

[\[Docs\]](#) [\[txt|pdf\]](#) [\[draft-ietf-mobile...\]](#) [\[Diff1\]](#) [\[Diff2\]](#) [\[IPR\]](#) [\[Errata\]](#)

Obsoleted by: [3344](#)

PROPOSED STANDARD

[Errata Exist](#)

Network Working Group  
Request for Comments: 3220  
Obsoletes: [2002](#)  
Category: Standards Track

C. Perkins, Ed.  
Nokia Research Center  
January 2002

## IP Mobility Support for IPv4

### Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

### Abstract

This document specifies protocol enhancements that allow transparent routing of IP datagrams to mobile nodes in the Internet. Each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. While situated away from its home, a mobile node is also associated with a care-of address, which provides information about its current point of attachment to the Internet. The protocol provides for registering the care-of address with a home agent. The home agent sends datagrams destined for the mobile node through a tunnel to the care-of address. After arriving at the end of the tunnel, each datagram is then delivered to the mobile node.

### Contents

1. Introduction	3
<a href="#">1.1.</a> Protocol Requirements	<a href="#">4</a>
<a href="#">1.2.</a> Goals	<a href="#">4</a>
<a href="#">1.3.</a> Assumptions	<a href="#">5</a>
<a href="#">1.4.</a> Applicability	<a href="#">5</a>
<a href="#">1.5.</a> New Architectural Entities	<a href="#">5</a>
<a href="#">1.6.</a> Terminology	<a href="#">6</a>
<a href="#">1.7.</a> Protocol Overview	<a href="#">6</a>
<a href="#">1.8.</a> Message Format and Protocol Extensibility	<a href="#">13</a>
1.9. Type-Length-Value Extension Format for Mobile IP	
Extensions	<a href="#">15</a>
<a href="#">1.10.</a> Long Extension Format	<a href="#">16</a>

1.11.	Short Extension Format . . . . .	16
2.	Agent Discovery . . . . .	17
2.1.	Agent Advertisement . . . . .	18
2.1.1.	Mobility Agent Advertisement Extension . . . . .	20
2.1.2.	Prefix-Lengths Extension . . . . .	22
2.1.3.	One-byte Padding Extension . . . . .	22
2.2.	Agent Solicitation . . . . .	23
2.3.	Foreign Agent and Home Agent Considerations . . . . .	23
2.3.1.	Advertised Router Addresses . . . . .	24
2.3.2.	Sequence Numbers and Rollover Handling . . . . .	24
2.4.	Mobile Node Considerations . . . . .	25
2.4.1.	Registration Required . . . . .	26
2.4.2.	Move Detection . . . . .	26
2.4.3.	Returning Home . . . . .	27
2.4.4.	Sequence Numbers and Rollover Handling . . . . .	28
3.	Registration . . . . .	28
3.1.	Registration Overview . . . . .	29
3.2.	Authentication . . . . .	30
3.3.	Registration Request . . . . .	30
3.4.	Registration Reply . . . . .	33
3.5.	Registration Extensions . . . . .	36
3.5.1.	Computing Authentication Extension Values . . . . .	36
3.5.2.	Mobile-Home Authentication Extension . . . . .	37
3.5.3.	Mobile-Foreign Authentication Extension . . . . .	37
3.5.4.	Foreign-Home Authentication Extension . . . . .	38
3.6.	Mobile Node Considerations . . . . .	38
3.6.1.	Sending Registration Requests . . . . .	40
3.6.2.	Receiving Registration Replies . . . . .	43
3.6.3.	Registration Retransmission . . . . .	46
3.7.	Foreign Agent Considerations . . . . .	46
3.7.1.	Configuration and Registration Tables . . . . .	47
3.7.2.	Receiving Registration Requests . . . . .	48
3.7.3.	Receiving Registration Replies . . . . .	51
3.8.	Home Agent Considerations . . . . .	53
3.8.1.	Configuration and Registration Tables . . . . .	54
3.8.2.	Receiving Registration Requests . . . . .	55
3.8.3.	Sending Registration Replies . . . . .	58
4.	Routing Considerations . . . . .	61
4.1.	Encapsulation Types . . . . .	61
4.2.	Unicast Datagram Routing . . . . .	61
4.2.1.	Mobile Node Considerations . . . . .	61
4.2.2.	Foreign Agent Considerations . . . . .	62
4.2.3.	Home Agent Considerations . . . . .	63
4.3.	Broadcast Datagrams . . . . .	65
4.4.	Multicast Datagram Routing . . . . .	65
4.5.	Mobile Routers . . . . .	66
4.6.	ARP, Proxy ARP, and Gratuitous ARP . . . . .	68
5.	Security Considerations . . . . .	72

5.1.	Message Authentication Codes . . . . .	72
5.2.	Areas of Security Concern in this Protocol . . . . .	72
5.3.	Key Management . . . . .	73
5.4.	Picking Good Random Numbers . . . . .	73
5.5.	Privacy . . . . .	73
5.6.	Ingress Filtering . . . . .	74
5.7.	Replay Protection for Registration Requests . . . . .	74
5.7.1.	Replay Protection using Timestamps . . . . .	74
5.7.2.	Replay Protection using Nonces . . . . .	76
6.	IANA Considerations . . . . .	76
6.1.	Mobile IP Message Types . . . . .	77
6.2.	Extensions to <a href="#">RFC 1256</a> Router Advertisement . . . . .	77
6.3.	Extensions to Mobile IP Registration Messages . . . . .	78
6.4.	Code Values for Mobile IP Registration Reply Messages . . . . .	78
7.	Acknowledgments . . . . .	79
A.	Patent Issues . . . . .	81
B.	Link-Layer Considerations . . . . .	81
C.	TCP Considerations . . . . .	82
C.1.	TCP Timers . . . . .	82
C.2.	TCP Congestion Management . . . . .	82
D.	Example Scenarios . . . . .	83
D.1.	Registering with a Foreign Agent Care-of Address . . . . .	83
D.2.	Registering with a Co-Located Care-of Address . . . . .	83
D.3.	Deregistration . . . . .	84
E.	Applicability of Prefix-Lengths Extension . . . . .	85
F.	Interoperability Considerations . . . . .	85
G.	Changes since <a href="#">RFC 2002</a> . . . . .	86
G.1.	Major Changes . . . . .	86
G.2.	Minor Changes . . . . .	88
G.3.	Changes since revision 04 of RFC2002bis . . . . .	90
H.	Example Messages . . . . .	91
H.1.	Example ICMP Agent Advertisement Message Format . . . . .	91
H.2.	Example Registration Request Message Format . . . . .	92
H.3.	Example Registration Reply Message Format . . . . .	93
References	. . . . .	93
Authors' Addresses	. . . . .	97
Full Copyright Statement	. . . . .	98

## 1. Introduction

IP version 4 assumes that a node's IP address uniquely identifies the node's point of attachment to the Internet. Therefore, a node must be located on the network indicated by its IP address in order to receive datagrams destined to it; otherwise, datagrams destined to the node would be undeliverable. For a node to change its point of attachment without losing its ability to communicate, currently one of the two following mechanisms must typically be employed:

- a) the node must change its IP address whenever it changes its point of attachment, or
- b) host-specific routes must be propagated throughout much of the Internet routing fabric.

Both of these alternatives are often unacceptable. The first makes it impossible for a node to maintain transport and higher-layer connections when the node changes location. The second has obvious and severe scaling problems, especially relevant considering the explosive growth in sales of notebook (mobile) computers.

A new, scalable, mechanism is required for accommodating node mobility within the Internet. This document defines such a mechanism, which enables nodes to change their point of attachment to the Internet without changing their IP address.

Changes between this revised specification for Mobile IP and the original specifications (see [[33](#), [32](#), [34](#), [43](#), [8](#)]) are detailed in the appendix section G.

### 1.1. Protocol Requirements

A mobile node must be able to communicate with other nodes after changing its link-layer point of attachment to the Internet, yet without changing its IP address.

A mobile node must be able to communicate with other nodes that do not implement these mobility functions. No protocol enhancements are required in hosts or routers that are not acting as any of the new architectural entities introduced in [Section 1.5](#).

All messages used to update another node as to the location of a mobile node must be authenticated in order to protect against remote redirection attacks.

### 1.2. Goals

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

### 1.3. Assumptions

The protocols defined in this document place no additional constraints on the assignment of IP addresses. That is, a mobile node can be assigned an IP address by the organization that owns the machine.

This protocol assumes that mobile nodes will generally not change their point of attachment to the Internet more frequently than once per second.

