

Network Working Group
Request for Comments: 3132
Category: Informational

J. Kempf
Sun Microsystems
June 2001

Dormant Mode Host Alerting ("IP Paging") Problem Statement

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 (2001). All Rights Reserved.

Abstract

This memo describes paging, assesses the need for IP paging, and presents a list of recommendations for Seamoby charter items regarding work on paging. The results are specifically directed toward the task undertaken by the design team, and are not meant to be the definitive word on paging for all time, nor to be binding on Seamoby or other working groups, should the situation with regard to IP mobility protocols or radio link support undergo a major change.

1.0 Introduction

The IESG has requested that the Seamoby Working Group develop a problem statement about the need for additional protocol work to support alerting of dormant mode mobile hosts, commonly known as IP paging, for seamless IP mobility. The paging design team interpreted this as direction to examine whether location of a mobile node in power saving mode can be supported by the existing Mobile IPv4 and Mobile IPv6 protocols given existing radio link protocols.

Many existing radio link protocols and mobile systems support location of and radio link establishment with mobile nodes that are in power saving mode and hence are not actively listening for delivery of IP packets all the time or are not listening on the radio channels normally associated with delivering IP traffic to mobile nodes. This alerting functionality allows mobile nodes to reduce power consumption and decreases signaling load on the network for tracking mobiles that are not actively participating in IP packet generation or reception.

When a mobile is in low power consumption mode, special steps need to be taken to locate the mobile and alert it. These steps differ depending on the radio link, but the generic name for this process is paging, a term that is commonly used in cellular telephony.

In this document, after some initial definitions and material related to more clearly explaining what paging is, we assess the need for paging in existing IP mobility protocols (namely Mobile IP [1] [2]). We then develop a list of work items for the Seamoby working group related to this need. Note that the discussion in this document and the conclusions regarding work items are directed toward existing IP mobility protocols and existing radio link protocols. Should a major change occur in radio link support or the available IP mobility protocols, such as the introduction of a micromobility protocol for IP, the issues examined in this document may need to be revisited.

2.0 Definitions

The following definitions are relevant with respect to clarifying the paging functionality:

Dormant Mode - A state in which the mobile restricts its ability to receive normal IP traffic by reducing monitoring of radio channels. This allows the mobile to save power and reduces signaling load on the network.

Time-slotted Dormant Mode - A dormant mode implementation in which the mobile alternates between periods of not listening for any radio traffic and listening for traffic. Time-slotted dormant mode implementations are typically synchronized with the network so the network can deliver traffic to the mobile during listening periods. Additionally, the mobile may be restricted to listening on specific signaling channels that, according to current practice, are not typically used to carry IP traffic.

Paging - As a consequence of a mobile-bound packet destined for a mobile currently in dormant mode, signaling by the network through radio access points directed to locating the mobile and alerting it to establish a last hop connection. This messaging is in addition to simply delivering the packet to the mobile, i.e., last hop routing of packets is NOT considered to be paging.

Paging Area - Collection of radio access points that are signaled to locate a dormant mode mobile node. A paging area does not necessarily correspond to an IP subnet. A dormant mode mobile node may be required to signal to the network when it crosses a paging area boundary, in order that the network can maintain a rough idea of where the mobile is located.

Paging Channel - A radio channel dedicated to signaling dormant mode mobiles for paging purposes. By current practice, the protocol used on a paging channel is usually dictated by the radio link protocol, although some paging protocols have provision for carrying arbitrary traffic (and thus could potentially be used to carry IP).

Traffic Channel - The radio channel on which IP traffic to an active mobile is typically sent. This channel is used by a mobile that is actively sending and receiving IP traffic, and is not continuously active in a dormant mode mobile. For some radio link protocols, this may be the only channel available.

Paging Area Registrations - Signaling from a dormant mode mobile node to the network when the mobile node crosses a paging area boundary to establish the mobile node's presence in the new paging area.

3.0 Discussion of Paging

Dormant mode is advantageous to a mobile node and the network for the following reasons:

- Power savings. By reducing the amount of time the mobile is required to listen to the radio interface, the drain on the mobile node's battery is reduced.
- Reduced signaling for location tracking. By requiring the mobile to only signal when it crosses a paging area boundary rather than when it switches between radio access points, the amount of signaling for tracking the mobile is reduced because paging areas typically contain many radio access points.

