

Mobile IP and the IETF

Charles E. Perkins

charles.perkins@eng.sun.com

SUN Microsystems, Mountain View, CA, USA

Introduction

IETF 41 was a time for consolidation of several existing efforts, and one new thread that shows good promise. Mobile IP has several new documents that have been accepted as Proposed Standards, and the IPv6 mobility support document has reached consensus on all major points within both the IPng working group and the Mobile IP working group. Within the Service Location Protocol areas of interest, there has been a new and hopefully final release of SLP version 2, and a Birds of Feather (BOF) about “wide-area” service location. Lastly, consideration of Mobile IP is spreading to other appropriate groups such as PPP-extensions, roamops, and ion.

Mobile IPv6

IPv6 mobility took big steps towards Proposed Standard in discussions within the Mobile IP working group and within the IPng working group. Dave Johnson ran through a lengthy list of design issues with Mobile IPv6, proposed solutions for them, and was rewarded with working group consensus on essentially every point. Thus, we expect that the document will be revised and sent along the way to working group last call before the end of May, and thus probably to IETF last call before the end of June. This could lead to Proposed Standard status *before* the next IETF meeting in August. The list of issues discussed at the working group meetings includes:

- Assignment of an anycast address for all home agents on the home network
- Creation of a new extension for Router Advertisement, enabling routers to specify how often advertisements are sent
- How home agents can send a list of possible home agents to a mobile node during Home Agent Discovery
- ‘H’ bit in Router Advertisement to indicate that the router is a home agent
- Discovery of routers’ global addresses
- Allow advertisements to come out more often than once every 3 seconds (actually, we raised this issue years ago)
- IANA considerations (getting the needed new option numbers)
- Dissemination of new requirements to IPv6 implementors

It appears that the Home Agent Anycast address will be ‘3’ in the low order bits of the address range defined by the home agent routing prefix. Moreover, the suggestion has been made that the first 255 such addresses be reserved for future anycast definition. Currently, ‘0’ is the anycast router address for a network, and ‘1’ and ‘2’ are reserved for point-to-point links. The new Router Advertisement extension enables a (mobile) node to know how long to wait before deciding that the router is no longer providing service. Specifying the ‘H’ bit is related to enabling Home Agents to build up a list of all home

agents on a network. When the list is available, then any home agent can send the list to a mobile node in response to a Home Agent discovery process. That way, the mobile node is likely to find a responsive home agent, even if the home agent that receives the mobile node’s request cannot itself perform the service. This might otherwise present problems because the same home agent might answer the discovery request every time, and if that home agent were “busy” then the mobile node would enter a very frustrating mode of operation. In point of fact, home agents are never supposed to be so “busy” to preclude offering their services, but hopefully this step will improve the robustness of the protocol.

In the Mobile IPv6 specification, mobile nodes are allowed to effect smooth handoffs from one network to another by sending a Binding Update to their previous router, indicating the new care-of address for address used at the mobile node’s previous point of attachment (which latter was itself used as a care-of address). This is only possible if the previous router is globally addressable, and (unfortunately – in fact I think mistakenly) the previous router is not required by protocol to have informed the mobile node of its global address. The problem is that Router Advertisements only contain the link-local address of the router, and the mobile has no guaranteed way to derive the global address from the link-local address. There is a trivial way that is *almost* guaranteed, but not an *absolutely guaranteed* way. The mistake, in the author’s opinion, is that the trivial way should be made guaranteed to work, and it is not being made so.

The last point, about disseminating requirements information to IPng implementors, arises because of the dual working-group nature of the Mobile IPv6 protocol specification. The work was largely done within the Mobile IP working group, and not so many IPng working group members attended the Mobile IP working group meetings. One notable exception has been Erik Nordmark, who is co-chair of the Mobile IP working group, and Erik has provided a great deal of guidance to the working group on IPv6 issues. Since so few IPng members were aware of the protocol requirements for mobility, the concern was raised that implementors may unwittingly neglect the needed logic for mobility support, and the discussion centered around how to provide the notification. It was agreed that announcements would be made to the mailing lists, and that the requirements would be posted to the appropriate Web pages. There is not time for a full-blown IPv6 “Host Requirements” document at this time.