This protocol assumes that IP unicast datagrams are routed based on the destination address in the datagram header (and not, for example, by source address).

### 1.4. Applicability

Mobile IP is intended to enable nodes to move from one IP subnet to another. It is just as suitable for mobility across homogeneous media as it is for mobility across heterogeneous media. That is, Mobile IP facilitates node movement from one Ethernet segment to another as well as it accommodates node movement from an Ethernet segment to a wireless LAN, as long as the mobile node's IP address remains the same after such a movement.

One can think of Mobile IP as solving the "macro" mobility management problem. It is less well suited for more "micro" mobility management applications -- for example, handoff amongst wireless transceivers, each of which covers only a very small geographic area. As long as node movement does not occur between points of attachment on different IP subnets, link-layer mechanisms for mobility (i.e., link-layer handoff) may offer faster convergence and far less overhead than Mobile IP.

### 1.5. New Architectural Entities

Mobile IP introduces the following new functional entities:

#### Mobile Node

A host or router that changes its point of attachment from one network or subnetwork to another. A mobile node may change its location without changing its IP address; it may continue to communicate with other Internet nodes at any location using its (constant) IP address, assuming link-layer connectivity to a point of attachment is available.

### Home Agent

A router on a mobile node's home network which tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information for the mobile node.

### Foreign Agent

A router on a mobile node's visited network which provides routing services to the mobile node while registered. The foreign agent detunnels and delivers datagrams to the mobile node that were tunneled by the mobile node's home agent. For datagrams sent by a mobile node, the foreign agent may serve as a default router for registered mobile nodes.

A mobile node is given a long-term IP address on a home network. This home address is administered in the same way as a "permanent" IP address is provided to a stationary host. When away from its home network, a "care-of address" is associated with the mobile node and reflects the mobile node's current point of attachment. The mobile node uses its home address as the source address of all IP datagrams that it sends, except where otherwise described in this document for datagrams sent for certain mobility management functions (e.g., as in [Section 3.6.1.1](#)).

## 1.6. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [4].

In addition, this document frequently uses the following terms:

#### Authorization-enabling extension

An authentication which makes a (registration) message acceptable to the ultimate recipient of the registration message. An authorization-enabling extension MUST contain an SPI.

In this document, all uses of authorization-enabling extension refer to authentication extensions that enable the Registration Request message to be acceptable to the home agent. Using additional protocol structures specified outside of this document, it may be possible for the mobile node to provide authentication of its registration to the

home agent, by way of another authenticating entity within the network that is acceptable to the home agent (for example, see [RFC 2794](#) [6]).

#### Agent Advertisement

An advertisement message constructed by attaching a special Extension to a router advertisement [[10](#)] message.

#### Authentication

The process of verifying (using cryptographic techniques, for all applications in this specification) the identity of the originator of a message.

#### Care-of Address

The termination point of a tunnel toward a mobile node, for datagrams forwarded to the mobile node while it is away from home. The protocol can use two different types of care-of address: a "foreign agent care-of address" is an address of a foreign agent with which the mobile node is registered, and a "co-located care-of address" is an externally obtained local address which the mobile node has associated with one of its own network interfaces.

#### Correspondent Node

A peer with which a mobile node is communicating. A correspondent node may be either mobile or stationary.

#### Foreign Network

Any network other than the mobile node's Home Network.

#### Gratuitous ARP

An ARP packet sent by a node in order to spontaneously cause other nodes to update an entry in their ARP cache [[45](#)]. See [section 4.6](#).

#### Home Address

An IP address that is assigned for an extended period of time to a mobile node. It remains unchanged regardless of where the node is attached to the Internet.

#### Home Network

A network, possibly virtual, having a network prefix matching that of a mobile node's home address. Note that standard IP routing mechanisms will deliver datagrams destined to a mobile node's Home Address to the mobile node's Home Network.

#### Link

A facility or medium over which nodes can communicate at the link layer. A link underlies the network layer.

#### Link-Layer Address

The address used to identify an endpoint of some communication over a physical link. Typically, the Link-Layer address is an interface's Media Access Control (MAC) address.

#### Mobility Agent

Either a home agent or a foreign agent.

#### Mobility Binding

The association of a home address with a care-of address, along with the remaining lifetime of that association.

#### Mobility Security Association

A collection of security contexts, between a pair of nodes, which may be applied to Mobile IP protocol messages exchanged between them. Each context indicates an authentication algorithm and mode ([Section 5.1](#)), a secret (a shared key, or appropriate public/private key pair), and a style of replay protection in use ([Section 5.7](#)).

#### Node

A host or a router.

#### Nonce

A randomly chosen value, different from previous choices, inserted in a message to protect against replays.



#### Security Parameter Index (SPI)

An index identifying a security context between a pair of nodes among the contexts available in the Mobility Security Association. SPI values 0 through 255 are reserved and MUST NOT be used in any Mobility Security Association.

#### Tunnel

The path followed by a datagram while it is encapsulated. The model is that, while it is encapsulated, a datagram is routed to a knowledgeable decapsulating agent, which decapsulates the datagram and then correctly delivers it to its ultimate destination.

#### Virtual Network

A network with no physical instantiation beyond a router (with a physical network interface on another network). The router (e.g., a home agent) generally advertises reachability to the virtual network using conventional routing protocols.

#### Visited Network

A network other than a mobile node's Home Network, to which the mobile node is currently connected.

#### Visitor List

The list of mobile nodes visiting a foreign agent.

### 1.7. Protocol Overview

The following support services are defined for Mobile IP:

#### Agent Discovery

Home agents and foreign agents may advertise their availability on each link for which they provide service. A newly arrived mobile node can send a solicitation on the link to learn if any prospective agents are present.

#### Registration

When the mobile node is away from home, it registers its care-of address with its home agent. Depending on its method of attachment, the mobile node will register either

directly with its home agent, or through a foreign agent which forwards the registration to the home agent.

silently discard

The implementation discards the datagram without further processing, and without indicating an error to the sender. The implementation SHOULD provide the capability of logging the error, including the contents of the discarded datagram, and SHOULD record the event in a statistics counter.

The following steps provide a rough outline of operation of the Mobile IP protocol:

- Mobility agents (i.e., foreign agents and home agents) advertise their presence via Agent Advertisement messages ([Section 2](#)). A mobile node may optionally solicit an Agent Advertisement message from any locally attached mobility agents through an Agent Solicitation message.
- A mobile node receives these Agent Advertisements and determines whether it is on its home network or a foreign network.
- When the mobile node detects that it is located on its home network, it operates without mobility services. If returning to its home network from being registered elsewhere, the mobile node deregisters with its home agent, through exchange of a Registration Request and Registration Reply message with it.
- When a mobile node detects that it has moved to a foreign network, it obtains a care-of address on the foreign network. The care-of address can either be determined from a foreign agent's advertisements (a foreign agent care-of address), or by some external assignment mechanism such as DHCP [[13](#)] (a co-located care-of address).
- The mobile node operating away from home then registers its new care-of address with its home agent through exchange of a Registration Request and Registration Reply message with it, possibly via a foreign agent ([Section 3](#)).
- Datagrams sent to the mobile node's home address are intercepted by its home agent, tunneled by the home agent to the mobile node's care-of address, received at the tunnel endpoint (either at a foreign agent or at the mobile node itself), and finally delivered to the mobile node ([Section 4.2.3](#)).

- In the reverse direction, datagrams sent by the mobile node are generally delivered to their destination using standard IP routing mechanisms, not necessarily passing through the home agent.

When away from home, Mobile IP uses protocol tunneling to hide a mobile node's home address from intervening routers between its home network and its current location. The tunnel terminates at the mobile node's care-of address. The care-of address must be an address to which datagrams can be delivered via conventional IP routing. At the care-of address, the original datagram is removed from the tunnel and delivered to the mobile node.

Mobile IP provides two alternative modes for the acquisition of a care-of address:

- a) A "foreign agent care-of address" is a care-of address provided by a foreign agent through its Agent Advertisement messages. In this case, the care-of address is an IP address of the foreign agent. In this mode, the foreign agent is the endpoint of the tunnel and, upon receiving tunneled datagrams, decapsulates them and delivers the inner datagram to the mobile node. This mode of acquisition is preferred because it allows many mobile nodes to share the same care-of address and therefore does not place unnecessary demands on the already limited IPv4 address space.
- b) A "co-located care-of address" is a care-of address acquired by the mobile node as a local IP address through some external means, which the mobile node then associates with one of its own network interfaces. The address may be dynamically acquired as a temporary address by the mobile node such as through DHCP [13], or may be owned by the mobile node as a long-term address for its use only while visiting some foreign network. Specific external methods of acquiring a local IP address for use as a co-located care-of address are beyond the scope of this document. When using a co-located care-of address, the mobile node serves as the endpoint of the tunnel and itself performs decapsulation of the datagrams tunneled to it.

