IPv6 Working Group
INTERNET DRAFT

Charles Perkins
IBM Corporation
David B. Johnson
Carnegie Mellon University
8 July 1995

Mobility Support in IPv6

<draft-perkins-ipv6-mobility-sup-02.txt>

Abstract

This document presents some suggestions for mobility support in IPv6, drawing on the experiences of the authors in our work with IPv4 mobility within the Mobile IP Working Group of the IETF. The development of IPv6 presents a rare opportunity to consider in what ways mobility could explicitly be taken into account in the design of IPv6, and in what ways the current work on mobility within IPv4 can or should be changed to take advantage of IPv6. We believe that the most important function needed to support mobility is the reliable and timely notification of a mobile node's current location to other nodes that need it: the home agent, the correspondent nodes, and its nearest router.

Status of This Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups and individuals may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

1. Introduction

A new version of the Internet Protocol, IPv6, is being developed with 128-bit addresses, which remedies many perceived flaws with the existing version (that is, IPv4). This document draws on the experiences of the authors during the design of a set of protocols for the operation of mobile computers for IPv4, in our work within the Mobile IP Working Group of the IETF [4]. Mobile computers are very likely to account for a substantial fraction of the future population of the Internet during the lifetime of IPv6. We expect that the combination of a projected need for mobile computing, and clearly specified features within IPv6 to enable it, should make the necessary operations essentially automatic and universally available.

The IETF Mobile IP Working Group's current protocol design for mobility in IPv4 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 [4]. 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 draft, therefore, considers how IPv6 can most naturally fulfill the support requirements for mobile nodes. and in what ways the IPv4 mobility design can or should be changed to take advantage of IPv6.

We believe that the most important function needed to support mobility is the reliable and timely notification of a mobile node's current location to other nodes that need it. The home agent needs this location information in order to forward intercepted packets from the home network to the mobile node, and correspondent nodes need this information in order to send their own packets directly to the mobile node.

In this document, we will first specify the way that the mobile node can send notifications about its current whereabouts, using mostly existing mechanisms available already in IPv6. Then we describe the mechanism by which a routing header can be used to deliver packets to the mobile node at its current whereabouts. We suppose that all IPv6 nodes and routers can support the operations required for mobility, since the additional overhead of doing so is not very high. This leads to dramatic simplifications in the required protocols. In this proposal, we preserve features analogous to all of the features available to mobile nodes using the IPv4 mobile-IP protocol.

2. Basic Operation

From the model of operation developed for enabling mobile networking for IPv4, we borrow the concepts of home network, home address, home agent, care-of address, and binding. Accordingly, mobile computers will use (at least) two IPv6 addresses whenever they are roaming away from their home network. One is permanent; the other is temporary.

In brief, using the IPv4 language, we have a basic model of operation in which a mobile node can always be reached by sending packets to its home (permanent) address. Assuming the mobile node is not present on its home network, packets arriving for it there will be intercepted by the home agent, and tunneled to a care-of address.

In the configurations described first in this document, the mobile node can itself receive packets addressed to the care-of address. Alternatively, a router will receive the tunneled packets, and deliver them directly to the mobile node.

Generally, mobile nodes will select one or more of the available care-of addresses, possibly by collecting offers of service from routers in the area, and make sure the home agent is aware of its currently valid care-of address(es). The method of reporting the binding to the home agent (i.e., the association between care-of address and home address) is substantially different than what is currently specified for IPv4. The method by which care-of addresses can be discovered will depend on the final form the Neighbor Discovery Protocol [7] for IPv6. The packets are preferentially delivered to mobile nodes by using routing headers instead of encapsulation.

3. Binding Updates

Also borrowed from existing work on route optimization for IPv4 mobility is the concept of a location cache for mobile node bindings. In IPv6, we specify that all IPv6 nodes be capable of caching the location of mobile nodes with which they want to communicate, and recommend that this location cache be integrated with the node's conventional routing table.

We view it as essential for scalability and performance that correspondent nodes be able to learn the location of a mobile node and to be able to cache this knowledge for use in sending future packets directly to the mobile node. By caching the location 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 also eliminates congestion at the home agent

and thus contributes significantly to the overall health of the Internet. Moreover, many communications between the mobile nodes and its correspondent nodes can be carried out 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 network, and thus a single point of failure for communications to nodes on that home network. Besides that, communications between the home agent and any mobile node depend on perhaps many 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 contact over the Internet.

