Mobility Support in IPv6

Charles E. Perkins
T. J. Watson Research Center
IBM Corporation
Hawthorne, NY 10532
perk@watson.ibm.com

David B. Johnson
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
dbj@cs.cmu.edu

Abstract

IP version 6 (IPv6) is being designed within the IETF as a replacement for the current version of the IP protocol used in the Internet (IPv4). We have designed protocol enhancements for IPv6, known as Mobile IPv6, that allow transparent routing of IPv6 packets to mobile nodes, taking advantage of the opportunities made possible by the design of a new version of IP. In Mobile IPv6, each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. While away from its home IP subnet, a mobile node is also associated with a care-of address, which indicates the mobile node's current location. Mobile IPv6 enables any IPv6 node to learn and cache the care-of address associated with a mobile node's home address, and then to send packets destined for the mobile node directly to it at this care-of address using an IPv6 Routing header.

David Johnson was supported in part by the National Science Foundation under CAREER Award NCR-9502725, by the Air Force Materiel Command (AFMC) and ARPA under contract number F196828-93-C-0193, and by the AT&T Foundation under a Special Purpose Grant in Science and Engineering. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of NSF, AFMC, ARPA, the AT&T Foundation, IBM, CMU, or the U.S. Government.

To appear in *Proceedings of the Second Annual International Conference on Mobile Computing and Networking (MobiCom'96)*, November 10–12, 1996, Rye, New York, USA. © 1996 ACM.

1 Introduction

In this paper, we describe the design of a new protocol for transparent routing of IPv6 packets to mobile IPv6 nodes operating in the Internet [16]. IP is the protocol which provides packet routing and delivery services for the Internet, and IP version 6 (IPv6) [9] is a new version of IP intended to replace the current version of IP (IPv4) [22]. IPv6 is in the final stages of design within several working groups of the Internet Engineering Task Force (IETF) [24, 5], the principle standards development body for the Internet.

Without specific support for mobility in IPv6, packets destined to a mobile node would not be able to reach it while the mobile node is away from its home IP subnet, since as in IPv4, routing is based on the network prefix in a packet's destination IP address. In order to continue communication in spite of its movement, a mobile node could change its IP address each time it moves to a new IP subnet, but the mobile node would then not be able to maintain transport and higher-layer connections when it changes location. Mobility support in IPv6 is particularly important, since mobile computers are likely to account for a majority or at least a substantial fraction of the population of the Internet during the lifetime of IPv6.

IPv6 is derived from IPv4 and is in many ways similar to it. As such, the IETF Mobile IP Working Group's current protocol design [17] for mobility of IPv4 nodes could be adapted for use in IPv6, with only the straightforward changes needed to accommodate differences between IPv4 and IPv6 such as the size of addresses. However, the development of IPv6 presents a rare opportunity, in that there is no existing installed base of IPv6 hosts or routers with which we must be compatible, and in that the design of IPv6 may still be adjusted to account for the few special needs of mobile nodes. This paper, therefore, considers how IPv6 can most naturally fulfill the support requirements for mobile nodes.

Each mobile node is assigned a (permanent) IP address in the same way as any other node, and this IP address is known as the mobile node's *home address*. A mobile node's home address remains unchanged regardless of where the node is attached to the Internet. The IP subnet indicated by this home address is the mobile node's *home subnet*, and standard IP routing mechanisms will deliver packets destined to a mobile node's home address only to the mobile node's home subnet. A mobile node is simply any node that may change its point of attachment from one IP subnet to another, while continuing to be addressed by its home address. Any node with which a mobile node is communicating we refer to here as a *correspondent node*, which itself may be either mobile or stationary.

A mobile node's current location while away from home is known as its *care-of address*, which is a globally-routable address acquired by the mobile node through IPv6 address autoconfiguration in the *foreign subnet* being visited by it. The association of a mobile node's home address with a care-of address, along with the remaining lifetime of that association, is known as a *binding*.

While away from its home subnet, a router on the mobile node's home subnet known as its *home agent* maintains a record of the current binding of the mobile node. The home agent then intercepts any packets on the home subnet addressed to the mobile node's home address and *tunnels* them to the mobile node at its current care-of address. This tunneling uses IPv6 encapsulation [8], and the path followed by a packet while it is encapsulated is known as a *tunnel*. Once a correspondent node has learned the mobile node's care-of address, it may cache it and route its own packets for the mobile node directly there using an IPv6 Routing header [9], bypassing the home agent completely.

The most important function needed to support mobility is the reliable and timely notification of a mobile node's current care-of address to those other nodes that need it, in order to avoid the routing anomaly known as *triangle routing*, as illustrated in Figure 1. In triangle routing, *all* packets sent *to* a mobile node must be routed first to the mobile node's home subnet and then forwarded to the mobile node at its current location by its home agent; packets sent *from* a mobile node are not forwarded in this way (unless they are destined to another mobile node), leading to this "triangular" combination of the two routes used for all communication between these two nodes.

Triangle routing, because of its poor route selection, has many attendant problems, including

- increased impact of possible network partitions,
- increased load on the network, and

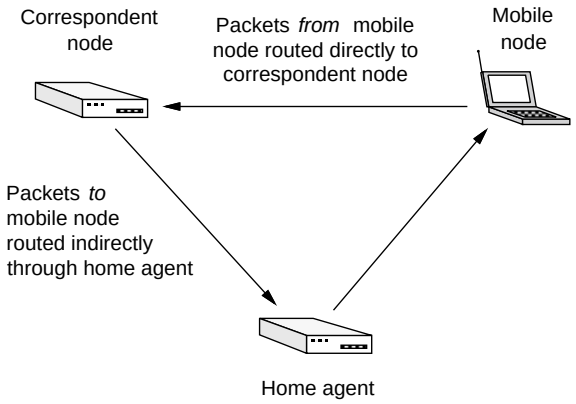

Figure 1 Triangle routing

- increased delay in delivering packets.

We believe it is reasonable to expect every IPv6 node to perform the extra steps of caching and using the care-of address for each mobile node with which it is communicating, since the additional overhead of doing so is quite small; the problems of triangle routing are then avoided by having each correspondent node route its own packets for a mobile node *directly* to its care-of address. Mobile IPv6 introduces a set of new *IPv6 Destination options*, called *Binding Update* and *Binding Acknowledgement*, to manage these cache entries as needed. The Binding Update and Binding Acknowledgement options conform to an option mechanism already present in IPv6 [9].

