAY compute a new CoA using the 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, 4, and 5 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. The Option-Code field in the New Access Point LLA option in this case 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 has been sent unsolicited. If a New CoA option is present following the New Router Prefix Information option, the MN MUST use the supplied NCoA and send an 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. This would be the case, for instance, when a number of access points are connected to the same router interface, or when network based mobility management mechanisms ensure that the specific mobile node always observes the same prefix regardless of whether there is a separate router attached to the target access point. 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.

A Proxy Router Advertisement with Code 5 means that the MN may use the new router information present for detecting movement to a new subnet, but the MN must perform DHCP [RFC3315] upon attaching to the NAR's link. The PAR and NAR will forward packets to the PCoA of the MN. The MN must still formulate an NCoA for transmitting FBU (using the information sent in this message), but that NCoA will not be used for forwarding packets.

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 an 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

ICMP Fields:

Type: 154

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 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 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 round trip time (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 the NAR to the 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.

ICMP Fields:

Type: 154

Code:

- 0: Handover Accepted, NCoA valid
- 1: Handover Accepted, NCoA not valid or in use
- 2: Handover Accepted, NCoA assigned (used in Assigned addressing)
- 3: Handover Accepted, use PCoA
- 4: Message sent unsolicited, usually to trigger an HI message
- 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 the PCoA (instead of accepting or assigning an NCoA), establish a host route entry for the PCoA, and set up a tunnel to the PAR to forward the MN's packets sent with the 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. The NAR indicates forwarding support for PCoA using Code value 3 in the HAck message. Subsequently, the PAR establishes a tunnel to the NAR in order to forward packets arriving for the PCoA.

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 [RFC3775]. The Fast Binding Update, Fast Binding Acknowledgment, and the (deprecated) Fast Neighbor Advertisement messages use the Mobility Header.

6.3.1. Fast Binding Update (FBU)

The Fast Binding Update message has a Mobility Header Type value of 8. The FBU is identical to the Mobile IPv6 Binding Update (BU) message. However, the processing rules are slightly different.

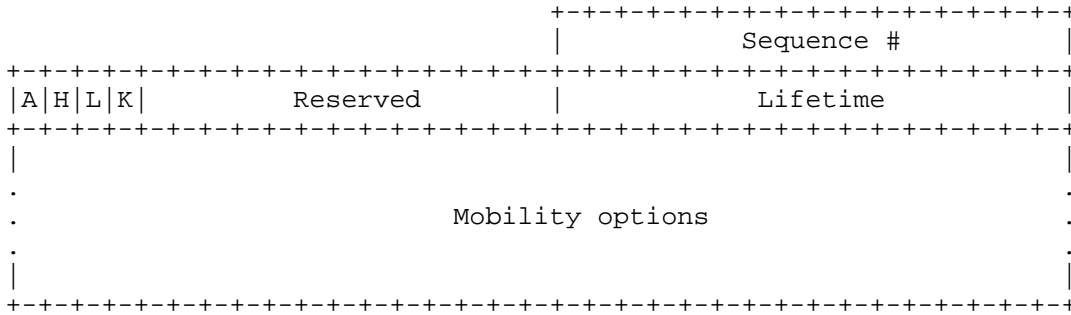


Figure 8: Fast Binding Update (FBU) Message

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 [RFC3775].

'L' flag: See [RFC3775].

'K' flag: See [RFC3775].

Reserved: This field is unused. MUST be set to zero.

Sequence Number: See [RFC3775].

Lifetime: The requested time in seconds for which the sender wishes to have a binding.

Mobility Options: MUST contain an alternate CoA option set to the NCoA when an FBU is sent from the PAR's link. MUST contain the Binding Authorization Data for the FMIP (BADF) option. See Section 6.5.4. MAY contain the Mobility Header LLA option (see Section 6.5.3).

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 the NCoA on the NAR according to Neighbor Discovery and IPv6 Address Configuration protocols. When the MN moves without having received a PrRtAdv message, it cannot transmit an UNA message upon attaching to the NAR's link.

