

July 2005

Simultaneous Bindings for Mobile IPv6 Fast Handovers
<[draft-elmalki-mobileip-bicasting-v6-06.txt](#)>

Status of this memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

By submitting this Internet-Draft, we accept the provisions of [Section 4 of RFC 3667](#) ([BCP 78](#)).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright Notice

Copyright (C) The Internet Society (2005). All Rights Reserved.

Abstract

Fast Handover for Mobile IPv6 [[1](#)] minimizes the amount of service disruption when performing layer-3 handovers. This draft extends the Fast Handover protocol with a simultaneous bindings function to minimize packet loss at the MN. Traffic for the MN is therefore bicast or n-cast for a short period to its current location and to one or more locations where the MN is expected to move to shortly.

This removes the timing ambiguity regarding when to start sending traffic for the MN to its new point of attachment following a Fast Handover and allows the decoupling of layer-2 and layer-3 handovers. It also saves the MN periods of service disruption in the case of ping-pong movement.

TABLE OF CONTENTS

- [1. Introduction.....](#)[2](#)
- [1.1 Terminology.....](#)[3](#)

- [2. Simultaneous Bindings.....](#)[3](#)

- [3. Fast Handovers in Mobile IPv6.....](#)[4](#)

- [4. Decoupling L3 Handovers from L2 handovers using Simultaneous Bindings](#)[5](#)

- [5. Avoiding service disruption due to ping-pong movement.....](#)[6](#)

- [6. Extensions to the Fast Handover Operations.....](#)[7](#)
- [6.1 MN Operation.....](#)[7](#)
- [6.2 HA/MAP/AR Operation.....](#)[7](#)

- [7. Simultaneous Bindings Flag in Fast Binding Update \(F-BU\) message.](#)[8](#)

- [8. Simultaneous Bindings option for Fast Binding Acknowledgement... \(F-BA\) message.....](#)[8](#)

- [9. Multiple copies of packets received at AR.....](#)[9](#)

- [10. Reception of multiple copies in the MN.....](#)[9](#)

- [11. References.....](#)[10](#)

- [12. Authors' Addresses.....](#)[10](#)

[1. Introduction](#)

Fast Handover for Mobile IPv6 (FMIPv6) describes a protocol to minimise the amount of service disruption when performing layer-3 handovers. This draft extends the Fast Handover protocol with a simultaneous bindings function to minimise packet loss at the MN. Traffic for the MN is therefore bicast or n-cast for a short period to its current location and to one or more locations where the MN is expected to move to shortly. This removes the timing ambiguity regarding when to start sending traffic for the MN to its new point of attachment following a Fast Handover and allows the decoupling of layer-2 and layer-3 handovers. It also saves the MN periods of service disruption in the case of ping-pong movement. [Appendix A](#)

contains some calculations illustrating how to achieve zero service disruption at L3 using FMIPv6 and bicasting.

1.1 Terminology

This section presents a few terms used throughout the document.

PAR $\hat{=}$ Previous Access Router.

NAR - New Access Router.

L2 handover - Movement of a MN's point of Layer 2 (L2) connection from one wireless access point to another.

L3 handover - Movement of a MN between ARs which involves changing the on-link care-of address at Layer 3 (L3).

L2 trigger - Information from L2 that informs L3 of particular events before and after L2 handover. 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.

Bicasting/n-casting - The splitting of a stream of packets destined for a MN into two or more streams, and the simultaneous transmission of the streams to PAR and one or more NARs. N/casting is a technique used to reduce packet loss during handover.

ping-ponging - Rapid back and forth movement of an MN between two wireless access points due to failure of L2 handover or frequent handovers. Ping-ponging 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.

2. Simultaneous Bindings

Simultaneous bindings are built into the Mobile IPv4 protocol [2]. To enable multiple simultaneous bindings using Mobile IPv4 the MN simply sends the first normal Registration Request for a care-of address and then sends other Registration messages for additional care-of addresses having the S bit set. The receiver of the Registration Requests (i.e. the HA) will then maintain all these care-of address bindings for the MN contemporarily rather than only servicing the MN at the care-of address in its most recent Registration Request (which would be the case had the S bit not been set). This results one copy of packets being sent to each of the registered care-of addresses (i.e. bicasting or n-casting of packets). This draft extends the

Mobile IPv6 protocol [3] with similar functionality and describes a new Simultaneous Bindings flag for the Fast Binding Update in [1].

