Informational   Alper E. Yegin,
                                                                 Editor
Internet Draft   Daichi Funato
Document: draft-manyfolks-l2-mobilereq-02.txt   Karim El Malki
Expires: December 2002   Youngjune Gwon
June 2002   James Kempf
                                                    Mattias Pettersson
                                                          Phil Roberts
                                                        Hesham Soliman
                                                     Atsushi Takeshita

Supporting Optimized Handover for IP Mobility - Requirements for Underlying Systems

Status of this Memo

Abstract

A critical factor in achieving good performance for IP mobility protocols is the design of L2 handover. Handover occurs when a Mobile Node moves from one radio Access Point to another. If the new radio Access Point is associated with a new subnet, a change in routing reachability may occur and require L3 protocol action on the part of the Mobile Node or Access Routers. If no change in subnet occurs, the Access Point may still need to take some action to inform the Access Router about a change in on-link reachability. In either case, prompt and timely information from L2 to the parties involved about the sequencing of handover can help optimize handover at the IP level. This draft discusses requirements for an L2 handover protocol or API to support optimized handover.

Table of Contents

1.0 Introduction

An important consideration in the design of IP mobility protocols is handover. A moving Mobile Node (MN) may irregularly need to change the terrestrial radio Access Point (AP) with which it is communicating. The change in L2 connectivity to a new AP may cause a change in IP routing reachability, and thus require either the MN or the Access Routers (ARs) to perform actions that update routing information for the MN. Even if no change in subnet occurs, the APs may still need to communicate the change in on-link reachability to the local AR. In order for handover to occur, candidate APs must be identified and a target AP must be selected [9]. Once this process has been complete, the handover process can begin.

Several protocol designs have been advanced for Mobile IP that seek to reduce the amount of handover latency at L3 [3] [4]. These protocols depend on obtaining timely information from the L2 protocol about the progress of handover. An additional beneficiary of timely handover progress information is context transfer [5]. Context transfer involves moving context information (QoS, header compression, authentication, etc.) from the old AR to the new. By moving such context information, the ARs can avoid requiring the MN to set up all the context information from scratch, considerably reducing the amount of time necessary to set up basic network service on the new subnet. If handover progress information is available from L2, context transfer can proceed more quickly.

This document discusses requirements on underlying systems for supporting optimized IP mobility, in particular, handover. While the document has been written with existing Mobile IP work in mind, it should be applicable to any protocol that can benefit from knowledge

about handover sequencing to facilitate mobility. Requirements for assisting in handover between two APs on the same subnet, between two ARs on different subnets, and for context transfer between ARs are discussed.

2.0 Terminology

The following terms are used in this document.

Access Point (AP)

A Layer 2 (L2) access entity, e.g. a radio transceiver station, that is connected to one or more Access Routers. Its primary function is to provide MNs an L2 wireless link via its specific air-interface technology.

Access Router (AR)

A Layer 3 (L3) IP router, residing in an access network and connected to one or more Access Points. An AR is the first hop router for a MN.

L2 Handover

Change of MN's link layer connection from one AP to another. No change in off-subnet routing reachability information is required if both APs are part of the same subnet.

L3 Handover

Change of MN's routable address from one AR to another. An L3 handover results in a change in the MN's routing reachability, that will require action on the part of the IP mobility protocol running in the MN and/or ARs.

3.0 L2 Trigger Definition

This section discusses defining L2 triggers that provide information on the sequencing of handover. An L2 trigger is not associated with any specific L2 but rather is based on the kind of L2 information that is or could be available from a wide variety of radio link protocols.

3.1. What is an L2 Trigger?

An L2 trigger is an abstraction of a notification from L2 (potentially including parameter information) that a certain event has happened or is about to happen. The trigger may be implemented in a variety of ways. Some examples are:

- The L2 driver may allow the IP stack to register a callback function that is called when the trigger fires. The parameters associated with the trigger are delivered to the callback.

- The operating system may allow a thread to call into a system call for the appropriate trigger or triggers. The system call blocks until a particular trigger has fired, then it returns with parameter information available in some way (return value of the system call, file descriptor, etc.).

- The trigger may consist of a protocol for transferring the trigger notification and parameter information at either L2 or L3 between the new AP or AR and the old AP or AR. The parameter information is included as part of the protocol. This allows the IP stack on a separate machine to react to the trigger. The IAPP protocol [6] is an example of such a protocol.

- The trigger information may be available within the operating system kernel to the IP stack from the driver as an out of band communication.

In any case, the implementation details of how the information involved in an L2 trigger are transferred to the IP mobility protocol are likely to color how the mobility protocol is implemented on top of that L2, but they should not influence the specification of the abstract L2 triggers themselves.