In existing radio link protocols, there is a clear distinction between those protocols that support dormant mode only and those that support dormant mode with paging. Radio link protocols that do not support paging have no paging areas, no dedicated paging channel, and no radio link protocol specifically directed towards locating a dormant mode mobile, while radio link protocols that do support paging have these features. Although generalizations always run the risk of being contradicted by specific exceptions, the following comparison of existing radio link protocol support for these two cases may be instructive.

3.1 Dormant Mode Support Only

In radio link protocols that only support dormant mode, a dormant mode mobile node typically operates in time slotted mode and there is only one radio channel available, namely the traffic channel. The mobile node periodically wakes up, and, synchronously, the radio access point in the network with which the mobile node is associated delivers any IP packets that have arrived while the mobile node was asleep. Radio access points are required to buffer incoming packets for dormant mode mobiles; exactly how many packets and how long they are buffered are implementation dependent.

If the mobile node happens to move out of range of the access point with which it was associated, while it is in dormant mode, it discovers this when it awakens and reassociates with a new access point. The new access point then contacts the old access point over the wired backbone, the old access point sends any buffered packets, and the new access point delivers them to the mobile.

Radio link protocols with dormant mode support only are typically wireless LAN protocols in unlicensed spectrum in which the mobile node is not charged for using a traffic channel, and hence there is no need for conserving spectrum usage.

3.2 Dormant Mode with Paging Support

In radio link protocols with support for paging, the radio link typically supports more than one channel. A dormant mode mobile node may operate in time slotted mode, periodically waking up to listen to the paging channel, or it may simply listen to the paging channel continuously. The important point is that the mobile does not listen to nor transmit on a traffic channel while in dormant mode.

The radio access points are grouped into paging areas, and the radio link protocol supports periodic signaling between the mobile and the network only when the mobile crosses a paging area boundary, for the purpose of giving the network a rough idea of the mobile's location (paging area registrations). Some deployments of paging do not even use paging area registrations. They use heuristics to determine where the mobile is located when a packet arrives, in which case, no signaling is required while the mobile is in dormant mode.

An incoming packet is directed to the paging area where the mobile last reported, or the paging area is determined by heuristics. The network performs a radio link page by sending out a signal on the paging channel. The signal may be repeated until the mobile answers or a timeout occurs. In the former case, the packet is delivered, in the latter, the mobile is assumed to be unreachable.

Radio link protocols with paging support tend to be in licensed spectrum where the network operator has an interest in reducing the amount of signaling over traffic channels. Such reduction frees traffic channel spectrum for revenue-producing use, and avoids charging the customer for signaling overhead.

4.0 Is IP Paging Necessary?

In this section, we consider whether IP paging support is necessary. We first consider radio link protocols that have no support for paging. We then examine radio link protocols that have paging support. As discussed in the introduction, the focus is on whether the existing IETF mobility protocol, namely Mobile IP, requires enhancement. We also briefly discuss the relationship between paging and a potential future micromobility protocol.

4.1 IP Paging for Dormant Mode Only Radio Links

One possible justification for IP paging is for radio links that do not support paging. The reasoning is that an IP paging protocol could allow location of a dormant mode mobile in radio networks that do not support paging in the radio protocol.

An important point to keep in mind when considering this possibility is that, for radio links that do support paging, paging is typically used to locate mobiles for which the network has a rough idea of where the mobile is located. More specifically, in order to conserve signaling between the network and the mobile and to reduce power drain on the mobile, the mobile only updates the network about its location when it crosses a paging area boundary (if even then), which is far less frequent than when it crosses a radio access point boundary. If IP paging is to be of any use to radio link protocols that do not support paging, it must also be the case that it allows the network to maintain a rough idea of where the mobile is, otherwise, the amount of signaling involved in tracking the mobile and power drain on the mobile is not reduced.

However, as the description in the previous section indicates, for radio links without paging support, the network always has an **exact** idea of where the mobile is located. When the mobile moves into range of a new radio access point, it re-registers with the access point in that cell allowing the new access point to contact the old and deliver any buffered traffic. Additionally, the new access point at that time may choose to deliver a foreign agent advertisement (for Mobile IPv4) or router advertisement (for Mobile IPv6) to the mobile if the mobile node has changed subnets, so that the mobile can perform Mobile IP re-registration in order to make sure its IP routing is current. There is absolutely no ambiguity in the mobile's

location as far as the network is concerned, and so the network can continue to route packets to the mobile node while the mobile is in dormant mode with assurance (modulo buffer overflows and timeouts at the radio access point) that the packets will be delivered to the mobile the next time it wakes up from dormant mode.

As a consequence, IP paging provides no advantages for radio link protocols in which the radio link does not have support for paging.