Binding updates should be considered a form of routing updates; thus, handled incorrectly, they could be a source of security problems and routing loops. We assume that in the deployed IPv6 systems there will be access to suitable authentication mechanisms which can authenticate binding updates.

3.1. Binding Update Option Format

We introduce a new IPv6 destination option by which a mobile node can transmit a binding update to another IPv6 node. A mobile node uses the Binding Update option to notify another node of its current care-of address. The binding update should be placed in the IPv6 packet after any routing header, since the binding update should only be processed by the destination node rather than by each hop along the path. The binding update is encoded as an option within the destination extension header. This alternative has the advantage that it does not require the allocation of any new protocol number, although there isn't any shortage yet of protocol numbers for IPv4 or IPv6. By encoding the binding update in this way, it can be included in any normal data packet or can be sent in a separate packet containing no data. The binding update should contain the mobile node's care-of address, an identification for the binding (to protect against attempts to replay the update), and possibly a lifetime for the binding. Note that the binding update is functionally similar to the previously suggested [9] "Remote Redirect", which was intended to facilitate the dissemination of mobility bindings to those correspondent hosts that need them.

This option format is adapted from that suggested in the IPv4 route optimization proposal [6]. Note that the home address is required to be the source address of IPv6 packet containing the binding update, and thus is not required to be located within the data of the destination option.

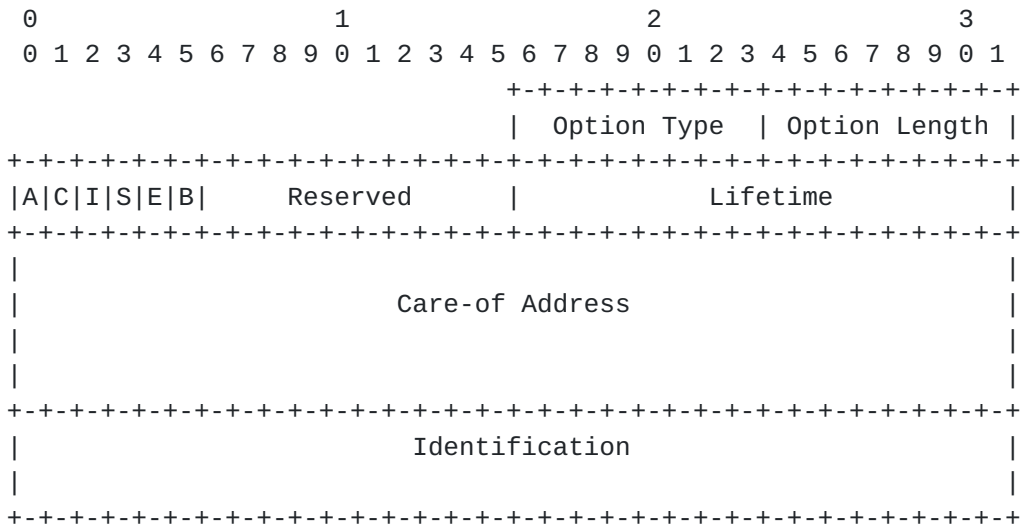

Option Type

8-bit identifier of the type of option. The first three bits of the option are 000, indicating first that a node receiving

the option may discard the option and still process the rest of the packet, and second that the option may not be modified enroute.

Option Length

8-bit unsigned integer. Length of the Option Data field of this option, in octets.

Acknowledge (A)

The Acknowledge (A) bit is set by a node if it wants a a Binding Acknowledge message to be returned upon receipt of the Binding Update Option.

Co-location (C)

The mobile node is itself the agent receiving datagrams at the care-of address.

Identification Present (I)

The (I) bit is set by the node sending the Binding Update option to indicate whether or not the Identification field is present.

Simultaneous (S)

The (S) bit is set by the mobile node if it wishes the receiver to maintain multiple simultaneous bindings for the mobile node.

Encapsulation (E)

The (E) bit is set by the mobile node to request that the receiving agent send encapsulated packets to the mobile node, instead of packets containing the care-of address in a routing header.

