

Low-Latency Handoffs in Mobile IPv4

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 IETF Trust (2007).

Abstract

Mobile IPv4 describes how a Mobile Node can perform IPv4-layer handoffs between subnets served by different Foreign Agents. In certain cases, the latency involved in these handoffs can be above the threshold required for the support of delay-sensitive or real-time services. The aim of this document is to present two methods to achieve low-latency Mobile IPv4 handoffs. In addition, a combination of these two methods is described. The described techniques allow greater support for real-time services on a Mobile IPv4 network by minimizing the period of time when a Mobile Node is unable to send or receive IPv4 packets due to the delay in the Mobile IPv4 Registration process.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. The Techniques	5
1.3. L2 Triggers	7
1.4. Requirements Language	9
2. Requirements	9
3. The PRE-REGISTRATION Handoff Method	10
3.1. Operation	11
3.2. Network-Initiated Handoff	13
3.3. Mobile-Initiated Handoff	15
3.4. Obtaining and Proxying nFA Advertisements	17
3.4.1. Inter-FA Solicitation	17
3.4.2. Tunneled nFA Advertisements	18
3.5. Caching Router Advertisements	19

3.6.	Movement Detection, MN, and FA Considerations	19
3.7.	L2 Address Considerations	21
3.8.	Applicability of PRE-REGISTRATION Handoff	21
4.	The POST-REGISTRATION Handoff Method	23
4.1.	Two-Party Handoff	24
4.2.	Three-Party Handoff	28
4.3.	Renewal or Termination of Tunneling Service	34
4.4.	When Will the MN Perform a Mobile IPv4 Registration?	34
4.5.	Handoff Request (HRqst) Message Format	36
4.6.	Handoff Reply (HRply) Message Format	38
4.7.	Handoff to Third (HTT) Message Format	40
4.8.	Applicability of POST-REGISTRATION Handoff Method	40
5.	Combined Handoff Method	41
6.	Layer 2 and Layer 3 Handoff Timing Considerations	42
7.	Reverse Tunneling Support	42
8.	Handoff Signaling Failure Recovery	43
8.1.	PRE-REGISTRATION Signaling Failure Recovery	43
8.1.1.	Failure of PrRtSol and PrRtAdv	43
8.1.2.	Failure of Inter-FA Solicitation and Advertisement	44
8.2.	POST-REGISTRATION Signaling Failure Recovery	44
8.2.1.	HRqst Message Dropped	44
8.2.2.	HRply Message Dropped	45
9.	Generalized Link Layer and IPv4 Address (LLA) Extension	46
9.1.	3GPP2 IMSI Link Layer Address and Connection ID Extension	47
9.2.	3GPP IMSI Link Layer Address Extension	48
9.3.	Ethernet Link Layer Address Extension	49
9.4.	IEEE 64-Bit Global Identifier (EUI-64) Address Extension ..	50
9.5.	Solicited IPv4 Address Extension	51
9.6.	Access Point Identifier Extension	52
9.7.	FA IPv4 Address Extension	53
10.	IANA Considerations	53
10.1.	New Extension Values	53
10.2.	Generalized Link Layer and IP Address Identifier (LLA) ..	54
10.3.	New Message Type and Code	54
11.	Security Considerations	55
12.	Acknowledgements	57
13.	References	57
13.1.	Normative References	57
13.2.	Informative References	58
	Appendix A - Gateway Foreign Agents.....	59
	Appendix B - Low Latency Handoffs for Multiple-Interface MNs.....	60
	Appendix C - PRE_REGISTRATION Message Summary.....	61

1. Introduction

Mobile IPv4 [1] describes how a Mobile Node (MN) can perform IPv4-layer handoff between subnets served by different Foreign Agents (FAs). In certain cases, the latency involved in handoff can be above the threshold required for the support of delay-sensitive or real-time services. The aim of this document is to present two techniques to achieve low-latency Mobile IPv4 handoff during movement between FAs. A further combination of these two techniques is also described. The presented techniques allow greater support for real-time services on a Mobile IPv4 network by minimizing the period of time during which an MN is unable to send or receive IPv4 packets due to the delay in the Mobile IPv4 Registration process. One or more of these techniques may be required to achieve fast Mobile IPv4 handoffs over different wireless technologies (e.g., WLAN, Cellular, WiMAX, Flash-OFDM, etc.). Each wireless technology has different layer 2 handoff procedures, and the best low-latency technique for each scenario should be used to optimize the handoff performance. Further deployment and experimentation are required to determine which technique is best suited to the wireless technologies in terms of implementation and performance. Therefore, the authors encourage further performance measurements and work on low-latency-over-foo specifications in collaboration with the appropriate wireless technology fora to describe the applicability to different wireless layer 2s.

In the rest of this section, terminology used throughout the document is presented, the handoff techniques are briefly described, and the use of link-layer information is outlined. In Section 2, a brief description of requirements is presented. Section 3 describes the details of the PRE-REGISTRATION handoff technique, and Section 4 describes the details of the POST-REGISTRATION handoff technique. In Section 5, a combined method using the two handoff techniques together is presented. Section 6 discusses layer 2 and layer 3 handoff timing considerations. Section 7 discusses reverse tunneling support, Section 8 describes mechanisms to recover from message failures, and Section 9 describes protocol extensions required by the handoff techniques. Sections 10 and 11 discuss IANA and security considerations. Finally, the three appendices discuss additional material related to the handoff techniques. Appendix A gives a short introduction to Regional Registrations [11], which can be used together with low-latency handoffs. Appendix B discusses low-latency handoff when an MN has multiple wireless L2 interfaces, in which case the techniques in this document may not be necessary. Appendix C provides a summary of the messages used in PRE-REGISTRATION.

1.1. Terminology

This section presents a few terms used throughout the document.

oFA - old Foreign Agent (FA), the FA involved in handling the care-of address (CoA) of a Mobile Node (MN) prior to a layer 3 (L3) handoff.

nFA - new Foreign Agent, the FA anticipated to be handling an MN's care-of address after completion of an L3 handoff.

aFA - anchor Foreign Agent, the FA that is currently handling the network end of the tunnel in POST-REGISTRATION.

L2 handoff - Movement of an MN's point of layer 2 (L2) connection from one wireless access point to another.

L3 handoff - Movement of an MN between FAs that involves changing the care-of address at Layer 3 (L3).

L2 trigger - Information from L2 that informs L3 of particular events before and after L2 handoff. The descriptions of L2 triggers in this document are not specific to any particular L2, but rather represent generalizations of L2 information available from a wide variety of L2 protocols.

L2-MT - An L2 trigger that occurs at the MN, informing of movement to a certain nFA (Mobile Trigger).

L2-ST or source trigger - An L2 trigger that occurs at oFA, informing the oFA that L2 handoff is about to occur.

L2-TT or target trigger - An L2 trigger that occurs at nFA, informing the nFA that an MN is about to be handed off to nFA.

L2-LU - An L2 trigger that occurs at the MN or nFA, informing that the L2 link between MN and nFA is established.

L2-LD - An L2 trigger that occurs at the oFA, informing the oFA that the L2 link between MN and oFA is lost.

low-latency handoff - L3 handoff in which the period of time during which the MN is unable to receive packets is minimized.

low-loss handoff - L3 handoff in which the number of packets dropped or delayed is minimized. Low-loss handoff is often called smooth handoff.

seamless handoff - L3 handoff that is both low latency and low loss.

bidirectional edge tunnel (BET) - A bidirectional tunnel established between two FAs for purposes of temporarily routing an MN's traffic to/from it on a new subnet without requiring the MN to change CoA.

ping-pong - Rapid back-and-forth movement between two wireless access points often due to failure of L2 handoff. Ping-pong can occur if radio conditions for both the old and new access points are about equivalent and less than optimal for establishing a good, low-error L2 connection.

network-initiated handoff - L3 handoff in which oFA or nFA initiates the handoff.

mobile-initiated handoff - L3 handoff in which the MN initiates the handoff.

MN or FA identifier - An IPv4 address of an MN or FA, or an L2 identifier that can be resolved to the IPv4 address of an MN or FA. If the identifier is an L2 identifier, it may be specific to the L2 technology.

1.2. The Techniques

Mobile IPv4 was originally designed without any assumptions about the underlying link layers over which it would operate so that it could have the widest possible applicability. This approach has the advantage of facilitating a clean separation between L2 and L3 of the protocol stack, but it has negative consequences for handoff latency. The strict separation between L2 and L3 results in the following built-in sources of delay:

- The MN may only communicate with a directly connected FA. This implies that an MN may only begin the registration process after an L2 handoff to nFA (new FA) has completed.
- The registration process takes some non-zero time to complete as the Registration Requests propagate through the network. During this period of time, the MN is not able to send or receive IPv4 packets.

This document presents techniques for reducing these built-in delay components of Mobile IPv4. The techniques can be divided into two general categories, depending on which of the above problems they are attempting to address:

- Allow the MN to communicate with the nFA while still connected to the oFA.
- Provide for data delivery to the MN at the nFA even before the formal registration process has completed.

The first category of techniques allows the MN to "pre-build" its registration state on the nFA prior to an underlying L2 handoff. The second category of techniques allows for service to continue uninterrupted while the handoff is being processed by the network without requiring the MN's involvement.

Three methods are presented in this document to achieve low-latency L3 handoff, one for each category described above and one as a combination of the two:

- PRE-REGISTRATION handoff method,
- POST-REGISTRATION handoff method, and
- combined handoff method.

The PRE-REGISTRATION handoff method allows the MN to be involved in an anticipated IPv4-layer handoff. The MN is assisted by the network in performing an L3 handoff before it completes the L2 handoff. The L3 handoff can be either network-initiated or mobile-initiated. Accordingly, L2 triggers are used both in the MN and in the FA to trigger particular L3 handoff events. The PRE-REGISTRATION method coupled with L2 mobility helps to achieve seamless handoffs between FAs. The basic Mobile IPv4 concept involving advertisement followed by registration is supported, and the PRE-REGISTRATION handoff method relies on Mobile IPv4 security. No new messages are proposed, except for an extension to the Agent Solicitation message in the mobile-initiated case.

The POST-REGISTRATION handoff method proposes extensions to the Mobile IPv4 protocol to allow the oFA (old FA) and nFA (new FA) to utilize L2 triggers to set up a bidirectional tunnel between oFA and nFA that allows the MN to continue using its oFA while on nFA's subnet. This enables a rapid establishment of service at the new point of attachment, which minimizes the impact on real-time applications. The MN must eventually perform a formal Mobile IPv4 Registration after L2 communication with the new FA is established, but this can be delayed as required by the MN or FA. Until the MN performs registration, the FAs will set up and move bidirectional tunnels as required to give the MN continued connectivity.

The combined method involves running a PRE-REGISTRATION and a POST-REGISTRATION handoff in parallel. If the PRE-REGISTRATION handoff can be performed before the L2 handoff completes, the combined method resolves to a PRE-REGISTRATION handoff. However, if the PRE-REGISTRATION handoff does not complete within an access technology dependent time period, the oFA starts forwarding traffic for the MN to the nFA as specified in the POST-REGISTRATION handoff method. This provides for a useful backup mechanism when completion of a PRE-REGISTRATION handoff cannot always be guaranteed before the L2 handoff completion.

It should be noted that the methods described in this document may be applied to MNs having a single interface (e.g., Wireless LAN interface) or multiple interfaces (e.g., one WLAN and one cellular interface). However, the case of multiply-interfaced MNs needs special consideration, since the handoff methods described in this document may not be required in all cases (see Appendix B).

1.3. L2 Triggers

An L2 trigger is a signal of an L2 event. In this document, the L2 events relate to the L2 handoff process. One possible event is early notice of an upcoming change in the L2 point of attachment of the mobile node to the access network. Another possible event is the completion of relocation of the mobile node's L2 point of attachment to a new L2 access point. This information may come explicitly from L2 in a solicited or unsolicited manner, or it may be derived from L2 messages. Although the protocols outlined in this document make use of specific L2 information, Mobile IPv4 should be kept independent of any specific L2. L2 triggers are an abstraction mechanism for a technology-specific trigger. Therefore, an L2 trigger that is made available to the Mobile IPv4 stack is assumed to be generic and technology independent. The precise format of these triggers is not covered in this document, but the information required to be contained in the L2 triggers for low-latency handoffs is specified.

In order to properly abstract from the L2, it is assumed that one of the three entities -- the MN, oFA, or nFA -- is made aware of the need for an L2 handoff and that the nFA or MN can optionally also be made aware that an L2 handoff has completed. A specific L2 will often dictate when a trigger is received and which entity will receive it. Certain L2s provide advance triggers on the network side, while others provide advance triggers on the MN. Also, the particular timing of the trigger with respect to the actual L2 handoff may differ from technology to technology. For example, some wireless links may provide such a trigger well in advance of the actual handoff. In contrast, other L2s may provide little or no information in anticipation of the L2 handoff.

An L2 trigger may be categorized according to whether it is received by the MN, oFA, or nFA. Table 1 gives such a categorization along with information contained in the trigger. The methods presented in this document operate based on different types of L2 triggers as shown in Table 1. Once the L2 trigger is received, the handoff processes described hereafter are initiated. The three triggers, L2-ST, L2-TT, and L2-MT, are independent of each other and are not expected to occur together since each will trigger a different type of handoff behaviour.

