

Network Working Group
Request for Comments: 3775
Category: Standards Track

D. Johnson
Rice University
C. Perkins
Nokia Research Center
J. Arkko
Ericsson
June 2004

Mobility Support in IPv6

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 (2004).

Abstract

This document specifies a protocol which allows nodes to remain reachable while moving around in the IPv6 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 the mobile node's current location. IPv6 packets addressed to a mobile node's home address are transparently routed to its care-of address. The protocol enables IPv6 nodes to cache the binding of a mobile node's home address with its care-of address, and to then send any packets destined for the mobile node directly to it at this care-of address. To support this operation, Mobile IPv6 defines a new IPv6 protocol and a new destination option. All IPv6 nodes, whether mobile or stationary, can communicate with mobile nodes.

Table of Contents

1.	Introduction	5
2.	Comparison with Mobile IP for IPv4	6
3.	Terminology.	7
3.1.	General Terms	8
3.2.	Mobile IPv6 Terms	10
4.	Overview of Mobile IPv6.	13
4.1.	Basic Operation	13
4.2.	New IPv6 Protocol	15
4.3.	New IPv6 Destination Option	17
4.4.	New IPv6 ICMP Messages.	17
4.5.	Conceptual Data Structure Terminology	17
4.6.	Site-Local Addressability	18
5.	Overview of Mobile IPv6 Security	18
5.1.	Binding Updates to Home Agents.	18
5.2.	Binding Updates to Correspondent Nodes.	20
5.2.1.	Node Keys	20
5.2.2.	Nonces.	20
5.2.3.	Cookies and Tokens.	21
5.2.4.	Cryptographic Functions	22
5.2.5.	Return Routability Procedure.	22
5.2.6.	Authorizing Binding Management Messages	27
5.2.7.	Updating Node Keys and Nonces	29
5.2.8.	Preventing Replay Attacks	30
5.3.	Dynamic Home Agent Address Discovery.	30
5.4.	Mobile Prefix Discovery	30
5.5.	Payload Packets	30
6.	New IPv6 Protocol, Message Types, and Destination Option	31
6.1.	Mobility Header	31
6.1.1.	Format.	32
6.1.2.	Binding Refresh Request Message	34
6.1.3.	Home Test Init Message.	35
6.1.4.	Care-of Test Init Message	36
6.1.5.	Home Test Message	37
6.1.6.	Care-of Test Message.	38
6.1.7.	Binding Update Message.	39
6.1.8.	Binding Acknowledgement Message	42
6.1.9.	Binding Error Message	44
6.2.	Mobility Options.	46
6.2.1.	Format.	46
6.2.2.	Pad1.	47
6.2.3.	PadN.	48
6.2.4.	Binding Refresh Advice.	48
6.2.5.	Alternate Care-of Address	49
6.2.6.	Nonce Indices	49
6.2.7.	Binding Authorization Data.	50
6.3.	Home Address Option	51

6.4.	Type 2 Routing Header	53
6.4.1.	Format.	54
6.5.	ICMP Home Agent Address Discovery Request Message .	55
6.6.	ICMP Home Agent Address Discovery Reply Message . .	56
6.7.	ICMP Mobile Prefix Solicitation Message Format. . .	57
6.8.	ICMP Mobile Prefix Advertisement Message Format . .	59
7.	Modifications to IPv6 Neighbor Discovery	61
7.1.	Modified Router Advertisement Message Format. . . .	61
7.2.	Modified Prefix Information Option Format	62
7.3.	New Advertisement Interval Option Format.	64
7.4.	New Home Agent Information Option Format.	65
7.5.	Changes to Sending Router Advertisements.	67
8.	Requirements for Types of IPv6 Nodes	69
8.1.	All IPv6 Nodes.	69
8.2.	IPv6 Nodes with Support for Route Optimization. . .	69
8.3.	All IPv6 Routers.	71
8.4.	IPv6 Home Agents.	71
8.5.	IPv6 Mobile Nodes	73
9.	Correspondent Node Operation	74
9.1.	Conceptual Data Structures.	74
9.2.	Processing Mobility Headers	75
9.3.	Packet Processing	76
9.3.1.	Receiving Packets with Home Address Option.	76
9.3.2.	Sending Packets to a Mobile Node.	77
9.3.3.	Sending Binding Error Messages.	78
9.3.4.	Receiving ICMP Error Messages	79
9.4.	Return Routability Procedure.	79
9.4.1.	Receiving Home Test Init Messages	80
9.4.2.	Receiving Care-of Test Init Messages.	80
9.4.3.	Sending Home Test Messages.	80
9.4.4.	Sending Care-of Test Messages	81
9.5.	Processing Bindings	81
9.5.1.	Receiving Binding Updates	81
9.5.2.	Requests to Cache a Binding	84
9.5.3.	Requests to Delete a Binding.	84
9.5.4.	Sending Binding Acknowledgements.	85
9.5.5.	Sending Binding Refresh Requests.	86
9.6.	Cache Replacement Policy.	86
10.	Home Agent Operation	87
10.1.	Conceptual Data Structures.	87
10.2.	Processing Mobility Headers	88
10.3.	Processing Bindings	88
10.3.1.	Primary Care-of Address Registration.	88
10.3.2.	Primary Care-of Address De-Registration	92
10.4.	Packet Processing	94
10.4.1.	Intercepting Packets for a Mobile Node.	94
10.4.2.	Processing Intercepted Packets.	95
10.4.3.	Multicast Membership Control.	96

	10.4.4. Stateful Address Autoconfiguration.	98
	10.4.5. Handling Reverse Tunneled Packets	98
	10.4.6. Protecting Return Routability Packets	99
10.5.	Dynamic Home Agent Address Discovery.	99
	10.5.1. Receiving Router Advertisement Messages	100
10.6.	Sending Prefix Information to the Mobile Node	102
	10.6.1. List of Home Network Prefixes	102
	10.6.2. Scheduling Prefix Deliveries.	102
	10.6.3. Sending Advertisements.	104
	10.6.4. Lifetimes for Changed Prefixes.	105
11.	Mobile Node Operation.	105
	11.1. Conceptual Data Structures.	105
	11.2. Processing Mobility Headers	107
	11.3. Packet Processing	107
	11.3.1. Sending Packets While Away from Home.	107
	11.3.2. Interaction with Outbound IPsec Processing.	110
	11.3.3. Receiving Packets While Away from Home.	112
	11.3.4. Routing Multicast Packets	114
	11.3.5. Receiving ICMP Error Messages	115
	11.3.6. Receiving Binding Error Messages.	116
	11.4. Home Agent and Prefix Management.	117
	11.4.1. Dynamic Home Agent Address Discovery.	117
	11.4.2. Sending Mobile Prefix Solicitations	118
	11.4.3. Receiving Mobile Prefix Advertisements.	118
	11.5. Movement.	120
	11.5.1. Movement Detection.	120
	11.5.2. Forming New Care-of Addresses	122
	11.5.3. Using Multiple Care-of Addresses.	123
	11.5.4. Returning Home.	124
	11.6. Return Routability Procedure.	126
	11.6.1. Sending Test Init Messages.	126
	11.6.2. Receiving Test Messages	127
	11.6.3. Protecting Return Routability Packets	128
	11.7. Processing Bindings	128
	11.7.1. Sending Binding Updates to the Home Agent	128
	11.7.2. Correspondent Registration.	131
	11.7.3. Receiving Binding Acknowledgements.	134
	11.7.4. Receiving Binding Refresh Requests.	136
	11.8. Retransmissions and Rate Limiting	137
12.	Protocol Constants	138
13.	Protocol Configuration Variables	138
14.	IANA Considerations.	139
15.	Security Considerations.	142
	15.1. Threats	142
	15.2. Features.	144
	15.3. Binding Updates to Home Agent	145
	15.4. Binding Updates to Correspondent Nodes.	148
	15.4.1. Overview.	149

15.4.2.	Achieved Security Properties.	149
15.4.3.	Comparison to Regular IPv6 Communications .	150
15.4.4.	Replay Attacks.	152
15.4.5.	Denial-of-Service Attacks	152
15.4.6.	Key Lengths	153
15.5.	Dynamic Home Agent Address Discovery.	154
15.6.	Mobile Prefix Discovery	155
15.7.	Tunneling via the Home Agent.	155
15.8.	Home Address Option	156
15.9.	Type 2 Routing Header	156
16.	Contributors	157
17.	Acknowledgements	157
18.	References	158
18.1.	Normative References.	158
18.2.	Informative References.	159
Appendix A.	Future Extensions	161
A.1.	Piggybacking.	161
A.2.	Triangular Routing.	161
A.3.	New Authorization Methods	161
A.4.	Dynamically Generated Home Addresses.	161
A.5.	Remote Home Address Configuration	162
A.6.	Neighbor Discovery Extensions	163
	Authors' Addresses.	164
	Full Copyright Statement.	165

1. Introduction

This document specifies a protocol which allows nodes to remain reachable while moving around in the IPv6 Internet. Without specific support for mobility in IPv6 [11], packets destined to a mobile node would not be able to reach it while the mobile node is away from its home link. In order to continue communication in spite of its movement, a mobile node could change its IP address each time it moves to a new link, but the mobile node would then not be able to maintain transport and higher-layer connections when it changes location. Mobility support in IPv6 is particularly important, as mobile computers are likely to account for a majority or at least a substantial fraction of the population of the Internet during the lifetime of IPv6.

The protocol defined in this document, known as Mobile IPv6, allows a mobile node to move from one link to another without changing the mobile node's "home address". Packets may be routed to the mobile node using this address regardless of the mobile node's current point of attachment to the Internet. The mobile node may also continue to communicate with other nodes (stationary or mobile) after moving to a

new link. The movement of a mobile node away from its home link is thus transparent to transport and higher-layer protocols and applications.

The Mobile IPv6 protocol is just as suitable for mobility across homogeneous media as for mobility across heterogeneous media. For example, Mobile IPv6 facilitates node movement from one Ethernet segment to another as well as it facilitates node movement from an Ethernet segment to a wireless LAN cell, with the mobile node's IP address remaining unchanged in spite of such movement.

One can think of the Mobile IPv6 protocol as solving the network-layer mobility management problem. Some mobility management applications -- for example, handover among wireless transceivers, each of which covers only a very small geographic area -- have been solved using link-layer techniques. For example, in many current wireless LAN products, link-layer mobility mechanisms allow a "handover" of a mobile node from one cell to another, re-establishing link-layer connectivity to the node in each new location.

Mobile IPv6 does not attempt to solve all general problems related to the use of mobile computers or wireless networks. In particular, this protocol does not attempt to solve:

- o Handling links with unidirectional connectivity or partial reachability, such as the hidden terminal problem where a host is hidden from only some of the routers on the link.
- o Access control on a link being visited by a mobile node.
- o Local or hierarchical forms of mobility management (similar to many current link-layer mobility management solutions).
- o Assistance for adaptive applications.
- o Mobile routers.
- o Service Discovery.
- o Distinguishing between packets lost due to bit errors vs. network congestion.

2. Comparison with Mobile IP for IPv4

The design of Mobile IP support in IPv6 (Mobile IPv6) benefits both from the experiences gained from the development of Mobile IP support in IPv4 (Mobile IPv4) [22, 23, 24], and from the opportunities provided by IPv6. Mobile IPv6 thus shares many features with Mobile

IPv4, but is integrated into IPv6 and offers many other improvements. This section summarizes the major differences between Mobile IPv4 and Mobile IPv6:

- o There is no need to deploy special routers as "foreign agents", as in Mobile IPv4. Mobile IPv6 operates in any location without any special support required from the local router.
- o Support for route optimization is a fundamental part of the protocol, rather than a nonstandard set of extensions.
- o Mobile IPv6 route optimization can operate securely even without pre-arranged security associations. It is expected that route optimization can be deployed on a global scale between all mobile nodes and correspondent nodes.
- o Support is also integrated into Mobile IPv6 for allowing route optimization to coexist efficiently with routers that perform "ingress filtering" [26].
- o The IPv6 Neighbor Unreachability Detection assures symmetric reachability between the mobile node and its default router in the current location.
- o Most packets sent to a mobile node while away from home in Mobile IPv6 are sent using an IPv6 routing header rather than IP encapsulation, reducing the amount of resulting overhead compared to Mobile IPv4.
- o Mobile IPv6 is decoupled from any particular link layer, as it uses IPv6 Neighbor Discovery [12] instead of ARP. This also improves the robustness of the protocol.
- o The use of IPv6 encapsulation (and the routing header) removes the need in Mobile IPv6 to manage "tunnel soft state".
- o The dynamic home agent address discovery mechanism in Mobile IPv6 returns a single reply to the mobile node. The directed broadcast approach used in IPv4 returns separate replies from each home agent.

3. Terminology

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [2].