3.2. Information in an L2 Trigger

There are three types of information involved in defining an L2 trigger:

1. The event that causes the L2 trigger to fire,

2. The IP entity that receives the trigger,

3. The parameters delivered with the trigger.

The IP entities that can receive the trigger depend on the particular IP mobility protocol in use. Here are some possible IP entities, based on work done with L2 triggers and Mobile IP:

MN   The MN may receive an L2 trigger allowing it to start or conclude a mobile controlled handover.

FA In Mobile IPv4, the Foreign Agent (FA) is located on the last hop before the wireless link. The last hop can be either an AP or AR or even a separate host. An FA can make use of triggers to start or conclude network controlled handover.

AR The AR can obtain an L2 trigger directly from the

wireless link if one of its interfaces is on the link (that is, the AR is also an AP), or it can obtain an L2 trigger indirectly by L2 or L3 protocol messages from the AP.

4.0 Requirements for L2 Triggers

L2 handover, L3 handover, and context transfer events and related protocols would benefit from a collection of L2 triggers. Some of these protocols directly rely on the existence of certain triggers, and perform better when others are available. As such, L2 triggers should be designed to meet these protocolsÆ needs.

4.1. L2 Handovers

On the face of it, specifying requirements for pure L2 handover (i.e. no change in IP routing reachability) might seem out of scope for IETF. Existing wireless networks typically have special L2 AP-AR interfaces with L2 address update built in. For these systems, L2 triggers are unnecessary.

However, current trends in wireless networking suggest that future wireless networks will consist of a variety of heterogeneous wireless APs bridged into the wired network, potentially on the same subnet. A change in wireless AP, either between an AP supporting one wireless link technology and an AP supporting another, or between two APs supporting the same wireless technology, necessarily results in a change in the on-subnet reachability. Packet delivery within the subnet can be optimized if this information can be propagated to the AR, so it can update its on-subnet L2 address to IP address mapping.

In addition, the old AP may benefit from a notification that the MN has moved in the event it is not involved in the handover (as is the case with some WLAN radio protocols), by allowing the old AP to more quickly de-allocate resources dedicated to the moved MN. Some radio link protocols already define IP-based L2 trigger protocols for this purpose [6]. When APs supporting multiple radio technologies on a single subnet are involved, however, interoperability suffers if there is no L2-independent way of reporting on-link movement.

4.2. L3 Handovers

Low latency handover protocol designs for Mobile IPv4 and Mobile IPv6 [3] [4] rely on the existence of certain L2 triggers. Either the MN or the AP/AR needs to receive an indication that the handoff is imminent for these L3 mobility protocols to work. This trigger must be received by the MN for mobile-controlled handovers, and received by the AP/AR for network-controlled handovers. Timely receipt of this trigger is needed as protocol signaling needs to take place in parallel with the handoff. Protocol signaling over the current link should be completed prior to loss of connectivity.

Additional triggers that indicate the link tear down and establishment can be used to indicate departure and arrival of a MN at AP/ARs. Such indications can replace L3 signal exchange and therefore expedite the process. An L2 that supports a collection of such triggers is a good candidate for a high performance Mobile IP implementation.

4.3. Context Transfers

Context transfer (CT) is a relatively new issue for supporting seamless mobility between two ARs or FAs that provide access to a mobile node. Although lacking a "de facto" CT protocol specification at this time, plausible approaches toward a CT framework are well described in [7] [8]. Conceptually, CT can take place before, during, or after handover. Exactly when and how CT takes place is highly dependent on the type of context being transferred.

L2 triggers are used to initiate the context transfer operation. Early notification of handovers is essential to having sufficient time to complete the required protocol signaling. Also link establishment trigger can be used for activating the state related to a context.

5.0 L2 Triggers

Based on the L2 handover, L3 handover, and context transfer protocolsÆ needs, underlying systems are required to implement basic L2 triggers as outlined in Table 1.

The description for a trigger contains the trigger name, the L2 handover event causing the trigger to fire, what entities receive the trigger, and parameters, if any. The recipient is qualified by the IP mobility protocol in which the recipient plays a role. If the recipient does not have AP functionality (i.e., the recipient does not have an interface directly on the wireless link), the trigger information must be conveyed from the AP where it occurs to the recipient by an L2 or L3 protocol [10].