L2 Trigger	Mobile Trigger (L2-MT)	Source Trigger (L2-ST)	
Recipient	MN	oFA	
Method	PRE mobile-initiated	PRE network-initiated	POST source trigger
When?	sufficiently before the L2 handoff so that MN can solicit PrRtAdv from oFA	sufficiently before L2 handoff for FA to send PrRtAdv to MN	sufficiently before L2 handoff for oFA & nFA to exchange HRqst/HRply
Parameters	nFA identifier	nFA identifier, MN identifier	

Table 1 - L2 Trigger
(continued on next page)

L2 Trigger	Target Trigger (L2-TT)		Link-Up (L2-LU)	Link-Down (L2-LD)
Recipient	nFA		MN or nFA	oFA
Method	PRE network-initiated	POST target trigger	PRE & POST	POST
When?	same as source trigger		when radio link between MN & nFA is established	when radio link between MN and oFA is lost
Parameters	oFA identifier MN identifier		@MN: nFA IPv4 or L2 addr. @nFA: MN IPv4 or L2 addr.	MN identifier

Table 1 - L2 Trigger

1.4. Requirements Language

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT" are to be interpreted as described in [2].

2. Requirements

The following requirements are applicable to low-latency handoff techniques and are supported by the methods in this document:

- to provide low-latency and low-loss handoff for real-time services,
- to have no dependence on a wireless L2 technology,
- to support inter- and intra-access technology handoffs, and
- to limit wireless bandwidth usage.

3. The PRE-REGISTRATION Handoff Method

The PRE-REGISTRATION handoff method is based on the normal Mobile IPv4 handoff procedure specified in [1], according to which:

- an advertisement for an FA is received by an MN,
- the advertisement allows the MN to perform movement detection, and
- the MN registers with the FA.

The basic messages specified in [1] are extended to carry information required to achieve fast handoffs. The PRE-REGISTRATION method allows both the MN and FA to initiate the layer 3 handoff and it can make use of L2 triggers on either the FA or MN side, depending on whether network-initiated or mobile-initiated handoff occurs.

PRE-REGISTRATION supports the normal Mobile IPv4 model [1] and optionally also the Regional Registration model [11]. There can be advantages in implementing [11] together with low-latency handoff mechanisms, in particular in cases where the Home Agent (HA) is at a distance (in terms of delay) from the nFA. The time required for the handoff procedure to complete can be reduced by using a closer local HA, called Gateway Foreign Agent (GFA) in [11]. However, implementation of [11] is not required by PRE-REGISTRATION. PRE-REGISTRATION also supports movement where a new Authentication, Authorization, and Accounting (AAA) transaction must occur to authenticate the MN with a new domain.

3.1. Operation

The PRE-REGISTRATION handoff mechanism is summarized in Figure 1.

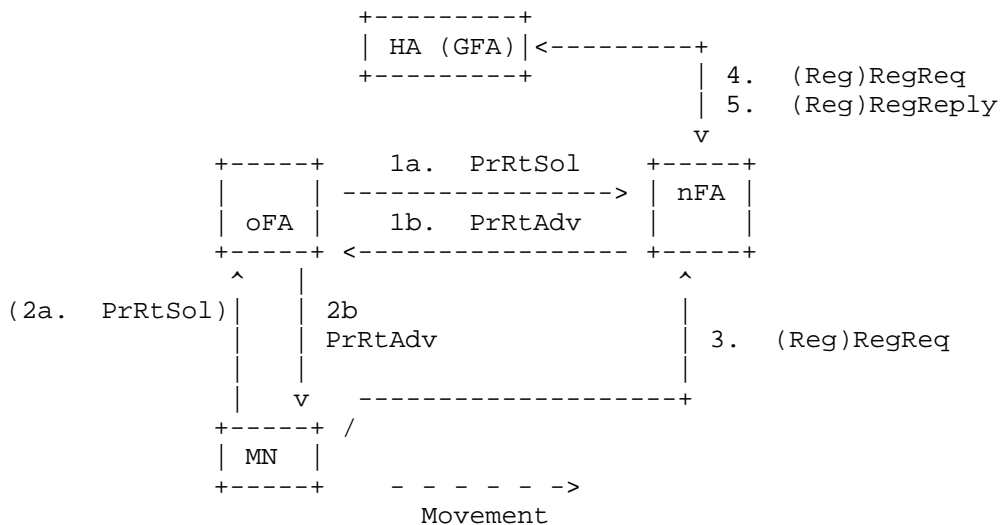


Figure 1 - PRE-REGISTRATION Handoff Protocol

The following steps provide more detail on the protocol:

1. Message 1a is a Proxy Router (Agent) Solicitation (PrRtSol) from oFA to nFA. It is a Mobile IP agent solicitation containing an identifier for the nFA (i.e., IP address or L2 address) in a Generalized Link Layer and IP Address Extension (see Section 9). When message 1a is received by the nFA containing nFA's correct identifier in the LLA extension, the nFA MUST return the Proxy Router Advertisement (Agent Advertisement) in message 1b. Message 1b is simply nFA's Agent Advertisement containing the nFA layer 2 address in a Generalized Link Layer and IP Address (LLA) Extension (see Section 9.3). Messages 1a and 1b SHOULD occur in advance of the PRE-REGISTRATION handoff in order not to delay the handoff. For this to occur, oFA SHOULD solicit and cache advertisements from neighboring nFAs using messages 1a and 1b, thus decoupling the timing of this exchange from the rest of the PRE-REGISTRATION handoff. When the L3 handoff is initiated by a target L2 trigger at nFA (L2-TT), message 1b equals message 2b and is sent unsolicited directly to MN (tunneled by nFA to MN through oFA) instead of being relayed by oFA.

2. Message 2a is a Proxy Router Solicitation (PrRtSol) from MN to oFA. It is different from a normal Router (Agent) Solicitation since it is soliciting an advertisement from a router different from the one receiving this message. It is a Mobile IP Agent Solicitation containing an identifier for the nFA (i.e., IP address or L2 address) in a Generalized Link Layer and IP Address Extension (see Section 9). The presence of message 2a indicates that the handoff is mobile-initiated and its absence means that the handoff is network-initiated. In mobile-initiated handoff, message 2a occurs if there is an L2 trigger in the MN to solicit for a Proxy Router Advertisement (PrRtAdv). When message 2a is received by the oFA, it MUST return the Proxy Router Advertisement (Agent Advertisement) in message 2b. This is simply nFA's Agent Advertisement containing the nFA layer 2 address in a Generalized Link Layer and IP Address (LLA) Extension (see Section 9.3). In network-initiated source-triggered handoff, the L2 trigger occurs at oFA, and oFA MUST relay the Agent Advertisement in message 2b without the need for the MN to solicit. Note that it is also possible for nFA to advertise directly to the MN in the network-initiated target-triggered case (see Section 3.2).
3. The MN performs movement detection upon receipt of a solicited or unsolicited Agent Advertisement and, if required, it sends a Registration Request (RegReq) message [1] in message 3 to nFA. When a local Gateway Foreign Agent (GFA) is present, this message can optionally be a Regional Registration Request (RegRegReq) [11]. Message 3 is routed through oFA since the MN is not directly connected to nFA prior to the L2 handoff.
4. Messages 4 and 5 complete the standard Mobile IPv4 Registration [1] or optionally Regional Registration [11] initiated with message 3. The Registration Request MUST contain the MN's layer 2 address in a Generalized Link Layer and IP Address Extension (see Sections 3.7 and 9). This identifier may be a plain Ethernet address or an identifier specific to the wireless technology. If the MN is not already connected to nFA, the Registration Reply in message 5 MUST be buffered by the nFA and unicast to the MN on-link as soon as the MN connects to nFA (i.e., L2-LU trigger at nFA, which can be implemented by the MN sending an Agent Solicitation or optionally using special layer 2 techniques, which are outside the scope of this document). This is necessary since the MN may have to detach from oFA, due to the wireless L2 connection, before it receives the reply. The MN's L2 address is obtained using the extensions in Section 9, as described in Section 3.7. Figures 2 and 3 illustrate this procedure.

5. If the registration is successful, packets for the MN are tunneled from the HA (or GFA) to the nFA and then to the MN.

PRE-REGISTRATION is not dependent on [11]. However, if the HA is at a distance (in terms of delay) from the nFA, the use of a local GFA may reduce the time required for the handoff procedure to complete.

The time at which the L2 trigger is received by the oFA or MN, thereby triggering the PRE-REGISTRATION handoff, compared to the time at which the actual L2 handoff occurs is important for the optimal performance of the low-latency handoff. That is, in the optimal case, the L2 trigger will be received and the four messaging steps of PRE-REGISTRATION described above will be completed (i.e., up to when the Registration Request is processed by HA or GFA) before the MN moves. Optimally, the Registration Reply and the first packet redirected by the HA (or GFA) to nFA will reach the MN at the moment in which the MN's L2 link to nFA is fully established. The MN would therefore not suffer any disruption due to the L3 handoff. This cannot always be guaranteed unless particular implementation techniques are used. To alleviate a part of this timing problem, the MN MAY set the S bit [1] in low-latency Registration Requests sent by the MN. This allows the MN to receive packets at both oFA and nFA during the short layer 2 handoff time. Other techniques may be required, such as L2 techniques or buffering, but these are outside the scope of this document. In addition, further handoff smoothing considerations may be required to prevent the loss of packets in-flight between HA (or GFA) and oFA while the MN performs a PRE-REGISTRATION handoff. These are also outside the scope of this document.

Figures 2, 3, and 4 contain message timing diagrams for the network-initiated and mobile-initiated PRE-REGISTRATION handoff procedures.

3.2. Network-Initiated Handoff

As described in Table 1, a PRE-REGISTRATION handoff can be initiated at oFA by a source trigger or at nFA by a target trigger. Figures 2 and 3 contain message timing diagrams for PRE-REGISTRATION network-initiated handoff for source and target triggers.

A source-triggered, network-initiated handoff occurs when an L2 trigger is received at the oFA informing it of a certain MN's upcoming movement from oFA to nFA. The L2 trigger contains information including the MN's identifier (i.e., the IPv4 address itself or an identifier that can be resolved to the IPv4 address) and the nFA's identifier. An identifier may be an IPv4 address or something specific to the wireless technology (e.g., Base Station or Access Point Identifier). A target-triggered, network-initiated

handoff occurs when an L2 trigger is received at the nFA informing it of a certain MN's upcoming movement from oFA. This type of trigger is also shown in Table 1 and contains information including the MN's and the oFA's identifier.

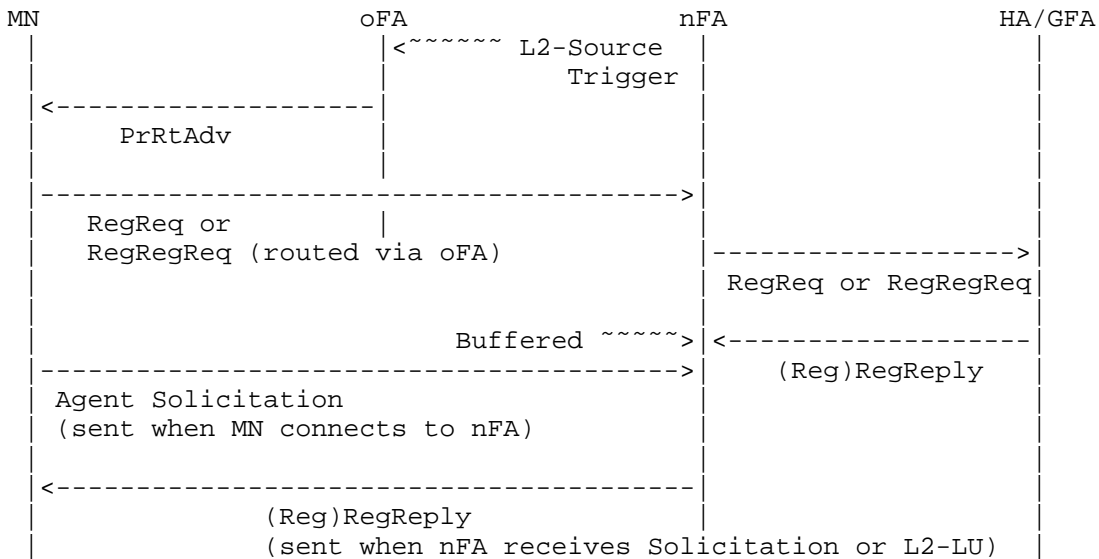


Figure 2 - PRE-REGISTRATION Handoff Message Timing Diagram (Network-Initiated, Source Trigger)

In a source-triggered handoff, when oFA receives the trigger (L2-ST), it MUST send message 2b, the Proxy Router Advertisement (PrRtAdv), to the MN. The PrRtAdv is nFA's Agent Advertisement [1] with one of the link-layer extensions described in Section 9. The use of the contents of this extension is described in Section 3.7. Messages 1a and 1b SHOULD be exchanged by oFA and nFA before the L2 trigger is received (see Section 3.4.1). Message 2a is not used.