Multiple simultaneous bindings and bicasting can be an important tool to decouple L3 handovers from L2 handovers and to reduce packet loss. This mechanism instructs the recipient of an F-BU [1] having the simultaneous bindings flag to make multiple copies of packets destined to the MN and send them to multiple MN care-of addresses before the MN actually moves there. This allows a smoothing of the L3 handover, meaning that packet loss is minimized or even eliminated. Simultaneous bindings are also useful to prevent service disruption due to ping-pong movement as described later.

3. Fast Handovers in Mobile IPv6

The mechanism to obtain fast L3 handovers for Mobile IPv6 is described in [1] and illustrated in Figure 1. This mechanism involves the use of L2 triggers which allow the L3 handover to be anticipated rather than being performed after the L2 handover completion as normal. Fast Handovers are required to ensure that the layer 3 (Mobile IP) handover delay is minimized, thus also minimizing and possibly eliminating the period of service disruption which normally occurs when a MN moves between two ARs. This period of service disruption usually occurs due to the time required by the MN to update its HA after it moves between ARs. During this time period the MN cannot resume or continue communications. Following is a short summary of the Fast Handover mechanism described in [1].

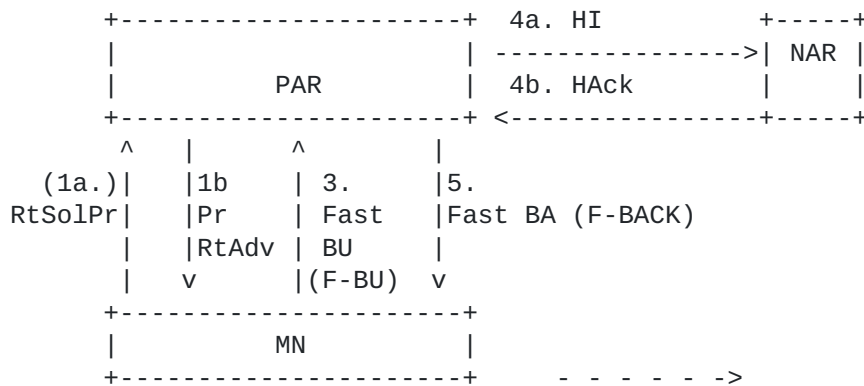


Figure 1 - Fast MIPv6 Handover Protocol

While the MN is connected to its Previous Access Router (PAR) and is about to move to a New Access Router (NAR), the Fast Handovers in Mobile IPv6 requires:

- the MN to obtain a new care-of address at the NAR while connected to the PAR the MN to send a Fast Binding Update (FBU) to its old anchor point (e.g. PAR) to update its binding cache with the MN's

new care-of address.

- the old anchor point (e.g. PAR) to start forwarding packets destined for the MN to NAR.

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

The MN obtains a new care-of address while connected to PAR by means of router advertisements containing information from the NAR (Proxy Router Advertisement, PrRtAdv, which may be sent due to a Proxy Router Solicitation, RtSolPr). The PAR will validate the MN's new COA by sending a Handover Initiate (HI) message to the NAR. Based on the response generated in the Handover Acknowledge (HACK) message, the PAR will either generate a tunnel to the MN's new COA (if the address was valid) or generate a tunnel to the NAR's address (if the address was already in use on the new subnet). If the address was already in use on the new subnet, the NAR will generate a host route for the MN using its old COA.

4. Decoupling L3 Handovers from L2 handovers using Simultaneous Bindings

The mechanisms described in [1] allow the anticipation of the layer 3 handover such that data traffic can be redirected to the MN's new location before it moves there. However it is not simple to determine the correct time to start forwarding between PAR and NAR, which has an impact on how smooth the handover will be. Packet loss will occur if this is performed too late or too early with respect to the time in which the MN detaches from PAR and attaches to NAR. Also, some measure is needed to support the case in which the MN moves quickly back-and-forth between ARs (ping-pong).

