

Mobile IPv6 Fast Handovers for 802.11 Networks

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document describes how a Mobile IPv6 Fast Handover could be implemented on link layers conforming to the 802.11 suite of specifications.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	2
2. Terminology	2
3. Deployment Architectures for Mobile IPv6 on 802.11	3
4. 802.11 Handovers in Detail	5
5. FMIPv6 Message Exchanges	7
6. Beacon Scanning and NAR Discovery	8
7. Scenarios	9
7.1. Scenario 1abcdef23456g	9
7.2. Scenario ab123456cdefg	10
7.3. Scenario 123456abcdefg	10
8. Security Considerations	10
9. Conclusions	12
10. References	13
10.1. Normative References	13
10.2. Informative References	13
11. Acknowledgements	13

1. Introduction

The Mobile IPv6 Fast Handover protocol [2] has been proposed as a way to minimize the interruption in service experienced by a Mobile IPv6 node as it changes its point of attachment to the Internet. Without such a mechanism, a mobile node cannot send or receive packets from the time that it disconnects from one point of attachment in one subnet to the time it registers a new care-of address from the new point of attachment in a new subnet. Such an interruption would be unacceptable for real-time services such as Voice-over-IP.

The basic idea behind a Mobile IPv6 fast handover is to leverage information from the link-layer technology to either predict or rapidly respond to a handover event. This allows IP connectivity to be restored at the new point of attachment sooner than would otherwise be possible. By tunneling data between the old and new access routers, it is possible to provide IP connectivity in advance of actual Mobile IP registration with the home agent or correspondent node. This allows real-time services to be reestablished without waiting for such Mobile IP registration to complete. Because Mobile IP registration involves time-consuming Internet round-trips, the Mobile IPv6 fast handover can provide for a smaller interruption in real-time services than an ordinary Mobile IP handover.

The particular link-layer information available, as well as the timing of its availability (before, during, or after a handover has occurred), differs according to the particular link-layer technology in use. This document gives a set of deployment examples for Mobile IPv6 Fast Handovers on 802.11 networks. We begin with a brief overview of relevant aspects of basic 802.11 [3]. We examine how and when handover information might become available to the IP layers that implement Fast Handover, both in the network infrastructure and on the mobile node. Finally, we trace the protocol steps for Mobile IPv6 Fast Handover in this environment.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1].

2. Terminology

This document borrows all of the terminology from Mobile IPv6 Fast Handovers [2], with the following additional terms from the 802.11 specification [3] (some definitions slightly modified for clarity):

Access Point (AP): Any entity that has station functionality and provides access to the distribution services, via the wireless medium (WM) for associated stations.

Association: The service used to establish access point/station (AP/STA) mapping and enable STA access to the Distribution System.

Basic Service Set (BSS): A set of stations controlled by a single coordination function, where the coordination function may be centralized (e.g., in a single AP) or distributed (e.g., for an ad hoc network). The BSS can be thought of as the coverage area of a single AP.

Distribution System (DS): A system used to interconnect a set of basic service sets (BSSs) and integrated local area networks (LANs) to create an extended service set (ESS).

Extended Service Set (ESS): A set of one or more interconnected basic service sets (BSSs) and integrated local area networks (LANs) that appears as a single BSS to the logical link control layer at any station associated with one of those BSSs. The ESS can be thought of as the coverage area provided by a collection of APs all interconnected by the Distribution System. It may consist of one or more IP subnets.

Station (STA): Any device that contains an IEEE 802.11 conformant medium access control (MAC) and physical layer (PHY) interface to the wireless medium (WM).

3. Deployment Architectures for Mobile IPv6 on 802.11

In this section, we describe the two most likely relationships between Access Points (APs), Access Routers (ARs), and IP subnets that are possible in an 802.11 network deployment. In this document, our focus is mainly on the infrastructure mode [3] of 802.11. Usually, a given STA is associated with one and only one AP at any given instant; however, implementations are possible [4] where multiple associations per STA may be maintained as long as the APs are connected to disjoint DSs. An STA may be in communication with an AP only when radio propagation conditions permit. Note that, as with any layer-2 technology, handover from one layer-2 point of attachment (AP) to another does not necessarily mean a change of AR or subnet.



Figure 1. An 802.11 deployment with relay APs.