The mode of using a co-located care-of address has the advantage that it allows a mobile node to function without a foreign agent, for example, in networks that have not yet deployed a foreign agent. It does, however, place additional burden on the IPv4 address space because it requires a pool of addresses within the foreign network to

be made available to visiting mobile nodes. It is difficult to efficiently maintain pools of addresses for each subnet that may permit mobile nodes to visit.

It is important to understand the distinction between the care-of address and the foreign agent functions. The care-of address is simply the endpoint of the tunnel. It might indeed be an address of a foreign agent (a foreign agent care-of address), but it might instead be an address temporarily acquired by the mobile node (a co-located care-of address). A foreign agent, on the other hand, is a mobility agent that provides services to mobile nodes. See Sections 3.7 and 4.2.2 for additional details.

For example, figure 1 illustrates the routing of datagrams to and from a mobile node away from home, once the mobile node has registered with its home agent. In figure 1, the mobile node is using a foreign agent care-of address, not a co-located care-of address.

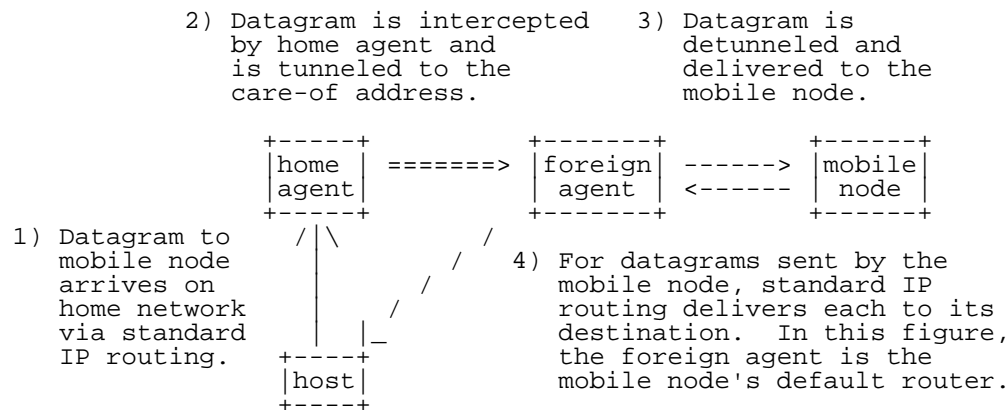


Figure 1: Operation of Mobile IPv4

A home agent MUST be able to attract and intercept datagrams that are destined to the home address of any of its registered mobile nodes. Using the proxy and gratuitous ARP mechanisms described in [Section 4.6](#), this requirement can be satisfied if the home agent has a network interface on the link indicated by the mobile node's home address. Other placements of the home agent relative to the mobile node's home location MAY also be possible using other mechanisms for intercepting datagrams destined to the mobile node's home address. Such placements are beyond the scope of this document.

Similarly, a mobile node and a prospective or current foreign agent MUST be able to exchange datagrams without relying on standard IP routing mechanisms; that is, those mechanisms which make forwarding decisions based upon the network-prefix of the destination address in the IP header. This requirement can be satisfied if the foreign agent and the visiting mobile node have an interface on the same link. In this case, the mobile node and foreign agent simply bypass their normal IP routing mechanism when sending datagrams to each other, addressing the underlying link-layer packets to their respective link-layer addresses. Other placements of the foreign agent relative to the mobile node MAY also be possible using other mechanisms to exchange datagrams between these nodes, but such placements are beyond the scope of this document.

If a mobile node is using a co-located care-of address (as described in (b) above), the mobile node MUST be located on the link identified by the network prefix of this care-of address. Otherwise, datagrams destined to the care-of address would be undeliverable.

### 1.8. Message Format and Protocol Extensibility

Mobile IP defines a set of new control messages, sent with UDP [[37](#)] using well-known port number 434. The following two message types are defined in this document:

- 1 Registration Request
- 3 Registration Reply

Up-to-date values for the message types for Mobile IP control messages are specified in the most recent "Assigned Numbers" [[40](#)].

In addition, for Agent Discovery, Mobile IP makes use of the existing Router Advertisement and Router Solicitation messages defined for ICMP Router Discovery [[10](#)].

Mobile IP defines a general Extension mechanism to allow optional information to be carried by Mobile IP control messages or by ICMP Router Discovery messages. Some extensions have been specified to be encoded in the simple Type-Length-Value format described in [Section 1.9](#).

Extensions allow variable amounts of information to be carried within each datagram. The end of the list of Extensions is indicated by the total length of the IP datagram.

Two separately maintained sets of numbering spaces, from which Extension Type values are allocated, are used in Mobile IP:

- The first set consists of those Extensions which may appear only in Mobile IP control messages (those sent to and from UDP port number 434). In this document, the following Types are defined for Extensions appearing in Mobile IP control messages:
  - 32 Mobile-Home Authentication
  - 33 Mobile-Foreign Authentication
  - 34 Foreign-Home Authentication
- The second set consists of those extensions which may appear only in ICMP Router Discovery messages [[10](#)]. In this document, the following Types are defined for Extensions appearing in ICMP Router Discovery messages:
  - 0 One-byte Padding (encoded with no Length nor Data field)
  - 16 Mobility Agent Advertisement
  - 19 Prefix-Lengths

Each individual Extension is described in detail in a separate section later in this document. Up-to-date values for these Extension Type numbers are specified in the most recent "Assigned Numbers" [[40](#)].

Due to the separation (orthogonality) of these sets, it is conceivable that two Extensions that are defined at a later date could have identical Type values, so long as one of the Extensions may be used only in Mobile IP control messages and the other may be used only in ICMP Router Discovery messages.

The type field in the Mobile IP extension structure can support up to 255 (skippable and not skippable) uniquely identifiable extensions. When an Extension numbered in either of these sets within the range 0 through 127 is encountered but not recognized, the message containing that Extension MUST be silently discarded. When an Extension numbered in the range 128 through 255 is encountered which is not recognized, that particular Extension is ignored, but the rest of the Extensions and message data MUST still be processed. The Length field of the Extension is used to skip the Data field in searching for the next Extension.

Unless additional structure is utilized for the extension types, new developments or additions to Mobile IP might require so many new extensions that the available space for extension types might run out. Two new extension structures are proposed to solve this problem. Certain types of extensions can be aggregated, using

subtypes to identify the precise extension, for example as has been done with the Generic Authentication Keys extensions [35]. In many cases, this may reduce the rate of allocation for new values of the type field.

Since the new extension structures will cause an efficient usage of the extension type space, it is recommended that new Mobile IP extensions follow one of the two new extension formats whenever there may be the possibility to group related extensions together.

The following subsections provide details about three distinct structures for Mobile IP extensions:

- The simple extension format
- The long extension format
- The short extension format

### 1.9. Type-Length-Value Extension Format for Mobile IP Extensions

The Type-Length-Value format illustrated in figure 2 is used for extensions which are specified in this document. Since this simple extension structure does not encourage the most efficient usage of the extension type space, it is recommended that new Mobile IP extensions follow one of the two new extension formats specified in sections 1.10 or 1.11 whenever there may be the possibility to group related extensions together.





Figure 2: Type-Length-Value extension format for Mobile IPv4

Type	Indicates the particular type of Extension.
Length	Indicates the length (in bytes) of the data field within this Extension. The length does NOT include the Type and Length bytes.
Data	The particular data associated with this Extension. This field may be zero or more bytes in length. The format and length of the data field is determined by the type and length fields.

**1.10. Long Extension Format**

This format is applicable for non-skippable extensions which carry information more than 256 bytes.



The Long Extension format requires that the following fields be specified as the first fields of the extension.

Type is the type, which describes a collection of extensions having a common data type.

Sub-Type is a unique number given to each member in the aggregated type.


Length indicates the length (in bytes) of the data field within this Extension. It does NOT include the Type, Length and Sub-Type bytes.

Data is the data associated with the subtype of this extension. This specification does not place any additional structure on the subtype data.

Since the length field is 16 bits wide, a the extension data can exceed 256 bytes in length.

**1.11. Short Extension Format**

This format is compatible with the skippable extensions defined in [section 1.9](#). It is not applicable for extensions which require more than 256 bytes of data; for such extensions, use the format described in [section 1.10](#).