L2 Trigger	Event	Recipient	Parameters
Link Up	When the L2 link comes up.	AP/AR/FA MN	MN L2 address to AP/AR/FA AP/AR/FA L2 address to MN
Link Down	When the L2 link goes down.	AP/AR/FA MN	MN L2 address to AP/AR/FA AP/AR/FA L2 address to MN Boolean cause (inadvertent/ deliberate)
Source Trigger	Sufficiently before L2 handover start for pre-handover L3 message exchange across the wired and/or wireless link	oAP/oAR/oFA	nAP/nAR/nFA L2 address that can be mapped to an IP address MN L2 address
Target Trigger	Sufficiently before L2 handover finish for pre-handover L3 message exchange across the wired	nAP/nAR/nFA	oAP/oAR/oFA L2 address that can be mapped to an IP address

	and/or wireless link.		MN L2 address
Mobile Trigger	Sufficiently before L2 handover start for pre-handover L3 message exchange across the wired and/or wireless link	MN	nAP/nAR/nFA L2 address that can be mapped to an IP address

Table 1. L2 trigger requirements

When a source trigger or target trigger is not followed by a link up or down trigger, this sequence of events can be interpreted as an indication of a failed handover.

6.0 Benefits of L2 Triggers for Other Systems

While the primary purpose of L2 triggers described in this draft is to aid L2 mobility optimization, L2 triggers can also benefit networks without Mobile IP or other IP mobility protocol support.

For example, IP addresses may change due to stateless or stateful address configuration whenever hosts are unplugged from the network or re-plugged into a different subnet.

Use of L2 triggers in such situations enables efficient state management in the AR. The AR can clean up the associated state as soon as it detects that a host has been disconnected through the L2 Link Down trigger, for example. State clean up includes removal of ARP or Neighbor Cache entries, and can save bandwidth by inhibiting incoming data on the link where the host was once connected.

Additionally, faster and more efficient router discovery is possible if the AR receives a L2 Link Up trigger for a host. When the AR receives the trigger, it can send an unsolicited unicast router advertisement to the host. The host can begin the process of establishing IP connectivity more quickly.

7.0 Example: 802.11

In this section, we give an example of how a subset of the L2 triggers could be implemented for the 802.11 wireless LAN protocol [11] and used by Mobile IPv6 [2].

The 802.11 protocol supports a MAC layer management frame called Reassociation.request. The Reassociation.request frame is sent by a MN to a peer acting as an AP when the MN is in infrastructure mode and the MN wishes to change its association from its current AP to a new AP. The MN determines that a new AP is available because it detects a beacon from the new AP. The MN sends the Reassociation.request because the bit error rate on the link with the old AP has become too high (the standard does not specify exactly how high is too high, however).

The 802.11 Reassociation.request message contains the MAC address of the MN's current AP, and it is sent to the MAC address of the MN's desired new AP. The MAC layer frame in addition contains the MN's MAC address. Upon receipt of the Reassociation.request, the AP determines if the MN may reassociate and replies with a Reassociation.reply message either allowing or denying the request.

The Reassociation.request and Reassociation.reply contain the material for the following L2 triggers:

- When the MN's 802.11 driver receives a Reassociation.reply from the new AP confirming reassociation, it can deliver a Link Up trigger to the Mobile IP stack (or a daemon that communicates with the Mobile IP stack and is monitoring the driver) containing the MAC address of the new AP.

- When the AP determines that it can send a positive Reassociation.reply to the MN, it can generate a Link Up trigger with the MN's MAC address and the MAC address of the MN's old AP.

On the MN, the Link Up trigger delivered to the Mobile IP stack triggers the stack to send a Solicitation for Router Advertisement [2] to the new AR. This causes the AR to reply with a Router Advertisement. The MN compares the subnet prefix advertised by the router with its current subnet prefix, and if it determines that it has moved into a new subnet, it begins the process of establishing a new Care of Address. In this case, the MN is running standard Mobile IP without any fast handover enhancements; but because of the L2 trigger, it is able to eliminate the latency involved in waiting for a Router Advertisement beacon from the AR, thus increasing handover performance.

On the AP, the disposition of the Link Up trigger depends on the relationship between the AP and the AR. If the AP is acting as a transparent Layer 2 bridge, then some type of protocol is needed to transfer the trigger from the AP to the AR. This could be an addition to the 802 protocol, or it could be an enhancement to IAPP, the 802.11 InterAccess Point Protocol [6], or it could be an IPv6 protocol enhancement [10]. Upon receipt of the trigger protocol, the AR's driver or Mobile IP stack disposes of it exactly as in the case of a trigger on the MN. If the AP is integrated with the AR, then the trigger is delivered programmatically to the Mobile IP stack.