The source IP address is the PCoA when the FBU is sent from the PAR's link, and the source IP address is the NCoA when the FBU sent from the NAR's link. When the source IP address is the PCoA, the MN MUST include the alternate CoA option set to NCoA. The PAR MUST process the FBU even though the address in the alternate CoA option is different from that in the source IP address, and ensure that the address in the alternate CoA option is used in the New CoA option in the HI message to the NAR.

The FBU MUST also include the Home Address Option set to PCoA. An FBU message MUST be protected so that the PAR is able to determine that the FBU message is sent by an MN that legitimately owns the PCoA.

6.3.2. Fast Binding Acknowledgment (FBack)

The Fast Binding Acknowledgment message has a Mobility Header Type value of 9. The FBack message is sent by the PAR to acknowledge receipt of a Fast Binding Update message in which the 'A' bit is set. If PAR sends an HI message to the NAR after processing an FBU, the FBack message SHOULD NOT be sent to the MN before the PAR receives a HAck message from the NAR. The PAR MAY send the FBack immediately in the reactive mode however. The Fast Binding Acknowledgment MAY also be sent to the MN on the old link.

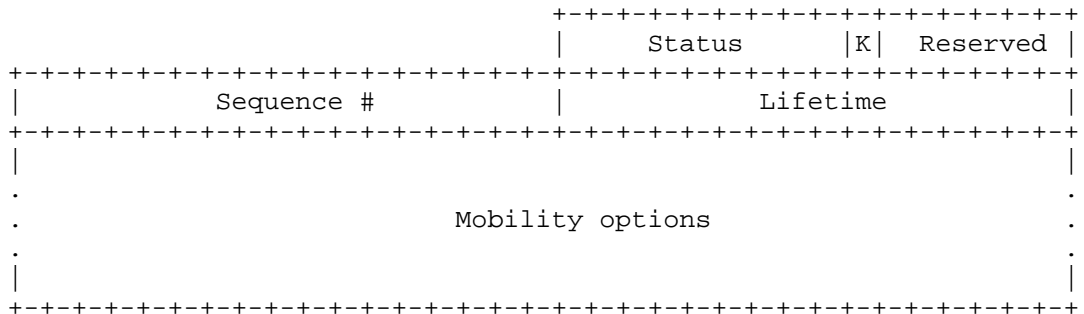


Figure 9: Fast Binding Acknowledgment (FBack) Message

IP Fields:

Source address: The IP address of the Previous Access Router

Destination Address: The NCoA, and optionally the PCoA

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 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 [RFC3775].

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. MUST contain the Binding Authorization Data for FMIP (BADF) option. See 6.4.5.

6.4. Unsolicited Neighbor Advertisement (UNA)

This is the same message as in [RFC4861] with the requirement that the 'O' bit is always set to zero. Since this is an unsolicited message, the 'S' bit is zero, and since this is sent by an MN, the 'R' bit is also zero.

If the NAR is proxying the NCoA (as a result of HI and HAcK exchange), then UNA processing has additional steps (see below). If the NAR is not proxying the NCoA (for instance, HI and HAcK exchange has not taken place), then UNA processing follows the same procedure as specified in [RFC4861]. Implementations MAY retransmit UNA subject to the specification in Section 7.2.6 of [RFC4861] while noting that the default RetransTimer value is large for handover purposes.

The Source Address in UNA MUST be the NCoA. The destination address is typically the all-nodes multicast address; however, some deployments may not prefer transmission to a multicast address. In such cases, the destination address SHOULD be the NAR's IP address.

The Target Address MUST include the NCoA, and the Target link-layer address MUST include the MN's LLA.

The MN sends an UNA message to the NAR, as soon as it regains connectivity on the new link. Arriving or buffered packets can be immediately forwarded. If the NAR is proxying the NCoA, it creates a neighbor cache entry in STALE state but forwards packets as it determines bidirectional reachability according to the standard Neighbor Discovery procedure. If there is an entry in INCOMPLETE state without a link-layer address, it sets it to STALE, again according to the procedure in [RFC4861].