The Short Extension format requires that the following fields be specified as the first fields of the extension:



Type is the type, which describes a collection of extensions having a common data type.

Sub-Type is a unique number given to each member in the aggregated type.

Length 8-bit unsigned integer. Length of the extension, in bytes, excluding the extension Type and the extension Length fields. This field MUST be set to 1 plus the total length of the data field.

Data is the data associated with this extension. This specification does not place any additional structure on the subtype data.

## 2. Agent Discovery

Agent Discovery is the method by which a mobile node determines whether it is currently connected to its home network or to a foreign network, and by which a mobile node can detect when it has moved from one network to another. When connected to a foreign network, the methods specified in this section also allow the mobile node to determine the foreign agent care-of address being offered by each foreign agent on that network.

Mobile IP extends ICMP Router Discovery [10] as its primary mechanism for Agent Discovery. An Agent Advertisement is formed by including a Mobility Agent Advertisement Extension in an ICMP Router Advertisement message (Section 2.1). An Agent Solicitation message is identical to an ICMP Router Solicitation, except that its IP TTL MUST be set to 1 (Section 2.2). This section describes the message formats and procedures by which mobile nodes, foreign agents, and home agents cooperate to realize Agent Discovery.

Agent Advertisement and Agent Solicitation may not be necessary for link layers that already provide this functionality. The method by which mobile nodes establish link-layer connections with prospective agents is outside the scope of this document (but see Appendix B). The procedures described below assume that such link-layer connectivity has already been established.

No authentication is required for Agent Advertisement and Agent Solicitation messages. They MAY be authenticated using the IP Authentication Header [22], which is unrelated to the messages described in this document. Further specification of the way in which Advertisement and Solicitation messages may be authenticated is outside of the scope of this document.

## 2.1. Agent Advertisement

Agent Advertisements are transmitted by a mobility agent to advertise its services on a link. Mobile nodes use these advertisements to determine their current point of attachment to the Internet. An Agent Advertisement is an ICMP Router Advertisement that has been extended to also carry a Mobility Agent Advertisement Extension ([Section 2.1.1](#)) and, optionally, a Prefix-Lengths Extension ([Section 2.1.2](#)), One-byte Padding Extension ([Section 2.1.3](#)), or other Extensions that might be defined in the future.

Within an Agent Advertisement message, ICMP Router Advertisement fields of the message are required to conform to the following additional specifications:

- Link-Layer Fields

- Destination Address

- The link-layer destination address of a unicast Agent Advertisement MUST be the same as the source link-layer address of the Agent Solicitation which prompted the Advertisement.

- IP Fields

- TTL           The TTL for all Agent Advertisements MUST be set to 1.

- Destination Address

- As specified for ICMP Router Discovery [[10](#)], the IP destination address of an multicast Agent Advertisement MUST be either the "all systems on this link" multicast address (224.0.0.1) [[11](#)] or the "limited broadcast" address (255.255.255.255). The subnet-directed broadcast address of the form <prefix>.<-1> cannot be used since mobile nodes will not generally know the prefix of the foreign network. When the Agent Advertisement is unicast to a mobile node, the IP home address of the mobile node SHOULD be used as the Destination Address.

- ICMP Fields

Code       The Code field of the agent advertisement is interpreted as follows:

- 0 The mobility agent handles common traffic -- that is, it acts as a router for IP datagrams not necessarily related to mobile nodes.
- 16 The mobility agent does not route common traffic. However, all foreign agents MUST (minimally) forward to a default router any datagrams received from a registered mobile node ([Section 4.2.2](#)).

Lifetime

The maximum length of time that the Advertisement is considered valid in the absence of further Advertisements.

Router Address(es)

See [Section 2.3.1](#) for a discussion of the addresses that may appear in this portion of the Agent Advertisement.

Num Adrs


The number of Router Addresses advertised in this message. Note that in an Agent Advertisement message, the number of router addresses specified in the ICMP Router Advertisement portion of the message MAY be set to 0. See [Section 2.3.1](#) for details.

If sent periodically, the nominal interval at which Agent Advertisements are sent SHOULD be no longer than 1/3 of the advertisement Lifetime given in the ICMP header. This interval MAY be shorter than 1/3 the advertised Lifetime. This allows a mobile node to miss three successive advertisements before deleting the agent from its list of valid agents. The actual transmission time for each advertisement SHOULD be slightly randomized [[10](#)] in order to avoid synchronization and subsequent collisions with other Agent

Advertisements that may be sent by other agents (or with other Router Advertisements sent by other routers). Note that this field has no relation to the "Registration Lifetime" field within the Mobility Agent Advertisement Extension defined below.

### 2.1.1.1. Mobility Agent Advertisement Extension

The Mobility Agent Advertisement Extension follows the ICMP Router Advertisement fields. It is used to indicate that an ICMP Router Advertisement message is also an Agent Advertisement being sent by a mobility agent. The Mobility Agent Advertisement Extension is defined as follows:



Type 16

Length (6 + 4\*N), where 6 accounts for the number of bytes in the Sequence Number, Registration Lifetime, flags, and reserved fields, and N is the number of care-of addresses advertised.

Sequence Number

The count of Agent Advertisement messages sent since the agent was initialized ([Section 2.3.2](#)).

Registration Lifetime

The longest lifetime (measured in seconds) that this agent is willing to accept in any Registration Request. A value of 0xffff indicates infinity. This field has no relation to the "Lifetime" field within the ICMP Router Advertisement portion of the Agent Advertisement.

R Registration required. Registration with this foreign agent (or another foreign agent on this link) is required even when using a co-located care-of address.

B Busy. The foreign agent will not accept registrations from additional mobile nodes.

H Home agent. This agent offers service as a home agent on the link on which this Agent Advertisement message is sent.

- F Foreign agent. This agent offers service as a foreign agent on the link on which this Agent Advertisement message is sent.
- M Minimal encapsulation. This agent implements receiving tunneled datagrams that use minimal encapsulation [34].
- G GRE encapsulation. This agent implements receiving tunneled datagrams that use GRE encapsulation [16].
- r Sent as zero; ignored on reception. SHOULD NOT be allocated for any other uses.
- T Foreign agent supports reverse tunneling [27].
- reserved Sent as zero; ignored on reception.

#### Care-of Address(es)


The advertised foreign agent care-of address(es) provided by this foreign agent. An Agent Advertisement MUST include at least one care-of address if the 'F' bit is set. The number of care-of addresses present is determined by the Length field in the Extension.

A home agent MUST always be prepared to serve the mobile nodes for which it is the home agent. A foreign agent may at times be too busy to serve additional mobile nodes; even so, it must continue to send Agent Advertisements, so that any mobile nodes already registered with it will know that they have not moved out of range of the foreign agent and that the foreign agent has not failed. A foreign agent may indicate that it is "too busy" to allow new mobile nodes to register with it, by setting the 'B' bit in its Agent Advertisements. An Agent Advertisement message MUST NOT have the 'B' bit set if the 'F' bit is not also set. Furthermore, at least one of the 'F' bit and the 'H' bit MUST be set in any Agent Advertisement message sent.

When a foreign agent wishes to require registration even from those mobile nodes which have acquired a co-located care-of address, it sets the 'R' bit to one. Because this bit applies only to foreign agents, an agent MUST NOT set the 'R' bit to one unless the 'F' bit is also set to one.

### 2.1.2. Prefix-Lengths Extension

The Prefix-Lengths Extension MAY follow the Mobility Agent Advertisement Extension. It is used to indicate the number of bits of network prefix that applies to each Router Address listed in the ICMP Router Advertisement portion of the Agent Advertisement. Note that the prefix lengths given DO NOT apply to care-of address(es) listed in the Mobility Agent Advertisement Extension. The Prefix-Lengths Extension is defined as follows:



Type        19 (Prefix-Lengths Extension)

Length      N, where N is the value (possibly zero) of the Num Addr field in the ICMP Router Advertisement portion of the Agent Advertisement.

Prefix Length(s)

The number of leading bits that define the network number of the corresponding Router Address listed in the ICMP Router Advertisement portion of the message. The prefix length for each Router Address is encoded as a separate byte, in the order that the Router Addresses are listed in the ICMP Router Advertisement portion of the message.

See [Section 2.4.2](#) for information about how the Prefix-Lengths Extension MAY be used by a mobile node when determining whether it has moved. See [Appendix E](#) for implementation details about the use of this Extension.