Broadcast (B)

The (B) bit is set by the mobile node to request that the home agent encapsulate and send broadcast packets to the mobile node at its care-of address. The (B) bit must only be used when sending binding updates to the home agent.

Reserved

Sent as 0; ignored on reception.

Lifetime

The number of seconds remaining before the location cache entry must be considered expired. A value of all ones indicates infinity. A value of zero indicates that the indicated location cache entry (or route table entry, in the case of a mobile node's previous router) for the mobile node should be deleted. The lifetime is typically equal to the remaining lifetime of the mobile node's binding with its care-of address.

Care-of Address

The current care-of address of the mobile node. When set equal to the home address of the mobile node, the Binding Update option instead indicates that no location cache entry for the mobile node should be created, and any existing location cache entry (and route table entry, in the case of a mobile node's previous router) for the mobile node should be deleted.

Identification

If present, a 64-bit number, used to assist in matching acknowledgements with binding updates, and in protecting against replay attacks.

The receiver of this message must be able to tell, say by employing whatever means adopted by the IPv6 working group for authenticating network-layer packets [1], that the mobile node is truly the agent which has generated the binding update.

4. Sending Binding Updates

After moving to a new location, the mobile node registers its new binding with its home agent by sending a packet containing a binding update to its home agent. This binding update MUST set the (A) bit, instructing the home agent to send an acknowledgement.

A binding update within an IPv6 header may also be included, when necessary, in any normal data packet sent to a correspondent node. For each correspondent node, an indication 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 any movement to a new care-of address has occurred. When a packet is sent to a correspondent node which hasn't been sent a fresh update, the mobile node includes the update within the packet's IPv6 header, and indicates that the update has been sent. Thus, correspondent nodes are generally kept updated and can send almost all data packets directly to the mobile node. Such binding updates are not generally required to be acknowledged. However, if the mobile node wants to be sure, an acknowledgment can be requested.

The binding update can also be sent in an otherwise empty packet whenever the mobile node wishes to update its correspondents. This would only be done if the mobile node suspects that its home agent is not operational, too far away, or that there may be an immediate need for the correspondent node to obtain the location information.

An IPv6 authentication header must be used so that the recipient can be assured that the routing information is authentic.

The mobile node achieves location privacy simply by limiting the correspondents to which it will send binding updates. No other IPv6 nodes are authorized to send binding updates on behalf of the mobile node.

No matter how binding updates are transmitted to correspondent nodes, some sort of back-off scheme must be implemented in the mobile node's software to avoid a rush of updates upon every movement to a new service area. Finally, some consideration should be made for the continued existence of IPv4 correspondent nodes, which are less likely to cache bindings.

5. Binding Acknowledgement Message

A Binding Acknowledge message is used to acknowledge receipt of a Binding Update (section 3.1) option, if that option has the Acknowledge (A) bit set.

Since the Binding Acknowledgement is mostly used by home agents and is not associated with any transmission of data packets, it is specified here as an informational ICMP message to the mobile node. However, all of the error conditions specified in the Registration Reply message of the IPv4 mobile-IP protocol may apply, so the general format and codes of that message are adapted here to fit the ICMP packet layout for IPv6 [2].

Nodes should send Binding Acknowledgement messages addressed to the mobile node originating the Binding Update, and if necessary use a routing header (routing type 0) containing the care-of address given in the Binding Update.

The acknowledgement message contains the necessary codes to inform the mobile node about the status of its binding. Additionally, the home agent MAY shorten the lifetime to be smaller than indicated in the original binding update. When the lifetime of the reply is greater than what was contained in the binding update, the excess time MUST be ignored. When the lifetime of the reply is smaller than the original request, another binding update SHOULD be sent before the lifetime expires.

If the mobile node is using a care-of address offered by a local router, the acknowledgement from the home agent will be sent to that care-of address and (presumably) relayed to the mobile node. Routers offering care-of addresses match incoming acknowledgements with previous routing entries by using the Identification field supplied with the binding acknowledgment.