Section 2 of this paper presents an overview of the important aspects of the operation of IPv6, and Section 3 presents an overview of the extensions for Mobile IPv6. In Sections 4 and 5, respectively, we then describe the new Binding Update and Binding Acknowledgement IPv6 Destination options used by Mobile IPv6. The operation of a mobile node is described in Section 6, Section 7 describes the operation of a correspondent node communicating with a mobile node, and Section 8 describes the operation of a home agent. Security issues are discussed in Section 9, and in Section 10, we present conclusions.

2 Overview of IPv6

In this section, we outline some of the basic characteristics of IP version 6 (IPv6) that are particularly relevant to our mobility protocol. The most visible difference is that IPv6 addresses are all 128 bits long, instead of 32 bits long as in IPv4. Within this huge address space, a tiny part is reserved for all current IPv4 addresses, and another tiny part is reserved for the *Link-Local* addresses, which are not routable but which are guaran-

teed to be unique on a link (i.e., on a local network). Nodes on the same link can communicate with each other even without any routers, by using their Link-Local addresses.

Nodes discover each other's presence, as well as each other's link-layer (i.e., MAC) addresses, by participating in the Neighbor Discovery protocol [?]; IPv6 nodes also discover local routers and network prefixes by means of Neighbor Discovery. The IPv6 Neighbor Discovery protocol can be characterized as a much improved version of two IPv4 protocols, the Address Resolution Protocol (ARP) [18] and the ICMP Router Discovery Protocol [10].

IPv6 defines several kinds of *extension headers*, which may be used to include additional information in the headers of an IPv6 packet. The defined IPv6 extension headers include:

- Destination Options header,
- Hop-by-Hop Options header,
- a Routing header, and
- an Authentication header.

The Destination Options header may be included in a packet to carry a sequence of one or more options to be processed only when the packet arrives at the final destination node. Similarly, the Hop-by-Hop Options header may be included to carry a sequence of one or more options, but these options are processed by every intermediate router which receives and forwards the packet as well as by the final destination node. In IPv4, every IP option [22] is treated essentially as a Hop-by-Hop option and thus causes performance degradation because of processing needed at every intermediate router, whether it pertains to that router or only to the final destination node.

The Routing header is particularly useful for our mobility protocol, and is similar to the Source Route options defined for IPv4. The IPv6 Routing header can serve both as a strict source route and a loose source route, although Mobile IPv6 uses it only as a loose source route. Unlike the IPv4 Source Route options, however, in IPv6, the Routing header is not examined or processed until it reaches the next node identified in the route. In addition, the destination node receiving a packet with a Routing header is under no obligation to reverse the route along which the packet was received, for routing packets back to the sender.

The Authentication header provides a means by which a packet can include optional authentication data, for example based on a one-way cryptographic hash (e.g., MD5 [14, ?, 23]) of the packet's contents. The inclusion of this authentication data allows the receiver to verify the authenticity of the packet sender, and also protects against modification of the packet while in transit, since a modified packet will be viewed by the receiver the same as a forged packet. The Authentication header may also be used to provide replay protection of packets, such that saved copies of an authenticated packet cannot later be be resent by an attacker [?]. The computation of the authentication data and use of replay protection are controlled by a "security association" that the sender of the packet must have established with the receiver [3]. Security associations may be manually configured or automatically established.

3 Overview of Mobile IPv6

3.1 Requirements, Goals, and Applicability

Mobile IPv6 is intended to enable IPv6 nodes to move from one IP subnet to another. It is just as suitable for mobility between subnets across homogeneous media as it is across heterogeneous media, although in the homogeneous case other solutions may also exist [1]. That is, Mobile IPv6 facilitates node movement from one Ethernet segment to another as well as it accommodates node movement from an Ethernet segment to a wireless LAN cell.

The protocol allows a mobile node to communicate with other nodes (stationary or mobile) after changing its link-layer point of attachment from one IP subnet to another, yet without changing the mobile node's IPv6 address. A mobile node is always addressable by its home address, and packets may be routed to it using this address regardless of the mobile node's current point of attachment to the Internet. The movement of a mobile node away from its home subnet is thus transparent to transport and higher-layer protocols and applications.

All packets used to inform another node about the location of a mobile node must be authenticated. Otherwise, a malicious host would be able to hijack traffic intended for a mobile node by the simple matter of causing the mobile node to seem to be elsewhere than its true location. Such hijacking attacks are called "remote redirection" attacks, since the malicious host, which may be operating at a network location far removed from the mobile node, nevertheless effectively redirects traffic away from the true location of the mobile node.

The link by which a mobile node is directly attached to the Internet may often be a wireless link. This link may thus have a substantially lower bandwidth and higher error rate than traditional wired networks. Moreover, mobile nodes are likely to be battery powered, and minimizing power consumption is important. Therefore, the number of administrative packets sent over

the link by which a mobile node is directly attached to the Internet should be minimized, and the size of these packets should be kept as small as is reasonably possible.

We assume that mobile nodes will generally not change their point of attachment to the Internet more frequently than once per second, although our protocol is likely to work even when the point of attachment changes more frequently. In addition, as is usual in the Internet today, we assume that IPv6 unicast packets are routed based on the destination address in the packet's IP header (and not, for example, influenced by the packet's IP source address).

3.2 Basic Operation of the Protocol

Mobile nodes will have assigned to their network interface(s) at least three IPv6 addresses whenever they are roaming away from their home subnet. One is its *home address*, which is permanently assigned to the mobile node in the same way as any IP node. The second address is the mobile node's current *Link-Local address*, as described in Section 2. Mobile IPv6 adds a third address, known as the mobile node's *care-of address*, which is associated with the mobile node only while visiting a particular foreign subnet. The network prefix of a mobile node's care-of address is equal to the network prefix of the foreign subnet being visited by the mobile node, and thus packets addressed to this care-of address will be routed by normal Internet routing mechanisms to the mobile node's location away from home.

Each time the mobile node moves its point of attachment from one IP subnet to another, the mobile node will configure its care-of address by stateless address autoconfiguration [?], or alternatively by some means of stateful address autoconfiguration such as DHCPv6 [6] or PPPv6 [12]. The decision about which manner of automatic address configuration to use is made according to the methods of IPv6 Neighbor Discovery [?]. A mobile node may have more than one care-of address at a time, for example if it is link-level attached to more than one (wireless) network at a time or if more than one IP network prefix is present on a network to which it is attached. As described in Section 1, the association between a mobile node's home address and its care-of address, along with the remaining lifetime of that association, is known as a *binding*. The central data structure used in Mobile IPv6 is a cache of mobile node bindings, maintained by each IPv6 node, known as a *Binding Cache*.