### 2.1.3. One-byte Padding Extension

Some IP protocol implementations insist upon padding ICMP messages to an even number of bytes. If the ICMP length of an Agent Advertisement is odd, this Extension MAY be included in order to make the ICMP length even. Note that this Extension is NOT intended to be a general-purpose Extension to be included in order to word- or long-align the various fields of the Agent Advertisement. An Agent Advertisement SHOULD NOT include more than one One-byte Padding Extension and if present, this Extension SHOULD be the last Extension in the Agent Advertisement.

Note that unlike other Extensions used in Mobile IP, the One-byte Padding Extension is encoded as a single byte, with no "Length" nor "Data" field present. The One-byte Padding Extension is defined as follows:

```

  0 1 2 3 4 5 6 7
  +---+---+---+---+---+---+
  |           |
  |   Type   |
  +---+---+---+---+---+---+

```

Type 0 (One-byte Padding Extension)

## 2.2. Agent Solicitation

An Agent Solicitation is identical to an ICMP Router Solicitation with the further restriction that the IP TTL Field MUST be set to 1.

## 2.3. Foreign Agent and Home Agent Considerations

Any mobility agent which cannot be discovered by a link-layer protocol MUST send Agent Advertisements. An agent which can be discovered by a link-layer protocol SHOULD also implement Agent Advertisements. However, the Advertisements need not be sent, except when the site policy requires registration with the agent (i.e., when the 'R' bit is set), or as a response to a specific Agent Solicitation. All mobility agents MUST process packets that they receive addressed to the Mobile-Agents multicast group, at address 224.0.0.11. A mobile node MAY send an Agent Solicitation to 224.0.0.11. All mobility agents SHOULD respond to Agent Solicitations.

The same procedures, defaults, and constants are used in Agent Advertisement messages and Agent Solicitation messages as specified for ICMP Router Discovery [10], except that:

- a mobility agent MUST limit the rate at which it sends broadcast or multicast Agent Advertisements; the maximum rate SHOULD be chosen so that the Advertisements do not consume a significant amount of network bandwidth, AND
- a mobility agent that receives a Router Solicitation MUST NOT require that the IP Source Address is the address of a neighbor (i.e., an address that matches one of the router's own addresses on the arrival interface, under the subnet mask associated with that address of the router).
- a mobility agent MAY be configured to send Agent Advertisements only in response to an Agent Solicitation message.

If the home network is not a virtual network, then the home agent for any mobile node SHOULD be located on the link identified by the mobile node's home address, and Agent Advertisement messages sent by the home agent on this link MUST have the 'H' bit set. In this way, mobile nodes on their own home network will be able to determine that they are indeed at home. Any Agent Advertisement messages sent by the home agent on another link to which it may be attached (if it is a mobility agent serving more than one link), MUST NOT have the 'H' bit set, unless the home agent also serves as a home agent (to other mobile nodes) on that other link. A mobility agent MAY use different settings for each of the 'R', 'H', and 'F' bits on different network interfaces.

If the home network is a virtual network, the home network has no physical realization external to the home agent itself. In this case, there is no physical network link on which to send Agent Advertisement messages advertising the home agent. Mobile nodes for which this is the home network are always treated as being away from home.

On a particular subnet, either all mobility agents MUST include the Prefix-Lengths Extension or all of them MUST NOT include this Extension. Equivalently, it is prohibited for some agents on a given subnet to include the Extension but for others not to include it. Otherwise, one of the move detection algorithms designed for mobile nodes will not function properly ([Section 2.4.2](#)).

#### **2.3.1. Advertised Router Addresses**

The ICMP Router Advertisement portion of the Agent Advertisement MAY contain one or more router addresses. An agent SHOULD only put its own addresses, if any, in the advertisement. Whether or not its own address appears in the Router Addresses, a foreign agent MUST route datagrams it receives from registered mobile nodes ([Section 4.2.2](#)).

#### **2.3.2. Sequence Numbers and Rollover Handling**

The sequence number in Agent Advertisements ranges from 0 to 0xffff. After booting, an agent MUST use the number 0 for its first advertisement. Each subsequent advertisement MUST use the sequence number one greater, with the exception that the sequence number 0xffff MUST be followed by sequence number 256. In this way, mobile nodes can distinguish a reduction in the sequence number that occurs after a reboot from a reduction that results in rollover of the sequence number after it attains the value 0xffff.



#### 2.4. Mobile Node Considerations

Every mobile node MUST implement Agent Solicitation. Solicitations SHOULD only be sent in the absence of Agent Advertisements and when a care-of address has not been determined through a link-layer protocol or other means. The mobile node uses the same procedures, defaults, and constants for Agent Solicitation as specified for ICMP Router Solicitation messages [10], except that the mobile node MAY solicit more often than once every three seconds, and that a mobile node that is currently not connected to any foreign agent MAY solicit more times than MAX\_SOLICITATIONS.

The rate at which a mobile node sends Solicitations MUST be limited by the mobile node. The mobile node MAY send three initial Solicitations at a maximum rate of one per second while searching for an agent. After this, the rate at which Solicitations are sent MUST be reduced so as to limit the overhead on the local link. Subsequent Solicitations MUST be sent using a binary exponential backoff mechanism, doubling the interval between consecutive Solicitations, up to a maximum interval. The maximum interval SHOULD be chosen appropriately based upon the characteristics of the media over which the mobile node is soliciting. This maximum interval SHOULD be at least one minute between Solicitations.

While still searching for an agent, the mobile node MUST NOT increase the rate at which it sends Solicitations unless it has received a positive indication that it has moved to a new link. After successfully registering with an agent, the mobile node SHOULD also increase the rate at which it will send Solicitations when it next begins searching for a new agent with which to register. The increased solicitation rate MAY revert to the maximum rate, but then MUST be limited in the manner described above. In all cases, the recommended solicitation intervals are nominal values. Mobile nodes MUST randomize their solicitation times around these nominal values as specified for ICMP Router Discovery [10].

Mobile nodes MUST process received Agent Advertisements. A mobile node can distinguish an Agent Advertisement message from other uses of the ICMP Router Advertisement message by examining the number of advertised addresses and the IP Total Length field. When the IP total length indicates that the ICMP message is longer than needed for the number of advertised addresses, the remaining data is interpreted as one or more Extensions. The presence of a Mobility Agent Advertisement Extension identifies the advertisement as an Agent Advertisement.











































































































instance as performed by the mobile registration protocol, is widely understood to be a security problem in the current Internet if not authenticated [2]. Moreover, the Address Resolution Protocol (ARP) is not authenticated, and can potentially be used to steal another host's traffic. The use of "Gratuitous ARP" ([Section 4.6](#)) brings with it all of the risks associated with the use of ARP.

### 5.3. Key Management

This specification requires a strong authentication mechanism (keyed MD5) which precludes many potential attacks based on the Mobile IP registration protocol. However, because key distribution is difficult in the absence of a network key management protocol, messages with the foreign agent are not all required to be authenticated. In a commercial environment it might be important to authenticate all messages between the foreign agent and the home agent, so that billing is possible, and service providers do not provide service to users that are not legitimate customers of that service provider.

### 5.4. Picking Good Random Numbers

The strength of any authentication mechanism depends on several factors, including the innate strength of the authentication algorithm, the secrecy of the key used, the strength of the key used, and the quality of the particular implementation. This specification requires implementation of keyed MD5 for authentication, but does not preclude the use of other authentication algorithms and modes. For keyed MD5 authentication to be useful, the 128-bit key must be both secret (that is, known only to authorized parties) and pseudo-random. If nonces are used in connection with replay protection, they must also be selected carefully. Eastlake, et al. [14] provides more information on generating pseudo-random numbers.

### 5.5. Privacy

Users who have sensitive data that they do not wish others to see should use mechanisms outside the scope of this document (such as encryption) to provide appropriate protection. Users concerned about traffic analysis should consider appropriate use of link encryption. If absolute location privacy is desired, the mobile node can create a tunnel to its home agent. Then, datagrams destined for correspondent nodes will appear to emanate from the home network, and it may be more difficult to pinpoint the location of the mobile node. Such mechanisms are all beyond the scope of this document.

## 5.6. Ingress Filtering

Many routers implement security policies such as "ingress filtering" [15] that do not allow forwarding of packets that have a Source Address which appears topologically incorrect. In environments where this is a problem, mobile nodes may use reverse tunneling [27] with the foreign agent supplied care-of address as the Source Address. Reverse tunneled packets will be able to pass normally through such routers, while ingress filtering rules will still be able to locate the true topological source of the packet in the same way as packets from non-mobile nodes.