Figure 1 depicts a typical 802.11 deployment with two IP subnets, each with three Access Routers and five Access Points. Note that the APs in this figure are acting as link-layer relays, which means that they transport Ethernet-layer frames between the wireless medium and the subnet. Note that APs do not generally implement any particular spanning tree algorithm, yet are more sophisticated than simple bridges that would relay all traffic; only traffic addressed to STAs known to be associated on a given AP would be forwarded. Each subnet is on top of a single LAN or VLAN, and we assume in this example that APs 6-10 cannot reach the VLAN on which Subnet 1 is implemented. Note that a handover from AP1 to AP2 does not require a change of AR (here we assume the STA will be placed on the same VLAN during such a handoff) because all three ARs are link-layer reachable from an STA connected to any AP1-5. Therefore, such handoffs would not require IP-layer mobility management, although some IP-layer signaling may be required to determine that connectivity to the existing AR is still available. However, a handover from AP5 to AP6 would require a change of AR, because AP6 cannot reach the VLAN on which Subnet 1 is implemented and therefore the STA would be attaching to a different subnet. An IP-layer handover mechanism would need to be invoked in order to provide low-interruption handover between the two ARs.




Figure 2. An 802.11 deployment with integrated APs/ARs.

Figure 2 depicts an alternative 802.11 deployment where each AP is integrated with exactly one AR on a disjoint VLAN. In this case, every change of AP would result in a necessary change of AR, which

would require some IP-layer handover mechanism to provide for low-interruption handover between the ARs. Also, the AR shares a MAC-layer identifier with its attached AP.

In the next section, we examine the steps involved in any 802.11 handover. Subsequent sections discuss how these steps could be integrated with an IP-layer handover mechanism in each of the above deployment scenarios.

4. 802.11 Handovers in Detail

An 802.11 handover takes place when an STA changes its association from one AP to another ("re-association"). This process consists of the following steps:

0. The STA realizes that a handoff is necessary due to degrading radio transmission environment for the current AP.
1. The STA performs a scan to see what APs are available. The result of the scan is a list of APs together with physical layer information, such as signal strength.
2. The STA chooses one of the APs and performs a join to synchronize its physical and MAC-layer timing parameters with the selected AP.
3. The STA requests authentication with the new AP. For an "Open System", such authentication is a single round-trip message exchange with null authentication.
4. The STA requests association or re-association with the new AP. A re-association request contains the MAC-layer address of the old AP, while a plain association request does not.
5. If operating in accordance with 802.11i [6], the STA and AP would execute 802.1X EAP-on-LAN procedures to authenticate the association (step 3 would have executed in "Open System" mode).
6. The new AP sends a Layer 2 Update frame on the local LAN segment to update the learning tables of any connected Ethernet bridges.

Although we preface step 1 with step 0 for illustration purposes, there is no standardized trigger for step 1. It may be performed as a result of decaying radio conditions on the current AP or at other times as determined by local implementation decisions. Some network interface cards (NICs) may do scanning in the background, interleaving scans between data packets. This decreases the time required to roam if the performance of the current AP proves

unsatisfactory, but it imposes more of a burden on the AP, since typically the STA places it in power-save mode prior to the scan, then once the scan is complete, returns to the AP channel in order to pick up queued packets. This can result in buffer exhaustion on the AP and attendant packet loss.

During step 2, the STA performs rate adjustment where it chooses the best available transmission rate. Rate adjustment can be quite time-consuming as well as unpredictable.

Note that in some existing 802.11 implementations, steps 1-4 are performed by firmware in rapid succession (note that even in these implementations step 3 is sometimes performed in a host driver, especially for newer implementations). This might make it impossible for the host to take any actions (including sending or receiving IP packets) before the handover is complete. In other 802.11 implementations, it is possible to invoke the scan (step 1) and join (step 2) operations independently. This would make it possible to, e.g., perform step 1 far in advance of the handover and perhaps in advance of any real-time traffic. This could substantially reduce the handover latency, as one study has concluded that the 802.11 beacon scanning function may take several hundred milliseconds to complete [8], during which time sending and receiving IP packets is not possible. However, scanning too far in advance may make the information out-of-date by the time of handover, which would cause the subsequent joint operation to fail if radio conditions have changed so much in the interim that the target AP is no longer reachable. So, a host may choose to do scanning based on, among other considerations, the age of the previously scanned information. In general, performing such subsequent scans is a policy issue that a given implementation of FMIPv6 over 802.11 must consider carefully.