The NAR MAY wish to provide a different IP address to the MN than the one in the UNA message. In such a case, the NAR MUST delete the proxy entry for the NCoA and send a Router Advertisement with the NAACK option containing the new IP address.

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

6.5. New Options

All the options, with the exception of Binding Data Authorization for FMIPv6 (BADF) discussed in Section 6.5.4, use Type, Length, and Option-Code format shown in Figure 10.

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 [RFC3775]. And, Option-Code provides additional information for each of the options (see individual options below).




Figure 10: Option Format

6.5.1. IP Address/Prefix Option

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

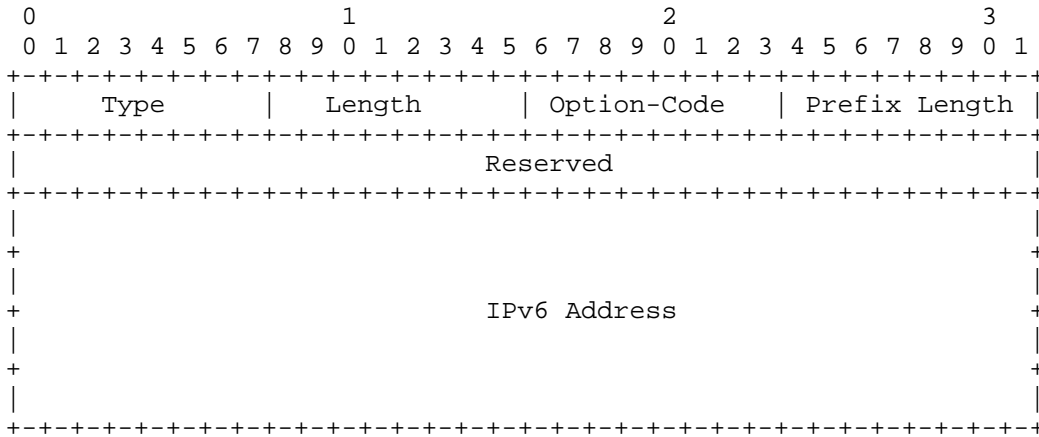


Figure 11: IPv6 Address/Prefix Option

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
- 4: NAR's Prefix, sent in PrRtAdv. 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.

Prefix Length: 8-bit unsigned integer that indicates the length of the IPv6 Address Prefix. The value ranges from 0 to 128.

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

IPv6 address: The IP address defined by the Option-Code field.

6.5.2. Link-Layer Address (LLA) Option




Figure 12: 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 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.

The LLA option does not have a length field for the LLA itself. The implementations must consult the specific link layer over which the protocol is run in order to determine the content and length of the LLA.

Depending on the size of 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 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.

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

This option is identical to the LLA option, but is carried in the Mobility Header messages, e.g., FBU. In the future, other Mobility Header messages may also make use of this option. The format of the option is shown in Figure 13. There are no alignment requirements for this option.




Figure 13: Mobility Header Link-Layer Address Option

Type: 7

Length: The size of this option in octets not including the Type and Length fields.

Option-Code: 2 Link-Layer Address of the MN.

LLA: The variable length link-layer address.

6.5.4. Binding Authorization Data for FMIPv6 (BADF)

This option MUST be present in FBU and FBack messages. The security association between the MN and the PAR is established by companion protocols [RFC5269]. This option specifies how to compute and verify a Message Authentication Code (MAC) using the established security association.

The format of this option is shown in Figure 14.