While away from home, a mobile node registers one of its bindings with a router in its home subnet, requesting this router to function as the *home agent* for the mobile node. The care-of address in this binding registered with its home agent is known as the mobile node's *primary care-of address*, and the mobile node's home agent retains this entry in its Binding Cache, marked as a "home registration," until its lifetime expires. While it has a home registration entry in its Binding Cache, the home agent uses proxy Neighbor Discovery [?] to intercept any IPv6 packets addressed to the mobile node's home address on the home subnet, and tunnels each intercepted packet to the mobile node's primary care-of address indicated in this Binding Cache entry. To tunnel the packet, the home agent encapsulates it using IPv6 encapsulation [8].

In addition, Mobile IPv6 provides a mechanism for IPv6 correspondent nodes communicating with a mobile node, to dynamically learn the mobile node's binding. The correspondent node adds this binding to its Binding Cache, although when space must be reclaimed in the Binding Cache, such a cache entry may be replaced at any time by any reasonable local cache replacement policy such as LRU. When sending a packet to any IPv6 destination, a node checks its Binding Cache for an entry for the packet's destination address, and if a cached binding for this address is found, the node routes the packet directly to the mobile node at the care-of address indicated in this binding; this routing uses an IPv6 Routing header [9] instead of IPv6 encapsulation, as this adds less overhead to the size of the packet. (The home agent cannot use a Routing header, since adding one to the packet at the home agent would invalidate the authentication in any IPv6 Authentication header included in the packet by the correspondent node.) If no Binding Cache entry is found, the node instead sends the packet normally (with no Routing header), and the packet is then intercepted and tunneled by the mobile node's home agent as described above.

This use of cached bindings for routing packets directly to a mobile node at its current care-of address is similar to the the existing work on "route optimization" [?] for IPv4 mobility [17]. However, since IPv6 is still being designed and there is no existing installed base of IPv6 nodes, we believe it is reasonable to require all IPv6 nodes to be capable of caching the binding of mobile nodes with which they are communicating. The additional burden on IPv6 nodes or implementations due to this additional functionality is negligible, since the operation of the Binding Cache is quite similar to the existing IPv6 Destination Cache [?] in the node and can easily be integrated with it.

Mobile IPv6 introduces two new IPv6 Destination options to allow a mobile node's home agent and correspondent nodes learn and cache the mobile node's binding. After configuring a new care-of address, a mobile node must send a *Binding Update* option containing that care-of address to its home agent, and to any

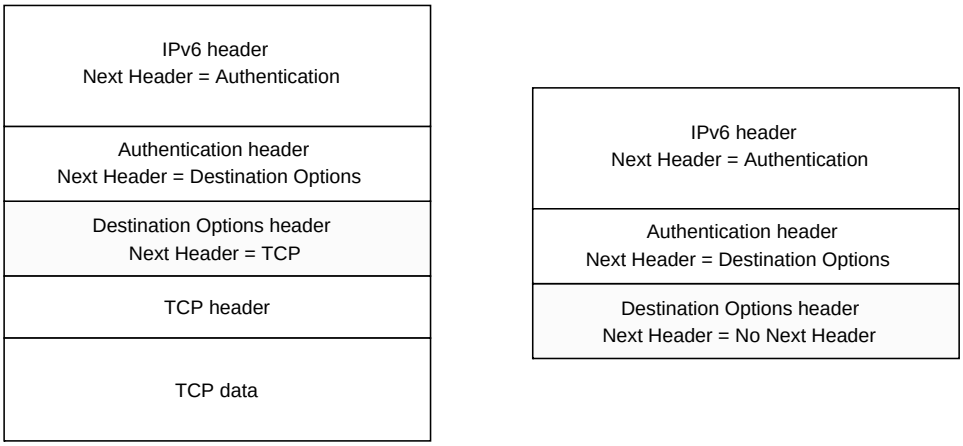

(a) Included within a TCP packet

(b) Sent alone in a separate packet

Figure 2 Sending Binding Updates and Acknowledgements as IPv6 Destination options

correspondent nodes that may have an out-of-date care-of address for the mobile node in their Binding Cache. Receipt of a Binding Update must be acknowledged using a *Binding Acknowledgement* option, if an acknowledgement was requested in the Binding Update.

Since an IPv6 Destination Options header containing one or more Destination options can appear in any IPv6 packet, any Mobile IPv6 option can be sent in either of two ways, as illustrated in Figure 2:

- A Binding Update or Binding Acknowledgement can be included within any IPv6 packet carrying any payload such as TCP [20] or UDP [19].
- A Binding Update or Binding Acknowledgement can be sent as a separate IPv6 packet containing no transport-level payload. In this case, the Next Header field in the Destination Options header is set to indicate "No Next Header" [9].

It is essential for scalability and minimizing network load that correspondent nodes be able to learn the care-of address of a mobile node, and that they cache this information for use in sending future packets to the mobile node's care-of address. By caching the care-of address of a mobile node, optimal routing of packets can be achieved between the correspondent node and the mobile node. Routing packets directly to the mobile node's care-of address also eliminates congestion at the home agent and thus contributes significantly to the overall health of the Internet.

Moreover, most packets sent by a correspondent node to a mobile node can be delivered with no assistance from the home agent. Thus, the impact of failure at the home agent can be drastically reduced; this is important because many administrative domains will have a single home agent to serve a particular home subnet, and thus a single point of failure for communications to nodes using that home agent. Furthermore, communications between the home agent and a mobile node may depend on a number of intervening networks; thus, there are many more ways that packets can fail to reach a mobile node when the home agent is required as an intermediate node. This would be particularly relevant on, say, trans-oceanic links between home agent and mobile node. Caching the binding of a mobile node at the correspondent node enables communication with the mobile nodes even if the home agent fails or is difficult to reach over the Internet.

4 The Binding Update Option

A mobile node sends a Binding Update to another node to inform it of its current binding. As noted in Section 3.2, since the Binding Update is sent as an IPv6 Destination option, the mobile node may include it in any existing packet (such as a TCP packet) that it is sending to that destination node, or it may send the Binding Update in a packet by itself. The Binding Update is used by a mobile node to register its primary care-of address with its home agent, and to notify correspondent nodes of its binding so that they may create or update entries in their Binding Cache for use in future communication with the mobile node.