Mobile IPv4

Two documents have been accepted as Proposed Standards, and one published as an Informational RFC after many months of review and improvement. The first two are “Mobile-IPv4 Configuration Option for PPP IPCP” and “Reverse Tunneling for Mobile IP”. The Informational document is “Firewall Support for Mobile IP”. The new PPP option enables a mobile node to use a NAS as a foreign agent, whenever the NAS sup-

ports Mobile IP and the new IPCP option. Reverse Tunneling is important in situations where the mobile node cannot transmit packets using its home address as the source IP address. The typical example in today's Internet occurs when border routers for a particular administrative domain perform what is known as "ingress filtering". That means that the border routers only allow packets to flow into the Internet if the packets have an allowable source IP address, typically an address known to reside on a network within the administrative domain. Since mobile nodes using Mobile IP have home addresses that would seem to lie outside the visited administrative domain, ingress filtering places severe constraints on the ability of mobile nodes to use their home address.

To get around this problem, a mobile node has to send its packets to correspondent nodes without using its home address as the source IP address, and one of the few methods known to accomplish this is by encapsulating the mobile node's outgoing packets within a tunnel. To satisfy the demands of ingress filtering, a local address is used as the source IP address of the outer IP header. The home agent is currently the only assured tunnel endpoint that can be chosen to receive the tunneled packets. This is a brief outline of the method specified in the Proposed Standard for Reverse Tunneling. This mode of operation has also been investigated as part of the MosquitoNet project at Stanford University.

Lastly, the Firewall Traversal specification gives some ideas about how to use a co-located care-of address and SKIP (a key-management protocol from Sun Microsystems) to enable a mobile node to tunnel past a firewall protecting its home administrative domain, when the mobile node is no longer attached to its home network (i.e., is traveling).

In related news, a new Internet Draft has been issued to deal with mobile nodes in a Non-Broadcast Multiple Access (NBMA) environment like ATM; look for draft-ietf-ion-nhrp-mobile-nhc-00.txt, in the "ion" (for IP over NBMA) working group. This specification deals with using mobile nodes with NHRP. For nodes that are dialing up into ISPs using RADIUS, a new draft has been prepared to specify how RADIUS can enable the use of Mobile IP. The ISP can use the information provided to determine whether or not the mobile node can use a co-located care-of address. For complete details, consult draft-ietf-roamops-mobileip-01.txt.

Service Location Protocol

Since the Service Location Protocol (SLP) was promulgated as a Proposed Standard (RFC 2165) last summer, there have been many developments. One purpose of the IETF process is to promote standards that are interoperable by using the combined experience of multiple implementation teams, typically from different vendors. The interoperability of SLP has been tested over the recent months in multiple venues, most recently at Connectathon'98. These experiences and the needs of product groups from the various vendors have indicated the need for refining the original Proposed Standard in several ways, most notably by reducing the number of protocol features and messages that are mandatory to implement. The reduced feature set is useful for devices (such as printers and fax machines) that wish to make their presence known on the network, but do not need to be burdened with unwanted features

that will never be used. Some of the other differences between SLP and SLPv2 were detailed in my previous column. Other differences approved since that time include revising the query language to match LDAP semantics, and to allow "service:" URLs for address families other than IP (notably, AppleTalk and IPX).

Discussing the final details of SLPv2 consumed most of the time in the working group, and by the middle of May the results are expected to be in Working Group Last Call. SLPv2 will need to go again to Proposed Standard instead of going to Draft Standard, because of the extent of the changes made to RFC 2165. The header formats are different, the scope model is different, the language and internationalization support is different, and the mandatory feature set is reduced. However, everyone agrees that the protocol has been greatly improved. The working group is set to go dormant while vendors release the new products and further experience is gained. Scheme documents for new service types, as well as the API and other SLP supporting drafts, are scheduled to go to Last Call and to Proposed Standard as soon as the revised protocol (SLPv2) once again is approved by the IESG.