4.2 IP Paging for Radio Links with Paging Support

In radio links that do support paging, there are two cases to consider: networks of radio links having a homogeneous radio technology and networks of radio links having heterogeneous radio technologies. We examine whether Mobile IP can support dormant mode location for both these cases.

4.2.1 Homogeneous Technology Networks

For homogeneous technology networks, the primary issue is whether signaling involved in Mobile IP is enough to provide support for locating dormant mode mobile nodes. Subnets constitute the unit of signaling for presence in IP. When a mobile node moves from one subnet to another, Mobile IP signaling is required to change the mobile's care-of address. This signaling establishes the mobile's presence in the new subnet. Paging areas constitute the unit of signaling for dormant mode mobile presence at the radio level. Paging area registrations or heuristics are used to establish a dormant mode mobile's presence in a particular paging area.

If paging area registrations can always serve to trigger Mobile IP registrations, there is no need for an IP paging protocol because the network (specifically the home or hierarchical agent) will always have an up-to-date picture of where the mobile is and can always route packets to the mobile. The key determining factor with regard to whether paging area registrations can be used in this fashion is how subnets are mapped into paging areas. If it is always possible to map the two such that a paging area registration can serve as a transport for a Mobile IP registration, or some other technique (such as network assisted handoff [3] [4]) can be used to transfer the Mobile IP registration, then no IP paging protocol is needed.

In general, the mapping between paging areas and subnets can be arbitrary, but we consider initially a smooth subset relationship, in which paging areas are subsets of subnets or vice versa. Network topologies in which one subnet is split between two or more paging areas are therefore eliminated. The restriction is arbitrary, but by

starting here, we can discover whether additional work is needed. We also consider a case where paging area registrations in the radio layer protocol are always done. This is also optimistic.

There are three cases:

- 1) The topological boundaries of the paging area and subnet are identical.
- 2) Multiple paging areas are part of the same subnet.
- 3) Multiple subnets are part of the same paging area.

Each case is considered in the following subsections.

4.2.1.1 Subnet and Paging Area Boundaries Identical

In the case where radio paging areas map one to one onto IP subnets (and hence Mobile IPv4 foreign agents or IPv6 access routers), it is possible to use radio link paging together with Mobile IP handoff techniques for the network to track the mobile's location. If the paging area update protocol supports sending arbitrary packet data over the paging channel, the access router or foreign agent can send a router advertisement or foreign agent advertisement to the mobile as part of the signal that the mobile has entered the new paging area, and the mobile can send a Mobile IP registration as part of the paging area update. For other cases, enhancements to Mobile IP network-assisted handoff techniques can allow the network to track the mobile as it moves from paging area (== subnet) to paging area. Other uses of the Mobile IP registration protocol are also possible depending on the level of paging support for packet data. As a consequence, the home or hierarchical agent has complete knowledge of routes to the mobile and can route packets to the foreign agent or access router. Radio layer paging may be needed at the foreign agent or access router in order to re-establish a traffic channel with the mobile, but no IP paging is required.

4.2.1.2 Multiple Paging Areas Map into One Subnet

The case where multiple radio paging areas map to a single IP subnet is the same as above, with the exception that the last hop Mobile IPv4 foreign agent or IPv6 access router for the subnet performs paging in multiple paging areas to locate the mobile.

4.2.1.3 Multiple Subnets Map into One Paging Area

In the case where a single radio paging area maps onto multiple IP subnets, it is not possible to directly use Mobile IP handoff between last hop access routers or foreign agents to track the mobile's location as it moves, because the mobile does not signal its location when it changes subnets. Within the set of subnets that span the paging area, the mobile's movement is invisible to the L2 paging system, so a packet delivered to the mobile's last known location may result in a page that is answered in a different subnet.

Consider the following example. Suppose we have a network in which there are two paging areas, PA(1) and PA(2). Within each, there are many subnets. Consider a mobile that moves from PA(1) to PA(2), and enters PA(2) at subnet X. Using the paging area registration, it signals the network that it has moved, and suppose that the paging area registration contains a Mobile IP registration. The agent handling the L2 paging protocol sends the registration to the home/hierarchical agent (or perhaps it simply gets routed). The home/hierarchical agent now knows that the mobile has a CoA in subnet X, as does the mobile. After the mobile has completed the paging area registration/Mobile IP registration, it goes back to sleep.

But the mobile does not stop in subnet X, it keeps moving while in dormant mode, when it is doing no signaling (L2, mobile IP or other) to the network. It moves from subnet X where it originally entered the paging area clear to the other side of the paging area, in a completely different subnet, subnet Y.