A Binding Update should be considered a form of routing update; handled incorrectly, a Binding Update could be a source of security problems due to the possibility of remote redirection attacks. Therefore, packets which carry a Binding Update must also include an IPv6 Authentication header, which provides authenti-

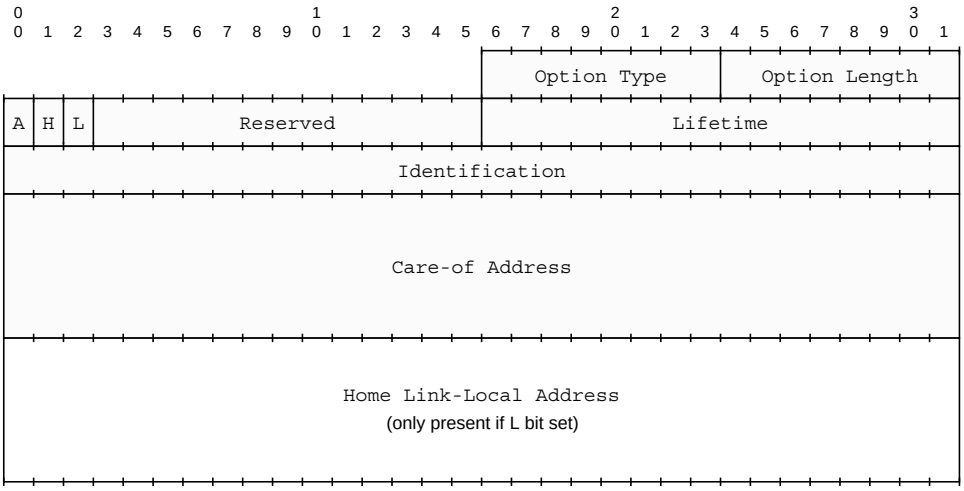

Figure 3 Binding Update option format

cation and replay protection for the Binding Update [2, ?].

The format of the Binding Update option is shown in Figure 3. Note that the sending mobile node's home address must be the source address in the IPv6 header of the packet containing the Binding Update, and thus need not be duplicated within the data of the Binding Update option.

The *Lifetime* field gives the number of seconds remaining before the binding must be considered expired. If the mobile node obtained its care-of address via some stateful address configuration service that assigns a lease time to the address, the Lifetime in the Binding Update could be set equal to the remaining lease time for the address. In any case, if the mobile node has an estimate for the length of time that it might be associated with the care-of address, it should use that estimate for the Lifetime of the Binding Update (subject to lease restrictions, if any). As special cases, a Lifetime value of all ones (0xffff) indicates infinity (the binding does not expire), and a value of zero indicates that the indicated binding for the mobile node should be deleted from the destination node's Binding Cache.

The *Identification* field is used to ensure that Binding Updates are not applied out of order at the destination in spite of varying network delays, and to match Binding Acknowledgements with outstanding Binding Updates at the mobile node. The mobile node increments a counter for each new Binding Update sent (*not* for retransmissions) and uses the value of the counter as the Identification field of the Binding Update.

The *Care-of Address* field gives the current care-of address of the mobile node. When the care-of address is equal to the home address of the mobile node, the Binding Update indicates that any existing entry for the mobile node in the destination node's Binding Cache should be deleted; no Binding Cache entry for the mobile node should be created.

When the *Home Agent (H)* bit is set in the Binding Update, the mobile node requests the destination of this Binding Update to serve as its home agent, with the binding contained in the Binding Update as the mobile node's primary care-of address. In this case, the mobile node may also include the Link-Local address it used when last on its home network, to cause the home agent to send proxy Neighbor Advertisements [?] for this address as well as for the mobile node's home address; the mobile node must set the the *Home Link-Local Address Present (L)* bit to indicate that the Home-Link Local Address field has been included in the Binding Update.

The mobile node may request a Binding Acknowledgement be returned for this Binding Update, by setting the *Acknowledge (A)* bit in the Update. In this case, if a mobile node fails to receive a Binding Acknowledgement within a specified period of time (nominally 1 second) after transmitting the Binding Update, it must retransmit the Binding Update with the same Identification value, and begin an exponential back-off process of retransmission. The time-out period is doubled upon each retransmission until a Binding Acknowledgement is received from the target of the Binding Update, or the time-out period reaches a maximum value (nominally 256 seconds). Typically, only Binding Updates to a mobile node's home agent with the Home Agent (H) bit set need to request an acknowledgement.

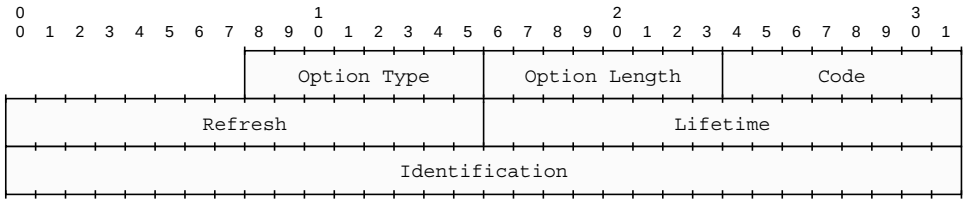

Figure 4 Binding Acknowledgement option format

5 The Binding Acknowledgement Option

The *Binding Acknowledgement* option is sent by a node to acknowledge receipt of a Binding Update (Section 4). When a node receives a Binding Update addressed to itself, in which the Acknowledge (A) bit is set, it must return a Binding Acknowledgement; other Binding Updates may also be acknowledged but need not be. The destination address in the IPv6 header of the packet carrying the Binding Acknowledgement must be the care-of address from the Binding Update, causing the Binding Acknowledgement to be returned directly to the mobile node sending the Binding Update. The format of the Binding Acknowledgement option is shown in Figure 4.

The *Code* field indicates the disposition of the Binding Update. Values of the Code field less than 128 indicate that the Binding Update was accepted by the receiving node, and values greater than or equal to 128 indicate that the Binding Update was rejected. A Binding Update may be rejected, for example, due to administrative reasons or lack of sufficient resources to create the Binding Cache entry, or due to a poorly formed Binding Update option. The Binding Update may also be rejected as part of the home agent discovery process, as outlined in Section 8.3.

The *Lifetime* field gives the lifetime for which the sending node will attempt to retain the binding being acknowledged in its Binding Cache. The Lifetime value must not be greater than that requested in the Binding Update, although the home agent may reduce the requested value. To provide some robustness if the home agent cannot store the binding in stable storage (so that it can survive a crash of the home agent), the *Refresh* field gives a suggested interval in seconds (possibly less than Lifetime) at which the mobile node should periodically send a new Binding Update to its home agent to reregister its primary care-of address.