3.1. General Terms

IP

Internet Protocol Version 6 (IPv6).

node

A device that implements IP.

router

A node that forwards IP packets not explicitly addressed to itself.

unicast routable address

An identifier for a single interface such that a packet sent to it from another IPv6 subnet is delivered to the interface identified by that address. Accordingly, a unicast routable address must have either a global or site-local scope (but not link-local).

host

Any node that is not a router.

link

A communication facility or medium over which nodes can communicate at the link layer, such as an Ethernet (simple or bridged). A link is the layer immediately below IP.

interface

A node's attachment to a link.

subnet prefix

A bit string that consists of some number of initial bits of an IP address.

interface identifier

A number used to identify a node's interface on a link. The interface identifier is the remaining low-order bits in the node's IP address after the subnet prefix.

link-layer address

A link-layer identifier for an interface, such as IEEE 802 addresses on Ethernet links.

packet

An IP header plus payload.

security association

An IPsec security association is a cooperative relationship formed by the sharing of cryptographic keying material and associated context. Security associations are simplex. That is, two security associations are needed to protect bidirectional traffic between two nodes, one for each direction.

security policy database

A database that specifies what security services are to be offered to IP packets and in what fashion.

destination option

Destination options are carried by the IPv6 Destination Options extension header. Destination options include optional information that need be examined only by the IPv6 node given as the destination address in the IPv6 header, not by routers in between. Mobile IPv6 defines one new destination option, the Home Address destination option (see Section 6.3).

routing header

A routing header may be present as an IPv6 header extension, and indicates that the payload has to be delivered to a destination IPv6 address in some way that is different from what would be carried out by standard Internet routing. In this document, use of the term "routing header" typically refers to use of a type 2 routing header, as specified in Section 6.4.

"|" (concatenation)

Some formulas in this specification use the symbol "|" to indicate bitwise concatenation, as in A | B. This concatenation requires that all of the octets of the datum A appear first in the result, followed by all of the octets of the datum B.

First (size, input)

Some formulas in this specification use a functional form "First (size, input)" to indicate truncation of the "input" data so that only the first "size" bits remain to be used.

3.2. Mobile IPv6 Terms

home address

A unicast routable address assigned to a mobile node, used as the permanent address of the mobile node. This address is within the mobile node's home link. Standard IP routing mechanisms will deliver packets destined for a mobile node's home address to its home link. Mobile nodes can have multiple home addresses, for instance when there are multiple home prefixes on the home link.

home subnet prefix

The IP subnet prefix corresponding to a mobile node's home address.

home link

The link on which a mobile node's home subnet prefix is defined.

mobile node

A node that can change its point of attachment from one link to another, while still being reachable via its home address.

movement

A change in a mobile node's point of attachment to the Internet such that it is no longer connected to the same link as it was previously. If a mobile node is not currently attached to its home link, the mobile node is said to be "away from home".

L2 handover

A process by which the mobile node changes from one link-layer connection to another. For example, a change of wireless access point is an L2 handover.

L3 handover

Subsequent to an L2 handover, a mobile node detects a change in an on-link subnet prefix that would require a change in the primary care-of address. For example, a change of access router subsequent to a change of wireless access point typically results in an L3 handover.

correspondent node

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

foreign subnet prefix

Any IP subnet prefix other than the mobile node's home subnet prefix.

foreign link

Any link other than the mobile node's home link.

care-of address

A unicast routable address associated with a mobile node while visiting a foreign link; the subnet prefix of this IP address is a foreign subnet prefix. Among the multiple care-of addresses that a mobile node may have at any given time (e.g., with different subnet prefixes), the one registered with the mobile node's home agent for a given home address is called its "primary" care-of address.

home agent

A router on a mobile node's home link with which the mobile node has registered its current care-of address. While the mobile node is away from home, the home agent intercepts packets on the home link destined to the mobile node's home address, encapsulates them, and tunnels them to the mobile node's registered care-of address.

binding

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

registration

The process during which a mobile node sends a Binding Update to its home agent or a correspondent node, causing a binding for the mobile node to be registered.

mobility message

A message containing a Mobility Header (see Section 6.1).

binding authorization

Correspondent registration needs to be authorized to allow the recipient to believe that the sender has the right to specify a new binding.

return routability procedure

The return routability procedure authorizes registrations by the use of a cryptographic token exchange.

correspondent registration

A return routability procedure followed by a registration, run between the mobile node and a correspondent node.

home registration

A registration between the mobile node and its home agent, authorized by the use of IPsec.

nonce

Nonces are random numbers used internally by the correspondent node in the creation of keygen tokens related to the return routability procedure. The nonces are not specific to a mobile node, and are kept secret within the correspondent node.

nonce index

A nonce index is used to indicate which nonces have been used when creating keygen token values, without revealing the nonces themselves.

cookie

A cookie is a random number used by a mobile node to prevent spoofing by a bogus correspondent node in the return routability procedure.

care-of init cookie

A cookie sent to the correspondent node in the Care-of Test Init message, to be returned in the Care-of Test message.

home init cookie

A cookie sent to the correspondent node in the Home Test Init message, to be returned in the Home Test message.

keygen token

A keygen token is a number supplied by a correspondent node in the return routability procedure to enable the mobile node to compute the necessary binding management key for authorizing a Binding Update.

care-of keygen token

A keygen token sent by the correspondent node in the Care-of Test message.

home keygen token

A keygen token sent by the correspondent node in the Home Test message.

binding management key (Kbm)

A binding management key (Kbm) is a key used for authorizing a binding cache management message (e.g., Binding Update or Binding Acknowledgement). Return routability provides a way to create a binding management key.

4. Overview of Mobile IPv6

4.1. Basic Operation

A mobile node is always expected to be addressable at its home address, whether it is currently attached to its home link or is away from home. The "home address" is an IP address assigned to the mobile node within its home subnet prefix on its home link. While a

mobile node is at home, packets addressed to its home address are routed to the mobile node's home link, using conventional Internet routing mechanisms.

While a mobile node is attached to some foreign link away from home, it is also addressable at one or more care-of addresses. A care-of address is an IP address associated with a mobile node that has the subnet prefix of a particular foreign link. The mobile node can acquire its care-of address through conventional IPv6 mechanisms, such as stateless or stateful auto-configuration. As long as the mobile node stays in this location, packets addressed to this care-of address will be routed to the mobile node. The mobile node may also accept packets from several care-of addresses, such as when it is moving but still reachable at the previous link.

The association between a mobile node's home address and care-of address is known as a "binding" for the mobile node. While away from home, a mobile node registers its primary care-of address with a router on its home link, requesting this router to function as the "home agent" for the mobile node. The mobile node performs this binding registration by sending a "Binding Update" message to the home agent. The home agent replies to the mobile node by returning a "Binding Acknowledgement" message. The operation of the mobile node is specified in Section 11, and the operation of the home agent is specified in Section 10.

Any node communicating with a mobile node is referred to in this document as a "correspondent node" of the mobile node, and may itself be either a stationary node or a mobile node. Mobile nodes can provide information about their current location to correspondent nodes. This happens through the correspondent registration. As a part of this procedure, a return routability test is performed in order to authorize the establishment of the binding. The operation of the correspondent node is specified in Section 9.

There are two possible modes for communications between the mobile node and a correspondent node. The first mode, bidirectional tunneling, does not require Mobile IPv6 support from the correspondent node and is available even if the mobile node has not registered its current binding with the correspondent node. Packets from the correspondent node are routed to the home agent and then tunneled to the mobile node. Packets to the correspondent node are tunneled from the mobile node to the home agent ("reverse tunneled") and then routed normally from the home network to the correspondent node. In this mode, the home agent uses proxy Neighbor Discovery to intercept any IPv6 packets addressed to the mobile node's home

address (or home addresses) on the home link. Each intercepted packet is tunneled to the mobile node's primary care-of address. This tunneling is performed using IPv6 encapsulation [15].

The second mode, "route optimization", requires the mobile node to register its current binding at the correspondent node. Packets from the correspondent node can be routed directly to the care-of address of the mobile node. When sending a packet to any IPv6 destination, the correspondent node checks its cached bindings for an entry for the packet's destination address. If a cached binding for this destination address is found, the node uses a new type of IPv6 routing header [11] (see Section 6.4) to route the packet to the mobile node by way of the care-of address indicated in this binding.

Routing packets directly to the mobile node's care-of address allows the shortest communications path to be used. It also eliminates congestion at the mobile node's home agent and home link. In addition, the impact of any possible failure of the home agent or networks on the path to or from it is reduced.

When routing packets directly to the mobile node, the correspondent node sets the Destination Address in the IPv6 header to the care-of address of the mobile node. A new type of IPv6 routing header (see Section 6.4) is also added to the packet to carry the desired home address. Similarly, the mobile node sets the Source Address in the packet's IPv6 header to its current care-of addresses. The mobile node adds a new IPv6 "Home Address" destination option (see Section 6.3) to carry its home address. The inclusion of home addresses in these packets makes the use of the care-of address transparent above the network layer (e.g., at the transport layer).

Mobile IPv6 also provides support for multiple home agents, and a limited support for the reconfiguration of the home network. In these cases, the mobile node may not know the IP address of its own home agent, and even the home subnet prefixes may change over time. A mechanism, known as "dynamic home agent address discovery" allows a mobile node to dynamically discover the IP address of a home agent on its home link, even when the mobile node is away from home. Mobile nodes can also learn new information about home subnet prefixes through the "mobile prefix discovery" mechanism. These mechanisms are described starting from Section 6.5.

4.2. New IPv6 Protocol

Mobile IPv6 defines a new IPv6 protocol, using the Mobility Header (see Section 6.1). This Header is used to carry the following messages:

Home Test Init

Home Test

Care-of Test Init

Care-of Test

These four messages are used to perform the return routability procedure from the mobile node to a correspondent node. This ensures authorization of subsequent Binding Updates, as described in Section 5.2.5.

Binding Update

A Binding Update is used by a mobile node to notify a correspondent node or the mobile node's home agent of its current binding. The Binding Update sent to the mobile node's home agent to register its primary care-of address is marked as a "home registration".

Binding Acknowledgement

A Binding Acknowledgement is used to acknowledge receipt of a Binding Update, if an acknowledgement was requested in the Binding Update, the binding update was sent to a home agent, or an error occurred.

Binding Refresh Request

A Binding Refresh Request is used by a correspondent node to request a mobile node to re-establish its binding with the correspondent node. This message is typically used when the cached binding is in active use but the binding's lifetime is close to expiration. The correspondent node may use, for instance, recent traffic and open transport layer connections as an indication of active use.

Binding Error

The Binding Error is used by the correspondent node to signal an error related to mobility, such as an inappropriate attempt to use the Home Address destination option without an existing binding.

4.3. New IPv6 Destination Option

Mobile IPv6 defines a new IPv6 destination option, the Home Address destination option. This option is described in detail in Section 6.3.

4.4. New IPv6 ICMP Messages

Mobile IPv6 also introduces four new ICMP message types, two for use in the dynamic home agent address discovery mechanism, and two for renumbering and mobile configuration mechanisms. As described in Section 10.5 and Section 11.4.1, the following two new ICMP message types are used for home agent address discovery:

- o Home Agent Address Discovery Request, described in Section 6.5.
- o Home Agent Address Discovery Reply, described in Section 6.6.

The next two message types are used for network renumbering and address configuration on the mobile node, as described in Section 10.6:

- o Mobile Prefix Solicitation, described in Section 6.7.
- o Mobile Prefix Advertisement, described in Section 6.8.

4.5. Conceptual Data Structure Terminology

This document describes the Mobile IPv6 protocol in terms of the following conceptual data structures:

Binding Cache

A cache of bindings for other nodes. This cache is maintained by home agents and correspondent nodes. The cache contains both "correspondent registration" entries (see Section 9.1) and "home registration" entries (see Section 10.1).

Binding Update List

This list is maintained by each mobile node. The list has an item for every binding that the mobile node has or is trying to establish with a specific other node. Both correspondent and home registrations are included in this list. Entries from the list are deleted as the lifetime of the binding expires. See Section 11.1.

Home Agents List

Home agents need to know which other home agents are on the same link. This information is stored in the Home Agents List, as described in more detail in Section 10.1. The list is used for informing mobile nodes during dynamic home agent address discovery.

4.6. Site-Local Addressability

This specification requires that home and care-of addresses MUST be unicast routable addresses. Site-local addresses may be usable on networks that are not connected to the Internet, but this specification does not define when such usage is safe and when it is not. Mobile nodes may not be aware of which site they are currently in, it is hard to prevent accidental attachment to other sites, and ambiguity of site-local addresses can cause problems if the home and visited networks use the same addresses. Therefore, site-local addresses SHOULD NOT be used as home or care-of addresses.

5. Overview of Mobile IPv6 Security

This specification provides a number of security features. These include the protection of Binding Updates both to home agents and correspondent nodes, the protection of mobile prefix discovery, and the protection of the mechanisms that Mobile IPv6 uses for transporting data packets.

Binding Updates are protected by the use of IPsec extension headers, or by the use of the Binding Authorization Data option. This option employs a binding management key, Kbm, which can be established through the return routability procedure. Mobile prefix discovery is protected through the use of IPsec extension headers. Mechanisms related to transporting payload packets - such as the Home Address destination option and type 2 routing header - have been specified in a manner which restricts their use in attacks.

5.1. Binding Updates to Home Agents

The mobile node and the home agent MUST use an IPsec security association to protect the integrity and authenticity of the Binding Updates and Acknowledgements. Both the mobile nodes and the home agents MUST support and SHOULD use the Encapsulating Security Payload (ESP) [6] header in transport mode and MUST use a non-NULL payload authentication algorithm to provide data origin authentication, connectionless integrity and optional anti-replay protection. Note that Authentication Header (AH) [5] is also possible but for brevity not discussed in this specification.