Even if steps 1 and 2 are performed in rapid succession, there is no guarantee that an AP found during step 1 will be available during step 2 because radio conditions can change dramatically from moment to moment. The STA may then decide to associate with a completely different AP. Often, this decision is implemented in firmware and the attached host would have no control over which AP is chosen. However, tools such as the host AP driver [10] offer full control over when and to which AP the host needs to associate. Operation as an Independent BSS (IBSS) or "ad-hoc mode" [3] may also permit the necessary control, although in this latter case attachment to an infrastructure AP would be impossible. Implementers can make use of such tools to obtain the best combination of flexibility and performance.

The coverage area of a single AP is known as a Basic Service Set (BSS). An Extended Service Set (ESS) is formed from a collection of APs that all broadcast the same ESSID. Note that an STA would send a re-association (which includes both the old and new AP addresses) only if the ESSID of the old and new APs are the same.

A change of BSS within an ESS may or may not require an IP-layer handover, depending on whether the APs can send packets to the same IP subnets. If an IP-layer handover is required, then FMIPv6 can decrease the overall latency of the handover. The main goal of this document is to describe the most reasonable scenarios for how the events of an 802.11 handover may interleave with the message exchanges in FMIPv6.

5. FMIPv6 Message Exchanges

An FMIPv6 handover nominally consists of the following messages:

- a. The mobile node (MN) sends a Router Solicitation for Proxy (RtSolPr) to find out about neighboring ARs.
- b. The MN receives a Proxy Router Advertisement (PrRtAdv) containing one or more [AP-ID, AR-Info] tuples.
- c. The MN sends a Fast Binding Update (FBU) to the Previous Access Router (PAR).
- d. The PAR sends a Handover Initiate (HI) message to the New Access Router (NAR).
- e. The NAR sends a Handover Acknowledge (HACK) message to the PAR.
- f. The PAR sends a Fast Binding Acknowledgement (FBack) message to the MS on the new link. The FBack is also optionally sent on the previous link if the FBU was sent from there.
- g. The MN sends Fast Neighbor Advertisement (FNA) to the NAR after attaching to it.

The MN may connect to the NAR prior to sending the FBU if the handover is unanticipated. In this case, the FNA (step g) would contain the FBU (listed as step c above) and then steps d, e, and f would take place from there.

6. Beacon Scanning and NAR Discovery

The RtSolPr message is used to request information about the router(s) connected to one or more APs. The APs are specified in the New Access Point Link-Layer Address option in the RtSolPr and associated IP-layer information is returned in the IP Address Option of the PrRtAdv [2]. In the case of an 802.11 link, the link-layer address is the BSSID of some AP.

Beacon scanning (step 1 from Section 4) produces a list of available APs along with signal strength information for each. This list would supply the necessary addresses for the New Access Point Link-Layer Address option(s) in the RtSolPr messages. To obtain this list, the host needs to invoke the MLME-SCAN.request primitive (see Section 10.3.2.1 of the 802.11 specification [3]). The BSSIDs returned by this primitive are the link-layer addresses of the available APs.

Because beacon scanning takes on the order of a few hundred milliseconds to complete, and because it is generally not possible to send and receive IP packets during this time, the MN needs to schedule these events with care so that they do not disrupt ongoing real-time services. For example, the scan could be performed at the time the MN attaches to the network prior to any real-time traffic. However, if the interval between scanning and handover is too long, the neighbor list may be out of date. For example, the signal strengths of neighboring APs may have dramatically changed, and a handover directed to the apparently best AP from the old list may fail. If the handover is executed in firmware, the STA may even choose a new target AP that is entirely missing from the old list (after performing its own scan). Both cases would limit the ability of the MN to choose the correct NAR for the FBU in step c during an anticipated handover. Ongoing work in the IEEE 802.11k task group may address extensions that allow interleaving beacon scanning with data transmission/reception along with buffering at APs to minimize packet loss.

Note that, aside from physical layer parameters such as signal strength, it may be possible to obtain all necessary information about neighboring APs by using the wildcard form of the RtSolPr message. This would cause the current access router to return a list of neighboring APs and would not interrupt ongoing communication with the current AP. This request could be made at the time the MN first attaches to the access router and periodically thereafter. This would enable the MN to cache the necessary [AP-ID, AR-Info] tuples and might enable it to react more quickly when a handover becomes necessary due to a changing radio environment. However, because the information does not include up-to-date signal strength, it would not enable the MN to predict accurately the next AP prior to a handover.