The *Identification* field is copied from the Binding Update, for use by the mobile node in matching the Binding Acknowledgement with an outstanding Binding Update.

6 Mobile Node Operation

Every IPv6 mobile node must be able to perform IPv6 decapsulation. Every mobile node must be able to send Binding Updates and to receive Binding Acknowledgements. Based on the Lifetime field in Binding Updates that it sends, every IPv6 mobile node must keep track of which other IPv6 nodes may need to receive a new Binding Update as a result of recent movement by the mobile node. Every IPv6 mobile node must also be able to send Binding Updates when it receives a packet from a correspondent node encapsulated to it by its home agent, rather than sent directly to it by the correspondent node using a Routing header.

6.1 Sending Binding Updates

After moving its link-layer point of attachment to a new IP subnet and configuring a new primary care-of address, the mobile node registers this new care-of address with its home agent by sending it a packet containing a Binding Update. This Binding Update must have the Home Agent (H) bit set to indicate this as a home registration, and must have the Acknowledge (A) bit set to request an acknowledgement of the Binding Update.

In the case in which the mobile node is returning to its home subnet, the Binding Update sent to its home agent must specify the mobile node's home address as the care-of address. The mobile node must also send the appropriate IPv6 Neighbor Advertisement messages with the Override flag set, so that its neighbors on its home subnet will update the link-layer address for the mobile node in their Neighbor Caches. The mobile node must do this for both its Link-Local address and its home address. The Neighbor Advertisement messages can be repeated a small number of times to guard against occasional loss of packets on the home subnet. Receipt of this Neighbor Advertisement message overrides the previous proxy Neighbor Discovery actions taken by the home agent when the mobile node earlier left its home subnet.

A Binding Update may also be included, whenever necessary, in any normal packet sent to a correspondent node. For each correspondent node, the remaining

lifetime of the binding sent to that node is kept by the mobile node to determine whether or not the correspondent node has been sent a fresh Binding Update since the last time the mobile node acquired a new care-of address. When a packet is to be sent to a correspondent node which has not been sent a fresh Binding Update, the mobile node should include the Update within the packet, and then record that the update has been sent. Thus, correspondent nodes are generally kept updated and can send almost every packet directly to the mobile node. Such Binding Updates are not generally required to be acknowledged, although the mobile node may request an acknowledgement if desired.

The Binding Update can also be sent alone in a separate packet, whenever the mobile node wishes to update its correspondents. This is typically done only if the mobile node suspects that its home agent is not operational or is too far away, a correspondent node is not sending the traffic to the proper care-of address, or there is an immediate need for the correspondent node to obtain the binding. A mobile node can detect that a correspondent node is not sending packets to the proper care-of address because in that case the packets arrive at the mobile node's care-of address by encapsulation (from the home agent) instead of by use of a Routing header in the packet.

The mobile node must not send Binding Updates too often (no more than once per second) to any correspondent node. After sending a number of consecutive Binding Updates to a particular correspondent node with the same care-of address, the mobile node must reduce its rate of sending Binding Updates to that correspondent node. The mobile node may continue to send Binding Updates at the slower rate indefinitely, in hopes that the correspondent node will finally be able to process one of them and begin to use the new route to the mobile node.

The mobile node may choose to keep its location private from certain correspondent nodes and need not send Binding Updates to those correspondents. No other IPv6 nodes are authorized to send Binding Updates on behalf of the mobile node. A mobile node can send a Binding Update to a correspondent node with its home address as the care-of address, instructing the correspondent node to delete any existing entry for the mobile node in its Binding Cache.

6.2 Movement Detection

A mobile node may use any combination of mechanisms available to it to detect when its link-level point of attachment has moved from one IP subnet to another. The primary movement detection mechanism for Mobile IPv6 uses the facilities of the IPv6 Neighbor Discovery protocol, including Router Discovery and Neighbor Unreachability Detection. The description here is based on the conceptual model of the organization and data structures defined by Neighbor Discovery [?].

Mobile nodes use Router Discovery to discover new routers and on-link network prefixes; a mobile node may send Router Solicitation messages, or may wait for unsolicited (periodic) Router Advertisement messages, as specified for Router Discovery. Based on received Router Advertisement messages, a mobile node (in the same way as any other node) maintains an entry in its Default Router List for each router, and an entry in its Prefix List for each network prefix, that it currently considers to be on-link. Each entry in these lists has an associated invalidation timer value (extracted from the Advertisement) used to expire the entry.

While away from home, a mobile node should select one router from its Default Router List to use as its default router, and one network prefix advertised by that router from its Prefix List to use as the network prefix in its primary care-of address. The mobile node configures a new care-of address using this prefix and registers it with its home agent as its primary care-of address, as described in Section 6.1. A mobile node may also have associated additional care-of addresses, using other network prefixes from its Prefix List.

While away from home and using some router as its default router, it is important for a mobile node to be able to quickly detect when that router becomes unreachable, so that it can switch to a new default router and to a new care-of address. Since some links (notably wireless) do not necessarily work equally well in both directions, it is likewise important for the mobile node to detect when it becomes unreachable to its default router, so that any correspondent nodes attempting to communicate with the mobile node can still reach it.

To detect when its default router becomes unreachable, a mobile node should use Neighbor Unreachability Detection. As specified in IPv6 Neighbor Discovery, while the mobile node is actively sending packets to (or through) its default router, the mobile node can detect that the router has become unreachable either through indications from upper layer protocols on the mobile node that a connection is not making "forward progress" (e.g., TCP timing out waiting for an acknowledgement after a number of retransmissions), or through the failure to receive a Neighbor Advertisement message form its default router in response to retransmitted explicit Neighbor Solicitation messages to it. No changes in IPv6 Neighbor Unreachability Detection are necessary for this aspect of movement detection in Mobile IPv6.

For a mobile node to detect when it has become unreachable to its default router, however, the mobile node

cannot efficiently rely on Neighbor Unreachability Detection alone, since the network overhead would be prohibitively high in many cases for a mobile node to continually probe its default router with Neighbor Solicitation messages even when it is not otherwise actively sending packets to it. Instead, a mobile node should consider receipt of any IPv6 packets from its current default router as an indication that it is still reachable from the router. Both packets from the router's IP address and (IPv6) packets from its link-layer address (e.g., those forwarded but not originated by the router) should be considered.