In order to protect messages exchanged between the mobile node and the home agent with IPsec, appropriate security policy database entries must be created. A mobile node must be prevented from using its security association to send a Binding Update on behalf of another mobile node using the same home agent. This MUST be achieved by having the home agent check that the given home address has been used with the right security association. Such a check is provided in the IPsec processing, by having the security policy database entries unequivocally identify a single security association for protecting Binding Updates between any given home address and home agent. In order to make this possible, it is necessary that the home address of the mobile node is visible in the Binding Updates and Acknowledgements. The home address is used in these packets as a source or destination, or in the Home Address Destination option or the type 2 routing header.

As with all IPsec security associations in this specification, manual configuration of security associations MUST be supported. The used shared secrets MUST be random and unique for different mobile nodes, and MUST be distributed off-line to the mobile nodes.

Automatic key management with IKE [9] MAY be supported. When IKE is used, either the security policy database entries or the Mobile IPv6 processing MUST unequivocally identify the IKE phase 1 credentials which can be used to authorize the creation of security associations for protecting Binding Updates for a particular home address. How these mappings are maintained is outside the scope of this specification, but they may be maintained, for instance, as a locally administered table in the home agent. If the phase 1 identity is a Fully Qualified Domain Name (FQDN), secure forms of DNS may also be used.

Section 11.3.2 discusses how IKE connections to the home agent need a careful treatment of the addresses used for transporting IKE. This is necessary to ensure that a Binding Update is not needed before the IKE exchange which is needed for securing the Binding Update.

When IKE version 1 is used with preshared secret authentication between the mobile node and the home agent, aggressive mode MUST be used.

The ID_IPV6_ADDR Identity Payload MUST NOT be used in IKEv1 phase 1.

Reference [21] contains a more detailed description and examples on using IPsec to protect the communications between the mobile node and the home agent.

5.2. Binding Updates to Correspondent Nodes

The protection of Binding Updates sent to correspondent nodes does not require the configuration of security associations or the existence of an authentication infrastructure between the mobile nodes and correspondent nodes. Instead, a method called the return routability procedure is used to assure that the right mobile node is sending the message. This method does not protect against attackers who are on the path between the home network and the correspondent node. However, attackers in such a location are capable of performing the same attacks even without Mobile IPv6. The main advantage of the return routability procedure is that it limits the potential attackers to those having an access to one specific path in the Internet, and avoids forged Binding Updates from anywhere else in the Internet. For a more in depth explanation of the security properties of the return routability procedure, see Section 15.

The integrity and authenticity of the Binding Updates messages to correspondent nodes is protected by using a keyed-hash algorithm. The binding management key, K_{bm}, is used to key the hash algorithm for this purpose. K_{bm} is established using data exchanged during the return routability procedure. The data exchange is accomplished by use of node keys, nonces, cookies, tokens, and certain cryptographic functions. Section 5.2.5 outlines the basic return routability procedure. Section 5.2.6 shows how the results of this procedure are used to authorize a Binding Update to a correspondent node.

5.2.1. Node Keys

Each correspondent node has a secret key, K_{cn}, called the "node key", which it uses to produce the keygen tokens sent to the mobile nodes. The node key MUST be a random number, 20 octets in length. The node key allows the correspondent node to verify that the keygen tokens used by the mobile node in authorizing a Binding Update are indeed its own. This key MUST NOT be shared with any other entity.

A correspondent node MAY generate a fresh node key at any time; this avoids the need for secure persistent key storage. Procedures for optionally updating the node key are discussed later in Section 5.2.7.

5.2.2. Nonces

Each correspondent node also generates nonces at regular intervals. The nonces should be generated by using a random number generator that is known to have good randomness properties [1]. A correspondent node may use the same K_{cn} and nonce with all the mobiles it is in communication with.

Each nonce is identified by a nonce index. When a new nonce is generated, it must be associated with a new nonce index; this may be done, for example, by incrementing the value of the previous nonce index, if the nonce index is used as an array pointer into a linear array of nonces. However, there is no requirement that nonces be stored that way, or that the values of subsequent nonce indices have any particular relationship to each other. The index value is communicated in the protocol, so that if a nonce is replaced by new nonce during the run of a protocol, the correspondent node can distinguish messages that should be checked against the old nonce from messages that should be checked against the new nonce. Strictly speaking, indices are not necessary in the authentication, but allow the correspondent node to efficiently find the nonce value that it used in creating a keygen token.

Correspondent nodes keep both the current nonce and a small set of valid previous nonces whose lifetime has not yet expired. Expired values **MUST** be discarded, and messages using stale or unknown indices will be rejected.

The specific nonce index values cannot be used by mobile nodes to determine the validity of the nonce. Expected validity times for the nonces values and the procedures for updating them are discussed later in Section 5.2.7.

A nonce is an octet string of any length. The recommended length is 64 bits.

5.2.3. Cookies and Tokens

The return routability address test procedure uses cookies and keygen tokens as opaque values within the test init and test messages, respectively.

- o The "home init cookie" and "care-of init cookie" are 64 bit values sent to the correspondent node from the mobile node, and later returned to the mobile node. The home init cookie is sent in the Home Test Init message, and returned in the Home Test message. The care-of init cookie is sent in the Care-of Test Init message, and returned in the Care-of Test message.
- o The "home keygen token" and "care-of keygen token" are 64-bit values sent by the correspondent node to the mobile node via the home agent (via the Home Test message) and the care-of address (by the Care-of Test message), respectively.

The mobile node should set the home init or care-of init cookie to a newly generated random number in every Home or Care-of Test Init message it sends. The cookies are used to verify that the Home Test or Care-of Test message matches the Home Test Init or Care-of Test Init message, respectively. These cookies also serve to ensure that parties who have not seen the request cannot spoof responses.

Home and care-of keygen tokens are produced by the correspondent node based on its currently active secret key (Kcn) and nonces, as well as the home or care-of address (respectively). A keygen token is valid as long as both the secret key (Kcn) and the nonce used to create it are valid.

5.2.4. Cryptographic Functions


In this specification, the function used to compute hash values is SHA1 [20]. Message Authentication Codes (MACs) are computed using HMAC_SHA1 [25, 20]. $HMAC_SHA1(K,m)$ denotes such a MAC computed on message m with key K .

5.2.5. Return Routability Procedure

The Return Routability Procedure enables the correspondent node to obtain some reasonable assurance that the mobile node is in fact addressable at its claimed care-of address as well as at its home address. Only with this assurance is the correspondent node able to accept Binding Updates from the mobile node which would then instruct the correspondent node to direct that mobile node's data traffic to its claimed care-of address.

This is done by testing whether packets addressed to the two claimed addresses are routed to the mobile node. The mobile node can pass the test only if it is able to supply proof that it received certain data (the "keygen tokens") which the correspondent node sends to those addresses. These data are combined by the mobile node into a binding management key, denoted K_{bm} .

The figure below shows the message flow for the return routability procedure.



The Home and Care-of Test Init messages are sent at the same time. The procedure requires very little processing at the correspondent node, and the Home and Care-of Test messages can be returned quickly, perhaps nearly simultaneously. These four messages form the return routability procedure.

Home Test Init

A mobile node sends a Home Test Init message to the correspondent node (via the home agent) to acquire the home keygen token. The contents of the message can be summarized as follows:

- * Source Address = home address
- * Destination Address = correspondent
- * Parameters:
 - + home init cookie

The Home Test Init message conveys the mobile node's home address to the correspondent node. The mobile node also sends along a home init cookie that the correspondent node must return later. The Home Test Init message is reverse tunneled through the home agent. (The headers and addresses related to reverse tunneling have been omitted from the above discussion of the message contents.) The mobile node remembers these cookie values to obtain some assurance that its protocol messages are being processed by the desired correspondent node.

Care-of Test Init

The mobile node sends a Care-of Test Init message to the correspondent node (directly, not via the home agent) to acquire the care-of keygen token. The contents of this message can be summarized as follows:

- * Source Address = care-of address
- * Destination Address = correspondent
- * Parameters:
 - + care-of init cookie

The Care-of Test Init message conveys the mobile node's care-of address to the correspondent node. The mobile node also sends along a care-of init cookie that the correspondent node must return later. The Care-of Test Init message is sent directly to the correspondent node.

Home Test

The Home Test message is sent in response to a Home Test Init message. It is sent via the home agent. The contents of the message are:

- * Source Address = correspondent
- * Destination Address = home address
- * Parameters:
 - + home init cookie
 - + home keygen token
 - + home nonce index

When the correspondent node receives the Home Test Init message, it generates a home keygen token as follows:

```
home keygen token :=  
  First (64, HMAC_SHA1 (Kcn, (home address | nonce | 0)))
```

where | denotes concatenation. The final "0" inside the HMAC_SHA1 function is a single zero octet, used to distinguish home and care-of cookies from each other.

The home keygen token is formed from the first 64 bits of the MAC. The home keygen token tests that the mobile node can receive were messages sent to its home address. Kcn is used in the production of home keygen token in order to allow the correspondent node to verify that it generated the home and care-of nonces, without forcing the correspondent node to remember a list of all tokens it has handed out.

The Home Test message is sent to the mobile node via the home network, where it is presumed that the home agent will tunnel the message to the mobile node. This means that the mobile node needs to already have sent a Binding Update to the home agent, so that the home agent will have received and authorized the new care-of address for the mobile node before the return routability procedure. For improved security, the data passed between the home agent and the mobile node is made immune to inspection and passive attacks. Such protection is gained by encrypting the home keygen token as it is tunneled from the home agent to the mobile node as specified in Section 10.4.6. The security properties of this additional security are discussed in Section 15.4.1.

The home init cookie from the mobile node is returned in the Home Test message, to ensure that the message comes from a node on the route between the home agent and the correspondent node.

The home nonce index is delivered to the mobile node to later allow the correspondent node to efficiently find the nonce value that it used in creating the home keygen token.

Care-of Test

This message is sent in response to a Care-of Test Init message. This message is not sent via the home agent, it is sent directly to the mobile node. The contents of the message are:

- * Source Address = correspondent
- * Destination Address = care-of address

* Parameters:

- + care-of init cookie
- + care-of keygen token
- + care-of nonce index

When the correspondent node receives the Care-of Test Init message, it generates a care-of keygen token as follows:

```
care-of keygen token :=
  First (64, HMAC_SHA1 (Kcn, (care-of address | nonce | 1)))
```

Here, the final "1" inside the HMAC_SHA1 function is a single octet containing the hex value 0x01, and is used to distinguish home and care-of cookies from each other. The keygen token is formed from the first 64 bits of the MAC, and sent directly to the mobile node at its care-of address. The care-of init cookie from the Care-of Test Init message is returned to ensure that the message comes from a node on the route to the correspondent node.

The care-of nonce index is provided to identify the nonce used for the care-of keygen token. The home and care-of nonce indices MAY be the same, or different, in the Home and Care-of Test messages.

When the mobile node has received both the Home and Care-of Test messages, the return routability procedure is complete. As a result of the procedure, the mobile node has the data it needs to send a Binding Update to the correspondent node. The mobile node hashes the tokens together to form a 20 octet binding key Kbm:

```
Kbm = SHA1 (home keygen token | care-of keygen token)
```


A Binding Update may also be used to delete a previously established binding (Section 6.1.7). In this case, the care-of keygen token is not used. Instead, the binding management key is generated as follows:

```
Kbm = SHA1(home keygen token)
```

Note that the correspondent node does not create any state specific to the mobile node, until it receives the Binding Update from that mobile node. The correspondent node does not maintain the value for the binding management key Kbm; it creates Kbm when given the nonce indices and the mobile node's addresses.

5.2.6. Authorizing Binding Management Messages

After the mobile node has created the binding management key (Kbm), it can supply a verifiable Binding Update to the correspondent node. This section provides an overview of this registration. The below figure shows the message flow.



Binding Update

To authorize a Binding Update, the mobile node creates a binding management key Kbm from the keygen tokens as described in the previous section. The contents of the Binding Update include the following:

- * Source Address = care-of address
- * Destination Address = correspondent
- * Parameters:
 - + home address (within the Home Address destination option if different from the Source Address)
 - + sequence number (within the Binding Update message header)
 - + home nonce index (within the Nonce Indices option)
 - + care-of nonce index (within the Nonce Indices option)
 - + First (96, HMAC_SHA1 (Kbm, (care-of address | correspondent | BU)))

The Binding Update contains a Nonce Indices option, indicating to the correspondent node which home and care-of nonces to use to recompute Kbm, the binding management key. The MAC is computed as described in Section 6.2.7, using the correspondent node's address as the destination address and the Binding Update message itself ("BU" above) as the MH Data.

Once the correspondent node has verified the MAC, it can create a Binding Cache entry for the mobile.

Binding Acknowledgement

The Binding Update is in some cases acknowledged by the correspondent node. The contents of the message are as follows:

- * Source Address = correspondent
- * Destination Address = care-of address
- * Parameters:
 - + sequence number (within the Binding Update message header)
 - + First (96, HMAC_SHA1 (Kbm, (care-of address | correspondent | BA)))

The Binding Acknowledgement contains the same sequence number as the Binding Update. The MAC is computed as described in Section 6.2.7, using the correspondent node's address as the destination address and the message itself ("BA" above) as the MH Data.