The ICMP packet is organized as follows:

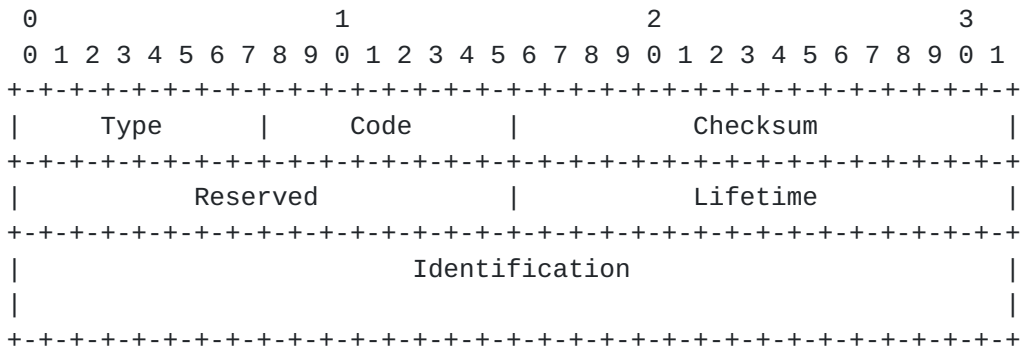

Type (provisionally) 192

Code One of the following codes:

0 service will be provided

1 service will be provided; simultaneous mobility bindings unsupported

Service denied by the owner of the care-of address:

16 reason unspecified
17 administratively prohibited
18 insufficient resources
19 mobile node failed authentication
20 home agent failed authentication
21 requested lifetime too long
22 home agent unreachable (ICMP error)
23 poorly formed binding update
24 poorly formed binding acknowledgement

Service denied by the home agent:

32 reason unspecified
33 administratively prohibited
34 insufficient resources
35 mobile node failed authentication
36 agent at care-of address failed authentication
37 identification mismatch
38 poorly formed binding update
39 too many simultaneous mobility bindings

Lifetime — The seconds remaining before the binding is considered expired. A value of zero confirms a request for removal of a binding. A value of all ones indicates infinity.

Identification — The acknowledgment identification is derived from the binding update message, for use by the mobile node in matching the acknowledgment with an outstanding update.

6. Delivering Packets to a Mobile Node

By default, the routing infrastructure of the Internet will route packets for a mobile node to its home network; this is true of any hierarchical routing and addressing scheme, whether provider-based or geographical. Since the mobile node's location is known on the home network (namely, by the home agent), packets can be addressed to the mobile node and intercepted by the home agent without the sender knowing that the node is mobile, and without requiring any special routing support for mobile nodes anywhere else in the Internet.

Placing the registry of a mobile node's current location at the home network also has the benefit of allowing each organization owning a home network to manage the home agent for the mobile nodes assigned to the organization's own network.

Correspondent nodes that have received a binding update for a mobile node, can send packets directly to that mobile node's current care-of address. There is already a routing header defined within the current IPv6 specification which is well-suited for this purpose, the routing header (routing type 0). To use the routing header for delivery of packets to a mobile node, a correspondent host just specifies the care-of address as the intermediate routing point and the mobile node as the (final) destination. When the packet arrives at the care-of address, normal processing of the routing header will ensure delivery to the mobile node.

The IPv6 routing header avoids the unfortunate semantics of the IPv4 loose source routing option which made it unsuitable for use with IPv4 mobility. In particular, it is fortunate that IPv6 routing headers do not carry the semantics which require reversal of source routes. Since the reversed source route will not be used by the mobile node, no additional security risks are introduced by using routing headers to deliver packets via the care-of address.

There is only one possible advantage afforded by the use of encapsulation, compared to the use of the existing routing header defined for IPv6. That only occurs when a mobile node uses a care-of address associated with a nearby router. If a mobile node has a link to a router over a low speed wireless link, and the router receives encapsulated packets for the mobile node, the encapsulation is stripped away before final delivery is made to the mobile node. In that case, fewer bytes are transmitted over the low-speed link, than would be the case for a normally processed routing header specifying the address of the nearby router (see also section 11.1).

Home agents are often unable to use routing headers to deliver packets to the mobile node, because they can't modify the packet and add to it in flight; therefore, we specify that they must always use encapsulation [8] for this purpose(section 9). It is unknown at this time whether there is sufficient reason to allow the use of alternative encapsulation protocols other than IPv6-within-IPv6, as is done in the mobile-IP specification for IPv4.