Since the router should be sending periodic multicast Router Advertisement messages, the mobile node will have frequent opportunity to check if it is still reachable to its default router, even in the absence of other packets to it from the router. On some types of network interfaces, the mobile node may also supplement this by at times setting its network interface into "promiscuous" receive mode, so that is able to receive all packets on the link, including those not link-level addressed to it. The mobile node will then be able to detect any packets sent by the router, in order to to detect reachability from the router.

If the above means do not provide indication that the mobile node is still reachable from its current default router (i.e., the mobile node receives no packets form the router for a period of time), then the mobile node should actively probe the router with Neighbor Solicitation messages, even if it is not otherwise actively sending packets to the router. If it receives a solicited Neighbor Advertisement message in response from the router, then the mobile node can deduce that it is still reachable. It is expected that the mobile node will in most cases be able to determine its reachability from the router by listening for packets from the router as described above, and thus, such extra Neighbor Unreachability Detection probes should rarely be necessary.

With some types of networks, it is possible that additional information can be obtained from lower-layer protocol or device driver software within the mobile node. For example, a mobile node may use wireless signal strength or signal quality information (with suitable hysteresis) for its link with the available default routers to decide when to switch to a new default router and primary care-of address. Even though the mobile node's current default router may still be reachable in terms of Neighbor Unreachability Detection, the mobile node may use such lower-layer information to determine that switching to a new default router would provide a better connection.

6.3 Smooth Handoffs with Overlapping Cells

In the case of wireless communications, a mobile node is often effectively attached to the Internet at multiple points of attachment. When a mobile node moves from one point of attachment towards another, the mobile node will, upon detection of the movement, likely configure a new care-of address for the new point of attachment, and report the new care-of address to its home agent (by way of a Binding Update).

For smooth handoffs, a mobile node should still accept packets at its previous care-of address even after reporting its new care-of address to its home agent. This is reasonable, since the mobile node could only receive packets at its previous care-of address if it were indeed still attached to the Internet at that care-of address. These considerations obtain whether or not the mobile node has configured its care-of address with stateless address autoconfiguration, or by way of a stateful address allocation authority such as DHCPv6 [6]. If the previous address is allocated by such a stateful address server, then such a mobile node may not wish to release the address immediately upon acquisition of a new care-of address. The stateful address server will allow mobile nodes to acquire new addresses while still using previously allocated addresses.

6.4 Router-Assisted Smooth Handoffs

Routers (just as any IPv6 node) must be able to accept authenticated Binding Updates for a mobile node and, subsequently, act on the cached binding by encapsulating packets for intermediate delivery to the care-of address specified in the binding. In cases in which a mobile node moves from one care-of address to another without being able to maintain simultaneous connectivity at both care-of addresses, the mobile node should send a Binding Update to the router servicing the previous care-of address, so that packets for the mobile node can be delivered to the new care-of address immediately. For example, a mobile node may move from one radio link to another on a different channel, and may be unable to monitor packets transmitted over both channels at once. In this example, the mobile node should send a Binding Update to the previous router, which is the entity delivering packets to the mobile node over the previous radio channel, so that those packets will instead be delivered via its new care-of address. This Binding Update associates the mobile node's (immediately) previous care-of address to the mobile node's new care-of address, and is authenticated using the IPv6 Authentication header with whatever security association

the previous router had with the mobile node's previous care-of address

Note that the previous router does not necessarily know anything about the mobile node's home address as part of this sequence of events; the previous router only knows about the care-of address used by the mobile node at its previous point of attachment. The mobile node in effect requests the previous router to serve as a temporary home agent for its own previous care-of address, with a primary care-of address of its new care-of address (through its new router). Thus, in the Binding Update to its previous router, the mobile node sets the Home Agent (H) bit to request the router to serve as a home agent, and sets the Acknowledge (A) bit to request a Binding Acknowledgement from the router.

The previous router then operates in the same way as when the mobile node's home agent (for its home address) receives a Binding Update from the mobile node. That is, the previous router must intercept any packets destined for the home address indicated in the Binding Update (the mobile node's previous care-of address), and tunnel any such intercepted packets to the care-of address indicated in the Binding Update (the mobile node's new care-of address). This tunneling is done using IPv6 encapsulation [8], in the same way as packets tunneled from any home agent. Once the mobile node receives the encapsulated packet, it can then typically follow the Routing header contained in the decapsulated packet (that the correspondent node used to route the packet to the mobile node's previous care-of address) and send a Binding Update to this correspondent node giving its new care-of address.

6.5 Renumbering the Home Subnet

IPv6 Neighbor Discovery specifies a mechanism by which all nodes on a subnet can gracefully autoconfigure new addresses, for example when the home subnet changes its Internet service to a different service provider and must change its network prefix (and thus the network prefix of all nodes on the home subnet). As currently specified, this mechanism works when the nodes are on the same link as the router issuing the necessary multicast Router Advertisement packets to advertise the new routing prefix(es) appropriate for the subnet.

However, for mobile nodes not currently attached to the same link as their home agent, special care must be taken to allow the mobile nodes to renumber gracefully along with the rest of its home subnet. The most direct method of correctly extending the renumbering to mobile nodes away from home is for the home agent to tunnel the multicast Router Advertisement packets used for renumbering, to the care-of address of each mobile node for which it is serving as the home agent. The rules for this are as follows:

- A mobile node assumes that its home network prefix has not changed unless it receives an authenticated Router Advertisement message from its home agent that the prefix has changed.
- When the mobile node is at home, the home agent does not tunnel Router Advertisements to it.
- When a home network prefix changes, the home agent tunnels the Router Advertisement messages to each mobile node which is currently away from home and using a home address with the affected network prefix. Such tunneled Router Advertisements must be authenticated using an IPv6 Authentication header [2].
- When a mobile node receives a tunneled Router Advertisement containing a new home network prefix, it must perform the standard autoconfiguration operation to create its new home address
- When a mobile node returns to its home subnet, it must again perform IPv6 Duplicate Address Detection at the earliest possible moment after it has registered with its home agent.
- A mobile node may send a Router Solicitation message to its home agent at any time, within the constraints imposed by rate control in IPv6 Neighbor Discovery [?].

7 Correspondent Node Operation

Every IPv6 node may at times be a correspondent node communicating with one or more mobile nodes. Thus, every IPv6 node must be able to process received Binding Updates, to send Binding Acknowledgements, and to maintain a Binding Cache.

An existing requirement in IPv6 is that every IPv6 node must be able to maintain security associations for use in IPv6 Authentication headers, as described in Section 2. Nodes receiving a Binding Update or Binding Acknowledgement must verify the authentication data contained in the Authentication header in the packet carrying the Update or Acknowledgement.