Also, if the scale of the network is such that a given access router is attached to many APs, then it is possible that there may not be room to list all APs in the PrRtAdv.

The time taken to scan for beacons is significant because it involves iteration through all 802.11 channels and listening on each one for active beacons. A more targeted approach would allow the STA to scan, e.g., only one or two channels of interest, which would provide for much shorter interruption of real-time traffic. However, such optimizations are currently outside the scope of 802.11 specifications.

7. Scenarios

In this section, we look at a few of the possible scenarios for using FMIPv6 in an 802.11 context. Each scenario is labeled by the sequence of events that take place, where the numbered events are from Section 4 and the lettered events are from Section 5. For example, "1abcde23456fg" represents step 1 from Section 4 followed by steps a-e from Section 5 followed by steps 2-6 from Section 4 followed by steps f and g from Section 5. This is the sequence where the MN performs a scan, then the MN executes the FMIPv6 messaging to obtain NAR information and send a binding update, then the PAR initiates HI/HACK exchange, then the 802.11 handover completes, and finally the HACK is received at the PAR and the MN sends an FNA.

Each scenario is followed by a brief description and discussion of the benefits and drawbacks.

7.1. Scenario 1abcdef23456g

This scenario is the predictive mode of operation from the FMIPv6 specification. In this scenario, the host executes the scan sometime prior to the handover and is able to send the FBU prior to handover. Only the FNA is sent after the handover. This mode of operation requires that the scan and join operations (steps 1 and 2) can be performed separately and under host control, so that steps a-f can be inserted between 1 and 2. As mentioned previously, such control may be possible in some implementations [10] but not in others.

Steps 1ab may be executed far in advance of the handover, which would remove them from the critical path. This would minimize the service interruption from beacon scanning and allow at least one RtSolPr/PrRtAdv exchange to complete so that the host has link-layer information about some NARs. Note that if steps ab were delayed until handover is imminent, there would be no guarantee that the RtSolPr/PrRtAdv exchange would complete especially in a radio environment where the connection to the old AP is deteriorating

rapidly. However, if there were a long interval between the scan and the handover, then the FBU (step c) would be created with out-of-date information. There is no guarantee that the MN will actually attach to the desired new AP after it has sent the FBU to the oAR, because changing radio conditions may cause NAR to be suddenly unreachable. If this were the case, then the handover would need to devolve into one of the reactive cases given below.

7.2. Scenario ab123456cdefg

This is the reactive mode of operation from the FMIPv6 specification. This scenario does not require host intervention between steps 1 and 2.

However, it does require that the MN obtain the link-layer address of NAR prior to handover, so that it has a link-layer destination address for outgoing packets (default router information). This would then be used for sending the FNA (with encapsulated FBU) when it reaches the new subnet.

7.3. Scenario 123456abcdefg

In this scenario, the MN does not obtain any information about the NAR prior to executing the handover. It is completely reactive and consists of soliciting a router advertisement after handover and then sending an FNA with encapsulated FBU immediately.

This scenario may be appropriate when it is difficult to learn the link-layer address of the NAR prior to handover. This may be the case, e.g., if the scan primitive is not available to the host and the wildcard PrRtAdv form returns too many results. It may be possible to skip the router advertisement/solicitation steps (ab) in some cases, if it is possible to learn the NAR's link-layer address through some other means. In the deployment illustrated in Figure 2, this would be exactly the new AP's MAC-layer address, which can be learned from the link-layer handover messages. However, in the case of Figure 1, this information must be learned through router discovery of some form. Also note that even in the case of Figure 2, the MN must somehow be made aware that it is in fact operating in a Figure 2 network and not a Figure 1 network.

8. Security Considerations

The security considerations applicable to FMIPv6 are described in the base FMIPv6 specification [2]. In particular, the PAR must be assured of the authenticity of the FBU before it begins to redirect user traffic. However, if the association with the new AP is not

protected using mutual authentication, it may be possible for a rogue AP to fool the MN into sending an FBU to the PAR when it is not in its best interest to do so.

Note that step 6 from Section 4 installs layer-2 forwarding state that can redirect user traffic and cause disruption of service if it can be triggered by a malicious node.