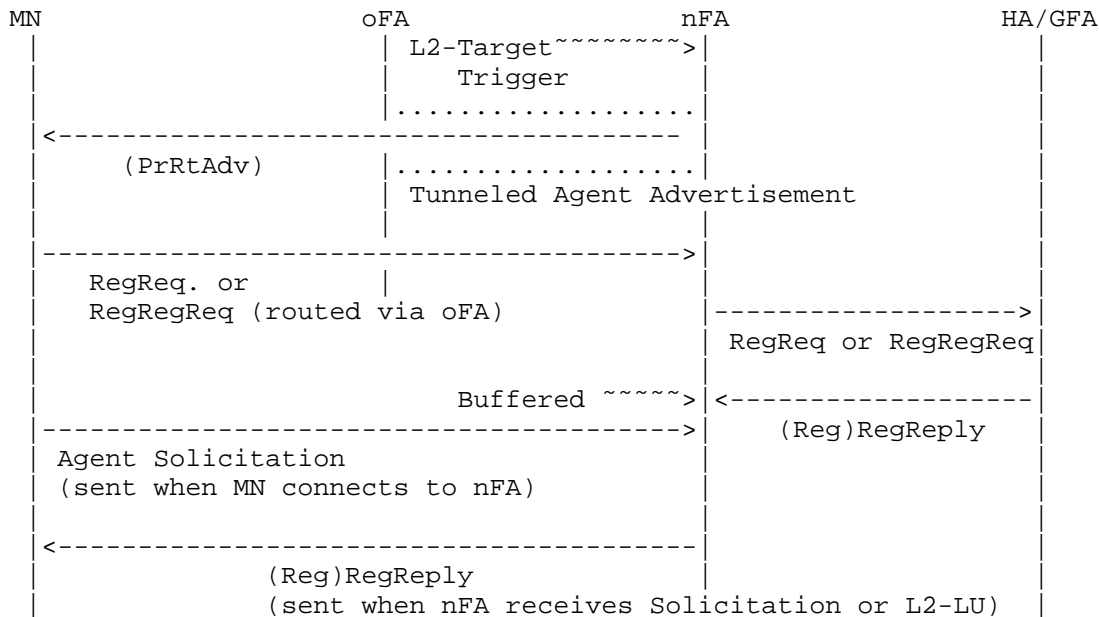


Figure 3 - PRE-REGISTRATION Handoff Message Timing Diagram (Network-Initiated, Target Trigger)

In a target-triggered handoff, when nFA receives the trigger (L2-TT), it MUST tunnel an Agent Advertisement to the MN through oFA to initiate the L3 handoff. The inner advertisement is unicast by nFA to MN, thus nFA treats the target trigger as a Router (Agent) Solicitation. This advertisement is tunneled to oFA, which functions as a normal router, decapsulating the advertisement and forwarding it to the MN. This message MUST be authenticated to prevent attacks (see Section 3.4.2).

3.3. Mobile-Initiated Handoff

As shown in Table 1, a mobile-initiated handoff occurs when an L2 trigger is received at the MN informing that it will shortly move to nFA. The L2 trigger contains information such as the nFA's identifier (i.e., nFA's IPv4 address or an identifier that can be resolved to the nFA's IPv4 address). As an example, a Wireless LAN MN may perform a scan to obtain the Base Station Identifier (BSSID) of the access point that is a potential handoff target (i.e., its signal is becoming stronger). The message timing diagram is shown in Figure 4.

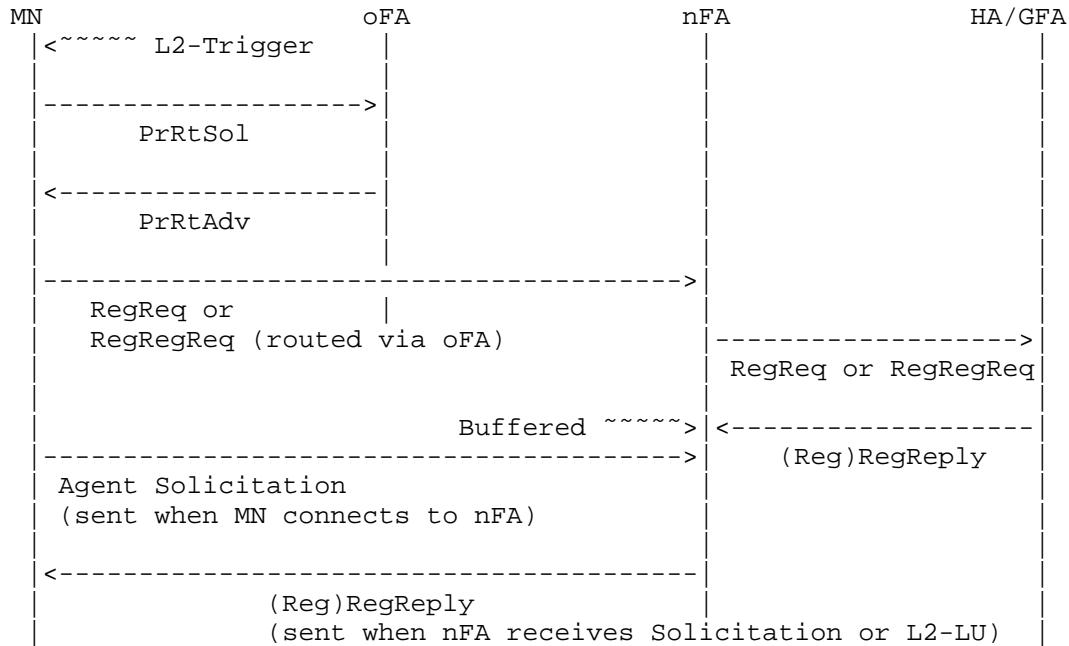


Figure 4 - PRE-REGISTRATION Handoff Message Timing Diagram
(Mobile-Initiated)

As a consequence of the L2 trigger (L2-MT), the MN MUST send message 1a, the Proxy Router Solicitation (PrRtSol). This message is a unicast Agent Solicitation to oFA for a Proxy Router Advertisement (PrRtAdv). This solicitation MUST have a TTL=1 as in [1]. The Proxy Router Advertisement Solicitation unicast to oFA is an Agent Solicitation with a special extension. The solicitation MUST have an extension containing an FA identifier (i.e., IPv4 address or L2 address contained in an LLA extension, see Section 9) because the MN is soliciting another specific FA's advertisement from the oFA. This specific FA will be the MN's nFA. The identifier is the IPv4 address of the nFA or another identifier that can be used by the oFA to resolve to nFA's IPv4 address. If the identifier is not an IPv4 address, it MAY be specific to the underlying wireless technology, for example, an access point or Base Station Identifier (e.g., WLAN BSSID) that can be mapped by oFA to the nFA IPv4 address as described in Section 3.4.1. The extension containing the identifier is a sub-type of the Generalized Link Layer Address Extension described in Section 9.

Two extension sub-types have been defined to contain the nFA's IPv4 address and an access point identifier. They are called the Solicited Agent IPv4 Address Extension and the Access Point

Identifier Extension, and are described in Sections 9.5 and 9.6. These two extensions SHOULD NOT be present in the same PrRtSol message.

When oFA receives the PrRtSol message, it MUST reply to the MN with the Proxy Router Advertisement (PrRtAdv, message 2b). The PrRtAdv is simply the Agent Advertisement for the requested nFA, proxied by oFA. In order to expedite the handoff, the actual nFA advertisement SHOULD be cached by the oFA following a previous exchange with nFA, shown in messages 1a and 1b, as specified in Section 3.5. The PrRtAdv message MUST contain the nFA's L2 address (using the LLA extension in Section 9.3). This is further described in Section 3.7.

3.4. Obtaining and Proxying nFA Advertisements

Since L2 triggers are involved in initiating PRE-REGISTRATION handoff, the trigger timing SHOULD be arranged such that a full L3 PRE-REGISTRATION handoff can complete before the L2 handoff process completes. That is, the L2 handoff should be completed after the MN's registration with the nFA is performed (message 3 in Figure 1). The registration MAY be transmitted in more than one copy (default recommendation: 2) to reduce the probability that it is lost due to errors on the wireless link. This would not apply to reliable wireless links where retransmissions are performed at layer 2 in case of error to guarantee packet delivery.

A PRE-REGISTRATION handoff in this case requires the MN to receive an Agent Advertisement from the nFA through the old wireless access point. How to achieve this is discussed in the following subsections. Messages exchanged between FAs MUST be authenticated to prevent impersonation attacks. The minimal requirement is that all FAs involved in low-latency handoffs MUST support manual pre-configuration of security associations with other neighboring FAs, involving shared keys and the default algorithms of [1] (see the Security Considerations of this document).

3.4.1. Inter-FA Solicitation

This applies to the network-initiated source-triggered (L2-ST) and mobile-initiated (L2-MT) cases only. Inter-FA solicitation assumes that oFA has access to the IPv4 address of the nFA. The IPv4 address of nFA is obtained by means of an L2 trigger at oFA in the network-initiated case (see Section 3.2) or by means of the extension to the Proxy Router Solicitation (PrRtSol) from the MN in the mobile-initiated case (see Section 3.3). This extension to the PrRtSol may contain an IPv4 address or another identifier, for example, an identifier of a Wireless Base Station such as the WLAN BSSID. In the latter case, the oFA must implement a mechanism to resolve the Base

Station Identifier to an IPv4 address. The default mechanism is to use a configured table of neighboring Base Station Identifiers (e.g., BSSID) to FA IPv4 address mappings in each FA. Other automated discovery mechanisms may also be used.

If oFA does not cache advertisements (see Section 3.5) once it receives an L2 trigger and obtains the address of the nFA for a specific MN, it MUST send a unicast Agent Solicitation (PrRtSol) to nFA. The nFA replies to the oFA by unicasting an Agent Advertisement with appropriate extensions (PrRtAdv). This method removes the TTL limitation of [1] for Mobile IPv4 messages (i.e., TTL=1 is not applicable here). The TTL limitation cannot be applied since oFA and nFA may be more than one hop away and since it is unnecessary for a secured unicast message. The ICMP solicitations and advertisements MUST be authenticated and integrity protected. These messages MUST be protected using Encapsulating Security Payload (ESP) [10] to prevent attacks (see the Security Considerations section of this document). An FA MUST NOT accept ICMP solicitations or advertisements from sources that are not authenticated.

As a practical matter, oFA SHOULD pre-solicit and cache advertisements from known neighboring FAs (see section 3.5) to avoid performing the solicitation during an actual handoff procedure.

3.4.2. Tunneled nFA Advertisements

This applies to the network-initiated target-triggered (L2-TT) case only. Following a target trigger (L2-TT) the nFA MUST send a tunneled Agent Advertisement to the MN through oFA. Tunneling nFA advertisements assumes that the nFA is aware of the IPv4 address for oFA and the MN. These IPv4 addresses are obtained by means of the FA and MN identifiers contained in an L2 trigger received at nFA in the network-initiated case (see Section 3.2). However, in [1] the TTL must be 1 on Agent Advertisements from the nFA. Therefore, tunneling advertisements is applicable if the TTL limitation of [1] is relaxed. For this purpose, a pre-established security association between oFA and nFA MUST be in place to authenticate this message and relax the TTL limitation. If the implementation requires this, a tunnel SHOULD be configured when the inter-FA security association is established. The tunneled ICMP advertisement MUST be secured using tunnel mode ESP [10] between nFA and oFA. An FA MUST NOT accept tunneled ICMP packets destined to it from sources that are not authenticated.

3.5. Caching Router Advertisements

In the mobile-initiated (L2-MT) case and the network-initiated source-triggered (L2-ST) case, the message exchange 1 in Figure 1 could impose an additional latency on the L3 handoff process if done as part of the handoff procedure. In order to remove this source of latency, the inter-FA Router (Agent) Solicitation and Advertisement exchange SHOULD be performed in advance of handoff. A process SHOULD be in place at the oFA to solicit its neighboring nFAs at a predefined time interval (MIN_SOLICITATION_INTERVAL). This interval SHOULD NOT be set too small to avoid unnecessary consumption of network bandwidth and nFA processing resources. The minimum value of MIN_SOLICITATION_INTERVAL is 1 second. If the FA Challenge/Response mechanism in [7] is used, then the MIN_SOLICITATION_INTERVAL MUST be set to a value smaller than the window of time in which a challenge remains valid so that the nFA challenge does not expire before the MN issues the Registration Request. Therefore, the recommended default value for the MIN_SOLICITATION_INTERVAL in oFA is $(0.5 * \text{nFA's CHALLENGE_WINDOW} * \text{nFA's Agent Advertisement interval})$. The CHALLENGE_WINDOW and Agent Advertisement interval are defined in [7] and [1] respectively. The minimum requirement is that the MIN_SOLICITATION_INTERVAL MUST be manually configurable, while possible autoconfiguration mechanisms are outside the scope of this document. To allow advertisement caching in certain implementations and in cases where the nFA advertisement interval is very small, it MAY be necessary for the implementation in nFA to allow different CHALLENGE_WINDOW and Agent Advertisement interval settings for its nFA-oFA interface.

The oFA SHOULD cache the most recent advertisement from its neighboring nFAs. This advertisement MUST be sent to the MN in message 2b with a TTL=1. The oFA SHOULD also have a mechanism in place to create a list of neighboring nFAs. The minimum requirement for each FA is that it SHOULD allow manual configuration of a list of nFA addresses that an MN could possibly perform an L3 handoff to. The FA addresses in this list will depend on deployment and radio coverage. It is also possible to specify another protocol to achieve nFA discovery, but this is outside the scope of this document.

3.6. Movement Detection, MN, and FA Considerations

When the MN receives an Agent Advertisement with a Mobility Agent extension, it performs actions according to the following movement detection mechanism: the MN SHOULD be "Eager" to perform new bindings. This means that the MN SHOULD perform registrations with any new FA from which it receives an advertisement (i.e., MN is Eager), as long as there are no locally-defined policies in the MN that discourage the use of the discovered FA. For example, the MN

could have a policy based on the cost of service. The method by which the MN determines whether the FA is a new FA is described in [1] and MAY use an FA-NAI extension [11]. By being "Eager" to perform registrations, the MN reduces latency times.

The MN also needs to change its default router from oFA to nFA. The MN MUST change its default router to nFA as soon as the PRE-REGISTRATION procedure has completed (i.e., Registration Reply is received by MN) as described in [1].

Overall, the MN behaves as described in [1] with the following changes: the specified movement detection mechanism mentioned above and the ability to use the L2-MT to initiate an Agent Solicitation with a special extension (PrRtSol). Also, when the MN receives an L2-LU trigger (i.e., new interface or link is up), it MUST immediately send an Agent Solicitation [1] on that interface. An nFA that receives an Agent Solicitation [1] will use it as an L2-LU trigger event, and according to [1] it will record the MN's IPv4/layer 2 addresses (i.e., the Address Resolution Protocol (ARP) entry). At that point, the nFA starts delivering data to the MN including the previously buffered Registration Reply. The nFA MAY also use other L2 mechanisms to detect earlier that the MN has attached to the new link and to start forwarding data to it. The MN SHOULD NOT attempt to retransmit a low-latency Registration Request (i.e., Registration Request containing an LLA extension described in Section 9.) when it does not receive the Registration Reply.