7.1 Delivering Packets to a Mobile Node

Before sending a packet to any destination address, a node must check its Binding Cache for a binding for

this address. If no Binding Cache entry is found, the node sends the packet normally (with no Routing header or encapsulation), and the packet thus will be routed by normal Internet routing mechanisms to the mobile node's home subnet; there, the packet will be intercepted by the mobile node's home agent and tunneled (using IPv6 encapsulation) to its current primary care-of address. Thus, packets can be sent by a correspondent node to the mobile node without the sender knowing that the node is mobile. Once the mobile node receives the encapsulated packet, it will send a Binding Update to the sender, as described in Section 6.1.

If, on the other hand, the correspondent node has a Binding Cache entry for the destination address of a packet it is sending, the correspondent node should send the packet directly to the care-of address indicated in the binding, using an IPv6 Routing header [9]. To use the Routing header for delivery of a packet to a mobile node, the correspondent node lists the care-of address in the packet's IPv6 header as the destination address, and lists the mobile node's home address (the original destination address of the packet) in the Routing header as the final destination of the packet. When the packet arrives at the care-of address, normal processing of the Routing header by the mobile node will cause the packet to be delivered to the upper-layer software in the mobile node using the mobile node's home address.

7.2 Handling Returned ICMP Errors

If a correspondent node receives persistent ICMP Host Unreachable or Network Unreachable messages after sending packets to a mobile node using its cached care-of address, it should delete the cache entry; it may then create a new Binding Cache entry for the mobile node when the mobile node's current care-of address becomes available again via a new Binding Update from the mobile node.

8 Home Agent Operation

Every IPv6 router must perform the mobility-related functions defined in Section 7 for correspondent nodes, but not necessarily the additional functions defined in Section 6 for mobile nodes.

In addition, every IPv6 router must be able to send Binding Acknowledgements in response to Binding Updates received from a mobile node. Every IPv6 router must also be able to encapsulate packets in order to tunnel them to a care-of address known for a mobile node for which it is serving as a home agent.

8.1 Delivering Packets to a Mobile Node

A home agent cannot use a Routing header to deliver intercepted packets to the mobile node for which it is serving as the home agent, because it cannot modify the packet or add to it in flight, as such modifications or additions would cause any IPv6 authentication [2] in the packet to fail at the receiver. Instead, a home agent must always use IPv6 encapsulation [8] for delivering intercepted packets to a mobile node.

When a home agent encapsulates a packet for delivery to the mobile node, the home agent uses the care-of address as the destination address in the outer IPv6 header, and uses its own address as the outer source address. Since the mobile node is presumed to be receiving packets at the care-of address, the delivery path from the care-of address to the mobile node's home address is then trivial and is entirely within the mobile node itself. The home agent is expected to be involved only rarely with the transmission of packets to the mobile node, because the mobile node will send Binding Updates as soon as possible to its correspondent nodes.

8.2 Proxy Neighbor Advertisements

When a mobile node first registers with its home agent after leaving its home subnet, the home agent must send onto the home subnet a gratuitous Neighbor Advertisement on behalf of the mobile node, with the Override flag set [?], giving its own link-layer address as the associated link-layer address for the mobile node's IPv6 home address. All nodes on the home subnet receiving this Neighbor Advertisement will then update their Destination Cache entry for the mobile node to contain the link-layer address of the home agent, allowing the home agent to intercept any packets sent by them to the mobile node. Some correspondent nodes on the home subnet may have established connections with the mobile node by using the mobile node's Link-Local address, and thus the home agent must also send a gratuitous Neighbor Advertisement for the mobile node's Link-Local address as well. Each of these gratuitous Neighbor Advertisements should be repeated a few times to increase reliability, although any node that does not receive one of these Advertisements will be able to detect the change in the mobile node's link-layer address using IPv6 Neighbor Unreachability Detection [?].

In addition, while the mobile node is away from home, the home agent should act as a proxy for the mobile node, replying to any received Neighbor Solicitation messages received for the mobile node's home address or Link-Local address. These proxy Neighbor Advertisement replies ensure that correspondent nodes on

the home subnet retain the home agent's link-layer address in their Destination Cache while the mobile node is away from home.

8.3 Home Agent Discovery

It is useful to be able to send a Binding Update to a mobile node's home agent without explicitly knowing the home agent's address. For example, since the mobile node was last at home, it may have become necessary to replace the node serving as its home agent due to the failure of the original node or due to reconfiguration of the home subnet. It thus may not always be possible or convenient for a mobile node to know the exact address of its own home agent.

Mobile nodes can dynamically discover the address of a home agent by sending a Binding Update to the IPv6 anycast address on their home subnet. A packet sent to an anycast address is received by exactly one router on the destination subnet. Any router on the home subnet which receives this Binding Update (responding to the anycast address) must reject the Binding Update and include its own (unicast) IPv6 address in the Binding Acknowledgement indicating the rejection. The mobile node will then repeat its Binding Update, sending it directly to the router that returned the rejection.

8.4 Handling Returned ICMP Errors

When sending a packet to a mobile node, it is important to correctly return to the original sender any ICMP error messages generated by the packet. Since it is expected that in most cases, the correspondent node will have a Binding Cache entry for the destination mobile node, most packets sent by to a mobile node will use a Routing header rather than encapsulation. In this case, the source address in the packet's IPv6 header along the entire path to the mobile node will be that of the correspondent node, allowing any ICMP error messages generated by the packet to be automatically returned directly to the correspondent node.

However, when the correspondent node has no Binding Cache entry for the destination mobile node, the packet will be intercepted by the mobile node's home agent, encapsulated, and tunneled to the mobile node's care-of address. In this case, the source address in the packet's (outer) IPv6 header between the home agent and the mobile node will be that of the home agent, not the correspondent node. In order to return the ICMP error message to the original sender of the packet (the correspondent node), IPv6 has defined tunneling using encapsulation to "relay" the ICMP error message to the original sender [8].

ICMP error messages contain as payload, a portion of the packet generating the error message. Whereas ICMP for IPv4 [21] generally returned only the first 8 bytes beyond the IP header of the packet generating the error, ICMP for IPv6 [7] returns as much of the original packet as will fit in the ICMP error message without causing the entire ICMP packet to exceed 576 bytes. This size is sufficient to allow the tunnel entry point node (the home agent in this case) to remove the IPv6 encapsulation header from the returned packet in the ICMP error message, and then to forward the modified ICMP message to the original sender of the packet causing the error (the correspondent node in this case).