Suppose a packet comes into the home/hierarchical agent for this mobile. Because the home/hierarchical agent believes the mobile is in subnet X, it sends the packet to the access router or foreign agent for subnet X. The packet gets to the access router or foreign agent, and the access router or foreign agent performs a radio page for the mobile in subnet X. Since the mobile isn't in subnet X, it wakes up in subnet Y because the radio page propagates throughout the paging area. It does a mobile IP re-registration because it sees that it is in a new subnet, but the packet at the access router or foreign agent in subnet X can't get to the mobile.

Without any further support, the access router or foreign agent in subnet X drops the packet. The only way to get the packet to the mobile node from the access router or foreign agent is for the mobile node to send a binding update to the access router or foreign agent when it wakes up in the new subnet. Once the access router or foreign agent has the new binding, it can forward the packet. Some smooth handoff techniques depend on sending binding updates to foreign agents [5], so arranging for the mobile node to send a

binding update would be possible. In IPv6, it becomes less attractive because of the need for security on the binding update. In either case, the result would be yet more Mobile IP signaling before the packet could be delivered, increasing the amount of latency experienced by the mobile.

While it may be possible with enhancements to Mobile IP to handle the case, the enhancements would probably introduce more latency and signaling into the initial connection between the mobile and the network when the mobile awakes from dormant mode. An IP paging protocol between the home or hierarchical agent and a paging agent in the paging area would serve to reduce the amount of latency involved in delivering the initial packet. With IP paging, the arrival of the packet at the home/hierarchical agent results in an IP page to a paging agent in the last reported paging area. The paging agent performs an L2 page to the mobile. The mobile answers the page with a mobile IP registration to the home/hierarchical agent and the home/hierarchical agent sends the packet. The home/hierarchical agent and the mobile already have a security association, so there is no need to negotiate one, and buffering of the first packet and any further incoming packets prior to the mobile IP registration is handled by the home/hierarchical agent rather than a router at the edge, so the edge routers can be simpler. Finally, the home/hierarchical agent can start routing to the mobile as soon as the registration comes in.

4.1.2.4 More Complex Homogeneous Network Cases

Up until now, the discussion has not identified any case where the problem of locating and delivering the first packet to a dormant mode mobile could not be handled by Mobile IP with enhancements. IP paging serves as a promising optimization in the multiple subnets to single paging area case, but in principle additional Mobile IP signaling (potentially lots in the case of IPv6 if a security association is needed) could handle the problem. However, the examples examined in the above sections are really best-case. In practice, the mapping of subnets to paging areas is likely to be far less clear cut, and the use of paging area registrations far less common than has been assumed in these cases.

Requiring network operators to make paging areas and subnets conform to a subset relationship that would allow mobile IP signaling to do double duty as paging area updates is unrealistic. In practice, paging areas often overlap and there is often not even a clear subset relationship between paging areas themselves. Some radio protocols, such as wCDMA [6], allow different mobile terminals in the same geographical area to have different paging area identifiers. Working through each case and trying to identify whether Mobile IP needs

enhancement would probably result in a much more complex result than having a simple IP paging protocol that allows a home/hierarchical agent to notify an L2 agent in the paging area when a new packet comes in.

Finally, requiring operators to always turn on paging area registrations is unacceptable, and using Mobile IP registrations won't work if paging area registrations are not done. The above description is ideal with regard to signaling between the mobile node in dormant mode and the network. Anecdotal evidence indicates that most operators do not turn on paging area registrations, they use heuristics to determine where to page for the mobile. If the operator does not turn on paging area registrations, there is no way for the mobile to report its position when it changes paging area, hence no L2 vehicle for potential dormant mode use of Mobile IP.

4.2.2 Heterogeneous Technology Networks

In a network composed of links with multiple technologies, the problems identified above become multiplied. Using Mobile IP becomes even more cumbersome, because the subnet to which the initial packet is delivered, besides not being in the same subnet on which the dormant mode mobile is located, may be on a radio network which the user would actually not prefer to use in their current location. This could happen, for example, if the mobile moved inside a building and radio coverage on one interface became weak or nonexistent, or if the user had a choice of a cheaper or higher bandwidth connection. The mobile may actually no longer be listening or reachable on the paging channel of the old network, so when the old access router or foreign agent pages on the old radio network, the mobile, which is now listening only for pages on the new network, may not answer, even though it is reachable on the new network. Arranging for pages in multiple radio networks is a possibility, but without an L3 paging protocol to abstract away from the L2 details, the details of each L2 protocol must be handled separately.