## 5.7. Replay Protection for Registration Requests

The Identification field is used to let the home agent verify that a registration message has been freshly generated by the mobile node, not replayed by an attacker from some previous registration. Two methods are described in this section: timestamps (mandatory) and "nonces" (optional). All mobile nodes and home agents MUST implement timestamp-based replay protection. These nodes MAY also implement nonce-based replay protection (but see [Appendix A](#)).

The style of replay protection in effect between a mobile node and its home agent is part of the mobile security association. A mobile node and its home agent MUST agree on which method of replay protection will be used. The interpretation of the Identification field depends on the method of replay protection as described in the subsequent subsections.

Whatever method is used, the low-order 32 bits of the Identification MUST be copied unchanged from the Registration Request to the Reply. The foreign agent uses those bits (and the mobile node's home address) to match Registration Requests with corresponding replies. of any Registration Reply are identical to the bits it sent in the Registration Request.

The Identification in a new Registration Request MUST NOT be the same as in an immediately preceding Request, and SHOULD NOT repeat while the same security context is being used between the mobile node and the home agent. Retransmission as in [Section 3.6.3](#) is allowed.

### 5.7.1. Replay Protection using Timestamps

The basic principle of timestamp replay protection is that the node generating a message inserts the current time of day, and the node receiving the message checks that this timestamp is sufficiently close to its own time of day. Unless specified differently in the security association between the nodes, a default value of 7 seconds

MAY be used to limit the time difference. This value SHOULD be greater than 3 seconds. Obviously the two nodes must have adequately synchronized time-of-day clocks. As with any messages, time synchronization messages may be protected against tampering by an authentication mechanism determined by the security context between the two nodes.

If timestamps are used, the mobile node MUST set the Identification field to a 64-bit value formatted as specified by the Network Time Protocol [26]. The low-order 32 bits of the NTP format represent fractional seconds, and those bits which are not available from a time source SHOULD be generated from a good source of randomness. Note, however, that when using timestamps, the 64-bit Identification used in a Registration Request from the mobile node MUST be greater than that used in any previous Registration Request, as the home agent uses this field also as a sequence number. Without such a sequence number, it would be possible for a delayed duplicate of an earlier Registration Request to arrive at the home agent (within the clock synchronization required by the home agent), and thus be applied out of order, mistakenly altering the mobile node's current registered care-of address.

Upon receipt of a Registration Request with an authorization-enabling extension, the home agent MUST check the Identification field for validity. In order to be valid, the timestamp contained in the Identification field MUST be close enough to the home agent's time of day clock and the timestamp MUST be greater than all previously accepted timestamps for the requesting mobile node. Time tolerances and resynchronization details are specific to a particular mobility security association.

If the timestamp is valid, the home agent copies the entire Identification field into the Registration Reply it returns the Reply to the mobile node. If the timestamp is not valid, the home agent copies only the low-order 32 bits into the Registration Reply, and supplies the high-order 32 bits from its own time of day. In this latter case, the home agent MUST reject the registration by returning Code 133 (identification mismatch) in the Registration Reply.

As described in [Section 3.6.2.1](#), the mobile node MUST verify that the low-order 32 bits of the Identification in the Registration Reply are identical to those in the rejected registration attempt, before using the high-order bits for clock resynchronization.

### 5.7.2. Replay Protection using Nonces

The basic principle of nonce replay protection is that node A includes a new random number in every message to node B, and checks that node B returns that same number in its next message to node A. Both messages use an authentication code to protect against alteration by an attacker. At the same time node B can send its own nonces in all messages to node A (to be echoed by node A), so that it too can verify that it is receiving fresh messages.

The home agent may be expected to have resources for computing pseudo-random numbers useful as nonces [14]. It inserts a new nonce as the high-order 32 bits of the identification field of every Registration Reply. The home agent copies the low-order 32 bits of the Identification from the Registration Request message into the low-order 32 bits of the Identification in the Registration Reply. When the mobile node receives an authenticated Registration Reply from the home agent, it saves the high-order 32 bits of the identification for use as the high-order 32 bits of its next Registration Request.

The mobile node is responsible for generating the low-order 32 bits of the Identification in each Registration Request. Ideally it should generate its own random nonces. However it may use any expedient method, including duplication of the random value sent by the home agent. The method chosen is of concern only to the mobile node, because it is the node that checks for valid values in the Registration Reply. The high-order and low-order 32 bits of the identification chosen SHOULD both differ from their previous values. The home agent uses a new high-order value and the mobile node uses a new low-order value for each registration message. The foreign agent uses the low-order value (and the mobile host's home address) to correctly match registration replies with pending Requests ([Section 3.7.1](#)).

If a registration message is rejected because of an invalid nonce, the Reply always provides the mobile node with a new nonce to be used in the next registration. Thus the nonce protocol is self-synchronizing.

## 6. IANA Considerations

Mobile IP specifies several new number spaces for values to be used in various message fields. These number spaces include the following:

- Mobile IP message types sent to UDP port 434, as defined in [section 1.8](#).

- types of extensions to Registration Request and Registration Reply messages (see sections [3.3](#) and [3.4](#), and also consult [27, 29, 6, 7, 12])
- values for the Code in the Registration Reply message (see [section 3.4](#), and also consult [27, 29, 6, 7, 12])
- Mobile IP defines so-called Agent Solicitation and Agent Advertisement messages. These messages are in fact Router Discovery messages [[10](#)] augmented with mobile-IP specific extensions. Thus, they do not define a new name space, but do define additional Router Discovery extensions as described below in [Section 6.2](#). Also see [Section 2.1](#) and consult [7, 12].

There are additional Mobile IP numbering spaces specified in [[7](#)].

Information about assignment of mobile-ip numbers derived from specifications external to this document is given by IANA at <http://www.iana.org/numbers.html>. From that URL, follow the hyperlinks to [M] within the "Directory of General Assigned Numbers", and subsequently to the specific section for "Mobile IP Numbers".

### 6.1. Mobile IP Message Types

Mobile IP messages are defined to be those that are sent to a message recipient at port 434 (UDP or TCP). The number space for Mobile IP messages is specified in [Section 1.8](#). Approval of new extension numbers is subject to Expert Review, and a specification is required [[30](#)]. The currently standardized message types have the following numbers, and are specified in the indicated sections.

Type	Name	Section
----	-----	-----
1	Registration Request	3.3
3	Registration Reply	3.4

### 6.2. Extensions to [RFC 1256](#) Router Advertisement

[RFC 1256](#) defines two ICMP message types, Router Advertisement and Router Solicitation. Mobile IP defines a number space for extensions to Router Advertisement, which could be used by protocols other than Mobile IP. The extension types currently standardized for use with Mobile IP have the following numbers.

Type	Name	Reference
0	One-byte Padding	2.1.3
16	Mobility Agent Advertisement	2.1.1
19	Prefix-Lengths	2.1.2

Approval of new extension numbers for use with Mobile IP is subject to Expert Review, and a specification is required [30].

### 6.3. Extensions to Mobile IP Registration Messages

The Mobile IP messages, specified within this document, and listed in sections 1.8 and 6.1, may have extensions. Mobile IP message extensions all share the same number space, even if they are to be applied to different Mobile IP messages. The number space for Mobile IP message extensions is specified within this document. Approval of new extension numbers is subject to Expert Review, and a specification is required [30].

Type	Name	Reference
0	One-byte Padding	
32	Mobile-Home Authentication	3.5.2
33	Mobile-Foreign Authentication	3.5.3
34	Foreign-Home Authentication	3.5.4

### 6.4. Code Values for Mobile IP Registration Reply Messages

The Mobile IP Registration Reply message, specified in section 3.4, has a Code field. The number space for the Code field values is also specified in Section 3.4. The Code number space is structured according to whether the registration was successful, or whether the foreign agent denied the registration request, or lastly whether the home agent denied the registration request, as follows:

0-8	Success Codes
9-63	No allocation guidelines currently exist
64-127	Error Codes from the Foreign Agent
128-192	Error Codes from the Home Agent
193-255	No allocation guidelines currently exist

Approval of new Code values requires Expert Review [30].

## 7. Acknowledgments

Special thanks to Steve Deering (Xerox PARC), along with Dan Duchamp and John Ioannidis (JI) (Columbia University), for forming the working group, chairing it, and putting so much effort into its early development. Columbia's early Mobile IP work can be found in [18, 19, 17].

Thanks also to Kannan Alaggapan, Greg Minshall, Tony Li, Jim Solomon, Erik Nordmark, Basavaraj Patil, and Phil Roberts for their contributions to the group while performing the duties of chairperson, as well as for their many useful comments.