8.0 Security Considerations

The L2 triggers convey information about the link state of the MN and this information can trigger IP layer changes in routing reachability. As such, the information in an L2 trigger, if misused by an adversary or fraudulently propagated, could result in denial of IP service to the MN or hijacking of the MN's packets to a hostile third party.

If the L2 trigger is implemented as an API on an AR or AP, then the operating system and API implementation are required to assure that only qualified users can call into the API. Normally this involves denying access through the API unless the process running the API client has the proper security credentials on the host. If the L2 trigger is implemented as an L2 or L3 protocol, the protocol is required to protect the trigger messages with the proper authentication. In particular, if the protocol is an IP-based protocol, it must include authenticators so the parties that use the protocol can authenticate each other. If the protocol is intended to be used on public data networks, the option of encrypting the traffic must be available, to grant some privacy over the MN movement information propagated by the protocol messages.

9.0 References

1 Perkins, C., ed., "IP Mobility Support," RFC 2002, October, 1996.

2 Johnson, D., and Perkins, C., "Mobility Support in IPv6," draft-ietf-mobileip-ipv6-17.txt, a work in progress.

3 El-Malki, K., et. al., "Low Latency Handoff in Mobile IPv4," draft-ietf-mobileip-lowlatencyhandoffs-v4-02.txt, a work in progress.

4 Dommety, G., et. al. "Fast Handovers for Mobile IPv6," draft-ietf-mobileip-fast-mipv6-04.txt, a work in progress.

5 Levkowetz, O.H., et. al., "Problem Description: Reasons for Performing Context Transfers Between Nodes in an IP Access Network," draft-ietf-seamoby-context-transfer-problem-stat-04.txt, a work in progress.

6 "Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation," IEEE Std

802.11f/D1, DRAFT.

7 Sayed, H., et. al., "Context Transfer Framework," draft-hamid-seamoby-ct-fwk-01.txt, a work in progress

8 Sayed, H., et. al., "General Requirements for a Context Transfer Framework," draft-hamid-seamoby-ct-reqs-02.txt, a work in progress

9 Trossen, D., et. al., "Issues in Candidate Access Router Discovery for Seamless IP Handoffs," draft-ietf-seamoby-CARdiscovery-issues-03.txt, a work in progress.

10 Yegin, A., "Link-layer Triggers Protocol", draft-yegin-l2-triggers-00.txt, a work in progress.

11 IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE Std. 802.11, 1999.

12 Krishnamurthi, et. al., "Requirements for CAR Discovery Protocols," draft-ietf-seamoby-card-requirements-00.txt, a work in progress.

10.0 Author's Addresses

Karim El Malki
Ericsson Radio Systems AB
LM Ericssons Vag. 8 Phone: +46 8 7195803
126 25 Stockholm Fax: +46 8 7190170
SWEDEN Email: Karim.El-Malki@era.ericsson.se

Daichi Funato
DoCoMo USA Labs
181 Metro Drive, Suite 300 Phone: +1 408 451 4736
San Jose, CA 95110 Fax: +1 408 573 1090
USA Email: funato@docomolabs-usa.com

Youngjune Gwon
DoCoMo USA Labs
181 Metro Drive, Suite 300 Phone: +1 408 451 4734
San Jose, CA 95110 Fax: +1 408 573 1090
USA Email: gyj@docomolabs-usa.com

James Kempf
DoCoMo USA Labs
181 Metro Drive, Suite 300 Phone: +1 408 451 4711
San Jose, CA 95110 Fax: +1 408 573 1090
USA Email: james@docomolabs-usa.com

Mattias Pettersson
Ericsson Radio Systems AB   Phone: +46 8 585 32 562
Torshamnsgatan 23   Fax: +46 8 404 70 20
SE-164 80 Stockholm   E-mail: Mattias.Pettersson@era.ericsson.se
SWEDEN

Phil Roberts
Megisto Systems
20251 Century Blvd, Suite 120 Email: proberts@megisto.com
Germantown, MD 20874-1191
USA

Hesham Soliman
Ericsson Radio Systems
Torshamnsgatan 29, Kista   Phone: +46 8 7578162
Stockholm   Fax: +46 8 4043630
SWEDEN   Email: Hesham.Soliman@era.ericsson.se

Atsushi Takeshita
DoCoMo USA Labs
181 Metro Drive, Suite 300   Phone: +1 408 451 4705
San Jose, CA 95110   Fax: +1 408 573 1090
USA   Email: takeshita@docomolabs-usa.com

Alper E. Yegin, Editor
DoCoMo USA Labs
181 Metro Drive, Suite 300   Phone: +1 408 451 4743
San Jose, CA 95110   Fax: +1 408 573 1090
USA   Email: alper@docomolabs-usa.com