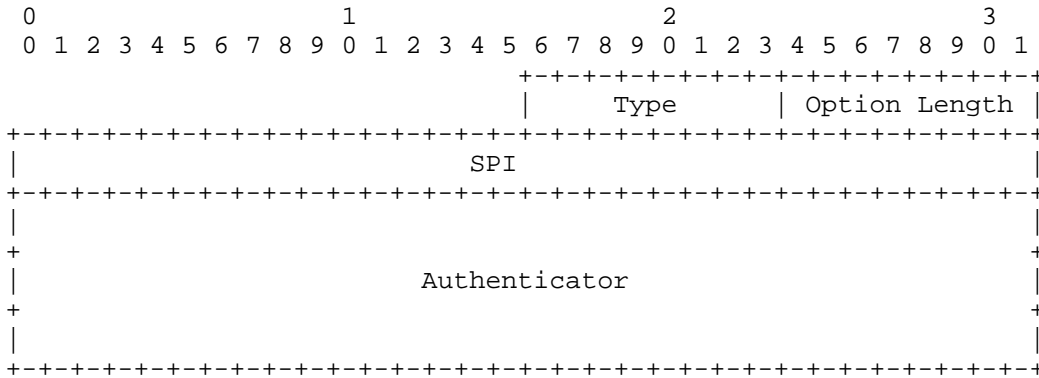


Figure 14: Binding Authorization Data for FMIPv6 (BADF) Option

Type: 21

Option Length: The length of the Authenticator in bytes

SPI: Security Parameter Index. SPI = 0 is reserved for the Authenticator computed using SEND-based handover keys.

Authenticator: Same as in RFC 3775, with "correspondent" replaced by the PAR's IP address, and Kbm replaced by the shared key between the MN and the PAR.

The default MAC calculation is done using HMAC_SHA1 with the first 96 bits used for the MAC. Since there is an Option Length field, implementations can use other algorithms such as HMAC_SHA256.

This option MUST be the last Mobility Option present.

6.5.5. Neighbor Advertisement Acknowledgment (NAACK)




Figure 15: Neighbor Advertisement Acknowledgment Option

Type: 20

Length: 8-bit unsigned integer. Length of the option, in 8 octets.

The length is 1 when a new CoA is not supplied. The length is 3 when a new CoA is present (immediately following the Reserved field)

Option-Code: 0

Status: 8-bit unsigned integer indicating the disposition of the Unsolicited Neighbor Advertisement message. The following Status values are currently defined:

- 1: NCoA is invalid, perform address configuration
- 2: NCoA is invalid, use the supplied NCoA. The supplied NCoA (in the form of an IP Address Option) MUST be present following the Reserved field.
- 3: NCoA is invalid, use NAR's IP address as NCoA in FBU
- 4: PCoA supplied, do not send FBU
- 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 UNA with the NAACK option to notify the MN to use a different NCoA than the one that the MN has used. If the NAR proposes a different NCoA, the Router Advertisement MUST use the source IP address in the UNA message as the destination address, and use the L2 address present in UNA. The MN MUST use the NCoA if it is supplied with the NAACK option. If the NAACK indicates that the Link-Layer Address is unrecognized, for instance, if the MN uses an LLA valid on PAR's link but the same LLA is not valid on NAR's link due to a different access technology, the MN MUST NOT use the NCoA or the PCoA and SHOULD start immediately the process of acquiring a different NCoA at the NAR.

In the future, new option types may be defined.

7. Related Protocol and Device Considerations

The protocol specified here, as a design principle, introduces no or minimal changes to related protocols. For example, no changes to the base Mobile IPv6 protocol are needed in order to implement this protocol. Similarly, no changes to the IPv6 stateless address auto-configuration protocol [RFC4862] and DHCP [RFC3315] are introduced. The protocol specifies an optional extension to Neighbor Discovery [RFC4861] in which an access router may send a router advertisement as a response to the UNA message (see Section 6.4). Other than this extension, the specification does not modify Neighbor Discovery behavior (including the procedures performed when attached to the PAR and when attaching to the NAR).

The protocol does not require changes to any intermediate Layer 2 device between an MN and its access router that supports this specification. This includes the wireless access points, switches, snooping devices, and so on.

8. Evolution from and Compatibility with RFC 4068

This document has evolved from [RFC4068]. Specifically, a new handover key establishment protocol (see [RFC5269]) has been defined to enable a security association between a mobile node and its access router. This allows the secure update of the routing of packets during a handover. In the future, new specifications may be defined to establish such security associations depending on the particular deployment scenario.