Thanks to the active members of the Mobile IP Working Group, particularly those who contributed text, including (in alphabetical order)

- Ran Atkinson (Naval Research Lab),
- Samita Chakrabarti (Sun Microsystems)
- Ken Imboden (Candlestick Networks, Inc.)
- Dave Johnson (Carnegie Mellon University),
- Frank Kastenholz (FTP Software),
- Anders Klemets (KTH),
- Chip Maguire (KTH),
- Alison Mankin (ISI)
- Andrew Myles (Macquarie University),
- Thomas Narten (IBM)
- Al Quirt (Bell Northern Research),
- Yakov Rekhter (IBM), and
- Fumio Teraoka (Sony).
- Alper Yegin (NTT DoCoMo)

Thanks to Charlie Kunzinger and to Bill Simpson, the editors who produced the first drafts for of this document, reflecting the discussions of the Working Group. Much of the new text in the later revisions preceding [RFC 2002](#) is due to Jim Solomon and Dave Johnson.

Thanks to Greg Minshall (Novell), Phil Karn (Qualcomm), Frank Kastenholz (FTP Software), and Pat Calhoun (Sun Microsystems) for their generous support in hosting interim Working Group meetings.

Sections [1.10](#) and [1.11](#), which specify new extension formats to be used with aggregatable extension types, were included from a specification document (entitled "Mobile IP Extensions Rationalization (MIER)", which was written by

- Mohamed M.Khalil, Nortel Networks
- Raja Narayanan, nVisible Networks
- Haseeb Akhtar, Nortel Networks
- Emad Qaddoura, Nortel Networks

Thanks to these authors, and also for the additional work on MIER, which was contributed by Basavaraj Patil, Pat Calhoun, Neil Justusson, N. Asokan, and Jouni Malinen.



## A. Patent Issues

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

## B. Link-Layer Considerations

The mobile node MAY use link-layer mechanisms to decide that its point of attachment has changed. Such indications include the Down/Testing/Up interface status [[24](#)], and changes in cell or administration. The mechanisms will be specific to the particular link-layer technology, and are outside the scope of this document.

The Point-to-Point-Protocol (PPP) [[42](#)] and its Internet Protocol Control Protocol (IPCP) [[25](#)], negotiates the use of IP addresses. The mobile node SHOULD first attempt to specify its home address, so that if the mobile node is attaching to its home network, the unrouted link will function correctly. When the home address is not accepted by the peer, but a transient IP address is dynamically assigned to the mobile node, and the mobile node is capable of supporting a co-located care-of address, the mobile node MAY register that address as a co-located care-of address. When the peer specifies its own IP address, that address MUST NOT be assumed to be a foreign agent care-of address or the IP address of a home agent.

PPP extensions for Mobile IP have been specified in [RFC 2290](#) [44]. Please consult that document for additional details for how to handle care-of address assignment from PPP in a more efficient manner.

## C. TCP Considerations

### C.1. TCP Timers

When high-delay (e.g. SATCOM) or low-bandwidth (e.g. High-Frequency Radio) links are in use, some TCP stacks may have insufficiently adaptive (non-standard) retransmission timeouts. There may be spurious retransmission timeouts, even when the link and network are actually operating properly, but just with a high delay because of the medium in use. This can cause an inability to create or maintain TCP connections over such links, and can also cause unneeded retransmissions which consume already scarce bandwidth. Vendors are encouraged to follow the algorithms in [RFC 2988](#) [31] when implementing TCP retransmission timers. Vendors of systems designed for low-bandwidth, high-delay links should consult RFCs 2757 and 2488 [28, 1]. Designers of applications targeted to operate on mobile nodes should be sensitive to the possibility of timer-related difficulties.

### C.2. TCP Congestion Management

Mobile nodes often use media which are more likely to introduce errors, effectively causing more packets to be dropped. This introduces a conflict with the mechanisms for congestion management found in modern versions of TCP [21]. Now, when a packet is dropped, the correspondent node's TCP implementation is likely to react as if there were a source of network congestion, and initiate the slow-start mechanisms [21] designed for controlling that problem. However, those mechanisms are inappropriate for overcoming errors introduced by the links themselves, and have the effect of magnifying the discontinuity introduced by the dropped packet. This problem has been analyzed by Caceres, et al. [5]. TCP approaches to the problem of handling errors that might interfere with congestion management are discussed in documents from the [pilc] working group [3, 9]. While such approaches are beyond the scope of this document, they illustrate that providing performance transparency to mobile nodes involves understanding mechanisms outside the network layer. Problems introduced by higher media error rates also indicate the need to avoid designs which systematically drop packets; such designs might otherwise be considered favorably when making engineering tradeoffs.

**D. Example Scenarios**

This section shows example Registration Requests for several common scenarios.

**D.1. Registering with a Foreign Agent Care-of Address**

The mobile node receives an Agent Advertisement from a foreign agent and wishes to register with that agent using the advertised foreign agent care-of address. The mobile node wishes only IP-in-IP encapsulation, does not want broadcasts, and does not want simultaneous mobility bindings:

## IP fields:

Source Address = mobile node's home address

Destination Address = copied from the IP source address of the Agent Advertisement

Time to Live = 1

## UDP fields:

Source Port = <any>

Destination Port = 434

## Registration Request fields:

Type = 1

S=0,B=0,D=0,M=0,G=0

Lifetime = the Registration Lifetime copied from the Mobility Agent Advertisement Extension of the Router Advertisement message

Home Address = the mobile node's home address

Home Agent = IP address of mobile node's home agent

Care-of Address = the Care-of Address copied from the Mobility Agent Advertisement Extension of the Router Advertisement message

Identification = Network Time Protocol timestamp or Nonce

## Extensions:

An authorization-enabling extension (e.g., the Mobile-Home Authentication Extension)

**D.2. Registering with a Co-Located Care-of Address**

The mobile node enters a foreign network that contains no foreign agents. The mobile node obtains an address from a DHCP server [13] for use as a co-located care-of address. The mobile node supports all forms of encapsulation (IP-in-IP, minimal encapsulation, and GRE), desires a copy of broadcast datagrams on the home network, and does not want simultaneous mobility bindings:

IP fields:  
Source Address = care-of address obtained from DHCP server  
Destination Address = IP address of home agent  
Time to Live = 64

UDP fields:  
Source Port = <any>  
Destination Port = 434

Registration Request fields:  
Type = 1  
S=0,B=1,D=1,M=1,G=1  
Lifetime = 1800 (seconds)  
Home Address = the mobile node's home address  
Home Agent = IP address of mobile node's home agent  
Care-of Address = care-of address obtained from DHCP server  
Identification = Network Time Protocol timestamp or Nonce

Extensions:  
The Mobile-Home Authentication Extension

### D.3. Deregistration

The mobile node returns home and wishes to deregister all care-of addresses with its home agent.

IP fields:  
Source Address = mobile node's home address  
Destination Address = IP address of home agent  
Time to Live = 1

UDP fields:  
Source Port = <any>  
Destination Port = 434

Registration Request fields:  
Type = 1  
S=0,B=0,D=0,M=0,G=0  
Lifetime = 0  
Home Address = the mobile node's home address  
Home Agent = IP address of mobile node's home agent  
Care-of Address = the mobile node's home address  
Identification = Network Time Protocol timestamp or Nonce

Extensions:  
An authorization-enabling extension (e.g., the Mobile-Home Authentication Extension)

### E. Applicability of Prefix-Lengths Extension

Caution is indicated with the use of the Prefix-Lengths Extension over wireless links, due to the irregular coverage areas provided by wireless transmitters. As a result, it is possible that two foreign agents advertising the same prefix might indeed provide different connectivity to prospective mobile nodes. The Prefix-Lengths Extension SHOULD NOT be included in the advertisements sent by agents in such a configuration.

Foreign agents using different wireless interfaces would have to cooperate using special protocols to provide identical coverage in space, and thus be able to claim to have wireless interfaces situated on the same subnetwork. In the case of wired interfaces, a mobile node disconnecting and subsequently connecting to a new point of attachment, may well send in a new Registration Request no matter whether the new advertisement is on the same medium as the last recorded advertisement. And, finally, in areas with dense populations of foreign agents it would seem unwise to require the propagation via routing protocols of the subnet prefixes associated with each individual wireless foreign agent; such a strategy could lead to quick depletion of available space for routing tables, unwarranted increases in the time required for processing routing updates, and longer decision times for route selection if routes (which are almost always unnecessary) are stored for wireless "subnets".