If a packet to the mobile node is encapsulated, it uses the care-of address as the destination address in the outer IPv6 header. Then, when the the encapsulated packet arrives at the care-of address, the encapsulation is stripped away and the packet delivered (if possible) to the mobile node. Of course, if the mobile node is

itself receiving packets addressed to the care-of address, the delivery path is trivial. In that case, however, it is more likely that the packet would have been delivered using the care-of address and a routing header.

6.1. Smooth Handoffs

As the mobile node moves from one place to another, in the case of wireless communications with the existing local area networks, it is probable that the mobile node will often reside within range of multiple wireless network points of attachment. If the mobile node obtains a new care-of address while it is still within range for delivery of packets to its old care-of address, then it is reasonable to expect that movement from one care-of address to the next can occur without dropping any packets.

It is likely that, when a mobile node obtains a new care-of address from an address allocation authority, it would explicitly deallocate the previous care-of address. For smooth handoffs, we specify that the mobile client must still accept packets at both addresses for a short time after configuring its newly allocated IPv6 address. If the previous address were allocated by a stateful address server, then the mobile client must not release the address immediately upon acquisition of a new care-of address, and the stateful address server must allow mobile clients to acquire multiple addresses.

7. Foreign Agents

In the IPv4 mobility protocol, packets for a mobile node are tunneled to the mobile node's current care-of address, for delivery to the mobile node. The care-of address must be an address associated with a router neighboring the mobile node, or the network being visited by the mobile node, so that the normal routing of the Internet will deliver the packet to the appropriate network. The care-of address may either be the address of a foreign agent in that network, or may be a temporary local address obtained by means such as DHCP [3].

One reason for favoring the use of a foreign agent in IPv4 is the preservation the limited IPv4 address space. To require each mobile node to acquire its own temporary local address within the network it is visiting would force possibly large portions of the address space to be left available for such dynamic allocation. Any network willing to have mobile nodes visit would need to leave a pool of available addresses, and the number of visiting mobile nodes would be limited to the size of that pool. The address space size is less a concern in IPv6, and so it is feasible to allow each mobile node to

obtain a new care-of address each time it enters a new area of mobile services.

In this section, we outline the advantages afforded by the use of care-of addresses associated with routers nearby the mobile node, and modifications to the previously outlined methods which are made necessary when a mobile node wishes to use such care-of addresses.

Many other operations, related to registration of the mobile node in a new service area, are likely to become important as mobile nodes become more prevalent. For instance, it may be required to:

- authenticate the identity of mobile clients
- charge for connection services
- produce or share a session key for use by new mobile clients (say, for encryption)
- negotiate a compression algorithm
- manage the resources of router's communications devices

These considerations are mostly outside the scope of this document. In all cases, though, we suggest that the need for performing such protocol actions, to satisfy additional requirements, must be indicated in extensions to the basic service advertisement protocol; this may depend on the form of neighbor discovery finally adopted by the IPv6 working group. The actual protocol actions performed in response to the extensions would be carried out at layers above IPv6 (e.g., UDP).

For instance, if a router wishes to authenticate the identity of its prospective clients, it should use an extension to the service advertisement message to indicate this. Then, the mobile host will satisfy the router's requirement by responding with the appropriate protocol operations (which are undefined here). Note that if the router can authenticate any binding update issued by the mobile node during operation, that authentication is likely to be good enough to also authenticate the identity of the mobile node.

If the mobile node is to be billed for services, then surely authentication will be needed. In addition, if billing is managed by the routers, then the billing agent should append a billing extension to its basic advertisement, so that the mobile node can select among competing services if they are available, and so that the mobile node can supply the information needed by the router to effect the financial transactions. Similarly, encryption and/or compression

services might be advertised by extensions to the basic service advertisements.

A nearby router can provide services to mobile nodes without requiring any registration transactions. It is also suggested that service advertisement messages issued by the router contain an indication whether additional resources are currently available, so that the mobile node does not have to waste time sending packets through an agent which cannot forward them. Otherwise, if a mobile node receives a broadcast or multicast advertisement for service that the router is not really equipped to provide, then the router might have to reject attempts by that mobile node to transmit data through its interfaces. The router could do that by sending an ICMP 'Resource Unavailable' message back to the mobile node.