9 Security Issues

The IPv6 and IP Security specifications require authentication to be implemented by all IPv6 nodes [3, 9]. Thus, whenever a mobile node is able to establish a security association with a correspondent node, the mobile node will be able to send authenticated Binding Updates to that correspondent node. Methods for distributing keys for use in computing the authentication data in the IPv6 Authentication header between nodes are not completely specified yet within the IETF but are receiving a tremendous amount of attention within the IETF's IP Security Working Group [4, 13, 11, 15].

9.1 Session Keys with Local Routers

In the IPv4 "route optimization" proposal [**?**], a mechanism is outlined whereby a session key can be established between a mobile node and its foreign agent, without requiring any pre-established security relationship between them. A similar mechanism will work in IPv6, to avoid the need for a possibly time-consuming negotiation between routers and mobile nodes for the purpose of obtaining the session key, which under many circumstances would only be used once. This mechanism, if needed, can be specified completely outside the Mobile IPv6 protocol and would amount to a way of creating a dynamic security association between two nodes which do not share a trust relationship, but which need to agree on a key for some limited purpose (here, allowing the future authentication of a single Binding Update to indicate that the mobile node has moved away from this router). Methods for key distribution for use by Internet hosts, being standardized now, should allow this function to be performed appropriately for mobile nodes, say based on a Diffie-Hellman key exchange.

9.2 Source Address Filtering by Firewalls

The current design of Mobile IPv6 does nothing to permit mobile nodes to send their packets through firewalls which filter out packets with the "wrong" source address in their IPv6 header. The mobile node's home address may be unlikely to fall within the ranges required to satisfy the firewall's criteria for allowing the packet through the firewall.

Firewalls are unlikely to disappear. However, any solution [25] to the firewall problem based on hiding the non-local source address outside the source address field of the IPv6 header is likely to fail. Any vendor or facilities administrator wanting to filter based on the address in the IPv6 source address field would also quickly begin also filtering on such hidden source addresses.

Assume, for the moment, that a mobile node is able to send packets through a firewall protecting the domain in which a correspondent node is located. The mobile node could then encapsulate its packet so that the outer IPv6 header was addressed to the firewall and used the mobile node's care-of address as the source address. When the firewall decapsulates, it would be able to authenticate the inner packet based on the mobile node's home address. After the authentication is performed, the firewall could forward the packet to the correspondent node as desired. This simple procedure has the feature that it requires the minimal amount of encapsulation, no assistance by routers or other agents, and that the firewall can establish a security relationship with the mobile node based on its home (i.e., permanent) address.

10 Conclusions and Current Status

We have presented an efficient and deployable protocol for handling mobility with the new IPv6 protocol, and suitable for use with the coming multitudes of mobile nodes. We believe our protocol is as lightweight as possible, given the need to be transparent to higher level protocols; among schemes which propagate updates to any agent on the home subnet, our proposal attempts to minimize the control traffic needed to effect mobility while nevertheless supplying the necessary information to all IPv6 nodes which need it, in an event-driven fashion. Thus, we expect that this protocol will minimize exposures to network partitions, congestion, and also minimize administrative network traffic between mobile nodes and correspondent nodes. This protocol has undergone a number of revisions in response to suggestions by members of the Mobile IP working group and the IPng (IP: The Next Generation) working group of the IETF. We hope to have this protocol progress to the next stage of standardization in the near future.

References

[1] Draft Standard, Wireless LAN MAC and PHY Specifications, Rev. D1. IEEE Document P802.11/D1-94/12, Dec 1994.

[2] R. Atkinson. IP Authentication Header. RFC 1826, August 1995.

[3] R. Atkinson. Security Architecture for the Internet Protocol. RFC 1825, August 1995.

[4] A. Aziz, T. Markson, and H. Prafullchandra. Simple Key-Management For Internet Protocols (SKIP). draft-ietf-ipsec-skip-07.txt, August 1996. (work in progress).

[5] Internet Architecture Board and Internet Engineering Steering Group. The Internet Standards Process – Revision 2. RFC 1602, March 1994.

[6] J. Bound and C. Perkins. Dynamic Host Configuration Protocol for IPv6. draft-ietf-dhc-dhcpv6-05.txt – work in progress, June 1996.

[7] A. Conta and S. Deering. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6). RFC 1885, December 1995.

[8] A. Conta and S. Deering. Generic Packet Tunneling in IPv6. draft-ietf-ipngwg-ipv6-tunnel-01.txt - work in progress, February 1996.

[9] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 1883, December 1995.

[10] Stephen E. Deering, editor. ICMP Router Discovery Messages. RFC 1256, September 1991.

[11] D. Harkins and D. Carrel. The resolution of ISAKMP with Oakley. draft-ietf-ipsec-isakmp-oakley-00.txt, June 1996. (work in progress).

[12] D. Haskin and E. Allen. IP Version 6 over PPP. draft-ietf-ipngwg-pppext-ipv6cp-03.txt - work in progress, June 1996.

[13] P. Karn and B. Simpson. The Photuris Session Key Management Protocol. draft-ietf-ipsec-photuris-08.txt, July 1996. (work in progress).

[14] M. Oehler and R. Glenn. HMAC-MD5 IP Authentication with Replay Prevention. draft-ietf-ipsec-ah-hmac-md5-01.txt, July 1996. (work in progress).

[15] H. Orman. The OAKLEY Key Determination Protocol. draft-ietf-ipsec-oakley-01.txt, May 1996.

(work in progress).

[16] C. Perkins and D. Johnson. Mobility Support in IPv6. draft-ietf-mobileip-ipv6-00.txt – work in progress, January 1996.

[17] C. Perkins, Editor. IPv4 Mobility Support. ietf-draft-mobileip-protocol-17.txt, May 1996. (work in progress).

[18] David C. Plummer. An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Addresses for Transmission on Ethernet Hardware. RFC 826, November 1982.

[19] J. B. Postel. User Datagram Protocol. RFC 768, August 1980.

[20] J. B. Postel, editor. Transmission Control Protocol. RFC 793, September 1981.

[21] J. B. Postel, Editor. Internet Control Message Protocol. RFC 792, September 1981.

[22] J. B. Postel, Editor. Internet Protocol. RFC 791, September 1981.

[23] Ronald L. Rivest. The MD5 Message-Digest Algorithm. RFC 1321, April 1992.

[24] IETF Secretariat and G. Malkin. The Tao of IETF. RFC 1718, November 1994.

[25] Fumio Teraoka. draft-teraoka-ipv6-mobility-sup-02.txt. Internet Draft – work in progress, January 1996.