In many wireless networks it is not possible to know in advance precisely when a MN will detach from the wireless link to PAR and attach to the one connected to NAR. Therefore determining the exact time when to start forwarding packets between PAR and NAR is not possible. Certain wireless technologies involve layer-2 messages which instruct the MN to handover immediately or simply identify that the MN has already detached/attached. Even if the ARs could extract this information, there may not be sufficient time for the PAR to detect the MN's detachment and start getting packets tunnelled over to NAR before the MN attached to NAR. This is because wireless layer-2 handover times are relatively small (i.e. range from 10's to 100's

ms). Thus a period of service disruption may occur due to this timing uncertainty unless further enhancements are made to the handover mechanism. If the L3 handover is anticipated and the PAR starts forwarding data to NAR upon receipt of the Fast BU in [1] then the period of service disruption will be according to the following:

A = L3 handover anticipation (time difference between the start of the L3 fast handover and the moment in which the L2 handover occurs)

h = L2 handover time (disconnection time due to L2 handover)

Approximate period before MN receives packets again = A + h

It is therefore necessary to decouple layer-3 handover timing from layer-2 handover timing. This can be solved by bicasting or n-casting packets destined to the MN for a short period from the old anchor point (e.g. PAR) to one or more potential future MN locations (e.g. NAR/s) before the MN actually moves there. This means that the handover procedure described previously would be enhanced by having the old anchor point (e.g. PAR) send one copy of packets to the MN's old on-link care-of address and another copy of the packets to the MN's new care-of address (or addresses) connected to NAR. The MN is thus able to receive traffic independently of the exact layer-2 handover timing during the period.

5. Avoiding service disruption due to ping-pong movement

It is possible that the layer-2 handover procedure may fail or terminate abruptly in wireless systems. Therefore a MN which expects to move between PAR and NAR may unexpectedly never complete the layer-2 handover and find itself connected to PAR. Another undesired effect is that the MN could ping-pong between ARs due to layer-2 mobility issues. Both these cases would leave the MN unable to resume communication and have to transmit a new F-BU in [1] before resuming communications.

This may be solved through the use of simultaneous bindings which allow the MN to maintain layer-3 connectivity with the PAR during the affected handover period, thus smoothing the handover. This eliminates the need for continuous transmission of Fast Binding Updates in [1]. It also prevents the period of service disruption from being extended due to the effect of the above link-layer issues on L3 handover.

6. Extensions to the Fast Handover Operations

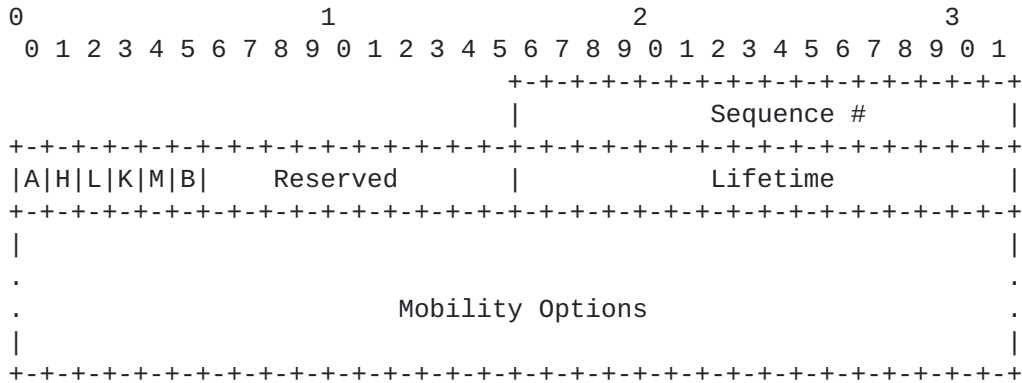
6.1 MN Operation

The MN operation in [1] is affected by the changes introduced by this document. The addition to [1] is that a MN with an existing active binding which receives a new router advertisement (PrRtAdv) MUST be "eager" to establish new bindings. When a MN has at least one existing binding and receives a new PrRtAdv it MUST send a Fast Binding Update (F-BU) with the Simultaneous Bindings flag set (B flag). The new flag is described in [section 8](#). In addition the MN MUST be able to process the new simultaneous bindings option in the Fast Binding Acknowledgement message described in [section 9](#). The lifetime field returned in this option MUST be used by the MN to identify the lifetime of the simultaneous binding requested. Two BU lifetime values will be returned: Bicasting lifetime (in the simultaneous bindings option) and new CoA lifetime (in the BA option) as described in the following sections. The new CoA lifetime (placed in the BA option as specified in [3]) runs in parallel with the Bicasting lifetime. Hence, when the bicasting lifetime ends, the MN will remove the special bicasting information from the Binding Update list and simply keep one entry for the new CoA with the remaining new CoA lifetime.