A paging protocol that unifies paging across multiple radio technologies therefore looks attractive. There may be commonalities in the corresponding radio paging protocols that allow a mapping to be established between the radio protocols and an abstract IP paging protocol. For example, assume we have a common paging area identifier defined at the IP layer that is mapped to each radio paging protocol by the access points. An IP paging message containing the identifier is sent to multiple access points, where the appropriate radio paging message is sent based on the particular technology implemented by the access points. The results are then returned by the radio paging responses, mapped back into IP by the access points, and delivered back to the origin of the page.

An additional case to consider is when a single subnet consists of multiple radio access technologies. A wireless access point usually provides L2 bridge behavior to the wired link with which it is connected. If two access points with incompatible technologies and non-overlapping cells are connected to the same subnet, a mobile node with interfaces to both technologies would need paging from both technologies. If reachability can be established simply by ARP or neighbor discovery, no IP paging is needed. However, note that ARP or neighbor discovery requires that a functional traffic channel be available to the mobile, since these protocols are typically implemented for wired networks in which a single channel exists on which all IP traffic is delivered. If the mobile is currently in the sleep phase of a time-slotted dormant mode, or if it is listening to a paging channel it will fail to respond to these requests. In this case, some means of triggering a radio page from IP is necessary to find the mobile. Modifying ARP or neighbor discovery to utilize a paging channel if available is a possible, if somewhat messy, alternative, but a dedicated location protocol may be somewhat cleaner.

4.3 Paging and Micromobility

If the Seamoby Working Group decides that an IP micromobility protocol is necessary, then the above analysis is no longer complete. A micromobility protocol may require some type of paging support. The design team does not want to include any further discussion of paging and micromobility at this point, because it is not clear whether micromobility will be pursued by Seamoby and hence such discussion would be premature.

5.0 What Exactly is the Problem?

While the above analysis has identified situations in which location of a mobile in dormant mode may require some action at the IP layer, it is important keep in mind what the problem is. The problem to be solved is the location of a mobile node because it has moved while in dormant mode. IP paging is one solution to the problem, there may be others.

6.0 Recommendations

The design group recommends the following charter items for Seamoby:

- 1) Since the design group has identified several network deployment scenarios where existing Mobile IP technology cannot find a mobile in dormant mode, protocol work is necessary to define a way for the network to find a mobile that is currently in dormant mode.

- 2) The work defined above should be pursued in a way that is maximally consistent with Mobile IP and other existing IETF protocols. The work should also generate recommendations about how to achieve the best match between existing radio paging protocols and IP.
- 3) If the Seamoby working group decides to pursue a micromobility protocol that requires paging, the Seamoby group should undertake the design of a new paging protocol within the context of that work.
- 4) There is some evidence that cellular operators' deployments of paging are highly variable, and may, in fact, be suboptimal in many cases with respect to supporting IP. The Seamoby working group should write a BCP which explains how to perform IP subnet to paging area mapping and which techniques to use when, so network designers in wireless networks have a guide when they are setting up their networks.

7.0 Acknowledgements

The editor would like to thank the Seamoby paging design team for helping formulate the first draft of the document. Jari Malinen contributed text to Section 4.2. Hesham Soliman, Karim El-Malki, and Behcet Sarikaya contributed critical commentary on the first draft, which was important in sharpening the reasoning about what can and can't be expected in the absence of radio layer paging support and how Mobile IP might be used to support dormant mode location.

8.0 References

- [1] Perkins, C., Editor, "IP Mobility Support", RFC 2002, October 1996.
- [2] Johnson, D., and C. Perkins, "Mobility Support in IPv6", Work in Progress.
- [3] El Malki, K. et. al., "Low Latency Handoff in Mobile IPv4", Work in Progress.
- [4] Tsirtsis, G., Editor, "Fast Handovers for Mobile IPv6", Work in Progress.
- [5] Perkins, C. and D. Johnson, "Route Optimization in Mobile IP", Work in Progress.
- [6] Holma, H. and A. Toskala, "WCDMA for UMTS: Radio Access for Third Generation Mobile Communication", John Wiley and Sons, New York, 2000.

9.0 Editor's Address

James Kempf
Sun Labs California
Sun Microsystems, Inc.
901 San Antonio Rd., UMPK15-214
Palo Alto, CA, 94303
USA

Phone: +1 650 786 5890
Fax: +1 650 786 6445
EMail: james.kempf@sun.com

10.0 Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