When moving from a PRE-REGISTRATION network to a normal Mobile IPv4 [1] network, the MN will no longer receive PrRtAdv messages (i.e., Agent Advertisements with the LLA extension). If the MN still receives L2-MTs, it will attempt to send PrRtSol messages. The normal FA will reply with a normal Agent Advertisement [1]. If the MN does not receive a PrRtAdv in reply to its PrRtSol, it MAY retransmit the PrRtSol message once after PRE_SOL_INTERVAL seconds and then for another PRE_SOL_ATTEMPTS times with exponential backoff of the transmission interval. If a PrRtAdv is not received within PRE_SOL_INTERVAL seconds after the last PrRtSol attempt, the MN MUST stop sending PrRtSol messages until after a registration with a new FA is performed. The default value for PRE_SOL_ATTEMPTS is 2, and for PRE_SOL_INTERVAL, it is 1 second. It should be noted that the performance of the movement detection mechanism mandated in PRE-REGISTRATION (i.e., eager to register) may have sub-optimal behaviour in a standard Mobile IPv4 [1] network. Therefore, standard movement detection mechanisms [1] should be used in plain Mobile IPv4 networks. Instead, when the MN moves from a normal Mobile IPv4 [1] network to a PRE-REGISTRATION network, the MN starts receiving L2-MT triggers or PrRtAdv messages. When the MN receives L2-MT triggers or PrRtAdv messages, it SHOULD follow the PRE-REGISTRATION procedure.

If there is uncertainty as to which mode to choose (e.g., MN receives messages from both PRE-REGISTRATION and normal FAs), the MN decides based on its registration status with the current FA. If the MN already has a valid normal Mobile IPv4 Registration [1] with the advertising FA, it SHOULD give priority to the PRE-REGISTRATION procedure. Otherwise it SHOULD give priority to normal Mobile IPv4 [1] Registration procedure. The MN SHOULD NOT attempt to perform PRE-REGISTRATION and standard Mobile IPv4 [1] Registrations in parallel.

3.7. L2 Address Considerations

Some special considerations should be taken with respect to the wireless system on which this handoff method is being implemented. Consider an Ethernet-like system such as IEEE 802.11, for example. In PRE-REGISTRATION, the MN is registering with an FA (nFA) that is not its current first-hop router; therefore, the L2 address of the Ethernet frame containing the MN's Registration Request reaching the nFA is not the MN's address. Therefore, the FA MUST NOT use the Ethernet address of the incoming Registration Request as the MN's L2 address as specified in [1]. This applies to the cases where the wireless access points are bridges or routers and independently of whether the FA is implemented in the wireless access points themselves. In this case, the MN's Registration Request (or Regional Registration Request) MUST use an L2 address extension to the registration message. Such an L2 address is either the same L2 address that remains constant as the MN moves, or it is the MN's L2 address at nFA. To communicate its L2 address, the MN includes a Generalized Link Layer and IP Address Extension (see Section 9) with its Registration Request (or Regional Registration Request) message. If this extension is present, the FA MUST use the L2 address contained in the extension to communicate with the MN. If a particular wireless L2 technology has defined a special interface to the wireless network that allows the FA to resolve the mapping between an MN's IPv4 address and its L2 address without the need to use the extension, the L2 address extension contents may be discarded. For the same reasons above, the MN MUST NOT use the source L2 address of the Agent Advertisement message (PrRtAdv) as its default router's L2 address. Therefore, the nFA MUST include a Generalized Link Layer and IP Address Extension (see Section 9.3) with its Agent Advertisement (PrRtAdv) messages.

3.8. Applicability of PRE-REGISTRATION Handoff

The PRE-REGISTRATION handoff method is applicable to scenarios where a period of service disruption due to layer 3 is not acceptable, for example, when performing real-time communications, and therefore where an anticipation of the layer 3 handoff is required. Security

for the PRE-REGISTRATION handoff method is based on the same security model as [1] including the use of AAA. A prerequisite for PRE-REGISTRATION is that the FA or MN is able to obtain an L2 trigger informing it of a pending L2 handoff procedure. The target of the L2 handoff is another access point or radio network that is in the coverage area of a new FA. The L2 trigger information may be in the form of identifiers that need to be resolved to IPv4 addresses using methods that may be specific to the wireless network and are not considered here. If, for example, the oFA or MN determines that the IPv4 address of the new FA matches oFA's address, then the PRE-REGISTRATION handoff SHOULD NOT be initiated.

The L2 trigger must allow enough time for the PRE-REGISTRATION handoff procedure to be performed. In many wireless L2 technologies, the L2 handoff procedure involves a number of message exchanges before the effective L2 handoff is performed. For such technologies, PRE-REGISTRATION handoff can be initiated at the beginning of the L2 handoff procedure and completed before the L2 handoff is completed. It is efficient to engineer the network such that this succession of events is ensured.

The PRE-REGISTRATION handoff method is applicable in the following cases:

- when the MN has locally defined policies that determine a preference for one access over another, for example, due to service cost within the same or different technology, and therefore where it is necessary to allow the MN to select the appropriate FA with which to connect.
- when L2 security between the MN and the FA is either not present or cannot be relied upon to provide adequate security.
- when the trigger to initiate the handoff is received at the MN.

In the first case, it is necessary to involve eventual local MN policies in the movement detection procedure as described in Section 3.6.

4. The POST-REGISTRATION Handoff Method

The POST-REGISTRATION handoff method uses bidirectional edge tunnels (BETs) or unidirectional tunnels to perform low-latency change in the L2 point of attachment for the MN without requiring any involvement by the MN. Figure 5 illustrates the basic POST-REGISTRATION handoff.

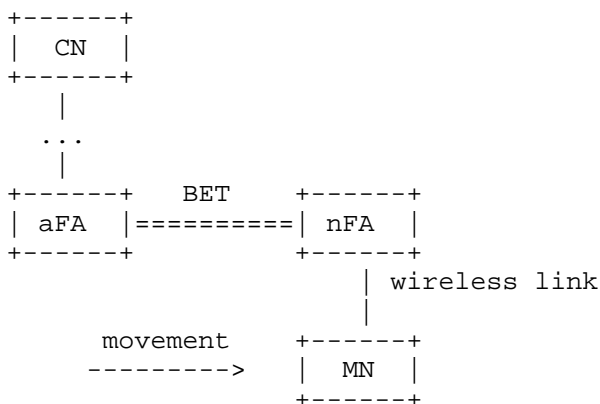


Figure 5 - POST-REGISTRATION Concept

Following a successful Mobile IPv4 Registration between MN and oFA, the oFA becomes the mobility anchor point for the MN, called the anchor FA (aFA). When the MN moves from oFA to nFA, rather than performing signaling over the wireless link to register with the nFA, the MN can defer the L3 handoff and continue to use its aFA (i.e., oFA in this case). If the MN moves to a third FA before registering with the nFA, in certain cases described later, the third FA signals aFA to move the wireless link end of the BET from nFA to it. The network end of the BET remains anchored at aFA until the MN performs the Mobile IPv4 Registration.

Messages between oFA/aFA and nFA MUST be authenticated. The minimal requirement is that all FAs involved in low-latency handoffs MUST support manual pre-configuration of security associations with other neighboring FAs, involving shared keys and the default algorithms of [1]. POST-REGISTRATION FAs MUST implement the inter-FA authentication extension (FA-FA authentication extension) specified in [11] and MAY additionally use other security mechanisms.

4.1. Two-Party Handoff

Two-party handoff occurs when the MN moves from oFA to nFA. Normally, this movement would result in a new Mobile IPv4 Registration at nFA. However, in POST-REGISTRATION, the MN and nFA MAY delay this but maintain connectivity using the BET (or alternatively unidirectional tunnel) between oFA and nFA. The protocol is shown in Figure 6.

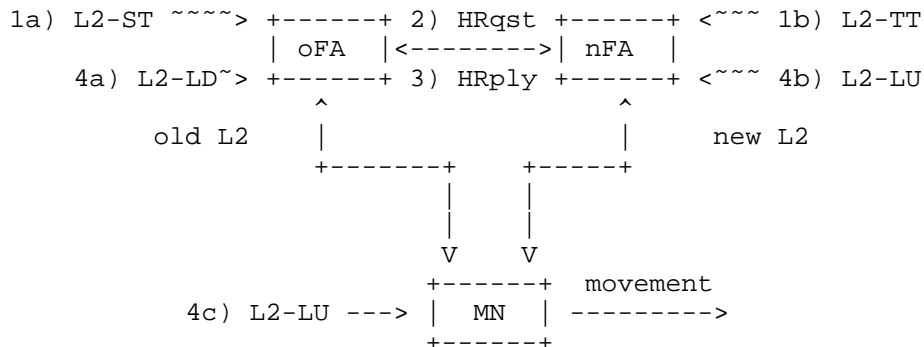


Figure 6 - Two-Party Handoff (POST-REGISTRATION)

The following describes the progress of a two-party handoff. The numbered items refer to steps in Figure 6. The source-triggered HRqst/HRply message for tunnel creation, the target-triggered HRqst/HRply message for tunnel creation, and the HRqst/HRply to extend or terminate a BET (or unidirectional tunnel) are identified using the suffixes (s), (t), and (r), respectively.

- 1) Either the oFA or nFA receives an L2 trigger informing it that a certain MN is about to move from oFA to nFA. The two cases are:
 - a) The L2 trigger is a source trigger (L2-ST) at oFA. The trigger contains the MN's L2 address and an identifier for the nFA (the IPv4 address itself or an L2 address that can be resolved to the IPv4 address of the nFA).
 - b) The L2 trigger is a target trigger (L2-TT) at nFA. The trigger contains the MN's L2 address and an identifier for the oFA (the IPv4 address itself or an L2 address that can be resolved to the IPv4 address of the oFA).
- 2) The FA receiving the trigger sends a Handoff Request (HRqst) to the other FA. There are two cases:

- a) If oFA is sending the HRqst, the H bit is set and the N bit is unset, indicating it is an HRqst(s). The HRqst(s) contains the lifetime of the tunnel the oFA is willing to support, the MN's IPv4 home address, the MN's HA address, and an LLA option with the MN's L2 address. If the lifetime is zero and the T bit is not set, the oFA is not willing to tunnel any packets for MN. A positive lifetime and a set T bit indicate that the oFA is willing to tunnel for the indicated time. Section 4.5 describes the HRqst(s) and Section 9 describes the LLA option.
 - b) If nFA is sending the HRqst, the N bit is set and the H bit is unset, indicating that it is an HRqst(t). If the T bit is set, nFA has requested a reverse tunnel and the HRqst(t) contains the lifetime of the tunnel the nFA is requesting. The HRqst(t) also contains an LLA option with the MN's L2 address. The MN's IPv4 home address and HA address are not sent, unless they are discovered by some means outside the scope of this document (for example, as part of the L2-TT). Section 4.5 describes the HRqst(t).
- 3) The FA receiving the HRqst sends a Handoff Reply (HRply) to the other FA. There are two cases:
- a) If oFA is sending the HRply, the N bit is set and the H and R bits are unset, indicating that the reply is in response to a HRqst(t), i.e., it is an HRply(t). If the T bit is set, the HRply(t) contains the tunnel lifetime the oFA is willing to provide; otherwise, the tunnel lifetime is set to zero indicating that the oFA is not willing to provide tunnel service. If both HRply(t) and HRqst(t) have the T bit set and non-zero lifetime, a BET is established. The HRply(t) also contains the MN's home subnet IPv4 address, the MN's HA address, and an LLA option containing the MN's L2 address. Section 4.6 describes the HRply(t).
 - b) If nFA sends the HRply, the H bit is set and the N and R bits are unset, indicating that this is a response to HRqst(s), i.e., it is an HRply(s). If the T bit is set, the nFA indicates that it requests a reverse tunnel, and the lifetime field is set with the requested tunnel lifetime. The T bit can be set in HRply only if the oFA had set the T bit in the corresponding HRqst or if the nFA is required to reverse tunnel incoming packets to oFA because ingress filtering is enabled on its network. This establishes a BET. The tunnel lifetime requested by the nFA must be less than or equal to the tunnel lifetime offered by oFA in the HRqst(s). Section 4.6 describes the HRply(s).

- 4) The point during the L2 handoff in which the MN is no longer connected on a given link is signaled by an L2-LD trigger at oFA and MN. Completion of L2 handoff is signaled by an L2-LU trigger at nFA and MN. The trigger is handled as follows:
 - a) When oFA receives the L2-LD trigger, it begins forwarding MN-bound packets through the forward tunnel to nFA.
 - b) When the nFA receives the L2-LU trigger, it begins delivering packets tunneled from oFA to MN and forwards outbound packets from MN using normal routing mechanisms or through a reverse tunnel to oFA or HA. The nFA at this point may not yet be the default router of the MN (see Section 4.4); therefore, to receive all outbound packets from the MN the nFA must send a unicast proxy ARP message (used in [1]) to the MN upon receiving an L2-LU trigger. This proxy ARP message is an ARP Reply [5] sent by the nFA on behalf of oFA, therefore supplying the nFA link-layer address in the Sender Hardware Address field and the oFA IPv4 address in the Target Protocol Address field.
 - c) When the MN receives the L2-LU, it MAY initiate the Mobile IPv4 Registration process by soliciting an Agent Advertisement as described in [1]. If the registration is successful, the nFA takes over the role of anchor FA (aFA) from the oFA. Alternatively, the MN MAY defer the Mobile IPv4 Registration (see Section 4.4).
- 5) The oFA becomes an aFA if the MN moves to a third FA before having performed a Mobile IPv4 Registration with nFA.
- 6) Should L2 handoff fail in Step 4 (due to L2 reasons) and a ping-pong situation arise, the oFA may be able to determine this case through the trigger mechanism (i.e., FA sees successive L2-ST/L2-TT followed by L2-LD and then L2-LU). The FA that originated the HRqst can in this case cancel the tunnel by sending an HRqst(r) to the other FA with lifetime zero. It will then simply continue delivering packets to MN exactly as if no handoff had been pending. Section 4.5 describes the HRqst(r).