6.2 HA/MAP/AR Operation

The HA [3], MAP [4] and AR [1] are the possible recipients of a F-BU message. Upon receiving a F-BU message having the B flag set (see [section 8](#)), the HA/MAP/PAR MUST create a new binding cache sub-entry (linked to the original entry for the old CoA) for the MN's new CoA. This sub-entry contains the same fields as normal binding cache entries but it MUST not replace any existing entries for the MN. The new sub-entry will have two lifetimes instead of one: the normal new CoA BU lifetime (sent in the BA) and a Bicasting lifetime set to SHORT_BINDING_LIFETIME (this value is sent in the BA option). The new CoA lifetime runs in parallel with the Bicasting lifetime. Until the Bicasting lifetime expires, the HA/MAP/PAR MUST send a copy of the data destined for the MN to the old CoA and to the new CoA/s in the linked binding cache sub-entry or sub-entries. When the Bicasting lifetime expires, the MAP/HA/PAR MUST remove the bicasting lifetime field and replace the old binding cache entry for the old CoA with the new CoA sub-entry. As a result, the HA/MAP/AR will end up with one entry for the MN's new CoA with the remaining new CoA lifetime. If the MAP/HA/AR receives a valid BU for the new COA without the S flag set before the expiration of the bicasting lifetime it must follow standard procedures of [1] and replace the existing binding cache entry.

7. Simultaneous Bindings Flag in Fast Binding Update (F-BU) message




Description of the flag added to the F-BU option already defined in [1]:

B When set indicates a request for bicasting all packets to both COAs of the MN (in the source address field and the alternate-CoA suboption). This BU will add another COA to the Binding Cache.

8. Simultaneous Bindings option for Fast Binding Acknowledgement

(F-BA) message



Option Type	TBD
Option Len	TBD
Status	Indicates success (0) or failure (128 and above).
Lifetime	The bicasting lifetime for the simultaneous binding requested in the F-BU. This value MUST be used by the MN to record the validity of this binding in its binding update list.

The alignment requirement for this option is 2n+2.

9. Multiple copies of packets received at AR

If the MN has simultaneous active bindings with HA/MAP/AR, it could (but preferably should not) receive multiple copies of the same traffic directed to it. The use of simultaneous bindings does not mean that the MN is receiving packets contemporarily from multiple sources. This depends on the characteristics of the access (L2) technology. The bicasting of packets involves sending a copy of the data to the AR which the MN is moving to (the NAR). Until the MN actually completes the L2 handover to the NAR and fully establishes the new L2 link, the NAR MAY receive packets for a MN to which it does not have a direct link layer connection. If the new AR is aware that the MN is performing a handover (due to earlier reception of the HI message) the AR MAY:

- drop all packets for the MN,
- drop some packets, based on local policies, or
- buffer packets for the MN.

The choice of which action to take may depend on the type of traffic involved (e.g. real-time or non real-time), but this is outside the scope of this document. The AR MAY also in parallel attempt to establish a link-layer connection with the MN. However an AR MUST NOT send ICMP Destination Unreachable messages if it drops packets or is unable to deliver the received IP packets due to unavailability of direct layer connection with the MN. This is because a copy of the packets would be dropped, but the MN is still receiving a copy of the packets through the PAR. Note that the MN may also select which flows need bicasting by adding a Flow movement option [7] to the simultaneous binding update. Therefore the simultaneous bindings mechanism may only be applied to traffic types that require this service.

10. Reception of multiple copies in the MN

In some scenarios it may be possible that the MN receives more than one copy of the same packet. Generally, Internet routing mechanisms cannot guarantee the delivery of a single copy of an IP packet to a node. However some TCP congestion avoidance implementations are known to react negatively to the reception of 3 duplicate acknowledgements. The Eifel detection and response algorithms in [5] and [6] address this problem. When using [5] and [6] bicasting should not cause any negative performance impacts for TCP.

11. References

Normative

- [1] R. Koodli (Editor) et al, "Fast Handovers for Mobile IPv6", [RFC 4068](#), July 2005.
- [3] D. Johnson, C. Perkins and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [4] H. Soliman, C. Castelluccia, K. El Malki and L. Bellier, "Hierarchical Mobile IPv6 mobility management (HMIPv6)", [draft-ietf-mipshop-hmipv6-04.txt](#), work in progress, December 2004.