7.1. Sending a Binding Update via a Foreign Agent

When a mobile node is attached to the network via a foreign agent, it is no longer possible to send binding updates directly to its home agent. The mobile node can still send the update to its home agent, through the foreign agent, by using a routing header and inserting the foreign agent as an intermediate router. In this case, however, the foreign agent has to remember that the mobile node, using its home address, is its neighbor. Otherwise, the foreign agent would not know how to deliver packets after decapsulation, or packets sent to the home address using the foreign agent's care-of address in a routing header. The foreign agent will already know the mobile node by the mobile node's local-use only address, and by the mobile node's globally routable or site-local address which was obtained from the Neighbor Discovery protocol, but those two addresses are logically different than the mobile node's home address.

We can insure this effect by by requiring the routing header used with the binding update to have a new routing type (routing type 1). With this routing type, the foreign agent will be required to transmit the expected binding acknowledgement back to the mobile node when it is received from the home agent.

Another strategy for ensuring that the foreign agent will associate the mobile node by its home address is to require the mobile node to deliver its binding update to the home agent through the foreign agent using encapsulation. The encapsulated binding update could itself have another binding update destination option, or some more economical means can be devised.

7.2. Smooth Handoffs between Routers

Given the ability to securely notify other IPv6 nodes of its current location, a mobile node can also facilitate a smooth transition between service from one router to another one simply by sending a binding update to its previous router. This binding update must be acknowledged by the previous router. Then, the router (acting as a cache agent) can forward packets to the new router for direct delivery to the mobile node. If a packet arrives at the previous router for the mobile node, the previous router encapsulates the packet and delivers it to the new router. If the previous router does not receive such a binding update, and the packet cannot be delivered to the mobile node, it encapsulates the packet for delivery to the home agent, using its own address as the source address in the outer header and the address of the mobile node as the destination address.

It is suggested that mobile nodes continue to receive packets at previous care-of addresses for as long as physically possible. This can be done with care-of addresses obtained by automatic or dynamic address configuration as well as those associated with routers. If wireless communications are continuously available in overlapping service areas, then mobile computers using such devices could thus reasonably expect to move between routers without dropping any packets.

8. After Decapsulating

After the router strips the encapsulation, it is no longer possible for the mobile node to determine whether the packet was encapsulated by the home agent, or by a correspondent node. If the packet was encapsulated by the home agent, then the correspondent node must have been unaware of the current location of the mobile node, and the mobile node should be advised to send its correspondent a binding update. This advice can be obtained in several ways, but perhaps the cleanest technique is for the router to re-encapsulate the packet (using essentially the same encapsulation protocol) before transmitting the packet to the mobile node. The extra transmission time from the router to the mobile node due to the encapsulation should not be an issue since this action will occur only rarely, compared to the flow of normal data packets. Re-encapsulation will be even rarer if the mobile node does a good job of including binding updates in the data packets it sends to its correspondent nodes. Note that, for implementation, no copying need occur for this operation; the encapsulating router may just replace the source address in the encapsulating header, and make other minor adjustments like resetting the hop limit, or the flow label.

A mobile node receiving such an encapsulated packet will send the correspondent node a binding update. Thus, the mobile node can quickly inform the correspondent node of its current care-of address, and the home agent will be relied upon for only a small percentage of the overall data traffic destined for the mobile node. Routers would only rarely have to perform this reencapsulation for the purposes of transmitting such advice. Another important advantage of this scheme is that the mobile node can send binding updates to its correspondent hosts without requiring any acknowledgement. Occasionally, the binding update might be lost, but in that case the mobile node will retransmit after a short timeout when it determines (say, with the help of its router) that the first attempt probably failed. Since the mobile host sends binding updates to its active correspondents soon after entering the service area of a new router, any delays due to stale or nonexistent location caches at correspondent nodes will be short-lived.

Just as a mobile node should avoid sending a rush of binding updates to its correspondent hosts when it migrates to a new care-of address, it would be advantageous if any nearby router could avoid sending a rush of advisory encapsulations to any of its newly acquired mobile clients.