If the oFA sets the B bit in HRqst/HRply and the nFA has not requested a reverse tunnel by setting the T bit, the nFA SHOULD tunnel outgoing packets from the MN to the HA because the MN has requested this service from the oFA. The nFA SHOULD offer this service only if no security between the nFA and the MN's HA is required, or if there is an existing nFA-HA security association.

The actual timing of BET or unidirectional tunnel placement depends on the available L2 triggers. The forward tunnel from oFA to nFA is constructed using one of the tunneling procedures described in [1] for the HA to FA tunnel with the difference that the ends of the tunnel are at the oFA and nFA, respectively. The reverse tunnel from nFA to oFA is constructed as described in [3] with the difference that the network end of the tunnel is at the oFA instead of the HA. If both forward and reverse tunnels are established, then a BET has been established. With optimal L2 trigger information, as described above, the FAs can set up the BET immediately when the L2 handoff is initiated, start tunneling MN-bound data when the link to the MN goes down, and the nFA can use the link-up trigger to start delivering packets. In the absence of optimal L2 trigger information, the HRply can act as the trigger to start tunneling MN-bound data, but in this case, the period of packet delivery disruption to the MN could still be present and additional measures may be required to provide uninterrupted service. Particular implementation and deployment scenarios could require techniques to smooth the handoff by providing a means to convey packets arriving during the L2 handoff. The exact techniques are outside the scope of this document.

Figures 7 and 8 show timing diagrams for source trigger (L2-ST) and target trigger (L2-TT) two-party handoffs, respectively.

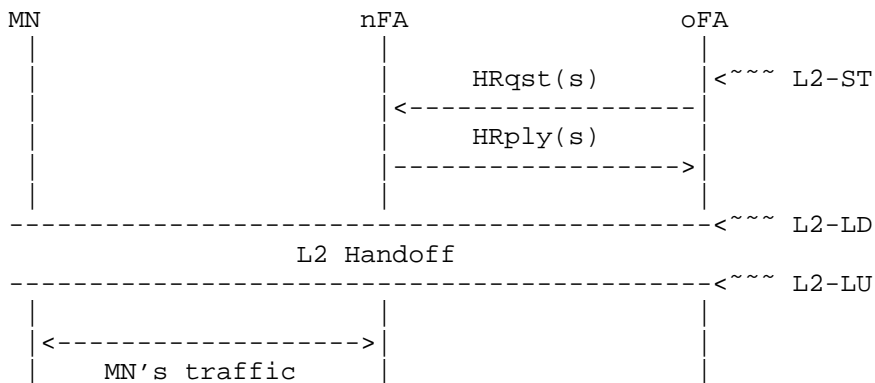


Figure 7 - Two-Party Source Trigger Handoff Timing

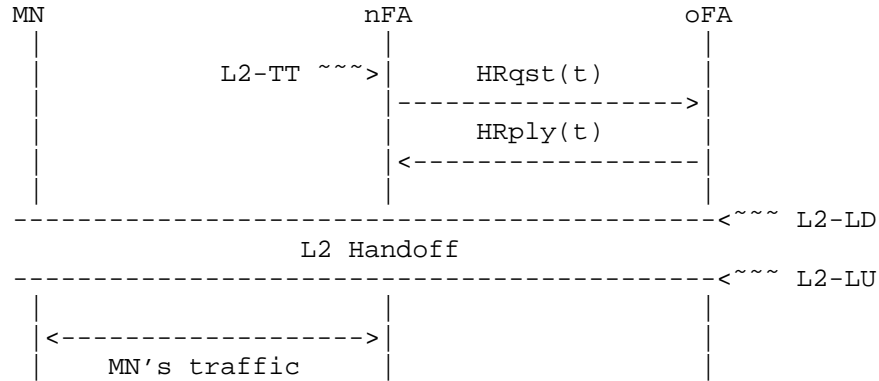


Figure 8 - Two-Party Target Trigger Handoff Timing

Once the tunnel between aFA and the current FA is in place, it is torn down by one of the following events:

- 1) The aFA decides to stop tunneling because the lifetime it sent expires and was not renewed, or the aFA or current FA decide to terminate tunnel service prematurely for some other reason (refer to Section 4.3).
- 2) The MN completes the process by performing a Mobile IPv4 Registration with the current FA. This may be initiated by the FA that sends an Agent Advertisement or by the MN that solicits for an Agent Advertisement as in [1].
- 3) The MN moves to a third FA (see Section 4.2)

4.2. Three-Party Handoff

Three-party handoff is applicable when an MN, which has already established an aFA and is receiving tunneled packets through its current FA, moves to a new FA without performing a Mobile IPv4 Registration.

The need for the three-party handoff function depends on the wireless system in which POST-REGISTRATION is being implemented. For radio L2 protocols in which it is possible for the MN to move so rapidly from one FA to another such that a probability exists that the Mobile IPv4 Registration with nFA will not complete before the MN moves on, HTT (Handoff to Third) SHOULD be implemented. Certain wireless systems and implementations do not allow such fast movement between FAs and may force the Mobile IPv4 Registration to occur soon after L2 handoff, in which case three-party handoff is not applicable. If this three-party handoff feature is not implemented, the FA SHOULD

send an Agent Advertisement to the MN after L2 handoff has completed (L2-LU at nFA) and/or the MN SHOULD solicit an Agent Advertisement after L2 handoff (L2-LU at MN).

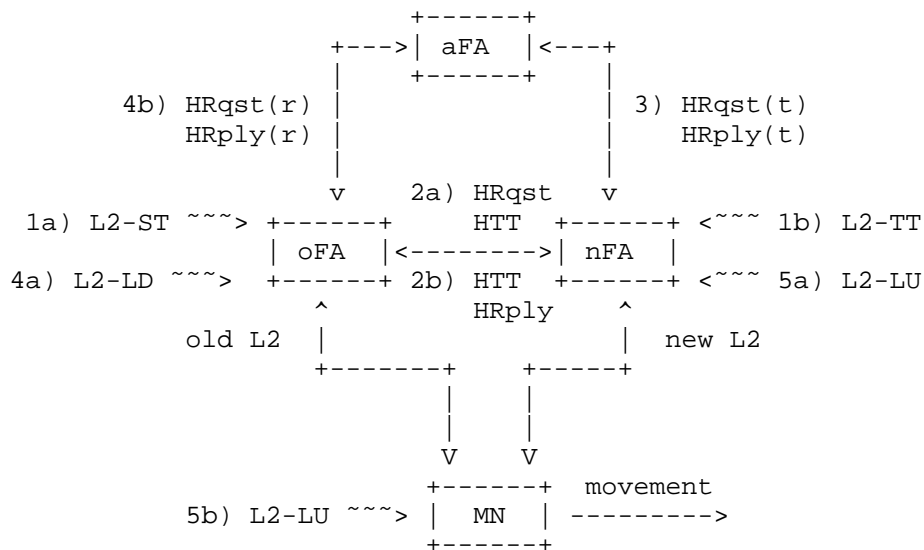


Figure 9 - Three-Party Handoff

The L3 handoff can be deferred either because of a decision by the MN/FA (i.e., MN does not send Agent Solicitations and FA does not send Agent Advertisements), or it may result from rapid movement between oFA and nFA that does not allow enough time for the registration to complete. This scenario is shown in Figure 9. In this case, oFA must inform nFA (i.e., the third FA) to contact aFA about moving the radio end of the tunnel. This is performed with the HTT message. The general idea behind the three-party handoff procedure is that the oFA supplies nFA with the same information it would have obtained via an L2-TT if handoff had occurred from aFA to nFA; then, the nFA performs an HRqst(t)/HRply(t) sequence with aFA in order to move the BET to nFA. When the L2 handoff is complete, oFA sends an HRqst(r) to aFA to terminate the previous BET.

The following describes the progress of a three-party handoff. The numbered items refer to steps in Figure 9.

- 1) Either the oFA or nFA receives an L2 trigger informing it that a certain MN is about to be moved. The two cases are:

- a) The L2 trigger is a source trigger (L2-ST) at oFA. The trigger contains the MN's L2 address and an identifier for the nFA (the IPv4 address itself or an L2 address that can be mapped to the IPv4 address of the nFA).
 - b) The L2 trigger is a target trigger (L2-TT) at nFA. The trigger contains the MN's L2 address and an identifier for the oFA (the IPv4 address itself or an L2 address that can be resolved to the IPv4 address of the oFA).
- 2) The oFA and nFA exchange an HTT/HRply or HRqst/HTT pair. HTT is indicated by setting both the H and N bits in the HRqst or HRply. The HTT message MUST NOT have any tunnel flag bits set, because the tunnel is negotiated between the aFA and nFA, not oFA and nFA. There are two cases:
- a) The L2 trigger is an L2-ST. The oFA sends HTT to nFA containing the MN's home IPv4 address, the MN's HA address, an LLA containing the aFA's IPv4 address, and an LLA containing the L2 address of the MN. This is enough information for nFA to perform a target-triggered handoff with aFA. The nFA responds with an HRply(s). Section 4.7 describes the HTT.
 - b) The L2 trigger is an L2-TT. The nFA sends HRqst(t) to oFA, exactly as if a two-party handoff were occurring. The oFA responds with HTT containing the same information as in a) above. This is enough information for nFA to perform a target-triggered handoff with aFA.
- 3) Upon receipt of the HTT, the nFA first checks its Visitor Cache to see whether it is already tunneling for MN. If so, Step 6 is performed. If not, nFA performs a target-triggered handoff with aFA, exactly as in Section 4.1, exchanging an HRqst(t)/HRply(t) pair. Because aFA receives no L2 trigger indicating when L2 handoff starts, it may start tunneling to nFA upon transmission of the HRply(t).
- 4) Once the L2 handoff is under way and the MN gets disconnected at L2, aFA and oFA exchange messages canceling tunnel service between aFA and oFA and allowing aFA to start the tunnel with nFA.
- a) The point in the L2 handoff process where the MN gets disconnected from oFA is signaled at oFA by L2-LD.

- b) The oFA exchanges an HRqst(r)/HRply(r) pair having lifetime zero with aFA. This cancels tunnel service between oFA and aFA. If aFA has not already established a tunnel to nFA, it must do so immediately upon receipt of the HRqst(r). The aFA provides tunneling service exactly as described in Section 4.1, Step 4a.
- 5) Completion of L2 handoff is signaled by an L2-LU trigger at nFA and/or MN. The nFA and MN handle the trigger as follows:
 - a) The nFA provides packet delivery service to the MN exactly as described in Section 4.1, Step 4b.
 - b) The MN either defers or initiates Mobile IPv4 Registration when it receives the L2-LU, as in Section 4.1.
- 6) In the special case where nFA and aFA are the same (i.e., the MN is moving back to the original anchor FA), aFA recognizes that it is tunneling to oFA when it checks its Visitor Cache in Step 3. In this case, there is no need for aFA to send the HRqst(t)/HRply(t) in Step 3. Upon receipt of the L2-LU trigger on handoff completion, the aFA begins routing packets to MN and the tunnel to nFA is torn down. The oFA still exchanges the HRqst(r)/HRply(r) with aFA in Step 4b because oFA cannot know a priori that aFA and nFA are the same, but they are redundant.

Figures 10 and 11 show timing diagrams for source trigger (L2-ST) and target trigger (L2-TT) three-party handoff, respectively.

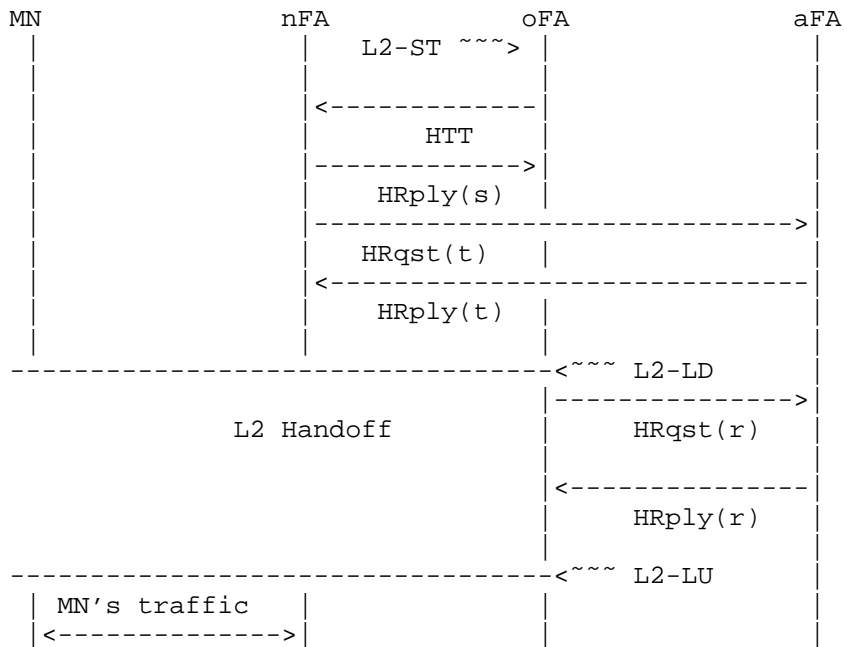


Figure 10 - Three-Party Source Trigger Handoff Timing

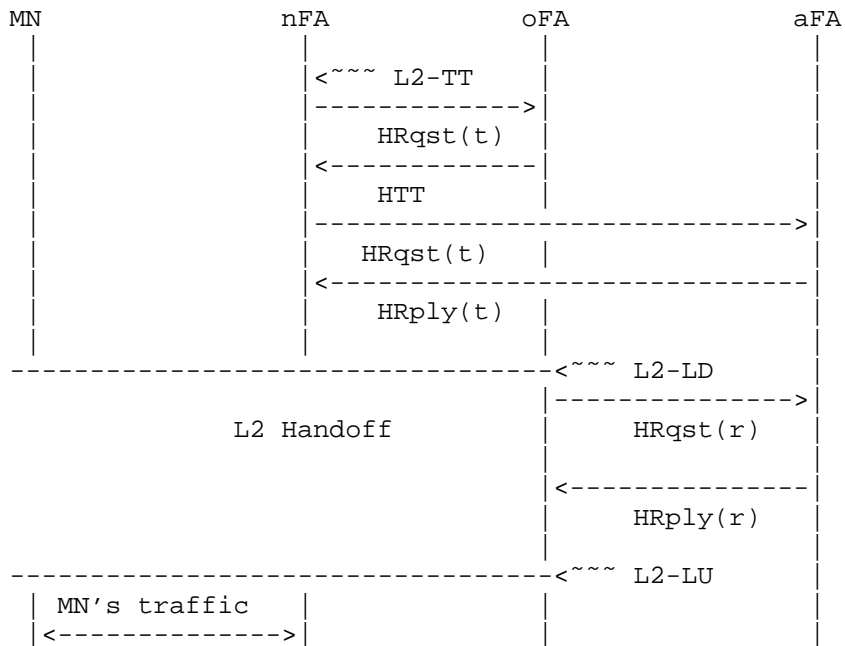


Figure 11 - Three-Party Target Trigger Handoff Timing

Unlike two-party handoff, the timing of BET establishment between aFA and nFA cannot fully depend on the availability of L2 trigger information because aFA does not receive an L2 trigger signaling L2 handoff. The two timing extremes at which aFA can place the BET with nFA are:

- 1) At the earliest, aFA MAY start tunneling packets using the BET to nFA after sending the HRply(t) to nFA in response to the request for target-triggered handoff.
- 2) At the latest, aFA MAY start tunneling packets using the BET to nFA and tear down the BET with oFA when receiving the HRqst(r) from oFA indicating that the MN has disconnected.

In addition, aFA MAY continue tunneling to oFA if 1) is selected, until the HRqst(r) is received. In this case, the result may be duplicated packets at the MN because the MN will receive packets through oFA on the old L2 until it disconnects (L2-LD). If 2) is selected, the additional latency will add to the MN's L3 service disruption period. Of course, aFA can choose to place the BET sometime between 1) and 2) if reliable bounds are available on the duration of time between L2-ST/L2-TT and the MN's disconnection (L2-LD). The exact selection of when to establish the BET is likely to

be influenced by network engineering and implementation considerations, including whether a handoff smoothing solution is used, and is beyond the scope of this specification.

4.3. Renewal or Termination of Tunneling Service

To prevent a BET from expiring when its lifetime runs out, the MN's current FA signals the aFA to either renew or terminate the BET. This may be the case when the MN defers Mobile IPv4 Registration. If no such signal is received, the aFA will terminate the BET when the lifetime expires. In addition, the current FA or aFA may need to terminate the BET prior to the lifetime expiring. In order to avoid error conditions in which tunnels do not expire even though the MN to which they apply is no longer reachable, FAs SHOULD set the tunnel lifetime field to some value other than 0xffff, which indicates "good until canceled".

Figure 12 illustrates the message exchange that occurs between the FA needing to terminate or extend the tunnel (designated FA(1) in the figure) and the other FA (designated FA(2) in the figure). The HRqst(r)/HRply(r) is indicated by setting the R bit in the HRqst/HRply messages. If the HRqst(r) is renewing a BET, then it contains a non-zero lifetime; otherwise, if the lifetime is set to zero, it indicates tunnel termination. The aFA has complete control over whether a tunnel is extended or terminated, and it MAY reply to a request for extension with a shorter lifetime than was requested.




Figure 12 - BET Renewal or Termination

4.4. When Will the MN Perform a Mobile IPv4 Registration?

The MN/FA have control over when to perform the Mobile IPv4 Registration. Although the MN/FA may decide to defer Mobile IPv4 Registration for a certain period, three possible events can lead to the need to terminate tunneling service. If this occurs, the MN MUST perform the Mobile IPv4 Registration. These events are:

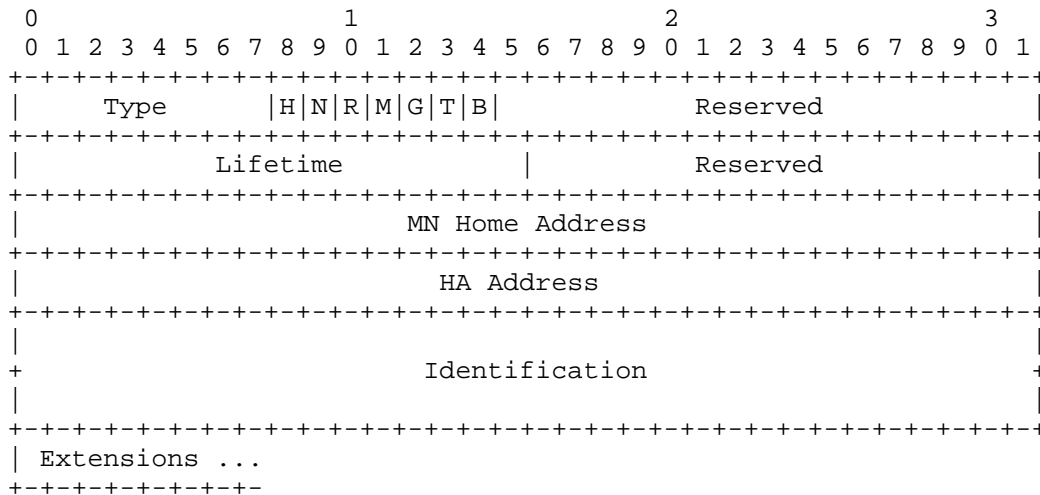
- 1) The end of life for the BET is pending and a request by the current FA to aFA for renewal has been denied, or alternatively the current FA or aFA needs to terminate the BET prematurely. The FA in this case MUST initiate the Mobile IPv4 Registration by sending an Agent Advertisement to the MN as in [1].

- 2) The MN itself decides to perform a Mobile IPv4 Registration and initiates it by sending an Agent Solicitation as in [1].
- 3) During a source-triggered handoff, the oFA attempts to perform BET handoff but nFA is not capable of performing it. The FA in this case MUST initiate the Mobile IPv4 Registration by sending the MN an Agent Advertisement as in [1]. Note that this situation will never arise during target-triggered handoff because an HRqst(t) will not be sent to oFA by an nFA that doesn't support POST-REGISTRATION.

Some detailed scenarios relating to case 2) will be described hereafter. According to [1], when using an FA care-of address, the MN MAY use the FA as its default router. Otherwise, it MUST choose its default router from those advertised in the ICMP Router Advertisement portion of the Agent Advertisement. Here we assume that the FA router is also the MN's default router. In POST-REGISTRATION, when a tunnel is established between oFA and nFA and the MN has moved to nFA, the oFA MUST NOT send Agent Advertisements to the MN. In this case, it is possible that the MN will not receive Agent Advertisements for extended periods of time. According to [8], hosts will remove default router entries if the lifetime of the Router Advertisement expires and no further advertisements are received. Note that the ICMP Router Advertisement lifetime is not related to the Registration Lifetime in the Mobility Agent Advertisement extension [1]. To avoid this disruption, the MN MUST solicit the default router (i.e., FA) before the lifetime of its active default router entry runs out, or alternatively, the FA MUST advertise as soon as the MN-nFA link is up (L2-LU). This effectively means that the MN will at most be able to defer Mobile IPv4 Registration for as long as the remaining lifetime of the active default router, as configured in the ICMP Router Advertisements. The MN MUST perform a Mobile IPv4 Registration [1] when it receives an Agent Advertisement following a POST-REGISTRATION handoff.

4.5. Handoff Request (HRqst) Message Format

This is a new Mobile IPv4 message carried on UDP (destination port 434) [1]. The UDP header is followed by the fields below.




- Type 16 (Handoff Request)
- H Source-triggered handoff request. When set and the N bit is unset, indicates that the request was the result of an L2-ST at oFA.
- N Target triggered handoff request. When set and the H bit is unset, indicates that the request was the result of an L2-TT at nFA.
- R Set if the request is an HRqst(r) (i.e., a request to renew the tunnel, H and N bits must be unset).
- M The FA issuing the HRqst will use Minimal Encapsulation as defined in [1,5] for the tunnel.
- G The FA issuing the HRqst will use Generic Routing Encapsulation (GRE) [4] as defined in [1,5] for the tunnel. Extensions of HRqst containing GRE type and key Fields are outside the scope of this document.

T	For an HRqst(s), indicates that the oFA is willing to support both forward and reverse tunnel service. For an HRqst(t), indicates that the nFA is requesting reverse tunnel service.
B	When sent in an HRqst(s), indicates that the MN has requested a reverse tunnel to the HA and that the nFA SHOULD use a reverse tunnel to the HA if it will not be reverse tunneling to the oFA.
Lifetime	The lifetime of the tunnel in seconds. If this is an HRqst(t), then the lifetime represents a request by nFA for a reverse tunnel. If this is an HRqst(s), then the lifetime represents the maximum amount of time that oFA is willing to maintain both forward and reverse tunnels. If this is an HRqst(r), then the lifetime represents a request for the amount of time to renew the tunnel's lifetime. A value of 0 on an HRqst(s) indicates that the oFA is unwilling to grant tunnel service. A value of 0 on an HRqst(t) indicates that the nFA does not require reverse tunnel service. A value of 0 on an HRqst(r) indicates that the tunnel should be terminated. A value of 0xffff indicates infinity.
MN Home Address	For HRqst(s), the home address of the MN.
HA Addr	For HRqst(s), the HA address of the mobile node.
Identification	As defined in [1].
Extensions	The message MUST include an LLA (see Section 9) containing the MN's L2 address and an L2 address that can be mapped to an IPv4 address for the FA. This message MUST contain the FA-FA Authentication Extension [11] that is used to secure the HRqst message.

4.6. Handoff Reply (HRply) Message Format

This is a new Mobile IPv4 message carried on UDP (destination port 434) [1]. The UDP header is followed by the fields below.



- Type 17 (Handoff Reply)
- Code A value indicating the result of the Handoff Request. Only two codes are currently supported, 0, indicating success, and 1, indicating that the handoff cannot be performed. The remaining values are for future use.
- Lifetime The lifetime, in seconds, for which the bidirectional tunnel for the MN will be maintained. If this is an HRply(s), then the lifetime represents a request by nFA, and it can be any value up to the maximum value sent in the HRqst(s). Larger values are assumed to default to oFA's maximum. If this is an HRply(t), then the lifetime represents the maximum amount of time that the oFA will grant to the nFA. If this is an HRply(r), then the lifetime represents the amount of time by which the tunnel life will be extended. If the Code field indicates that handoff failed, the Lifetime field will be ignored and SHOULD be set to zero. A value of 0 on an HRply(t) indicates that the oFA is unwilling to grant service. A value of 0 on an HRply(s) indicates that the nFA does not

require service. A value of 0 on an HRply(r) indicates that the tunnel lifetime will be terminated. A value of 0xffff indicates an infinite lifetime.

- H Source-triggered handoff reply. When set and the N bit is unset, indicates that the reply is in response to an HRqst(s).
- N Target-triggered handoff reply. When set and the H bit is unset, indicates that the reply is in response to an HRqst(t).
- R Set if the reply is an HRply(r). Neither the H nor the N bit are set.
- M The FA issuing the HRqst will use Minimal Encapsulation as defined in [1,5] for the tunnel.
- G The FA issuing the HRqst will use GRE [4] Encapsulation as defined in [1,5] for the tunnel. When this flag bit is set, the HRply may require extensions containing the GRE type and key fields, but they are outside the scope of this document.
- T For an HRply(s), indicates that the nFA is requesting to reverse tunnel service. For an HRply(t), indicates that the oFA is willing to provide both forward and reverse tunnel service.
- B When sent in an HRply(t), indicates that the MN has requested a reverse tunnel to the HA and that the nFA SHOULD use a reverse tunnel to the HA if it will not be reverse tunneling to the oFA. It can be set in HRply(t) only if the T bit was unset in the corresponding HRqst(t).
- MN Home Address For HRply(t), the home IPv4 address of the MN.
- HA Addr For HRply(t), the HA IPv4 address of the MN.
- Identification As defined in [1].
- Extensions This Message MUST contain the FA-FA Authentication Extension [11] that is used to secure the HRply message.

4.7. Handoff to Third (HTT) Message Format

The Handoff to Third message has the same format as the Handoff Request and Handoff Reply messages, except both the H and N bits are set. If the HTT message is in response to an L2-ST and is sent to initiate a handoff, then, with the exception of the H and N bits, the message has the same fields set and includes the same extensions as an HRqst(s). If the HTT message is sent in response to an HRqst(t), then, with the exception of the H and N bits, the message has the same fields set and includes the same extensions as an HRply(t). The tunnel bits MUST NOT be set in the HTT message because BET construction is not negotiated between oFA and nFA; it is negotiated between nFA and aFA in the ensuing HRqst(t)/HRply(t).

In addition, the HTT MUST contain the following extensions in the specified order:

Solicited IPv4 Address Option: containing aFA's Address

LLA Option: containing the L2 address of the MN.

4.8. Applicability of POST-REGISTRATION Handoff Method

The POST-REGISTRATION handoff approach allows FAs to communicate directly about a pending handoff, and does not require any IPv4-layer messages to be sent to or from an MN prior to the L2 handoff event. Therefore, it eliminates a possible source of handoff latency. This may be required when the link layer imposes hard deadlines on the time at which a handoff must occur, such as when an MN is rapidly moving out of a radio coverage area. Consequently, POST-REGISTRATION is primarily of interest in handoff between FAs that support the same radio access technology. Handoff between heterogeneous radio technologies will, of necessity, require interaction between the MN and the network, and so is not a domain of applicability for POST-REGISTRATION.