The protocol has improved from the experiences in implementing [RFC4068], and from experimental usage. The input has improved the specification of parameter fields (such as lifetime, codepoints, etc.) as well as inclusion of new parameter fields in the existing messages. As of this writing, there are two publicly available implementations, [fmipv6] and [tarzan], and multiple proprietary implementations. Some experience suggests that the protocol meets the delay and packet loss requirements when used appropriately with particular radio access protocols. For instance, see [RFC5184] and [mip6-book]. Nevertheless, it is important to recognize that handover performance is a function of both IP layer operations, which this protocol specifies, and the particular radio access technology itself, which this protocol relies upon but does not modify.

An existing implementation of [RFC4068] needs to be updated in order to support this specification. The primary addition is the establishment of a security association between an MN and its access router (i.e., MN and PAR). One way to establish such a security association is specified in [RFC5269]. An implementation that complies with the specification in this document is likely to also work with [RFC4068], except for the Binding Authorization Data for FMIPv6 option (see Section 6.5.4) that can only be processed when security association is in place between a mobile node and its access router. This specification deprecates the Fast Neighbor Advertisement (FNA) message. However, it is acceptable for a NAR to process this message from a mobile node as specified in [RFC4068].

9. Configurable Parameters

Mobile nodes rely on configuration parameters shown in the table below. Each mobile node MUST have a configuration mechanism to adjust the parameters. Such a configuration mechanism may be either local (such as a command line interface) or based on central management of a number of mobile nodes.

Parameter Name	Default Value	Definition
RTSOLPR_RETRIES	3	Section 6.1.1
MAX_RTSOLPR_RATE	3	Section 6.1.1
FBU_RETRIES	3	Section 6.3.1
PROXY_ND_LIFETIME	1.5 seconds	Section 6.2.2
HI_RETRIES	3	Section 6.2.1

10. Security Considerations

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

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 relies on a companion protocol [RFC5269] to establish such a security association. Using the shared handover key from [RFC5269], the Authenticator in BADF option (see Section 6.5.4) MUST be computed, and the BADF option included in FBU and FBack messages.

Secure FBU, malicious or inadvertent redirection: in this case, the FBU is secured, but the target of binding happens to be an unsuspecting node either due to inadvertent operation or due to 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. In order 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 UNA. Therefore, any inadvertent or malicious redirection to a host is avoided. It is still possible to jam a NAR's buffer with redirected traffic. However, since a NAR's handover state corresponding to an NCoA has a finite (and short) lifetime corresponding to a small multiple of anticipated handover latency, the extent of this vulnerability is arguably small.

Sending an FBU from a NAR's link: A malicious node may send an FBU from a NAR's link providing an unsuspecting node's address as an NCoA. This is similar to base Mobile IP where the MN can provide some other node's IP address as its CoA to its Home Agent; here the PAR acts like a "temporary Home Agent" having a security association with the Mobile Node and providing forwarding support for the handover traffic. As in base Mobile IP, this misdelivery is traceable to the MN that has a security association with the router. So, it is possible to isolate such an MN if it continues to misbehave. Similarly, an MN that has a security association with the PAR may provide the LLA of some other node on NAR's link, which can cause misdelivery of packets (meant for the NCoA) to an unsuspecting node. It is possible to trace the MN in this case as well.

Apart from the above, the RtSolPr (Section 6.1.1) and PrRtAdv (Section 6.1.2) messages inherit the weaknesses of Neighbor Discovery protocol [RFC4861]. Specifically, when its access router is compromised, the MN's RtSolPr message may be answered by an attacker that provides a rogue router as the resolution. Should the MN attach to such a rogue router, its communication can be compromised. Similarly, a network-initiated PrRtAdv message (see Section 3.3) from an attacker could cause an MN to handover to a rogue router. Where these weaknesses are a concern, a solution such as Secure Neighbor Discovery (SEND) [RFC3971] SHOULD be considered.

The protocol provides support for buffering packets during an MN's handover. This is done by securely exchanging the Handover Initiate (HI) and Handover Acknowledgment (HACK) messages in response to the FBU message from an MN. It is possible that an MN may fail, either inadvertently or purposely, to undergo handover to the NAR, which typically provides buffering support. This can cause the NAR to waste its memory containing the buffered packets, and in the worst case, could create resource exhaustion concerns. Hence, implementations must limit the size of the buffer as a local policy configuration, which may consider parameters such as the average handover delay, expected size of packets, and so on.