Bindings established with correspondent nodes using keys created by way of the return routability procedure MUST NOT exceed MAX_RR_BINDING_LIFETIME seconds (see Section 12).

The value in the Source Address field in the IPv6 header carrying the Binding Update is normally also the care-of address which is used in the binding. However, a different care-of address MAY be specified by including an Alternate Care-of Address mobility option in the Binding Update (see Section 6.2.5). When such a message is sent to the correspondent node and the return routability procedure is used as the authorization method, the Care-of Test Init and Care-of Test messages MUST have been performed for the address in the Alternate Care-of Address option (not the Source Address). The nonce indices and MAC value MUST be based on information gained in this test.

Binding Updates may also be sent to delete a previously established binding. In this case, generation of the binding management key depends exclusively on the home keygen token and the care-of nonce index is ignored.

5.2.7. Updating Node Keys and Nonces

Correspondent nodes generate nonces at regular intervals. It is recommended to keep each nonce (identified by a nonce index) acceptable for at least `MAX_TOKEN_LIFETIME` seconds (see Section 12) after it has been first used in constructing a return routability message response. However, the correspondent node **MUST NOT** accept nonces beyond `MAX_NONCE_LIFETIME` seconds (see Section 12) after the first use. As the difference between these two constants is 30 seconds, a convenient way to enforce the above lifetimes is to generate a new nonce every 30 seconds. The node can then continue to accept tokens that have been based on the last 8 (`MAX_NONCE_LIFETIME / 30`) nonces. This results in tokens being acceptable `MAX_TOKEN_LIFETIME` to `MAX_NONCE_LIFETIME` seconds after they have been sent to the mobile node, depending on whether the token was sent at the beginning or end of the first 30 second period. Note that the correspondent node may also attempt to generate new nonces on demand, or only if the old nonces have been used. This is possible, as long as the correspondent node keeps track of how long a time ago the nonces were used for the first time, and does not generate new nonces on every return routability request.

Due to resource limitations, rapid deletion of bindings, or reboots the correspondent node may not in all cases recognize the nonces that the tokens were based on. If a nonce index is unrecognized, the correspondent node replies with an error code in the Binding Acknowledgement (either 136, 137, or 138 as discussed in Section 6.1.8). The mobile node can then retry the return routability procedure.

An update of `Kcn` **SHOULD** be done at the same time as an update of a nonce, so that nonce indices can identify both the nonce and the key. Old `Kcn` values have to be therefore remembered as long as old nonce values.

Given that the tokens are normally expected to be usable for `MAX_TOKEN_LIFETIME` seconds, the mobile node **MAY** use them beyond a single run of the return routability procedure until `MAX_TOKEN_LIFETIME` expires. After this the mobile node **SHOULD NOT** use the tokens. A fast moving mobile node **MAY** reuse a recent home keygen token from a correspondent node when moving to a new location, and just acquire a new care-of keygen token to show routability in the new location.

While this does not save the number of round-trips due to the simultaneous processing of home and care-of return routability tests, there are fewer messages being exchanged, and a potentially long round-trip through the home agent is avoided. Consequently, this optimization is often useful. A mobile node that has multiple home addresses, MAY also use the same care-of keygen token for Binding Updates concerning all of these addresses.

5.2.8. Preventing Replay Attacks

The return routability procedure also protects the participants against replayed Binding Updates through the use of the sequence number and a MAC. Care must be taken when removing bindings at the correspondent node, however. Correspondent nodes must retain bindings and the associated sequence number information at least as long as the nonces used in the authorization of the binding are still valid. Alternatively, if memory is very constrained, the correspondent node MAY invalidate the nonces that were used for the binding being deleted (or some larger group of nonces that they belong to). This may, however, impact the ability to accept Binding Updates from mobile nodes that have recently received keygen tokens. This alternative is therefore recommended only as a last measure.

5.3. Dynamic Home Agent Address Discovery

No security is required for dynamic home agent address discovery.

5.4. Mobile Prefix Discovery

The mobile node and the home agent SHOULD use an IPsec security association to protect the integrity and authenticity of the Mobile Prefix Solicitations and Advertisements. Both the mobile nodes and the home agents MUST support and SHOULD use the Encapsulating Security Payload (ESP) header in transport mode with a non-NULL payload authentication algorithm to provide data origin authentication, connectionless integrity and optional anti-replay protection.

5.5. Payload Packets

Payload packets exchanged with mobile nodes can be protected in the usual manner, in the same way as stationary hosts can protect them. However, Mobile IPv6 introduces the Home Address destination option, a routing header, and tunneling headers in the payload packets. In the following we define the security measures taken to protect these, and to prevent their use in attacks against other parties.

This specification limits the use of the Home Address destination option to the situation where the correspondent node already has a Binding Cache entry for the given home address. This avoids the use of the Home Address option in attacks described in Section 15.1.

Mobile IPv6 uses a Mobile IPv6 specific type of a routing header. This type provides the necessary functionality but does not open vulnerabilities discussed in Section 15.1.

Tunnels between the mobile node and the home agent are protected by ensuring proper use of source addresses, and optional cryptographic protection. The mobile node verifies that the outer IP address corresponds to its home agent. The home agent verifies that the outer IP address corresponds to the current location of the mobile node (Binding Updates sent to the home agents are secure). The home agent identifies the mobile node through the source address of the inner packet. (Typically, this is the home address of the mobile node, but it can also be a link-local address, as discussed in Section 10.4.2. To recognize the latter type of addresses, the home agent requires that the Link-Local Address Compatibility (L) was set in the Binding Update.) These measures protect the tunnels against vulnerabilities discussed in Section 15.1.

For traffic tunneled via the home agent, additional IPsec ESP encapsulation MAY be supported and used. If multicast group membership control protocols or stateful address autoconfiguration protocols are supported, payload data protection MUST be supported.

6. New IPv6 Protocol, Message Types, and Destination Option

6.1. Mobility Header


The Mobility Header is an extension header used by mobile nodes, correspondent nodes, and home agents in all messaging related to the creation and management of bindings. The subsections within this section describe the message types that may be sent using the Mobility Header.

Mobility Header messages MUST NOT be sent with a type 2 routing header, except as described in Section 9.5.4 for Binding Acknowledgement. Mobility Header messages also MUST NOT be used with a Home Address destination option, except as described in Section 11.7.1 and Section 11.7.2 for Binding Update. Binding Update List or Binding Cache information (when present) for the destination MUST NOT be used in sending Mobility Header messages. That is, Mobility Header messages bypass both the Binding Cache check described in Section 9.3.2 and the Binding Update List check described in Section

11.3.1 which are normally performed for all packets. This applies even to messages sent to or from a correspondent node which is itself a mobile node.

6.1.1. Format

The Mobility Header is identified by a Next Header value of 135 in the immediately preceding header, and has the following format:



Payload Proto

8-bit selector. Identifies the type of header immediately following the Mobility Header. Uses the same values as the IPv6 Next Header field [11].

This field is intended to be used by a future extension (see Appendix B.1).

Implementations conforming to this specification SHOULD set the payload protocol type to IPPROTO_NONE (59 decimal).

Header Len

8-bit unsigned integer, representing the length of the Mobility Header in units of 8 octets, excluding the first 8 octets.

The length of the Mobility Header MUST be a multiple of 8 octets.

MH Type

8-bit selector. Identifies the particular mobility message in question. Current values are specified in Section 6.1.2 and onward. An unrecognized MH Type field causes an error indication to be sent.

Reserved

8-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Checksum

16-bit unsigned integer. This field contains the checksum of the Mobility Header. The checksum is calculated from the octet string consisting of a "pseudo-header" followed by the entire Mobility Header starting with the Payload Proto field. The checksum is the 16-bit one's complement of the one's complement sum of this string.

The pseudo-header contains IPv6 header fields, as specified in Section 8.1 of RFC 2460 [11]. The Next Header value used in the pseudo-header is 2. The addresses used in the pseudo-header are the addresses that appear in the Source and Destination Address fields in the IPv6 packet carrying the Mobility Header.

Note that the procedures of calculating upper layer checksums while away from home described in Section 11.3.1 apply even for the Mobility Header. If a mobility message has a Home Address destination option, then the checksum calculation uses the home address in this option as the value of the IPv6 Source Address field. The type 2 routing header is treated as explained in [11].

The Mobility Header is considered as the upper layer protocol for the purposes of calculating the pseudo-header. The Upper-Layer Packet Length field in the pseudo-header MUST be set to the total length of the Mobility Header.

For computing the checksum, the checksum field is set to zero.

Message Data

A variable length field containing the data specific to the indicated Mobility Header type.

Mobile IPv6 also defines a number of "mobility options" for use within these messages; if included, any options MUST appear after the fixed portion of the message data specified in this document. The presence of such options will be indicated by the Header Len field within the message. When the Header Len value is greater than the length required for the message specified here, the remaining octets are interpreted as mobility options. These options include padding options that can be used to ensure that other options are aligned


properly, and that the total length of the message is divisible by 8. The encoding and format of defined options are described in Section 6.2.

Alignment requirements for the Mobility Header are the same as for any IPv6 protocol Header. That is, they MUST be aligned on an 8-octet boundary.

6.1.2. Binding Refresh Request Message

The Binding Refresh Request (BRR) message requests a mobile node to update its mobility binding. This message is sent by correspondent nodes according to the rules in Section 9.5.5. When a mobile node receives a packet containing a Binding Refresh Request message it processes the message according to the rules in Section 11.7.4.

The Binding Refresh Request message uses the MH Type value 0. When this value is indicated in the MH Type field, the format of the Message Data field in the Mobility Header is as follows:



Reserved

16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Mobility Options


Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The encoding and format of defined options are described in Section 6.2. The receiver MUST ignore and skip any options which it does not understand.

There MAY be additional information, associated with this Binding Refresh Request message that need not be present in all Binding Refresh Request messages sent. Mobility options allow future extensions to the format of the Binding Refresh Request message to be defined. This specification does not define any options valid for the Binding Refresh Request message.

If no actual options are present in this message, no padding is necessary and the Header Len field will be set to 0.

6.1.3. Home Test Init Message

A mobile node uses the Home Test Init (HoTI) message to initiate the return routability procedure and request a home keygen token from a correspondent node (see Section 11.6.1). The Home Test Init message uses the MH Type value 1. When this value is indicated in the MH Type field, the format of the Message Data field in the Mobility Header is as follows:



Reserved

16-bit field reserved for future use. This value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Home Init Cookie

64-bit field which contains a random value, the home init cookie.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The receiver


MUST ignore and skip any options which it does not understand. This specification does not define any options valid for the Home Test Init message.

If no actual options are present in this message, no padding is necessary and the Header Len field will be set to 1.

This message is tunneled through the home agent when the mobile node is away from home. Such tunneling SHOULD employ IPsec ESP in tunnel mode between the home agent and the mobile node. This protection is indicated by the IPsec security policy database. The protection of Home Test Init messages is unrelated to the requirement to protect regular payload traffic, which MAY use such tunnels as well.

6.1.4. Care-of Test Init Message

A mobile node uses the Care-of Test Init (CoTI) message to initiate the return routability procedure and request a care-of keygen token from a correspondent node (see Section 11.6.1). The Care-of Test Init message uses the MH Type value 2. When this value is indicated in the MH Type field, the format of the Message Data field in the Mobility Header is as follows:



Reserved

16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Care-of Init Cookie

64-bit field which contains a random value, the care-of init cookie.

Care-of Init Cookie

64-bit field which contains the care-of init cookie.

Care-of Keygen Token

This field contains the 64 bit care-of keygen token used in the return routability procedure.

Mobility Options


Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The receiver MUST ignore and skip any options which it does not understand. This specification does not define any options valid for the Care-of Test message.

If no actual options are present in this message, no padding is necessary and the Header Len field will be set to 2.

6.1.7. Binding Update Message

The Binding Update (BU) message is used by a mobile node to notify other nodes of a new care-of address for itself. Binding Updates are sent as described in Section 11.7.1 and Section 11.7.2.

The Binding Update uses the MH Type value 5. When this value is indicated in the MH Type field, the format of the Message Data field in the Mobility Header is as follows:



Acknowledge (A)

The Acknowledge (A) bit is set by the sending mobile node to request a Binding Acknowledgement (Section 6.1.8) be returned upon receipt of the Binding Update.

Home Registration (H)

The Home Registration (H) bit is set by the sending mobile node to request that the receiving node should act as this node's home agent. The destination of the packet carrying this message MUST be that of a router sharing the same subnet prefix as the home address of the mobile node in the binding.

Link-Local Address Compatibility (L)

The Link-Local Address Compatibility (L) bit is set when the home address reported by the mobile node has the same interface identifier as the mobile node's link-local address.

Key Management Mobility Capability (K)

If this bit is cleared, the protocol used for establishing the IPsec security associations between the mobile node and the home agent does not survive movements. It may then have to be rerun. (Note that the IPsec security associations themselves are expected to survive movements.) If manual IPsec configuration is used, the bit MUST be cleared.

This bit is valid only in Binding Updates sent to the home agent, and MUST be cleared in other Binding Updates. Correspondent nodes MUST ignore this bit.