9. Home Agent Considerations

When the home agent receives a packet for the mobile node, it encapsulates the packet and delivers it to the mobile node. Methods by which the home agent could insert a routing header, or modify the destination address of the mobile node, may be unavailable because of the expected prevalence and operation of IPv6 authentication mechanisms [1].

If the home agent receives such an encapsulated packet for the mobile node, it decapsulates and re-encapsulates the packet (some optimization may be available here) for delivery to the mobile node's new care-of address. In this situation, the home agent must check that it is not trying to deliver the packet back to the same care-of address from which it came; otherwise, routing loops might develop. If the home agent determines that it does not have a valid binding for the mobile node, it may deliver the unencapsulated packet onto the home network, and should discontinue any proxy ARP operations it may be performing for the mobile node.

It is useful to be able to send a packet to a mobile node's home agent without explicitly knowing the home agent's address. For example, a mobile node must communicate with its home agent to send it a binding update; but 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 network. It thus may not always be possible or convenient for a mobile node to know the exact address of its own home agent.

In IPv4, one method for accomplishing this is for the mobile node to use the directed broadcast address for its home IP subnet. When the packet reaches the nearest router on the home network, a copy will be broadcast onto the local subnet, thus reaching the home agent, although all other nodes on the home network will also receive a copy of the packet and must ignore it. Then, any home agent on the home network which chooses to respond will inform the mobile node of its address, and the mobile node can subsequently perform the IPv4 registration procedure with the newly discovered home agent address.

In the current IPv6 specification [5], no directed broadcast technique is available. The anycast addresses proposed in IPv6 provide a similar functionality, but are restricted to addressing only the nearest of those routers at the boundary of those nodes identified by a common routing prefix. The home agent, though, may not be the nearest boundary router or may not be a boundary router at all.

Since all routers on the home network are assumed to process binding updates, we can be assured at least that any binding update sent to the anycast address will be processed correctly. Other routers on the home network must be instructed to forward packets to the current router which is serving as the mobile node's home agent. This can be done using the same proxy mechanisms already made available in Neighbor Discovery. The current home agent multicasts the equivalent of a Proxy ARP onto the home network, and subsequently the other routers on the home network will forward packets destined to the mobile node to the particular router which is serving as the home agent for that mobile node.

10. Compatibility with ICMP

When sending a packet to a mobile node, it is important to correctly return to the original sender any ICMP error messages generated by this packet. Since in most cases such packets use a routing header containing the care-of address, this is usually not a problem.

However, when a packet encapsulated at the home agent encounters such an error condition, returning the ICMP error message to mobile nodes away from home is more complicated: the ICMP error message should travel back to the original sender along the same path as

the original packet, and must be in a form that makes sense to the original sender when it gets there.

For example, if the original sender did not know that the mobile node is mobile, the original packet would have been routed to the mobile node's home agent, which then should have tunneled the packet to the mobile node's care-of address. If an ICMP error message were generated along the path between the home agent and the care-of address, the ICMP error message should have been returned first to the home agent, so that it could process the message and possibly attempt to recover from the error. If appropriate for the type of error, the ICMP error message should then be forwarded back to the original sender so that it may also process the error. Since the home agent added the tunneling to the original packet, it should remove this from the copy of the returned packet in the ICMP error message before returning it to the original sender.

In IPv4, this handling of returned ICMP error messages was complicated by the definition of the ICMP protocol. Originally, ICMP was specified to return only the first 8 data octets of the packet in error, and even though this has been changed in the current ``Host Requirements'' RFC to specify the return of AT LEAST the first 8 octets, many implementations still return only 8 octets. The problem is that no matter how tunneling is encoded in the packet to the mobile node, returning only 8 data octets form the packet cannot return both the tunneling information and a portion of the original data of the packet.

ICMP for IP version 6 has been specified to return as much of the original packet as will fit in the ICMP error message without the ICMP packet exceeding 576 octets [2]. This size should be sufficient for correctly returning ICMP error messages backwards along the tunnel, as long as the original sender does not expect to get this full size returned. Since the tunneling information is removed from the original packet by the home agent, the length of the ICMP packet will in this case be less than 576 octets, and correspondingly less of the original packet will be returned in the ICMP error message.