### F. Interoperability Considerations

This document specifies revisions to [RFC 2002](#) that are intended to improve interoperability by resolving ambiguities contained in the earlier text. Implementations that perform authentication according to the new more precisely specified algorithm would be interoperable with earlier implementations that did what was originally expected for producing authentication data. That was a major source of non-interoperability before.

However, this specification does have new features that, if used, would cause interoperability problems with older implementations. All features specified in [RFC 2002](#) will work with the new implementations, except for V-J compression [20]. The following list details some of the possible areas of compatibility problems that may be experienced by nodes conforming to this revised specification, when attempting to interoperate with nodes obeying [RFC 2002](#).

- A client that expects some of the newly mandatory features (like reverse tunneling) from a foreign agent would still be interoperable as long as it pays attention to the 'T' bit.

- Mobile nodes that use the NAI extension to identify themselves would not work with old mobility agents.
- Mobile nodes that use a zero home address and expect to receive their home address in the Registration Reply would not work with old mobility agents.
- Mobile nodes that attempt to authenticate themselves without using the Mobile-Home authentication extension will be unable to successfully register with their home agent.

In all of these cases, a robust and well-configured mobile node is very likely to be able to recover if it takes reasonable actions upon receipt of a Registration Reply with an error code indicating the cause for rejection. For instance, if a mobile node sends a registration request that is rejected because it contains the wrong kind of authentication extension, then the mobile node could retry the registration with a mobile-home authentication extension, since the foreign agent and/or home agent in this case will not be configured to demand the alternative authentication data.

## G. Changes since [RFC 2002](#)

This section details differences between the original Mobile IP base specification ([RFC 2002](#) and ff.) that have been made as part of this revised protocol specification for Mobile IP.

### G.1. Major Changes

- Specification for Destination IP address of Registration Reply transmitted from Foreign Agent, to avoid any possible transmission to IP address 0.0.0.0.
- Specification of two new formats for Mobile IP extensions, according to the ideas contained in MIER.
- Specification that the SPI of the MN-HA authentication extension is to be used as part of the data over which the authentication algorithm must be computed.
- Eliminated Van-Jacobson Compression feature
- Specification that foreign agents MAY send advertisements at a rate faster than once per second, but chosen so that the advertisements do not burden the capacity of the local link. For simplicity, the foreign agent now MAY send advertisements at an interval less than 1/3 the advertised ICMP Lifetime.

- Specification that foreign agents SHOULD support reverse tunneling, and home agents MUST support decapsulation of reverse tunnels.
- Changed the preconfiguration requirements in [section 3.6](#) for the mobile node to reflect the capability, specified in [RFC 2794](#) [6], for the mobile node to identify itself by using its NAI, and then getting a home address from the Registration Reply.
- Changed [section 3.7.3.1](#) so that a foreign agent is not required to discard Registration Replies that have a Home Address field that does not match any pending Registration Request.
- Allowed registrations to be authenticated by use of a security association between the mobile node and a suitable authentication entity acceptable to the home agent. Defined "Authorization-enabling Extension" to be an authentication extension that makes a registration message acceptable to the recipient. This is needed according to specification in [6].
- Mandated that HMAC-MD5 be used instead of the "prefix+suffix" mode of MD5 as originally mandated in [RFC 2002](#).
- Specified that the mobile node SHOULD take the first care-of address in a list offered by a foreign agent, and MAY try each subsequent advertised address in turn if the attempted registrations are rejected by the foreign agent
- Clarification that a mobility agent SHOULD only put its own addresses into the initial (i.e., not mobility-related) list of routers in the mobility advertisement. [RFC 2002](#) suggests that a mobility agent might advertise other default routers.
- Specification that a mobile node MUST ignore reserved bits in Agent Advertisements, as opposed to discarding such advertisements. In this way, new bits can be defined later, without affecting the ability for mobile nodes to use the advertisements even when the newly defined bits are not understood. Furthermore, foreign agents can set the `R' bit to make sure that new bits are handled by themselves instead of some legacy mobility agent.
- Specification that the foreign agent checks to make sure that the indicated home agent address does not belong to any of its network interfaces before relaying a Registration Request. If

the check fails, and the foreign agent is not the mobile node's home agent, then the foreign agent rejects the request with code 136 (unknown home agent address).

- Specification that, while they are away from the home network, mobile nodes MUST NOT broadcast ARP packets to find the MAC address of another Internet node. Thus, the (possibly empty) list of Router Addresses from the ICMP Router Advertisement portion of the message is not useful for selecting a default router, unless the mobile node has some means not involving broadcast ARP and not specified within this document for obtaining the MAC address of one of the routers in the list. Similarly, in the absence of unspecified mechanisms for obtaining MAC addresses on foreign networks, the mobile node MUST ignore redirects to other routers on foreign networks.
- Specification that a foreign agent MUST NOT use broadcast ARP for a mobile node's MAC address on a foreign network. It may obtain the MAC address by copying the information from an Agent Solicitation or a Registration Request transmitted from a mobile node.
- Specification that a foreign agent's ARP cache for the mobile node's IP address MUST NOT be allowed to expire before the mobile node's visitor list entry expires, unless the foreign agent has some way other than broadcast ARP to refresh its MAC address associated to the mobile node's IP address.
- At the end of [section 4.6](#), clarified that a home agent MUST NOT make any changes to the way it performs proxy ARP after it rejects an invalid deregistration request.
- In [section 4.2.3](#), specification that multihomed home agents MUST use the the address sent to the mobile node in the home agent field of the registration reply as the source address in the outer IP header of the encapsulated datagram.
- Inserted 'T' bit into its proper place in the Registration Request message format ([section 3.3](#)).

## G.2. Minor Changes

- Allowed registration replies to be processed by the mobile node, even in the absence of any Mobile-Home Authentication extension, when containing rejection code by the foreign agent.



- Specification that the foreign agent MAY configure a maximum number of pending registrations that it is willing to maintain (typically 5). Additional registrations SHOULD then be rejected by the foreign agent with code 66. The foreign agent MAY delete any pending Registration Request after the request has been pending for more than 7 seconds; in this case, the foreign agent SHOULD reject the Request with code 78 (registration timeout).
- Relaxation of the requirement that, when a mobile node has joined a multicast group at the router on the foreign network, the mobile node MUST use its home address as the source IP address for multicast packets,
- Clarification that a mobility agent MAY use different settings for each of the 'R', 'H', and 'F' bits on different network interfaces.
- Replacement of the terminology "recursive tunneling" by the terminology "nested tunneling".
- Specification that the mobile node MAY use the IP source address of an agent advertisement as its default router address.
- Clarification that keys with arbitrary binary values MUST be supported as part of mobility security associations.
- Specification that the default value may be chosen as 7 seconds, for allowable time skews between a home agent and mobile node using timestamps for replay protection. Further specification that this value SHOULD be greater than 3 seconds.
- Specification that Registration Requests with the 'D' bit set to 0, and specifying a care-of address not offered by the foreign agent, MUST be rejected with code 77 (invalid care-of address).
- Clarification that the foreign agent SHOULD consider its own maximum value when handling the Lifetime field of the Registration Reply.
- Clarification that the home agent MUST ignore the 'B' bit (as opposed to rejecting the Registration Request) if it does not support broadcasts.

- Advice about the impossibility of using dynamic home agent discovery in the case when routers change the IP destination address of a datagram from a subnet-directed broadcast address to 255.255.255.255 before injecting it into the destination subnet.
- Clarified that when an Agent Advertisement is unicast to a mobile node, the specific IP home address of a mobile node MAY be used as the destination IP address.
- Included a reference to [RFC 2290](#) within [appendix B](#), which deals with PPP operation.
- Created IANA Considerations section
- In [section 3.8.3](#), clarified that a home agent SHOULD arrange the selection of a home address for a mobile node when the Registration Reply contains a zero Home Address.

### G.3. Changes since revision 04 of RFC2002bis

This section lists the changes between this version (...-06.txt) and the previous version (...-04.txt) of the document. This section can be deleted by the RFC editor.

- Noted that HMAC-MD5 should be considered for use in place of the "prefix+suffix" mode of MD5 as originally mandated in [RFC 2002](#).
- Included a reference to [RFC 2290](#) within [appendix B](#), which deals with PPP operation.
- Revamped IANA Considerations section
- Revamped Changes section
- Replaced Patents section with wording mandated from [RFC 2026](#).
- Updated citations.

















## Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

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

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

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Acknowledgement

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