Reserved

These fields are unused. They MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Sequence

A 16-bit unsigned integer used by the receiving node to sequence Binding Updates and by the sending node to match a returned Binding Acknowledgement with this Binding Update.

Lifetime

16-bit unsigned integer. The number of time units remaining before the binding MUST be considered expired. A value of zero indicates that the Binding Cache entry for the mobile node MUST be deleted. (In this case the specified care-of address MUST also be set equal to the home address.) One time unit is 4 seconds.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The encoding and format of defined options are described in Section 6.2. The receiver MUST ignore and skip any options which it does not understand.

The following options are valid in a Binding Update:

- * Binding Authorization Data option (this option is mandatory in Binding Updates sent to a correspondent node)
- * Nonce Indices option.
- * Alternate Care-of Address option

If no options are present in this message, 4 octets of padding are necessary and the Header Len field will be set to 1.

The care-of address is specified either by the Source Address field in the IPv6 header or by the Alternate Care-of Address option, if present. The care-of address MUST be a unicast routable address. IPv6 Source Address MUST be a topologically correct source address. Binding Updates for a care-of address which is not a unicast routable address MUST be silently discarded. Similarly, the Binding Update MUST be silently discarded if the care-of address appears as a home address in an existing Binding Cache entry, with its current location creating a circular reference back to the home address specified in the Binding Update (possibly through additional entries).

The deletion of a binding can be indicated by setting the Lifetime field to 0 and by setting the care-of address equal to the home address. In deletion, the generation of the binding management key depends exclusively on the home keygen token, as explained in Section 5.2.5. (Note that while the senders are required to set both the Lifetime field to 0 and the care-of address equal to the home address, Section 9.5.1 rules for receivers are more liberal, and interpret either condition as a deletion.)

Correspondent nodes SHOULD NOT delete the Binding Cache entry before the lifetime expires, if any application hosted by the correspondent node is still likely to require communication with the mobile node. A Binding Cache entry that is de-allocated prematurely might cause subsequent packets to be dropped from the mobile node, if they contain the Home Address destination option. This situation is recoverable, since a Binding Error message is sent to the mobile node

Status

8-bit unsigned integer indicating the disposition of the Binding Update. Values of the Status field less than 128 indicate that the Binding Update was accepted by the receiving node. Values greater than or equal to 128 indicate that the Binding Update was rejected by the receiving node. The following Status values are currently defined:

- 0 Binding Update accepted
- 1 Accepted but prefix discovery necessary
- 128 Reason unspecified
- 129 Administratively prohibited
- 130 Insufficient resources
- 131 Home registration not supported
- 132 Not home subnet
- 133 Not home agent for this mobile node
- 134 Duplicate Address Detection failed
- 135 Sequence number out of window
- 136 Expired home nonce index
- 137 Expired care-of nonce index
- 138 Expired nonces
- 139 Registration type change disallowed

Up-to-date values of the Status field are to be specified in the IANA registry of assigned numbers [19].

Sequence

The Sequence Number in the Binding Acknowledgement is copied from the Sequence Number field in the Binding Update. It is used by the mobile node in matching this Binding Acknowledgement with an outstanding Binding Update.

Lifetime

The granted lifetime, in time units of 4 seconds, for which this node SHOULD retain the entry for this mobile node in its Binding Cache.

The value of this field is undefined if the Status field indicates that the Binding Update was rejected.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The encoding and format of defined options are described in Section 6.2. The receiver MUST ignore and skip any options which it does not understand.

There MAY be additional information, associated with this Binding Acknowledgement that need not be present in all Binding Acknowledgements sent. Mobility options allow future extensions to the format of the Binding Acknowledgement to be defined. The following options are valid for the Binding Acknowledgement:

- * Binding Authorization Data option (this option is mandatory in Binding Acknowledgements sent by a correspondent node, except where otherwise noted in Section 9.5.4)
- * Binding Refresh Advice option

If no options are present in this message, 4 octets of padding are necessary and the Header Len field will be set to 1.

6.1.9. Binding Error Message

The Binding Error (BE) message is used by the correspondent node to signal an error related to mobility, such as an inappropriate attempt to use the Home Address destination option without an existing binding; see Section 9.3.3 for details.

Mobility Options

Variable-length field of such length that the complete Mobility Header is an integer multiple of 8 octets long. This field contains zero or more TLV-encoded mobility options. The receiver MUST ignore and skip any options which it does not understand.

There MAY be additional information, associated with this Binding Error message that need not be present in all Binding Error messages sent. Mobility options allow future extensions to the format of the format of the Binding Error message to be defined. The encoding and format of defined options are described in Section 6.2. This specification does not define any options valid for the Binding Error message.

If no actual options are present in this message, no padding is necessary and the Header Len field will be set to 2.


6.2. Mobility Options

Mobility messages can include zero or more mobility options. This allows optional fields that may not be needed in every use of a particular Mobility Header, as well as future extensions to the format of the messages. Such options are included in the Message Data field of the message itself, after the fixed portion of the message data specified in the message subsections of Section 6.1.

The presence of such options will be indicated by the Header Len of the Mobility Header. If included, the Binding Authorization Data option (Section 6.2.7) MUST be the last option and MUST NOT have trailing padding. Otherwise, options can be placed in any order.

6.2.1. Format

Mobility options are encoded within the remaining space of the Message Data field of a mobility message, using a type-length-value (TLV) format as follows:



Option Type

8-bit identifier of the type of mobility option. When processing a Mobility Header containing an option for which the Option Type value is not recognized by the receiver, the receiver MUST quietly ignore and skip over the option, correctly handling any remaining options in the message.

Option Length

8-bit unsigned integer, representing the length in octets of the mobility option, not including the Option Type and Option Length fields.

Option Data

A variable length field that contains data specific to the option.

The following subsections specify the Option types which are currently defined for use in the Mobility Header.

Implementations MUST silently ignore any mobility options that they do not understand.

Mobility options may have alignment requirements. Following the convention in IPv6, these options are aligned in a packet so that multi-octet values within the Option Data field of each option fall on natural boundaries (i.e., fields of width n octets are placed at an integer multiple of n octets from the start of the header, for $n = 1, 2, 4, \text{ or } 8$) [11].

6.2.2. Pad1

The Pad1 option does not have any alignment requirements. Its format is as follows:

```


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

NOTE! the format of the Pad1 option is a special case - it has neither Option Length nor Option Data fields.

The Pad1 option is used to insert one octet of padding in the Mobility Options area of a Mobility Header. If more than one octet of padding is required, the PadN option, described next, should be used rather than multiple Pad1 options.

6.2.3. PadN


The PadN option does not have any alignment requirements. Its format is as follows:



The PadN option is used to insert two or more octets of padding in the Mobility Options area of a mobility message. For N octets of padding, the Option Length field contains the value N-2, and the Option Data consists of N-2 zero-valued octets. PadN Option data MUST be ignored by the receiver.

6.2.4. Binding Refresh Advice

The Binding Refresh Advice option has an alignment requirement of 2n. Its format is as follows:



The Binding Refresh Advice option is only valid in the Binding Acknowledgement, and only on Binding Acknowledgements sent from the mobile node's home agent in reply to a home registration. The Refresh Interval is measured in units of four seconds, and indicates remaining time until the mobile node SHOULD send a new home registration to the home agent. The Refresh Interval MUST be set to indicate a smaller time interval than the Lifetime value of the Binding Acknowledgement.


The Nonce Indices option is valid only in the Binding Update message sent to a correspondent node, and only when present together with a Binding Authorization Data option. When the correspondent node authorizes the Binding Update, it needs to produce home and care-of keygen tokens from its stored random nonce values.

The Home Nonce Index field tells the correspondent node which nonce value to use when producing the home keygen token.

The Care-of Nonce Index field is ignored in requests to delete a binding. Otherwise, it tells the correspondent node which nonce value to use when producing the care-of keygen token.

6.2.7. Binding Authorization Data

The Binding Authorization Data option does not have alignment requirements as such. However, since this option must be the last mobility option, an implicit alignment requirement is 8n + 2. The format of this option is as follows:



The Binding Authorization Data option is valid in the Binding Update and Binding Acknowledgement.

The Option Length field contains the length of the authenticator in octets.

The Authenticator field contains a cryptographic value which can be used to determine that the message in question comes from the right authority. Rules for calculating this value depends on the used authorization procedure.

For the return routability procedure, this option can appear in the Binding Update and Binding Acknowledgements. Rules for calculating the Authenticator value are the following:

```
Mobility Data = care-of address | correspondent | MH Data
Authenticator = First (96, HMAC_SHA1 (Kbm, Mobility Data))
```

Where | denotes concatenation. "Care-of address" is the care-of address which will be registered for the mobile node if the Binding Update succeeds, or the home address of the mobile node if this option is used in de-registration. Note also that this address might be different from the source address of the Binding Update message, if the Alternative Care-of Address mobility option is used, or when the lifetime of the binding is set to zero.

The "correspondent" is the IPv6 address of the correspondent node. Note that, if the message is sent to a destination which is itself mobile, the "correspondent" address may not be the address found in the Destination Address field of the IPv6 header; instead the home address from the type 2 Routing header should be used.


"MH Data" is the content of the Mobility Header, excluding the Authenticator field itself. The Authenticator value is calculated as if the Checksum field in the Mobility Header was zero. The Checksum in the transmitted packet is still calculated in the usual manner, with the calculated Authenticator being a part of the packet protected by the Checksum. Kbm is the binding management key, which is typically created using nonces provided by the correspondent node (see Section 9.4). Note that while the contents of a potential Home Address destination option are not covered in this formula, the rules for the calculation of the Kbm do take the home address in account. This ensures that the MAC will be different for different home addresses.

The first 96 bits from the MAC result are used as the Authenticator field.

6.3. Home Address Option

The Home Address option is carried by the Destination Option extension header (Next Header value = 60). It is used in a packet sent by a mobile node while away from home, to inform the recipient of the mobile node's home address.

The Home Address option is encoded in type-length-value (TLV) format as follows:



Option Type

201 = 0xC9

Option Length

8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields. This field MUST be set to 16.

Home Address

The home address of the mobile node sending the packet. This address MUST be a unicast routable address.

The alignment requirement [11] for the Home Address option is 8n+6.

The three highest-order bits of the Option Type field are encoded to indicate specific processing of the option [11]; for the Home Address option, these three bits are set to 110. This indicates the following processing requirements:

- o Any IPv6 node that does not recognize the Option Type must discard the packet, and if the packet's Destination Address was not a multicast address, return an ICMP Parameter Problem, Code 2, message to the packet's Source Address. The Pointer field in the ICMP message SHOULD point at the Option Type field. Otherwise, for multicast addresses, the ICMP message MUST NOT be sent.
- o The data within the option cannot change en route to the packet's final destination.

The Home Address option MUST be placed as follows:

- o After the routing header, if that header is present
- o Before the Fragment Header, if that header is present
- o Before the AH Header or ESP Header, if either one of those headers are present

For each IPv6 packet header, the Home Address Option MUST NOT appear more than once. However, an encapsulated packet [15] MAY contain a separate Home Address option associated with each encapsulating IP header.

The inclusion of a Home Address destination option in a packet affects the receiving node's processing of only this single packet. No state is created or modified in the receiving node as a result of receiving a Home Address option in a packet. In particular, the presence of a Home Address option in a received packet MUST NOT alter the contents of the receiver's Binding Cache and MUST NOT cause any changes in the routing of subsequent packets sent by this receiving node.

6.4. Type 2 Routing Header

Mobile IPv6 defines a new routing header variant, the type 2 routing header, to allow the packet to be routed directly from a correspondent to the mobile node's care-of address. The mobile node's care-of address is inserted into the IPv6 Destination Address field. Once the packet arrives at the care-of address, the mobile node retrieves its home address from the routing header, and this is used as the final destination address for the packet.

The new routing header uses a different type than defined for "regular" IPv6 source routing, enabling firewalls to apply different rules to source routed packets than to Mobile IPv6. This routing header type (type 2) is restricted to carry only one IPv6 address. All IPv6 nodes which process this routing header MUST verify that the address contained within is the node's own home address in order to prevent packets from being forwarded outside the node. The IP address contained in the routing header, since it is the mobile node's home address, MUST be a unicast routable address. Furthermore, if the scope of the home address is smaller than the scope of the care-of address, the mobile node MUST discard the packet (see Section 4.6).

6.4.1. Format

The type 2 routing header has the following format:

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Next Header | Hdr Ext Len=2 | Routing Type=2 | Segments Left=1 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Reserved                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Home Address                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Next Header

8-bit selector. Identifies the type of header immediately following the routing header. Uses the same values as the IPv6 Next Header field [11].

Hdr Ext Len

2 (8-bit unsigned integer); length of the routing header in 8-octet units, not including the first 8 octets.

Routing Type

2 (8-bit unsigned integer).

Segments Left

1 (8-bit unsigned integer).

Reserved

32-bit reserved field. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Home Address


The Home Address of the destination Mobile Node.

For a type 2 routing header, the Hdr Ext Len MUST be 2. The Segments Left value describes the number of route segments remaining; i.e., number of explicitly listed intermediate nodes still to be visited before reaching the final destination. Segments Left MUST be 1. The ordering rules for extension headers in an IPv6 packet are described in Section 4.1 of RFC 2460 [11]. The type 2 routing header defined for Mobile IPv6 follows the same ordering as other routing headers. If both a type 0 and a type 2 routing header are present, the type 2 routing header should follow the other routing header. A packet containing such nested encapsulation should be created as if the inner (type 2) routing header was constructed first and then treated as an original packet by the outer (type 0) routing header construction process.

In addition, the general procedures defined by IPv6 for routing headers suggest that a received routing header MAY be automatically "reversed" to construct a routing header for use in any response packets sent by upper-layer protocols, if the received packet is authenticated [6]. This MUST NOT be done automatically for type 2 routing headers.

6.5. ICMP Home Agent Address Discovery Request Message

The ICMP Home Agent Address Discovery Request message is used by a mobile node to initiate the dynamic home agent address discovery mechanism, as described in Section 11.4.1. The mobile node sends the Home Agent Address Discovery Request message to the Mobile IPv6 Home-Agents anycast address [16] for its own home subnet prefix. (Note that the currently defined anycast addresses may not work with all prefix lengths other than those defined in RFC 2373 [3, 35].)



Type
144
Code
0

Checksum

The ICMP checksum [14].

Identifier

An identifier to aid in matching Home Agent Address Discovery Reply messages to this Home Agent Address Discovery Request message.


Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

The Source Address of the Home Agent Address Discovery Request message packet is typically one of the mobile node's current care-of addresses. At the time of performing this dynamic home agent address discovery procedure, it is likely that the mobile node is not registered with any home agent. Therefore, neither the nature of the address nor the identity of the mobile node can be established at this time. The home agent MUST then return the Home Agent Address Discovery Reply message directly to the Source Address chosen by the mobile node.

6.6. ICMP Home Agent Address Discovery Reply Message

The ICMP Home Agent Address Discovery Reply message is used by a home agent to respond to a mobile node that uses the dynamic home agent address discovery mechanism, as described in Section 10.5.



Type

145

Code

0

Checksum

The ICMP checksum [14].

Identifier

The identifier from the invoking Home Agent Address Discovery Request message.

Reserved


This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Home Agent Addresses

A list of addresses of home agents on the home link for the mobile node. The number of addresses presented in the list is indicated by the remaining length of the IPv6 packet carrying the Home Agent Address Discovery Reply message.

6.7. ICMP Mobile Prefix Solicitation Message Format

The ICMP Mobile Prefix Solicitation Message is sent by a mobile node to its home agent while it is away from home. The purpose of the message is to solicit a Mobile Prefix Advertisement from the home agent, which will allow the mobile node to gather prefix information about its home network. This information can be used to configure and update home address(es) according to changes in prefix information supplied by the home agent.



IP Fields:

Source Address

The mobile node's care-of address.

Destination Address

The address of the mobile node's home agent. This home agent must be on the link that the mobile node wishes to learn prefix information about.

Hop Limit

Set to an initial hop limit value, similarly to any other unicast packet sent by the mobile node.

Destination Option:

A Home Address destination option MUST be included.

ESP header:

IPsec headers MUST be supported and SHOULD be used as described in Section 5.4.

ICMP Fields:

Type

146

Code

0

Checksum

The ICMP checksum [14].

Identifier

An identifier to aid in matching a future Mobile Prefix Advertisement to this Mobile Prefix Solicitation.


Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

The Mobile Prefix Solicitation messages may have options. These options MUST use the option format defined in RFC 2461 [12]. This document does not define any option types for the Mobile Prefix Solicitation message, but future documents may define new options. Home agents MUST silently ignore any options they do not recognize and continue processing the message.

6.8. ICMP Mobile Prefix Advertisement Message Format

A home agent will send a Mobile Prefix Advertisement to a mobile node to distribute prefix information about the home link while the mobile node is traveling away from the home network. This will occur in response to a Mobile Prefix Solicitation with an Advertisement, or by an unsolicited Advertisement sent according to the rules in Section 10.6.



IP Fields:

Source Address

The home agent's address as the mobile node would expect to see it (i.e., same network prefix).

Destination Address

If this message is a response to a Mobile Prefix Solicitation, this field contains the Source Address field from that packet. For unsolicited messages, the mobile node's care-of address SHOULD be used. Note that unsolicited messages can only be sent if the mobile node is currently registered with the home agent.

Routing header:

A type 2 routing header MUST be included.

ESP header:

IPsec headers MUST be supported and SHOULD be used as described in Section 5.4.

ICMP Fields:

Type

147

Code

0

Checksum

The ICMP checksum [14].

Identifier

An identifier to aid in matching this Mobile Prefix Advertisement to a previous Mobile Prefix Solicitation.

M

1-bit Managed Address Configuration flag. When set, hosts use the administered (stateful) protocol for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration. The use of this flag is described in [12, 13].

O

1-bit Other Stateful Configuration flag. When set, hosts use the administered (stateful) protocol for autoconfiguration of other (non-address) information. The use of this flag is described in [12, 13].

Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

The Mobile Prefix Advertisement messages may have options. These options MUST use the option format defined in RFC 2461 [12]. This document defines one option which may be carried in a Mobile Prefix Advertisement message, but future documents may define new options. Mobile nodes MUST silently ignore any options they do not recognize and continue processing the message.

Prefix Information

Each message contains one or more Prefix Information options. Each option carries the prefix(es) that the mobile node should use to configure its home address(es). Section 10.6 describes which prefixes should be advertised to the mobile node.

The Prefix Information option is defined in Section 4.6.2 of RFC 2461 [12], with modifications defined in Section 7.2 of this specification. The home agent MUST use this modified Prefix Information option to send home network prefixes as defined in Section 10.6.1.

If the Advertisement is sent in response to a Mobile Prefix Solicitation, the home agent MUST copy the Identifier value from that message into the Identifier field of the Advertisement.


The home agent MUST NOT send more than one Mobile Prefix Advertisement message per second to any mobile node.

The M and O bits MUST be cleared if the Home Agent DHCPv6 support is not provided. If such support is provided then they are set in concert with the home network's administrative settings.

7. Modifications to IPv6 Neighbor Discovery

7.1. Modified Router Advertisement Message Format

Mobile IPv6 modifies the format of the Router Advertisement message [12] by the addition of a single flag bit to indicate that the router sending the Advertisement message is serving as a home agent on this link. The format of the Router Advertisement message is as follows:



This format represents the following changes over that originally specified for Neighbor Discovery [12]:

Home Agent (H)

The Home Agent (H) bit is set in a Router Advertisement to indicate that the router sending this Router Advertisement is also functioning as a Mobile IPv6 home agent on this link.

Reserved


Reduced from a 6-bit field to a 5-bit field to account for the addition of the above bit.

7.2. Modified Prefix Information Option Format

Mobile IPv6 requires knowledge of a router's global address in building a Home Agents List as part of the dynamic home agent address discovery mechanism.

However, Neighbor Discovery [12] only advertises a router's link-local address, by requiring this address to be used as the IP Source Address of each Router Advertisement.

Mobile IPv6 extends Neighbor Discovery to allow a router to advertise its global address, by the addition of a single flag bit in the format of a Prefix Information option for use in Router Advertisement messages. The format of the Prefix Information option is as follows:



This format represents the following changes over that originally specified for Neighbor Discovery [12]:

Router Address (R)

1-bit router address flag. When set, indicates that the Prefix field contains a complete IP address assigned to the sending router. The indicated prefix is the first Prefix Length bits of the Prefix field. The router IP address has the same scope and conforms to the same lifetime values as the advertised prefix. This use of the Prefix field is compatible with its use in advertising the prefix itself, since Prefix Advertisement uses only the leading bits. Interpretation of this flag bit is thus independent of the processing required for the On-Link (L) and Autonomous Address-Configuration (A) flag bits.

Reserved1

Reduced from a 6-bit field to a 5-bit field to account for the addition of the above bit.


In a Router Advertisement, a home agent MUST, and all other routers MAY, include at least one Prefix Information option with the Router Address (R) bit set. Neighbor Discovery specifies that, if including all options in a Router Advertisement causes the size of the Advertisement to exceed the link MTU, multiple Advertisements can be sent, each containing a subset of the options [12]. Also, when sending unsolicited multicast Router Advertisements more frequently

than the limit specified in RFC 2461 [12], the sending router need not include all options in each of these Advertisements. However, in both of these cases the router SHOULD include at least one Prefix Information option with the Router Address (R) bit set in each such advertisement, if this bit is set in some advertisement sent by the router.

In addition, the following requirement can assist mobile nodes in movement detection. Barring changes in the prefixes for the link, routers that send multiple Router Advertisements with the Router Address (R) bit set in some of the included Prefix Information options SHOULD provide at least one option and router address which stays the same in all of the Advertisements.

7.3. New Advertisement Interval Option Format

Mobile IPv6 defines a new Advertisement Interval option, used in Router Advertisement messages to advertise the interval at which the sending router sends unsolicited multicast Router Advertisements. The format of the Advertisement Interval option is as follows:



Type

7

Length

8-bit unsigned integer. The length of the option (including the type and length fields) is in units of 8 octets. The value of this field MUST be 1.

Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Advertisement Interval


32-bit unsigned integer. The maximum time, in milliseconds, between successive unsolicited Router Advertisement messages sent by this router on this network interface. Using the conceptual router configuration variables defined by Neighbor Discovery [12], this field MUST be equal to the value MaxRtrAdvInterval, expressed in milliseconds.

Routers MAY include this option in their Router Advertisements. A mobile node receiving a Router Advertisement containing this option SHOULD utilize the specified Advertisement Interval for that router in its movement detection algorithm, as described in Section 11.5.1.

This option MUST be silently ignored for other Neighbor Discovery messages.

7.4. New Home Agent Information Option Format