The Handover Initiate (HI) and Handover Acknowledgement (HACK) messages exchanged between the PAR and NAR MUST be protected using end-to-end security association(s) offering integrity and data origin authentication.

The PAR and the NAR MUST implement IPsec [RFC4301] for protecting the HI and HACK messages. IPsec Encapsulating Security Payload (ESP) [RFC4303] in transport mode with mandatory integrity protection SHOULD be used for protecting the signaling messages. Confidentiality protection of these messages is not required.

The security associations can be created by using either manual IPsec configuration or a dynamic key negotiation protocol such as Internet Key Exchange Protocol version 2 (IKEv2) [RFC4306]. If IKEv2 is used, the PAR and the NAR can use any of the authentication mechanisms, as specified in RFC 4306, for mutual authentication. However, to ensure a baseline interoperability, the implementations MUST support shared secrets for mutual authentication. The following sections describe the Peer Authorization Database (PAD) and Security Policy Database (SPD) entries specified in [RFC4301] when IKEv2 is used for setting up the required IPsec security associations.

10.1. Peer Authorization Database Entries when Using IKEv2

This section describes PAD entries on the PAR and the NAR. The PAD entries are only example configurations. Note that the PAD is a logical concept and a particular PAR or NAR implementation can implement the PAD in any implementation specific manner. The PAD state may also be distributed across various databases in a specific implementation.

PAR PAD:

```
- IF remote_identity = nar_identity_1
  THEN authenticate (shared secret/certificate/EAP) and authorize
  CHILD_SA for remote address nar_address_1
```

NAR PAD:

```
- IF remote_identity = par_identity_1
  THEN authenticate (shared secret/certificate/EAP) and authorize
  CHILD_SAs for remote address par_address_1
```

The list of authentication mechanisms in the above examples is not exhaustive. There could be other credentials used for authentication stored in the PAD.

10.2. Security Policy Database Entries

This section describes the security policy entries on the PAR and the NAR required to protect the HI and HAcK messages. The SPD entries are only example configurations. A particular PAR or NAR implementation could configure different SPD entries as long as they provide the required security.

In the examples shown below, the identity of the PAR is assumed to be `par_1`, the address of the PAR is assumed to be `par_address_1`, and the address of the NAR is assumed to be `nar_address_1`.

PAR SPD-S:

```
- IF local_address = par_address_1 & remote_address =
  nar_address_1 & proto = ICMPv6 & local_icmpv6_type = HI &
  remote_icmpv6_type = HAcK
  THEN use SA ESP transport mode Initiate using IDi = par_1 to
  address nar_address_1
```

NAR SPD-S:

```
- IF local_address = nar_address_1 & remote_address =
  par_address_1 & proto = ICMPv6 & local_icmpv6_type = HAcK &
  remote_icmpv6_type = HI
  THEN use SA ESP transport mode
```

11. IANA Considerations

This document defines a new ICMPv6 message, which has been allocated from the ICMPv6 Type registry.

154 FMIPv6 Messages

This document creates a new registry for the 'Subtype' field in the above ICMPv6 message, called the "FMIPv6 Message Types". IANA has assigned the following values.

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

The values '0' and '1' are reserved. The upper limit is 255. An RFC is required for new message assignment.

The document defines a new Mobility Option that has received Type assignment from the Mobility Options Type registry.

1. Binding Authorization Data for FMIPv6 (BADF) option, described in Section 6.5.4

The document has received Type assignments for the following (see [RFC4068]):

The document defines the following Neighbor Discovery [RFC4861] options that have received Type assignment from IANA.

Type	Description	Reference
17	IP Address/Prefix Option	Section 6.5.1
19	Link-layer Address Option	Section 6.5.2
20	Neighbor Advertisement Acknowledgment Option	Section 6.5.5

The document defines the following Mobility Header messages that have received Type allocation from the Mobility Header Types registry.