There was a Wide Area Service Location (wasrv) BOF held, with the goal to assess interest and determine what the charter might be for a new working group. This might turn out to be a fantastic area for mobile networking, if a mobile computer could count on finding network services in the wide area. One service frequently mentioned was associated with Internet telephony; presumably a mobile user would like to find the closest "gateway" to support voice-over-IP service wherever the user happens to be attached to the Internet. Other working groups are defining the basic service types and attributes for such services now. Two proposals were described during the BOF. The first proposal outlined the use of multicast for advertising and discovering wide-area services. Another proposal (by yours truly) described a possible way to use existing Web search engines to collect services of a particular type, and then in conjunction with that a way for "Wide Area Directory Agents" to extract the lists of services from the search engines. To make this work, the wide-area services would make themselves available to the search engines by hyperlinks from certain well-known locations that would certainly be combed by the search engines on a regular basis (e.g., www.big-company.com).

DHCP options for Service Location Protocol

Lastly, the specification document for DHCP options 78 and 79 is finally heading for Last Call. These options allow a SLP user agent (or service agent) to acquire the address of a directory agent, and/or the name of a scope from which services may be requested. With the addition of the DHCP options, a complete strategy for scalable and low-administrative operation for user agents is at last available. The sequence of events would be as follows:

1. A user agent requests an IP address and the address of a directory agent from DHCP.
2. The user agent then can resolve service needs dynamically as needed.

By this procedure, a newly attached network node would be able to access network services without any preconfiguration, which is a far cry from the situation in today's typical networks. Of course, the realization of this goal relies heavily on local administrative policies. Note also that the user agent on a mobile node would typically request a new IP address and directory agent at every new point of attachment. The availability to mobile nodes of such local configuration will depend also upon local administrative policies, and network administrators may not be willing to make all such services available to all mobile nodes without a great deal of deployment and experience. Security, as so often happens, is an interesting design problem in such scenarios.



Call for Participation

MMT'98



IEEE
Communications
Society

Multiaccess, Mobility and Teletraffic for Wireless Communications

October 21-23, 1998, George Washington Univ., Washington, DC, USA

Organized by IEEE Communications Society in Co-operation with ACM SIGMOBILE (pending)

Advisory Board:

Norman Abramson, ALOHA Networks, US
Hamid Aghvami, King's College, UK
Donald Cox, Stanford Univ., US
Anthony Ephremides, UMD, US
David Everitt, Univ. of Melbourne, Australia
Robert Gallager, MIT, US
Philippe Godlewski, ENST, France
Bijan Jabbari, GMU, US
Jim Massey, ETH, Switzerland
Raj Pandya, BNR, Canada
Raymond Pickholtz, GWU, US
Stephen Rappaport, SUNY, US
Raymond Steele, MAC Ltd., UK
Andrew Viterbi, Qualcomm, US

General Workshop Co-chairs:

Raymond Pickholtz, George Washington University,
Washington, DC, US
pickholt@seas.gwu.edu

Bijan Jabbari, George Mason University,
Fairfax, VA, US
bjabbari@gmu.edu

Technical Program Co-chairs:

Kin K. Leung, AT&T Labs, New Jersey, US
kkleung@research.att.com

Branimir Vojcic, George Washington University,
Washington D.C., US
vojcic@seas.gwu.edu

The focus of this workshop is to identify, present and discuss the theoretical and implementation issues critical to the design of land and satellite-based mobile cellular and microcellular, wireless personal communications as well as wireless local area networks. Topics of interest but not limited to:

Multi-user channels, receiver design and channel modeling issues
Antenna techniques
Channel equalization
Spread spectrum techniques
Multiaccess methods and their performance analysis
Analytical techniques and capacity evaluation of wireless networks
Teletraffic issues
Mobility management and call control
Handoff strategies and implementation
Power control algorithms and implementations
Dynamic channel assignment
Network control and signaling
Routing schemes and optimization
Universal personal telecommunications issues
Mobile computing
Internet access in wireless networks

The workshop will feature distinguished speakers and a number of high quality technical presentations. You are cordially invited to participate at the workshop. For registration and other information pertinent to the workshop, see <http://mozart.gmu.edu/MMT98>, or contact:

MMT'98
Attention: Kin K. Leung
Tel: 732-345-3153, Fax: 732-345-3038
Email: MMT98@research.att.com (preferred) or kkleung@research.att.com.