Mobile IPv6 defines a new Home Agent Information option, used in Router Advertisements sent by a home agent to advertise information specific to this router's functionality as a home agent. The format of the Home Agent Information option is as follows:



Type

8

Length

8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 octets. The value of this field MUST be 1.

Reserved

This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Home Agent Preference

16-bit unsigned integer. The preference for the home agent sending this Router Advertisement, for use in ordering the addresses returned to a mobile node in the Home Agent Addresses field of a Home Agent Address Discovery Reply message. Higher values mean more preferable. If this option is not included in a Router Advertisement in which the Home Agent (H) bit is set, the preference value for this home agent MUST be considered to be 0. Greater values indicate a more preferable home agent than lower values.

The manual configuration of the Home Agent Preference value is described in Section 8.4. In addition, the sending home agent MAY dynamically set the Home Agent Preference value, for example basing it on the number of mobile nodes it is currently serving or on its remaining resources for serving additional mobile nodes; such dynamic settings are beyond the scope of this document. Any such dynamic setting of the Home Agent Preference, however, MUST set the preference appropriately, relative to the default Home Agent Preference value of 0 that may be in use by some home agents on this link (i.e., a home agent not including a Home Agent Information option in its Router Advertisements will be considered to have a Home Agent Preference value of 0).

Home Agent Lifetime

16-bit unsigned integer. The lifetime associated with the home agent in units of seconds. The default value is the same as the Router Lifetime, as specified in the main body of the Router Advertisement. The maximum value corresponds to 18.2 hours. A value of 0 MUST NOT be used. The Home Agent Lifetime applies only to this router's usefulness as a home agent; it does not apply to information contained in other message fields or options.

Home agents MAY include this option in their Router Advertisements. This option MUST NOT be included in a Router Advertisement in which the Home Agent (H) bit (see Section 7.1) is not set. If this option is not included in a Router Advertisement in which the Home Agent (H) bit is set, the lifetime for this home agent MUST be considered to be the same as the Router Lifetime in the Router Advertisement. If multiple Advertisements are being sent instead of a single larger unsolicited multicast Advertisement, all of the multiple Advertisements with the Router Address (R) bit set MUST include this option with the same contents, otherwise this option MUST be omitted from all Advertisements.

This option MUST be silently ignored for other Neighbor Discovery messages.

If both the Home Agent Preference and Home Agent Lifetime are set to their default values specified above, this option SHOULD NOT be included in the Router Advertisement messages sent by this home agent.

7.5. Changes to Sending Router Advertisements

The Neighbor Discovery protocol specification [12] limits routers to a minimum interval of 3 seconds between sending unsolicited multicast Router Advertisement messages from any given network interface (limited by MinRtrAdvInterval and MaxRtrAdvInterval), stating that:

"Routers generate Router Advertisements frequently enough that hosts will learn of their presence within a few minutes, but not frequently enough to rely on an absence of advertisements to detect router failure; a separate Neighbor Unreachability Detection algorithm provides failure detection."

This limitation, however, is not suitable to providing timely movement detection for mobile nodes. Mobile nodes detect their own movement by learning the presence of new routers as the mobile node moves into wireless transmission range of them (or physically connects to a new wired network), and by learning that previous routers are no longer reachable. Mobile nodes MUST be able to quickly detect when they move to a link served by a new router, so that they can acquire a new care-of address and send Binding Updates to register this care-of address with their home agent and to notify correspondent nodes as needed.

One method which can provide for faster movement detection, is to increase the rate at which unsolicited Router Advertisements are sent. Mobile IPv6 relaxes this limit such that routers MAY send unsolicited multicast Router Advertisements more frequently. This method can be applied where the router is expecting to provide service to visiting mobile nodes (e.g., wireless network interfaces), or on which it is serving as a home agent to one or more mobile nodes (who may return home and need to hear its Advertisements).

Routers supporting mobility SHOULD be able to be configured with a smaller MinRtrAdvInterval value and MaxRtrAdvInterval value to allow sending of unsolicited multicast Router Advertisements more often. The minimum allowed values are:

- o MinRtrAdvInterval 0.03 seconds
- o MaxRtrAdvInterval 0.07 seconds

In the case where the minimum intervals and delays are used, the mean time between unsolicited multicast router advertisements is 50 ms. Use of these modified limits MUST be configurable (see also the configuration variable MinDelayBetweenRas in Section 13 which may also have to be modified accordingly). Systems where these values are available MUST NOT default to them, and SHOULD default to values specified in RFC 2461. Knowledge of the type of network interface and operating environment SHOULD be taken into account in configuring these limits for each network interface. This is important with some wireless links, where increasing the frequency of multicast beacons can cause considerable overhead. Routers SHOULD adhere to the intervals specified in RFC 2461 [12], if this overhead is likely to cause service degradation.

Additionally, the possible low values of MaxRtrAdvInterval may cause some problems with movement detection in some mobile nodes. To ensure that this is not a problem, Routers SHOULD add 20 ms to any Advertisement Intervals sent in RAs, which are below 200 ms, in order to account for scheduling granularities on both the MN and the Router.

Note that multicast Router Advertisements are not always required in certain wireless networks that have limited bandwidth. Mobility detection or link changes in such networks may be done at lower layers. Router advertisements in such networks SHOULD be sent only when solicited. In such networks it SHOULD be possible to disable unsolicited multicast Router Advertisements on specific interfaces. The MinRtrAdvInterval and MaxRtrAdvInterval in such a case can be set to some high values.

Home agents MUST include the Source Link-Layer Address option in all Router Advertisements they send. This simplifies the process of returning home, as discussed in Section 11.5.4.

Note that according to RFC 2461 [12], AdvDefaultLifetime is by default based on the value of MaxRtrAdvInterval. AdvDefaultLifetime is used in the Router Lifetime field of Router Advertisements. Given that this field is expressed in seconds, a small MaxRtrAdvInterval value can result in a zero value for this field. To prevent this, routers SHOULD keep AdvDefaultLifetime in at least one second, even if the use of MaxRtrAdvInterval would result in a smaller value.

8. Requirements for Types of IPv6 Nodes

Mobile IPv6 places some special requirements on the functions provided by different types of IPv6 nodes. This section summarizes those requirements, identifying the functionality each requirement is intended to support.

The requirements are set for the following groups of nodes:

- o All IPv6 nodes.
- o All IPv6 nodes with support for route optimization.
- o All IPv6 routers.
- o All Mobile IPv6 home agents.
- o All Mobile IPv6 mobile nodes.

It is outside the scope of this specification to specify which of these groups are mandatory in IPv6. We only describe what is mandatory for a node that supports, for instance, route optimization. Other specifications are expected to define the extent of IPv6.

8.1. All IPv6 Nodes

Any IPv6 node may at any time be a correspondent node of a mobile node, either sending a packet to a mobile node or receiving a packet from a mobile node. There are no Mobile IPv6 specific MUST requirements for such nodes, and basic IPv6 techniques are sufficient. If a mobile node attempts to set up route optimization with a node with only basic IPv6 support, an ICMP error will signal that the node does not support such optimizations (Section 11.3.5), and communications will flow through the home agent.

An IPv6 node MUST NOT support the Home Address destination option, type 2 routing header, or the Mobility Header unless it fully supports the requirements listed in the next sections for either route optimization, mobile node, or home agent functionality.

8.2. IPv6 Nodes with Support for Route Optimization

Nodes that implement route optimization are a subset of all IPv6 nodes on the Internet. The ability of a correspondent node to participate in route optimization is essential for the efficient operation of the IPv6 Internet, for the following reasons:

- o Avoidance of congestion in the home network, and enabling the use of lower-performance home agent equipment even for supporting thousands of mobile nodes.
- o Reduced network load across the entire Internet, as mobile devices begin to predominate.
- o Reduction of jitter and latency for the communications.
- o Greater likelihood of success for QoS signaling as tunneling is avoided and, again, fewer sources of congestion.
- o Improved robustness against network partitions, congestion, and other problems, since fewer routing path segments are traversed.

These effects combine to enable much better performance and robustness for communications between mobile nodes and IPv6 correspondent nodes. Route optimization introduces a small amount of additional state for the peers, some additional messaging, and up to 1.5 roundtrip delays before it can be turned on. However, it is believed that the benefits far outweigh the costs in most cases. Section 11.3.1 discusses how mobile nodes may avoid route optimization for some of the remaining cases, such as very short-term communications.

The following requirements apply to all correspondent nodes that support route optimization:

- o The node **MUST** be able to validate a Home Address option using an existing Binding Cache entry, as described in Section 9.3.1.
- o The node **MUST** be able to insert a type 2 routing header into packets to be sent to a mobile node, as described in Section 9.3.2.
- o Unless the correspondent node is also acting as a mobile node, it **MUST** ignore type 2 routing headers and silently discard all packets that it has received with such headers.
- o The node **SHOULD** be able to interpret ICMP messages as described in Section 9.3.4.
- o The node **MUST** be able to send Binding Error messages as described in Section 9.3.3.
- o The node **MUST** be able to process Mobility Headers as described in Section 9.2.

- o The node MUST be able to participate in a return routability procedure (Section 9.4).
- o The node MUST be able to process Binding Update messages (Section 9.5).
- o The node MUST be able to return a Binding Acknowledgement (Section 9.5.4).
- o The node MUST be able to maintain a Binding Cache of the bindings received in accepted Binding Updates, as described in Section 9.1 and Section 9.6.
- o The node SHOULD allow route optimization to be administratively enabled or disabled. The default SHOULD be enabled.

8.3. All IPv6 Routers

All IPv6 routers, even those not serving as a home agent for Mobile IPv6, have an effect on how well mobile nodes can communicate:

- o Every IPv6 router SHOULD be able to send an Advertisement Interval option (Section 7.3) in each of its Router Advertisements [12], to aid movement detection by mobile nodes (as in Section 11.5.1). The use of this option in Router Advertisements SHOULD be configurable.
- o Every IPv6 router SHOULD be able to support sending unsolicited multicast Router Advertisements at the faster rate described in Section 7.5. If the router supports a faster rate, the used rate MUST be configurable.
- o Each router SHOULD include at least one prefix with the Router Address (R) bit set and with its full IP address in its Router Advertisements (as described in Section 7.2).
- o Routers supporting filtering packets with routing headers SHOULD support different rules for type 0 and type 2 routing headers (see Section 6.4) so that filtering of source routed packets (type 0) will not necessarily limit Mobile IPv6 traffic which is delivered via type 2 routing headers.

8.4. IPv6 Home Agents

In order for a mobile node to operate correctly while away from home, at least one IPv6 router on the mobile node's home link must function as a home agent for the mobile node. The following additional requirements apply to all IPv6 routers that serve as a home agent:

- o Every home agent MUST be able to maintain an entry in its Binding Cache for each mobile node for which it is serving as the home agent (Section 10.1 and Section 10.3.1).
- o Every home agent MUST be able to intercept packets (using proxy Neighbor Discovery [12]) addressed to a mobile node for which it is currently serving as the home agent, on that mobile node's home link, while the mobile node is away from home (Section 10.4.1).
- o Every home agent MUST be able to encapsulate [15] such intercepted packets in order to tunnel them to the primary care-of address for the mobile node indicated in its binding in the home agent's Binding Cache (Section 10.4.2).
- o Every home agent MUST support decapsulating [15] reverse tunneled packets sent to it from a mobile node's home address. Every home agent MUST also check that the source address in the tunneled packets corresponds to the currently registered location of the mobile node (Section 10.4.5).
- o The node MUST be able to process Mobility Headers as described in Section 10.2.
- o Every home agent MUST be able to return a Binding Acknowledgement in response to a Binding Update (Section 10.3.1).
- o Every home agent MUST maintain a separate Home Agents List for each link on which it is serving as a home agent, as described in Section 10.1 and Section 10.5.1.
- o Every home agent MUST be able to accept packets addressed to the Mobile IPv6 Home-Agents anycast address [16] for the subnet on which it is serving as a home agent, and MUST be able to participate in dynamic home agent address discovery (Section 10.5).
- o Every home agent SHOULD support a configuration mechanism to allow a system administrator to manually set the value to be sent by this home agent in the Home Agent Preference field of the Home Agent Information Option in Router Advertisements that it sends (Section 7.4).
- o Every home agent SHOULD support sending ICMP Mobile Prefix Advertisements (Section 6.8), and SHOULD respond to Mobile Prefix Solicitations (Section 6.7). If supported, this behavior MUST be configurable, so that home agents can be configured to avoid sending such Prefix Advertisements according to the needs of the network administration in the home domain.

- o Every home agent MUST support IPsec ESP for protection of packets belonging to the return routability procedure (Section 10.4.6).
- o Every home agent SHOULD support the multicast group membership control protocols as described in Section 10.4.3. If this support is provided, the home agent MUST be capable of using it to determine which multicast data packets to forward via the tunnel to the mobile node.
- o Home agents MAY support stateful address autoconfiguration for mobile nodes as described in Section 10.4.4.

8.5. IPv6 Mobile Nodes

Finally, the following requirements apply to all IPv6 nodes capable of functioning as mobile nodes:

- o The node MUST maintain a Binding Update List (Section 11.1).
- o The node MUST support sending packets containing a Home Address option (Section 11.3.1), and follow the required IPsec interaction (Section 11.3.2).
- o The node MUST be able to perform IPv6 encapsulation and decapsulation [15].
- o The node MUST be able to process type 2 routing header as defined in Section 6.4 and Section 11.3.3.
- o The node MUST support receiving a Binding Error message (Section 11.3.6).
- o The node MUST support receiving ICMP errors (Section 11.3.5).
- o The node MUST support movement detection, care-of address formation, and returning home (Section 11.5).
- o The node MUST be able to process Mobility Headers as described in Section 11.2.
- o The node MUST support the return routability procedure (Section 11.6).
- o The node MUST be able to send Binding Updates, as specified in Section 11.7.1 and Section 11.7.2.
- o The node MUST be able to receive and process Binding Acknowledgements, as specified in Section 11.7.3.

- o The node MUST support receiving a Binding Refresh Request (Section 6.1.2), by responding with a Binding Update.
- o The node MUST support receiving Mobile Prefix Advertisements (Section 11.4.3) and reconfiguring its home address based on the prefix information contained therein.
- o The node SHOULD support use of the dynamic home agent address discovery mechanism, as described in Section 11.4.1.
- o The node MUST allow route optimization to be administratively enabled or disabled. The default SHOULD be enabled.
- o The node MAY support the multicast address listener part of a multicast group membership protocol as described in Section 11.3.4. If this support is provided, the mobile node MUST be able to receive tunneled multicast packets from the home agent.
- o The node MAY support stateful address autoconfiguration mechanisms such as DHCPv6 [29] on the interface represented by the tunnel to the home agent.

9. Correspondent Node Operation

9.1. Conceptual Data Structures

IPv6 nodes with route optimization support maintain a Binding Cache of bindings for other nodes. A separate Binding Cache SHOULD be maintained by each IPv6 node for each of its unicast routable addresses. The Binding Cache MAY be implemented in any manner consistent with the external behavior described in this document, for example by being combined with the node's Destination Cache as maintained by Neighbor Discovery [12]. When sending a packet, the Binding Cache is searched before the Neighbor Discovery conceptual Destination Cache [12].

Each Binding Cache entry conceptually contains the following fields:

- o The home address of the mobile node for which this is the Binding Cache entry. This field is used as the key for searching the Binding Cache for the destination address of a packet being sent.