1. Fast Binding Update, described in Section 6.3.1
2. Fast Binding Acknowledgment, described in Section 6.3.2

The document defines the following Mobility Option that has received Type assignment from the Mobility Options Type registry.

1. Mobility Header Link-Layer Address option, described in Section 6.5.3

12. Acknowledgments

The editor would like to thank all those who have provided feedback on this specification, but can only mention a few here: Vijay Devarapalli, Youn-Hee Han, Emil Ivov, Syam Madanapalli, Suvidh Mathur, Andre Martin, Javier Martin, Koshiro Mitsuya, Gabriel Montenegro, Takeshi Ogawa, Sun Peng, YC Peng, Alex Petrescu, Domagoj Premec, Subba Reddy, K. Raghav, Ranjit Wable, and Jonathan Wood. Behcet Sarikaya and Frank Xia are acknowledged for the feedback on operation over point-to-point links. The editor would like to acknowledge a contribution from James Kempf to improve this

specification. Vijay Devarapalli provided text for the security configuration between access routers in Section 10. Thanks to Jari Arkko for the detailed AD Review, which has improved this document. 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.

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5269] Kempf, J. and R. Koodli, "Distributing a Symmetric Fast Mobile IPv6 (FMIPv6) Handover Key Using SEcure Neighbor Discovery (SEND)", RFC 5269, June 2008.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

13.2. Informative References

- [fmipv6] "fmipv6.org : Home Page", <<http://fmipv6.org>>.
- [mip6-book] Koodli, R. and C. Perkins, "Mobile Internetworking with IPv6, Chapter 22, John Wiley & Sons.", July 2007.
- [RFC3290] Bernet, Y., Blake, S., Grossman, D., and A. Smith, "An Informal Management Model for Diffserv Routers", RFC 3290, May 2002.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SECure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC4068] Koodli, R., Ed., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.
- [RFC5184] Teraoka, F., Gogo, K., Mitsuya, K., Shibui, R., and K. Mitani, "Unified Layer 2 (L2) Abstractions for Layer 3 (L3)-Driven Fast Handover", RFC 5184, May 2008.
- [tarzan] "Nautilus6 - Tarzan",
<<http://software.nautilus6.org/TARZAN/>>.

Appendix A. Contributors

This document has its origins in the fast handover design team in the erstwhile [mobile ip] working group. The members of this design team in alphabetical order were; Gopal Dommetty, Karim El-Malki, Mohammed Khalil, Charles Perkins, Hesham Soliman, George Tsirtsis, and Alper Yegin.

Appendix B. Changes since RFC 4068

Following are the major changes and clarifications:

- o Specified security association between the MN and its Access Router in the companion document [RFC5269].
- o Specified Binding Authorization Data for Fast Handovers (BADF) option to carry the security parameters used for verifying the authenticity of FBU and FBack messages. The handover key used for computing the Authenticator is specified in companion documents.
- o Specified the security configuration for inter - access router signaling (HI, HAcK).
- o Added a section on prefix management between access routers and illustrated protocol operation over point-to-point links.
- o Deprecated FNA, which is a Mobility Header message. In its place, the Unsolicited Neighbor Advertisement (UNA) message from RFC 4861 is used.
- o Combined the IPv6 Address Option and IPv6 Prefix Option.
- o Added description of DAD requirement on NAR when determining NCoA uniqueness in Section 4, "Protocol Details".
- o Added a new code value for gratuitous HAcK message to trigger a HI message.
- o Added Option-Code 5 in PrRtAdv message to indicate NETLMM usage.
- o Clarified protocol usage when DHCP is used for NCoA formulation (Sections 6.1.2, 3.1, and 5.2). Added a new Code value (5) in PrRtAdv (Section 6.1.2).
- o Clarified that IPv6 Neighbor Discovery operations are a must in Section 7, "Related Protocol and Device Considerations".

- o Clarified "PAR = temporary HA" for FBUs sent by a genuine MN to an unsuspecting CoA.

Editor's Address

Rajeev Koodli
Starent Networks
USA

EMail: rkoodli@starentnetworks.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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

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

Intellectual Property

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

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

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