Because a POST-REGISTRATION handoff is triggered by an unspecified mechanism that informs the oFA or nFA that an L2 handoff is pending, the POST-REGISTRATION approach is only applicable to networks where such a mechanism is available. For example, an L2 may provide indications of radio signal quality that cause the oFA or nFA to send the POST-REGISTRATION handoff messages. Any such indications must also provide each FA involved in the handoff with the identity of the other, so that messages can be sent to the right place. This may involve mapping L2 information onto FA IPv4 addresses. Also, the FAs involved in a handoff must have pre-provisioned security arrangements so that the POST-REGISTRATION messages can be authenticated. If a handoff is to be completed as a result of the POST-REGISTRATION

messaging, any L2 handoff indications must also be securely authenticated so that traffic to the old point of attachment is not improperly halted.

POST-REGISTRATION handoff is appropriate in the following cases:

- L2 triggers are available on the network to indicate that L2 handoff is pending.
- Pre-provisioned security mechanisms are in place to allow fast and secure messaging between the FAs and between the MN and an FA.
- Access point choice by the MN is not a concern or the choice requires user intervention and therefore is not on the critical path for handoff.

5. Combined Handoff Method

The combined method uses both PRE-REGISTRATION and POST-REGISTRATION handoff. If PRE-REGISTRATION does not complete prior to the expiration of a timer on the nFA, the POST-REGISTRATION mechanism is used to create the tunnel between oFA and nFA. This protects the MN from delays caused by errors such as loss of the Mobile IPv4 Registration Reply message involved in PRE-REGISTRATION for the mobile-initiated and network-initiated source-triggered cases. It also protects the MN from delays caused by errors or the loss of any of the Mobile IPv4 messages involved in PRE-REGISTRATION for the network-initiated target-triggered case.

When the nFA receives a target trigger, it will follow the PRE-REGISTRATION procedure. When the combined method is used, the nFA MUST also start a timer when it receives a target trigger. The timer should be set to a small value (default for target trigger case: 1 second).

According to PRE-REGISTRATION, the nFA will receive the Registration Request from the MN. When the combined method is used, this Registration Request sent by the MN MUST contain the IPv4 address of the oFA in an FA IPv4 address LLA extension (see Section 9.7). This same Registration Request message will contain multiple LLA extensions, since it will also contain the MN's layer 2 address in an LLA extension as described for PRE-REGISTRATION (see Sections 3.7 and 9). When the nFA has not started the handoff procedure using a target trigger (i.e., mobile-initiated or network-initiated target-triggered cases), the nFA MUST start a timer as soon as it receives the low-latency Registration Request from the MN. This timer should be set to a small value (default: 1 second).

In all cases, the timer MUST be reset when the Registration Reply message is received by nFA. If the timer expires before the Registration Reply is received, the nFA MUST initiate the POST-REGISTRATION procedure. The nFA utilizes the oFA IPv4 address (previously received in the extension to the Registration Request message) as the destination of the POST-REGISTRATION HRqst message to create the tunnel between nFA and oFA. The nFA MAY tear down this tunnel when it receives and forwards a successful Registration Reply for that MN.

6. Layer 2 and Layer 3 Handoff Timing Considerations

In the optimal cases considered in the PRE-REGISTRATION and POST-REGISTRATION handoffs, it was assumed that a timely L2 trigger would be received in such a way that packets could be delivered to the MN via its nFA immediately upon connection. In this way, the MN does not suffer disruption due to the L3 handoff. However, such precise timing of the L2 trigger and handoff mechanism with respect to the actual L2 handoff event will not be possible in all wireless systems and may depend on particular implementation techniques. Therefore, some uncertainty may exist at L3 as to exactly when the L2 connection between the MN and the nFA becomes fully established and can be used for L3 traffic. It is possible that in certain implementations traffic will be re-routed too early or too late with respect to the moment when the connection between the MN and the nFA becomes fully established. The techniques that may solve this problem and allow the MN to receive traffic independently of the timing of the L2 handoff event include buffering and simultaneous bindings (i.e., bicasting: setting the S bit [1] in Registration Requests). However, these are optional and are not mandated.

7. Reverse Tunneling Support

The handoff methods all support reverse tunneling. The MN may request reverse tunneling [3] by setting the T bit in its Registration Request. In the case of POST-REGISTRATION, if the MN had requested reverse tunneling previously at oFA, the handoff message from oFA (see Section 4) includes the T bit enabled to inform nFA to establish a BET for the visitor entry. Typically, the T bit will always be set to ensure that any delays in the MN receiving its new care-of address do not result in any delay in uplink packet transmission from the MN, but local policies and particular L2 technologies may allow the reverse tunnel to be turned off.

8. Handoff Signaling Failure Recovery

In general and to a greater extent in wireless networks, packets carrying handoff signaling may be dropped or lost due to errors on the link. In this section, we consider mechanisms for recovery from handoff signaling failures.

8.1. PRE-REGISTRATION Signaling Failure Recovery

Failure of PRE-REGISTRATION signaling breaks down into three cases:

- 1) Loss of messages PrRtSol and PrRtAdv on the air link.
- 2) Loss of the solicitation by an FA to obtain another neighboring FA's Advertisement or loss of the neighboring FA's advertisement.
- 3) Failure of the standard Mobile IPv4 Registration.

Of these, case 3) is handled by standard Mobile IPv4 mechanisms described in [1]. Case 2) is expected to be a rare event because spontaneous packet drop rates on the fixed network are caused by congestion or router failure. Since bit error rates on wireless links are higher than on fixed links, case 1) is more likely to occur. In the following subsections, cases 1) and 2) are considered.

8.1.1. Failure of PrRtSol and PrRtAdv

PRE-REGISTRATION handoff can fail in network-initiated handoff when the PrRtAdv sent by oFA in response to the source trigger (L2-ST) or the advertisement sent by nFA in response to the target trigger (L2-TT) fails to reach the MN. PRE-REGISTRATION handoff can also fail in mobile-initiated handoff when either the PrRtSol sent from the MN or return PrRtAdv sent from the oFA is dropped. To reduce the probability that PrRtAdv and PrRtSol are lost, the MN and FA MAY transmit multiple copies of these messages. Should these messages fail anyway, in both cases the MN connects to the nFA without having received any prior signaling. In this case, the MN solicits an FA Advertisement when it connects to nFA at L2 (L2-LU), as described in Section 3.6, and performs a standard Mobile IPv4 Registration with the nFA as specified in [1].

8.1.2. Failure of Inter-FA Solicitation and Advertisement

The solicitation from an FA to another neighboring FA may fail or the corresponding advertisement from the neighboring FA may be lost. To reduce the probability that these messages are lost, the FAs MAY transmit multiple copies of these messages. If a failure occurs anyway, the FA soliciting the Agent Advertisement is unable to send a PrRtAdv in response to a source trigger or to a mobile-initiated PrRtSol. In these cases, when the MN does not receive a notification or confirmation of a PRE-REGISTRATION handoff, the MN MUST perform a standard Mobile IPv4 Registration as soon as it connects to the nFA (L2-LU) as described in Section 8.1.1.

8.2. POST-REGISTRATION Signaling Failure Recovery

Failure occurs in POST-REGISTRATION when either the HRqst or HRply message is dropped. The effects of the failure and the recovery procedure depend on which message is dropped, and whether the handoff is source or target triggered. Since all of the POST-REGISTRATION signaling is going over the fixed network, it can be expected that spontaneous dropping of packets in the absence of congestion and router failure should be a relatively rare event. Nevertheless, failure recovery mechanisms SHOULD be implemented.

8.2.1. HRqst Message Dropped

If the HRqst message is dropped, the effect is the same for both source- and target-triggered handoffs. In either case, the FA to which the HRqst was destined will never respond with an HRply message. If the handoff is source triggered, then the nFA never learns of the handoff, and the oFA never receives confirmation. If the handoff is target-triggered, then the oFA never learns of the handoff, and the nFA never receives confirmation.

The recovery procedure in this case is as follows: the oFA MUST NOT construct a forward tunnel when the MN moves off-link (L2-LD) if the handoff is source-triggered, and the nFA MUST NOT construct a reverse tunnel if the handoff is target triggered. If the nFA was not informed of the handoff by an HRqst message (corresponding to failure of source-triggered handoff) or if the handoff was not confirmed by an HRply message (corresponding to failure of target-triggered handoff), the nFA MUST unicast an Agent Advertisement to the MN as soon as its L2 connection is established (L2-LU at nFA).

8.2.2. HRply Message Dropped


If the HRply message is dropped, the FA sending the HRply will assume that the handoff has been confirmed, but the FA that is expecting to receive the HRply does not receive confirmation. In this case, the failure recovery procedure is different for source-triggered and target-triggered handoffs.

In a target-triggered handoff, the oFA assumes that the handoff is confirmed because it has sent the HRply, but the nFA has not received it so it does not have confirmation. The oFA starts tunneling packets to the nFA when the MN moves from its link (L2-LD). The nFA MUST send an FA Advertisement to the MN as soon as its L2 link is up (L2-LU at nFA) and MAY drop the tunneled packets. The nFA SHOULD send an ICMP Destination Unreachable [9] message to the oFA. When the oFA receives this message, it will terminate the tunnel and stop forwarding packets. If reverse tunneling was requested, the nFA MUST NOT reverse tunnel because it has not received handoff confirmation.

In source-triggered handoff, the nFA assumes that the handoff is confirmed because it has sent the HRply, but the oFA has not received it so it does not have confirmation. Without failure recovery, the MN could move to the nFA without the oFA being able to start the forward tunnel for the MN's packets, and the MN would not be able to initiate a Mobile IPv4 Registration because it does not know that the handoff has failed. In this situation, the oFA MUST send out an HRqst message to the nFA with lifetime zero as soon as the MN leaves its link (L2-LD). The oFA SHOULD continue to retransmit the HRqst message, with exponential backoff for CONFIG-HFAIL seconds or until it receives an HRply acknowledging the request to cancel the tunnel. The default value for CONFIG-HFAIL is 10 seconds. When the nFA receives the HRqst, it MUST immediately send an Agent Advertisement to the MN, as is the case whenever a tunnel is canceled. In addition, the oFA MUST also drop any packets received through the reverse tunnel from the nFA. The oFA SHOULD NOT send the ICMP Destination Unreachable message to the nFA because the nFA has been informed by the HRqst message to cancel the tunnel. However, if the nFA receives an ICMP Destination Unreachable message for the tunnel prior to receiving the HRqst canceling the tunnel, it MUST send an FA Advertisement to the MN and cancel the tunnel.

9. Generalized Link Layer and IPv4 Address (LLA) Extension

This section defines the Generalized Link Layer and IPv4 Address (LLA) Extension, used by any node that needs to communicate link layer and IPv4 addresses. The format of the extension relies on sub-types, where each sub-type defines its own sub-structure. This document defines six sub-types. Future RFCs should allocate their own sub-type and define their own address formats.



Type

- 138 (skippable) [1] - when used in Registration Requests
- 140 (skippable) [1] - when used in Agent Advertisements

Length

The length of the Link Layer Address + the one-octet Sub-Type field

Sub-Type

This field contains the Link Layer sub-type identifier

LLA

Contains the Link Layer Address


In this document, seven sub-types are defined:

- 1 3GPP2 International Mobile Station Identity and Connection ID [13]
- 2 3GPP International Mobile Subscriber Identity [15]
- 3 Ethernet 48-bit MAC address [5]
- 4 64-bit Global ID, EUI-64 [6]
- 5 Solicited IPv4 Address
- 6 Access Point Identifier
- 7 FA IPv4 Address

The following subsections describe the extensions.

9.1. 3GPP2 IMSI Link Layer Address and Connection ID Extension

The IMSI Link Layer Address Extension contains the International Mobile Station Identity (IMSI).



Type

1 (skippable) [1]

Length

The length of the IMSI field + the one-octet Sub-Type field

Sub-Type

1

IMSI


Contains the IMSI, in the form:

<IMSI>:<Connection Id>

Where the <IMSI> is an ASCII-based representation of the International Mobile Station Identity, most significant digit first, ":" is ASCII 0x3a, and the Connection ID is the ASCII representation of a small, decimal number used for distinguishing different link-layer connections from the same mobile device.

9.2. 3GPP IMSI Link Layer Address Extension

The IMSI Link Layer Address Extension contains the International Mobile Station Identity.



Type

2 (skippable) [1]

Length

The length of the IMSI field + the one-octet Sub-Type field

Sub-Type


2

IMSI

Contains the IMSI, a number composed of 15 digits or less, coded as described in [15].

9.3. Ethernet Link Layer Address Extension

The Ethernet Link Layer Address Extension contains the 48-bit Ethernet MAC Address, as defined in [5].



Type

3 (skippable) [1]

Length

7 (includes the Sub-Type field)

Sub-Type


3

MAC

Contains the 48-bit Ethernet MAC Address.

9.4. IEEE 64-Bit Global Identifier (EUI-64) Address Extension

The 64-bit Global Identifier (EUI-64) Address Extension contains the 64-bit address, as defined in [6].



Type

4 (skippable) [1]

Length

9 (includes the Sub-Type field)

Sub-Type


4

MAC

Contains the 64-bit Global Identifier Address.

9.5. Solicited IPv4 Address Extension

The 32-bit Solicited IPv4 Address Extension contains the IPv4 address of the agent (FA) being solicited. This extension MAY be present in an ICMP Agent Solicitation as explained in Section 3.3.



Type

5 (skippable) [1]

Length

5 (includes the Sub-Type field)

Sub-Type


5

IPv4 Address

Contains the 32-bit IPv4 Address of the solicited node.