Informative

- [2] C. Perkins (Editor), "IP Mobility Support for IPv4", [RFC 3220](#), Jan 2002.
- [5] R. Ludwig, "The Eifel Detection Algorithm for TCP", [RFC 3522](#), April 2003.
- [6] R. Ludwig and A. Gurtov, "The Eifel response algorithm for TCP", [RFC 4015](#), February 2005.

12. Authors' Addresses

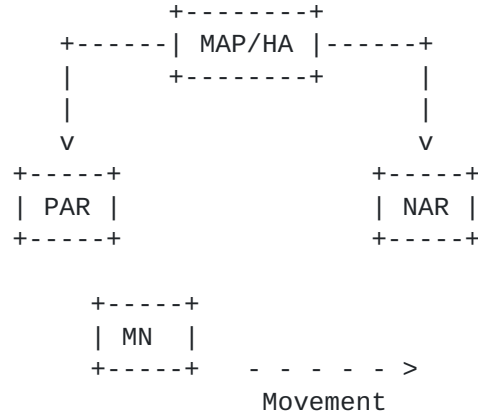
The authors may be contacted at the addresses below:

Karim El Malki
Ericsson AB
Phone: +46 8 7195803
E-mail: Karim.El-Malki@ericsson.com

Hesham Soliman
Flarion
E-mail: H.Soliman@flarion.com

[Appendix A](#) - Timing Calculations for bicasting

Example 1



This is the case specified by [1] with the extension of using the MAP from [4].

- A = anticipation time (F-BU is sent from MN at time t-A, where t is the time when the MN actually hands-off at L2)
- h = handover time (L2 only)
- D1 = MN to MAP delay (through PAR)
- D2 = MN to MAP delay (through NAR)
- p = F-BU and routing table processing time in the MAP and MN

To achieve zero L3 service disruption it is necessary for the time period between starting the fast handover and the MN completing the L2 handover to be greater than or equal to the time it takes for traffic to reach the MN at its new link (through NAR). This is represented by the following formula:

$$(A+h) \geq ((D1+D2)+p)$$

Assuming that $p \ll (D1+D2)$ this can be simplified to:

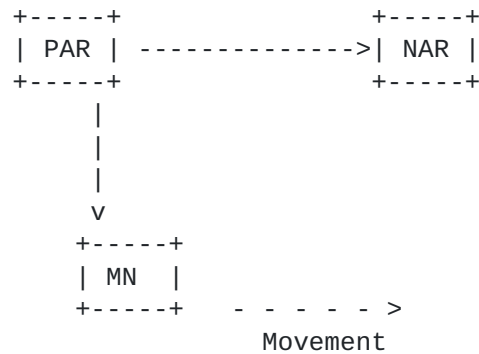
$$(A+h) \geq (D1+D2)$$

To achieve maximum performance from simultaneous bindings it is necessary for the above relation to hold.

The Anticipation time (A) is important and needs to be calculated appropriately for the link-layer being used. Depending on the L2 this may need engineering to synchronize the L2 and L3 handovers.

Once the MN has moved to NAR, it will be receiving traffic delayed by (D2-D1) with respect to when it was attached to PAR. To smooth this delay variation (jitter), which may be a problem for real-time services, it may be necessary to implement a smoothing buffer at NAR.

Example 2



When the MAP/HA/PAR are one entity (as considered in [1]), the following calculations apply.

- A = anticipation time (F-BU is sent from MN at time t-A, where t is the time when the MN actually hands-off at L2)
- h = handover time (L2 only)
- d = MN to AR delay (assume constant as MN moves ARs)
- L = PAR to NAR delay

As previously, the following must be true for the simultaneous bindings to yield zero L3 disruption:

$$(A+h) \geq (d+L+d)$$

$$\Rightarrow (A+h) \geq (2d+L)$$

The Anticipation time (A) is important and needs to be calculated appropriately for the link-layer being used. Depending on the L2 this may need engineering to synchronise the L2 and L3 handovers.

Once the MN has moved to NAR, it will be receiving traffic delayed by an amount L with respect to when it was attached to PAR. To smooth this delay variation (jitter), which may be a problem for real-time services, it may be necessary to implement a smoothing buffer at NAR.

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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

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

Intellectual Property

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

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

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

Acknowledgement

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