Note that step 3 from Section 4 could potentially provide some security; however, due to the identified weaknesses in Wired Equivalent Privacy (WEP) shared key security [9] this should not be relied upon. Instead, the Robust Security Network [6] will require the STA to undergo 802.1X Port-Based Network Access Control [5] before proceeding to steps 5 or 6. 802.1X defines a way to encapsulate Extensible Authentication Protocol (EAP) on 802 networks (EAPOL, for "EAP over LANs"). With this method, the client and AP participate in an EAP exchange that itself can encapsulate any of the various EAP authentication methods. The EAPOL exchange can output a Master Session Key (MSK) and Extended Master Session Key (EMSK), which can then be used to derive transient keys, which in turn can be used to encrypt/authenticate subsequent traffic. It is possible to use 802.1X pre-authentication [6] between an STA and a target AP while the STA is associated with another AP; this would enable authentication to be done in advance of handover, which would allow faster resumption of service after roaming. However, because EAPOL frames carry only MAC-layer instead of IP-layer addresses, this is currently only specified to work within a single VLAN, where IP-layer handover mechanisms are not necessarily needed anyway. In the most interesting case for FMIPv6 (roaming across subnet boundaries), the 802.1X exchange would need to be performed after handover to the new AP. This would introduce additional handover delay while the 802.1X exchange takes place, which may also involve round-trips to RADIUS or Diameter servers. The EAP exchange could be avoided if a preexisting Pairwise Master Key (PMK) is found between the STA and the AP, which may be the case if the STA has previously visited that AP or one that shares a common back-end infrastructure.

Perhaps faster cross-subnet authentication could be achieved with the use of pre-authentication using an IP-layer mechanism that could cross subnet boundaries. To our knowledge, this sort of work is not currently under way in the IEEE. The security considerations of these new approaches would need to be carefully studied.

9. Conclusions

The Mobile IPv6 Fast Handover specification presents a protocol for shortening the period of service interruption during a change in link-layer point of attachment. This document attempts to show how this protocol may be applied in the context of 802.11 access networks.

Implementation of FMIPv6 must be done in the context of a particular link-layer implementation, which must provide the triggers for the FMIPv6 message flows. For example, the host must be notified of such events as degradation of signal strength or attachment to a new AP.

The particular implementation of the 802.11 hardware and firmware may dictate how FMIPv6 is able to operate. For example, to execute a predictive handover, the scan request primitive must be available to the host and the firmware must execute join operations only under host control [10], not autonomously in response to its own handover criteria. Obtaining the desired PrRtAdv and sending an FBU immediately prior to handover requires that messages be exchanged over the wireless link during a period when connectivity is degrading. In some cases, the scenario given in Section 7.1 may not complete successfully or the FBU may redirect traffic to the wrong NAR. However, in these cases the handover may devolve to the scenario from Section 7.2 or the scenario from Section 7.3. Ultimately, falling back to basic Mobile IPv6 operation [7] and sending a Binding Update directly to the Home Agent can be used to recover from any failure of the FMIPv6 protocol.

10. References

10.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Koodli, R., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.
- [3] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", ANSI/IEEE Std 802.11, 1999 Edition.
- [4] Bahl, P., Bahl, P., and Chandra, R., "MultiNet: Enabling Simultaneous Connections to Multiple Wireless Networks Using a Single Radio", Microsoft Tech Report, MSR-TR-2003-46, June 2003.
- [5] "Port-Based Network Access Control", IEEE Std 802.1X-2004, July 2004.
- [6] "Medium Access Control (MAC) Security Enhancements", IEEE Std 802.11i-2004, July 2004.
- [7] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

10.2. Informative References

- [8] Mitra, A., Shin, M., and Arbaugh, W., "An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process", CS-TR-4395, University of Maryland Department of Computer Science, September 2002.
- [9] Borisov, N., Goldberg, I., and Wagner, D., "Intercepting Mobile Communications: The Insecurity of 802.11", Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking, July 2001, pp. 180-188.
- [10] Malinen, J., "Host AP driver for Intersil Prism2/2.5/3 and WPA Supplicant", <http://hostap.epitest.fi/>, July 2004.

11. Acknowledgements

Thanks to Bob O'Hara for providing explanation and insight on the 802.11 standards. Thanks to James Kempf, Erik Anderlind, Rajeev Koodli, and Bernard Aboba for providing comments on earlier versions.

Author's Address

Pete McCann
Lucent Technologies
Rm 9C-226R
1960 Lucent Lane
Naperville, IL 60563

Phone: +1 630 713 9359
Fax: +1 630 713 1921
EMail: mccap@lucent.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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

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

Intellectual Property

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

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

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

Acknowledgement

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