9.6. Access Point Identifier Extension

The 32-bit Access Point Identifier Extension contains an identifier of the access point to which the MN will move. This may be a wireless L2 identifier. The MN is able to solicit an advertisement from the FA servicing a certain access point by using this extension with Agent Solicitations as explained in Section 3.3.



Type

6 (skippable) [1]

Length

5 (includes the Sub-Type field)

Sub-Type


6

AP ID

Contains the 32-bit Access Point Identifier.

9.7. FA IPv4 Address Extension

The 32-bit FA IPv4 Address Extension contains the IPv4 address of the agent (FA). This extension MAY be present in a Registration Request message to identify the oFA as explained in Section 5.



Type

7 (skippable) [1]

Length

5 (includes the Sub-Type field)

Sub-Type

7

IPv4 Address

Contains the 32-bit IPv4 Address of the FA node.

10. IANA Considerations

This document defines one new extension to Mobile IPv4 Control messages and one new extension to Mobile IPv4 Router Discovery messages already maintained by IANA. This document also defines a new Mobile IPv4 Control message type to be used between FAs. To ensure correct interoperation based on this specification, IANA must reserve values in the Mobile IPv4 number space for two new extensions and one new message type. IANA must also manage the numbering spaces created by the two new extensions, the message type, and its associated Code field.

10.1. New Extension Values

Section 9 introduces two extensions.

Generalized Link Layer and IPv4 Address (LLA) Extension for Router Discovery messages: A new Mobile IPv4 extension that follows after Mobile IPv4 ICMP Router Discovery messages (e.g., Mobile IP Agent Advertisements). The type value of this extension belongs to the

Mobile IPv4 number space for Router Discovery messages maintained by IANA. The value assigned by IANA is 140. This new extension uses the Link Layer and IPv4 Address Identifier (LLA) sub-type numbering space that requires IANA management (see Section 10.2).

Generalized Link Layer and IPv4 Address (LLA) Extension for Mobile IP Control messages: A new Mobile IPv4 extension appended to Mobile IP Control messages (e.g., Registration Request). The type value of this extension belongs to the Mobile IPv4 number space for extensions to Mobile IPv4 Control messages maintained by IANA. It MUST be in the skippable (128-255) range as defined in [1]. The value assigned is 138 by IANA. This new extension uses the Link Layer and IP Address Identifier (LLA) sub-type numbering space that requires IANA management (see Section 10.2).

10.2. Generalized Link Layer and IP Address Identifier (LLA) Sub-type Values

This section describes the sub-type values that are applicable to both the Generalized LLA Extensions for Mobile IP Control and Router Discovery messages. This specification makes use of the sub-type values 1-7, and all other values other than zero (reserved) are available for assignment via IETF consensus [14]. The seven sub-type values defined in this specification are:

1	3GPP2 International Mobile Station Identity and Connection ID [13]
2	3GPP International Mobile Subscriber Identity [15]
3	Ethernet 48-bit MAC address [5]
4	64-bit Global ID, EUI-64 [6]
5	Solicited IPv4 Address
6	Access Point Identifier
7	FA IPv4 Address

10.3. New Message Type and Code

Sections 4.5 and 4.6 define two new Mobile IPv4 message types: Handoff Request and Handoff Reply. These require two type numbers to be assigned by IANA from the Mobile IPv4 Control message type address space. The Handoff Reply message also introduces its own Code field that requires IANA to manage a new Code address space. This specification makes use of the Code values 0-1, where 0 identifies a successful handoff and 1 defines a generic handoff failure. All other values are available for assignment via IETF consensus [14].

Code Values for Mobile IP Handoff Reply Messages

0	Successful Handoff
1	Generic Handoff Failure
2-255	Unallocated

11. Security Considerations

For the PRE-REGISTRATION method, as discussed in Section 3.8, the oFA and nFA MUST share a security association to authenticate and integrity protect messages transported between them. In addition, oFA must be authorized to solicit nFA based on the security association. The minimal requirement to establish a security association between FAs is that both FAs support manual pre-configuration of security associations involving shared keys. Other mechanisms to establish security associations using IKE [16] based on shared secrets or public keys may also be used. The inter-FA ICMP messages (solicitations and advertisements) MUST be authenticated and integrity protected using ESP [10]. The default ESP authentication algorithm for use in this specification is HMAC-SHA1-96 [12]. The absence of this security would allow denial-of-service attacks from malicious nodes at any distance from the FA. To secure Registration Request and Reply messages, PRE-REGISTRATION uses the security mechanisms already described in [1] and optionally [11].

POST-REGISTRATION introduces a new change to Mobile IPv4, which is the possibility that an MN may receive packets from an FA with which it has not yet performed a Mobile IPv4 Registration. It is not recommended that the MN drop packets from unknown FAs since it would effectively eliminate the advantages of POST-REGISTRATION. From a security viewpoint, dropping packets from unknown FAs would not provide significant protection for an MN from any attack. This is because any malicious host may use the MN's home address to send packets to the MN through its current known FA; therefore, processing packets received from unknown FAs would not provide worse security than with normal Mobile IPv4.

In a similar way to PRE-REGISTRATION, in POST-REGISTRATION, oFA and nFA MUST share a security association required to protect the Handoff Request and Reply messages. The minimal requirement to establish a security association between FAs is that the FAs support manual pre-configuration of security associations involving shared keys. Other mechanisms to establish security associations using IKE [16] based on shared secrets or public keys may also be used. The Handoff Request and Reply messages MUST be authenticated using the FA-FA authentication extension [11] that uses the default algorithm specified in [7]. The absence of this security would allow impersonation attacks and denial-of-service attacks.

The minimal requirement is that all FAs involved in low latency handoffs MUST support manual pre-configuration of peer-to-peer security associations with neighboring FAs, involving shared secrets and are already required to support the default algorithms of [1]. Other mechanisms to establish security associations using IKE [16] based on shared or public keys may also be used.

Since the techniques outlined in this document depend on particular L2 information (triggers) to optimize performance, some level of L2 security is assumed. Both PRE- and POST-REGISTRATION techniques depend on L2 triggers, but the L2 security implications are different for the two techniques.

In particular, in POST-REGISTRATION, the L2 triggers initiate the establishment of tunnels that route IPv4 packets for the MN to its new location. Therefore, the L2 triggers MUST be secured against any tampering by malicious nodes, either mobile or within the wired network. The L2 addresses or IPv4 addresses for the MN and the FAs that appear in the L2 triggers MUST correspond to the actual nodes that are participating in the handoff. If there is any possibility that tampering may occur, the recipient of an L2 trigger MUST have some way of authenticating the L2 information. Wireless networks that do not provide such features will be subject to impersonation attacks, where malicious nodes could cause FAs to believe that an MN has moved to other service areas or to allow a bogus MN to obtain unauthorized service from an FA prior to performing a Mobile IPv4 Registration. In POST-REGISTRATION, the L2 triggers would typically be sent between a wireless base station and the FA. No standard protocol exists at this time to communicate the L2 trigger information, but it is important that any future protocol used for this purpose provides adequate security. If the wireless base station and FA were integrated, then this security threat would not apply. Also the layer 2 control messages on the wireless link must be secured appropriately to prevent a malicious node from running impersonation attacks and causing unwanted L2 triggers to be generated. Integrity and replay protection would be required to avoid impersonation threats and resource consumption threats where a malicious node replays old messages to cause resource consumption. This depends on the type of L2 security of the wireless link. For example, in cellular technologies, the control messages are secured, although the type of security varies depending on the cellular standard, but this is not typically the case in WLAN IEEE 802.11 networks.

In PRE-REGISTRATION, the security of L2 triggers has different implications. The PRE-REGISTRATION technique depends on Mobile IPv4 security between MN and FA, so the same security considerations in [1] apply. Should malicious nodes be able to generate or modify L2

trigger information (i.e., L2-ST or L2-TT), this would cause advertisements to be sent to the MN. They would consume wireless resources and processing in the MN, but would not allow an impersonation attack. In order to prevent such denial-of-service attacks, there should be a limit on the number of advertisements that an FA (oFA) will relay to the MN as a result of the reception of L2 triggers. This number will depend on the L2 technology, and the default limit is 10 per second.

12. Acknowledgements

The authors want to thank Lennart Bang, Bryan Hartwell, Joel Hortelius, Gianluca Verin, and Jonathan Wood for valuable comments and suggestions on the whole document. The authors also thank the Mobile IPv4 WG chairs, Phil Roberts and Basavaraj Patil, for their input.

13. References

13.1. Normative References

- [1] Perkins, C., Ed., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Montenegro, G., Ed., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [4] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [5] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, November 1982.
- [6] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", <http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>, March 1997.
- [7] Perkins, C., Calhoun, P., and J. Bharatia, "Mobile IPv4 Challenge/Response Extensions (Revised)", RFC 4721, January 2007.

- [8] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [9] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.
- [10] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [11] Fogelstroem, E., Jonsson, A., and C. Perkins, "Mobile IPv4 Regional Registration", RFC 4857, June 2007.
- [12] Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", RFC 2404, November 1998.

13.2. Informative References

- [13] TIA/EIA/IS-2000.
- [14] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [15] 3GPP TS 23.003 (www.3gpp.org).
- [16] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

Appendix B - Low-Latency Handoffs for Multiple-Interface MNs

For MNs that have two wireless network interfaces, either on the same wireless network or on wireless networks having different wireless L2 technologies, the techniques discussed in this document may be unnecessary if the Mobile IPv4 stack on the MN allows switching an IPv4 address binding between interfaces. This Appendix discusses how multiple wireless interfaces can aid low-latency handoff.




Figure B.1 - Network Model for Mobile IPv4 with Multi-Access

Figure B.1 illustrates the normal and hierarchical MIPv4 models. As shown in the figure, assume that the MN is connected to Radio Network 1 (RN1) and is registered with oFA through which it is receiving traffic. Suppose MN enters the coverage area of RN2 and nFA and that it prefers connectivity to this network for reasons beyond the scope of this document (e.g., user preferences, cost, QoS available, etc.). The MN activates the interface to RN2 but continues communicating through RN1. The MN may solicit advertisements from nFA through the interface connected to RN1 to speed up the handoff process, provided there is no TTL restriction, or it can solicit advertisements through the interface connected to RN2 if it has been configured for IPv4 traffic.

Once the MN is registered with nFA and is successfully receiving and transmitting through the new network, it tears down the interface to RN1. If the MN has enough time to complete this procedure without incurring degraded service or disconnection, the MN would experience a seamless multi-access handoff, but it may not be possible in all

cases, due to network coverage or for other reasons. Should multiple interface handoff be possible, then the low-latency methods described in this document are not necessary.

In order to support the possible failure of the connectivity with the new network (RN2/nFA) in the short period following handoff, the MN may use the S bit in its Mobile IPv4 Registration Request to maintain simultaneous bindings with both its existing (HA or GFA) binding with oFA and a new binding with nFA.

Appendix C - PRE-REGISTRATION Message Summary

This appendix contains a quick reference for IPv4 and layer 2 addresses to be used in PRE-REGISTRATION messages.

Proxy Router Advertisement (PrRtAdv)

This is a standard Router/Agent Advertisement [1] with the following characteristics:

- Source IPv4 Address: nFA IPv4 Address
- Source Layer 2 Address: oFA L2 Address
- Destination IPv4 Address: MN IPv4 Address (from PrRtSol)
- Destination Layer 2 Address: MN L2 Address (from PrRtSol)
- LLA Extension (defined in this spec): containing nFA Layer 2 Address.

Proxy Router Solicitation (PrRtSol)

This is a standard Router/Agent Solicitation [1] with the following characteristics:

- Source IPv4 Address: MN Address
- Source Layer 2 Address: MN Address
- Destination IPv4 Address: oFA Address (from source address of previous Router Advertisement or PrRtAdv)
- Destination Layer 2 Address: oFA Address (from source address of previous Router Advertisement or PrRtAdv LLA)
- LLA Extension (defined in this spec): depends on the layer 2 technology (e.g., typically for WLAN, this would be the BSSID of the new WLAN Access Point)

Registration Request (as seen on MN-oFA link)

This is a Mobile IPv4 Registration Request message [1] with the following characteristics:

- Source IPv4 Address: MN Address
- Source Layer 2 Address: MN Address
- Destination IPv4 Address: nFA Address (from source addr of PrRtAdv)

Destination Layer 2 Address: Default Router (i.e., oFA Address)
LLA Extension (defined in this spec): containing the MN's L2 address that must be used by nFA. This will typically be an Ethernet MAC address but other types can be used as specified in Section 9 of this document.

Although this is not mandated, an MN implementation may set the S bit (see Section 6) in Registration Request messages to improve the handoff and avoid problems due to failed layer 2 handoffs and layer 2 ping-pong effects between two base stations.

Registration Reply (as seen on oFA-MN link)

This is a Mobile IPv4 Registration Reply message [1] with the following characteristics:

Source IPv4 Address: nFA Address
Source Layer 2 Address: oFA Address
Destination IPv4 Address: MN Address (from source of Registration Request)
Destination Layer 2 Address: MN Address (from source of Registration Request)

Contributing Authors

Pat Calhoun
Cisco Systems
EMail: pcalhoun@cisco.com

Tom Hiller
Lucent Technologies
EMail: tom.hiller@lucent.com

James Kempf
NTT DoCoMo USA Labs
EMail: kempf@docomolabs-usa.com

Peter J. McCann
Motorola Labs
EMail: pete.mccann@motorola.com

Ajoy Singh
Motorola
EMail: asinghl@email.mot.com

Hesham Soliman
Elevate Technologies
EMail: Hesham@elevatemobile.com

Sebastian Thalanany
US Cellular
EMail: Sebastian.thalanany@uscellular.com

Editor's Address

Karim El Malki
Athonet
EMail: karim@athonet.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

Acknowledgement

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