11. Summary

In this protocol document, we have proposed a departure from the mobile-IP protocol for IPv4 in several important ways:

- We propose the use of routing headers instead of encapsulation for common traffic destined to a mobile node, since the problems with loose source routing in IPv4 are no longer present in IPv6

- We build upon the commonality between IPv4 registration and route optimization protocols to permit a cleaner and smaller mechanism for accomplishing the same things

- We emphasize the need for distributing bindings to the entities that need them, in order to reduce drastically the role of the home agent in handling traffic destined for the mobile node

- We build on the expected capability for mobile nodes to receive datagrams at several IPv6 addresses, to suggest the reduced need for external care-of addresses in some common circumstances.

- We specify that all IPv6 routers and nodes accept binding updates, and thus that all IPv6 routers be able to operate as home agents, affording major simplifications and optimization at reasonable implementation cost.

11.1. Open issues

The common use of binding updates as Destination Options is architecturally very clean, but the IPv4 registration request has a more extensible mechanism for adding future functionality via Mobility Extensions. If the need for additional flexibility becomes important soon enough to affect this specification, then we will suggest a return to the IPv4 registration requests and replies for delivering binding updates to the home agents.

When a mobile node migrates away from an area in which it had its own care-of address, there won't be any way for "straggling packets" to be redirected back to the home agent, or to the new care-of address for the mobile nodes. This probably isn't any worse than IPv4 traffic delivery anyway, as far as we know now. Nevertheless, mobile nodes can alleviate this problem if the areas of service overlap sufficiently, and they continue to receive packets at their previous care-of address. How important is it to save such straggling packets? Should we emphasize the need to overlap service areas, or build additional protocol requirements to minimize the problems resulting from non-overlapped service?

Should alternative encapsulation techniques be defined for use with these protocols? Should a minimal encapsulation be defined and specified as the default? Perhaps this would be better taken care of by defining something like TCP header compression over the link from the router to the mobile node.

Another alternative would be to provide another type of routing header (routing type == 2, say) which would allow an intermediate

node to delete itself from the list instead of just rearranging the list. This would completely eliminate the need for encapsulation for normal datagrams from correspondent host to mobile node.

In the IPv4 route optimization proposal, a mechanism is outlined whereby a session key can be established between foreign agents and mobile clients, without requiring any pre-established security relationship between them. It could very well be the case that a similar mechanism should be defined for 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.

References

[1] R. Atkinson. IPv6 Authentication Header. draft-atkinson-ipng-auth-02.txt, work in progress, June 1995.

[2] A. Conta and S. Deering. ICMP for the Internet Protocol Version 6 (IPv6). draft-ietf-ipngwg-icmp-02.txt -- work in progress, June 1995.

[3] R. Droms. Dynamic Host Configuration Protocol. RFC 1541, October 1993.

[4] IETF Mobile-IP Working Group. IPv4 Mobility Support. ietf-draft-mobileip-protocol-10.txt -- work in progress, May 1995.

[5] Bob Hinden. Internet Protocol, Version 6 (IPv6) Specification. Internet draft -- work in progress, October 1994.

[6] David B. Johnson and Charles E. Perkins. Route Optimization in Mobile-IP. Internet Draft -- work in progress, January 1995.

[7] T. Narten, E. Nordmark, and W. Simpson. IPv6 Neighbor Discovery. draft-ietf-ipngwg-discovery-00.txt -- work in progress, June 1995.

[8] C. Perkins. IPv6 Mobility Support. draft-perkins-mobility-ipv6-00.txt -- work in progress, May 1995.

[9] Bill Simpson. IPv6 Neighbor Discover -- ICMP Message Formats. draft-simpson-ipv6-discov-formats-01.txt -- work in progress, November 1994.

Authors' Addresses

Charles Perkins
Room J1-A25
T. J. Watson Research Center
IBM Corporation
30 Saw Mill River Rd.
Hawthorne, NY 10532

Work: +1 914 789-7350
Fax: +1 914 784-7007
E-mail: perk@watson.ibm.com

David B. Johnson
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3891

Work: +1 412 268-7399
Fax: +1 412 268-5576
E-mail: dbj@cs.cmu.edu