- o The care-of address for the mobile node indicated by the home address field in this Binding Cache entry.

- o A lifetime value, indicating the remaining lifetime for this Binding Cache entry. The lifetime value is initialized from the Lifetime field in the Binding Update that created or last modified this Binding Cache entry.
- o A flag indicating whether or not this Binding Cache entry is a home registration entry (applicable only on nodes which support home agent functionality).
- o The maximum value of the Sequence Number field received in previous Binding Updates for this home address. The Sequence Number field is 16 bits long. Sequence Number values MUST be compared modulo 2^{16} as explained in Section 9.5.1.
- o Usage information for this Binding Cache entry. This is needed to implement the cache replacement policy in use in the Binding Cache. Recent use of a cache entry also serves as an indication that a Binding Refresh Request should be sent when the lifetime of this entry nears expiration.

Binding Cache entries not marked as home registrations MAY be replaced at any time by any reasonable local cache replacement policy but SHOULD NOT be unnecessarily deleted. The Binding Cache for any one of a node's IPv6 addresses may contain at most one entry for each mobile node home address. The contents of a node's Binding Cache MUST NOT be changed in response to a Home Address option in a received packet.

9.2. Processing Mobility Headers

Mobility Header processing MUST observe the following rules:

- o The checksum must be verified as per Section 6.1. Otherwise, the node MUST silently discard the message.
- o The MH Type field MUST have a known value (Section 6.1.1). Otherwise, the node MUST discard the message and issue a Binding Error message as described in Section 9.3.3, with Status field set to 2 (unrecognized MH Type value).
- o The Payload Proto field MUST be IPPROTO_NONE (59 decimal). Otherwise, the node MUST discard the message and SHOULD send ICMP Parameter Problem, Code 0, directly to the Source Address of the packet as specified in RFC 2463 [14]. Thus no Binding Cache information is used in sending the ICMP message. The Pointer field in the ICMP message SHOULD point at the Payload Proto field.

- o The Header Len field in the Mobility Header MUST NOT be less than the length specified for this particular type of message in Section 6.1. Otherwise, the node MUST discard the message and SHOULD send ICMP Parameter Problem, Code 0, directly to the Source Address of the packet as specified in RFC 2463 [14]. (The Binding Cache information is again not used.) The Pointer field in the ICMP message SHOULD point at the Header Len field.

Subsequent checks depend on the particular Mobility Header.

9.3. Packet Processing

This section describes how the correspondent node sends packets to the mobile node, and receives packets from it.

9.3.1. Receiving Packets with Home Address Option

Packets containing a Home Address option MUST be dropped if the given home address is not a unicast routable address.

Mobile nodes can include a Home Address destination option in a packet if they believe the correspondent node has a Binding Cache entry for the home address of a mobile node. Packets containing a Home Address option MUST be dropped if there is no corresponding Binding Cache entry. A corresponding Binding Cache entry MUST have the same home address as appears in the Home Address destination option, and the currently registered care-of address MUST be equal to the source address of the packet. These tests MUST NOT be done for packets that contain a Home Address option and a Binding Update.

If the packet is dropped due the above tests, the correspondent node MUST send the Binding Error message as described in Section 9.3.3. The Status field in this message should be set to 1 (unknown binding for Home Address destination option).

The correspondent node MUST process the option in a manner consistent with exchanging the Home Address field from the Home Address option into the IPv6 header and replacing the original value of the Source Address field there. After all IPv6 options have been processed, it MUST be possible for upper layers to process the packet without the knowledge that it came originally from a care-of address or that a Home Address option was used.

The use of IPsec Authentication Header (AH) for the Home Address option is not required, except that if the IPv6 header of a packet is covered by AH, then the authentication MUST also cover the Home Address option; this coverage is achieved automatically by the definition of the Option Type code for the Home Address option, since

it indicates that the data within the option cannot change en route to the packet's final destination, and thus the option is included in the AH computation. By requiring that any authentication of the IPv6 header also cover the Home Address option, the security of the Source Address field in the IPv6 header is not compromised by the presence of a Home Address option.

When attempting to verify AH authentication data in a packet that contains a Home Address option, the receiving node MUST calculate the AH authentication data as if the following were true: The Home Address option contains the care-of address, and the source IPv6 address field of the IPv6 header contains the home address. This conforms with the calculation specified in Section 11.3.2.

9.3.2. Sending Packets to a Mobile Node

Before sending any packet, the sending node SHOULD examine its Binding Cache for an entry for the destination address to which the packet is being sent. If the sending node has a Binding Cache entry for this address, the sending node SHOULD use a type 2 routing header to route the packet to this mobile node (the destination node) by way of its care-of address. However, the sending node MUST not do this in the following cases:

- o When sending an IPv6 Neighbor Discovery [12] packet.
- o Where otherwise noted in Section 6.1.

When calculating authentication data in a packet that contains a type 2 routing header, the correspondent node MUST calculate the AH authentication data as if the following were true: The routing header contains the care-of address, the destination IPv6 address field of the IPv6 header contains the home address, and the Segments Left field is zero. The IPsec Security Policy Database lookup MUST be based on the mobile node's home address.

For instance, assuming there are no additional routing headers in this packet beyond those needed by Mobile IPv6, the correspondent node could set the fields in the packet's IPv6 header and routing header as follows:

- o The Destination Address in the packet's IPv6 header is set to the mobile node's home address (the original destination address to which the packet was being sent).

- o The routing header is initialized to contain a single route segment, containing the mobile node's care-of address copied from the Binding Cache entry. The Segments Left field is, however, temporarily set to zero.

The IP layer will insert the routing header before performing any necessary IPsec processing. Once all IPsec processing has been performed, the node swaps the IPv6 destination field with the Home Address field in the routing header, sets the Segments Left field to one, and sends the packet. This ensures the AH calculation is done on the packet in the form it will have on the receiver after advancing the routing header.

Following the definition of a type 2 routing header in Section 6.4, this packet will be routed to the mobile node's care-of address, where it will be delivered to the mobile node (the mobile node has associated the care-of address with its network interface).

Note that following the above conceptual model in an implementation creates some additional requirements for path MTU discovery since the layer that decides the packet size (e.g., TCP and applications using UDP) needs to be aware of the size of the headers added by the IP layer on the sending node.

If, instead, the sending node has no Binding Cache entry for the destination address to which the packet is being sent, the sending node simply sends the packet normally, with no routing header. If the destination node is not a mobile node (or is a mobile node that is currently at home), the packet will be delivered directly to this node and processed normally by it. If, however, the destination node is a mobile node that is currently away from home, the packet will be intercepted by the mobile node's home agent and tunneled to the mobile node's current primary care-of address.

9.3.3. Sending Binding Error Messages

Section 9.2 and Section 9.3.1 describe error conditions that lead to a need to send a Binding Error message.

A Binding Error message is sent directly to the address that appeared in the IPv6 Source Address field of the offending packet. If the Source Address field does not contain a unicast address, the Binding Error message MUST NOT be sent.

The Home Address field in the Binding Error message MUST be copied from the Home Address field in the Home Address destination option of the offending packet, or set to the unspecified address if no such option appeared in the packet.

Note that the IPv6 Source Address and Home Address field values discussed above are the values from the wire, i.e., before any modifications possibly performed as specified in Section 9.3.1.

Binding Error messages SHOULD be subject to rate limiting in the same manner as is done for ICMPv6 messages [14].

9.3.4. Receiving ICMP Error Messages

When the correspondent node has a Binding Cache entry for a mobile node, all traffic destined to the mobile node goes directly to the current care-of address of the mobile node using a routing header. Any ICMP error message caused by packets on their way to the care-of address will be returned in the normal manner to the correspondent node.

On the other hand, if the correspondent node has no Binding Cache entry for the mobile node, the packet will be routed through the mobile node's home link. Any ICMP error message caused by the packet on its way to the mobile node while in the tunnel, will be transmitted to the mobile node's home agent. By the definition of IPv6 encapsulation [15], the home agent MUST relay certain ICMP error messages back to the original sender of the packet, which in this case is the correspondent node.

Thus, in all cases, any meaningful ICMP error messages caused by packets from a correspondent node to a mobile node will be returned to the correspondent node. If the correspondent node receives persistent ICMP Destination Unreachable messages after sending packets to a mobile node based on an entry in its Binding Cache, the correspondent node SHOULD delete this Binding Cache entry. Note that if the mobile node continues to send packets with the Home Address destination option to this correspondent node, they will be dropped due to the lack of a binding. For this reason it is important that only persistent ICMP messages lead to the deletion of the Binding Cache entry.

9.4. Return Routability Procedure

This subsection specifies actions taken by a correspondent node during the return routability procedure.

9.4.1. Receiving Home Test Init Messages

Upon receiving a Home Test Init message, the correspondent node verifies the following:

- o The packet MUST NOT include a Home Address destination option.

Any packet carrying a Home Test Init message which fails to satisfy all of these tests MUST be silently ignored.

Otherwise, in preparation for sending the corresponding Home Test Message, the correspondent node checks that it has the necessary material to engage in a return routability procedure, as specified in Section 5.2. The correspondent node MUST have a secret Kcn and a nonce. If it does not have this material yet, it MUST produce it before continuing with the return routability procedure.

Section 9.4.3 specifies further processing.

9.4.2. Receiving Care-of Test Init Messages

Upon receiving a Care-of Test Init message, the correspondent node verifies the following:

- o The packet MUST NOT include a Home Address destination option.

Any packet carrying a Care-of Test Init message which fails to satisfy all of these tests MUST be silently ignored.

Otherwise, in preparation for sending the corresponding Care-of Test Message, the correspondent node checks that it has the necessary material to engage in a return routability procedure in the manner described in Section 9.4.1.

Section 9.4.4 specifies further processing.

9.4.3. Sending Home Test Messages

The correspondent node creates a home keygen token and uses the current nonce index as the Home Nonce Index. It then creates a Home Test message (Section 6.1.5) and sends it to the mobile node at the latter's home address.

9.4.4. Sending Care-of Test Messages

The correspondent node creates a care-of keygen token and uses the current nonce index as the Care-of Nonce Index. It then creates a Care-of Test message (Section 6.1.6) and sends it to the mobile node at the latter's care-of address.

9.5. Processing Bindings

This section explains how the correspondent node processes messages related to bindings. These messages are:

- o Binding Update
- o Binding Refresh Request
- o Binding Acknowledgement
- o Binding Error

9.5.1. Receiving Binding Updates

Before accepting a Binding Update, the receiving node MUST validate the Binding Update according to the following tests:

- o The packet MUST contain a unicast routable home address, either in the Home Address option or in the Source Address, if the Home Address option is not present.
- o The Sequence Number field in the Binding Update is greater than the Sequence Number received in the previous valid Binding Update for this home address, if any.

If the receiving node has no Binding Cache entry for the indicated home address, it MUST accept any Sequence Number value in a received Binding Update from this mobile node.

This Sequence Number comparison MUST be performed modulo 2^{16} , i.e., the number is a free running counter represented modulo 65536. A Sequence Number in a received Binding Update is considered less than or equal to the last received number if its value lies in the range of the last received number and the preceding 32768 values, inclusive. For example, if the last received sequence number was 15, then messages with sequence numbers 0 through 15, as well as 32783 through 65535, would be considered less than or equal.

When the Home Registration (H) bit is not set, the following are also required:

- o A Nonce Indices mobility option MUST be present, and the Home and Care-of Nonce Index values in this option MUST be recent enough to be recognized by the correspondent node. (Care-of Nonce Index values are not inspected for requests to delete a binding.)
- o The correspondent node MUST re-generate the home keygen token and the care-of keygen token from the information contained in the packet. It then generates the binding management key Kbm and uses it to verify the authenticator field in the Binding Update as specified in Section 6.1.7.
- o The Binding Authorization Data mobility option MUST be present, and its contents MUST satisfy rules presented in Section 5.2.6. Note that a care-of address different from the Source Address MAY have been specified by including an Alternate Care-of Address mobility option in the Binding Update. When such a message is received and the return routability procedure is used as an authorization method, the correspondent node MUST verify the authenticator by using the address within the Alternate Care-of Address in the calculations.
- o The Binding Authorization Data mobility option MUST be the last option and MUST NOT have trailing padding.

If the Home Registration (H) bit is set, the Nonce Indices mobility option MUST NOT be present.

If the mobile node sends a sequence number which is not greater than the sequence number from the last valid Binding Update for this home address, then the receiving node MUST send back a Binding Acknowledgement with status code 135, and the last accepted sequence number in the Sequence Number field of the Binding Acknowledgement.

If a binding already exists for the given home address and the home registration flag has a different value than the Home Registration (H) bit in the Binding Update, then the receiving node MUST send back a Binding Acknowledgement with status code 139 (registration type change disallowed). The home registration flag stored in the Binding Cache entry MUST NOT be changed.

If the receiving node no longer recognizes the Home Nonce Index value, Care-of Nonce Index value, or both values from the Binding Update, then the receiving node MUST send back a Binding Acknowledgement with status code 136, 137, or 138, respectively.

Packets carrying Binding Updates that fail to satisfy all of these tests for any reason other than insufficiency of the Sequence Number, registration type change, or expired nonce index values, MUST be silently discarded.

If the Binding Update is valid according to the tests above, then the Binding Update is processed further as follows:

- o The Sequence Number value received from a mobile node in a Binding Update is stored by the receiving node in its Binding Cache entry for the given home address.
- o If the Lifetime specified in the Binding Update is nonzero and the specified care-of address is not equal to the home address for the binding, then this is a request to cache a binding for the home address. If the Home Registration (H) bit is set in the Binding Update, the Binding Update is processed according to the procedure specified in Section 10.3.1; otherwise, it is processed according to the procedure specified in Section 9.5.2.
- o If the Lifetime specified in the Binding Update is zero or the specified care-of address matches the home address for the binding, then this is a request to delete the cached binding for the home address. In this case, the Binding Update MUST include a valid home nonce index, and the care-of nonce index MUST be ignored by the correspondent node. The generation of the binding management key depends then exclusively on the home keygen token (Section 5.2.5). If the Home Registration (H) bit is set in the Binding Update, the Binding Update is processed according to the procedure specified in Section 10.3.2; otherwise, it is processed according to the procedure specified in Section 9.5.3.

The specified care-of address MUST be determined as follows:

- o If the Alternate Care-of Address option is present, the care-of address is the address in that option.
- o Otherwise, the care-of address is the Source Address field in the packet's IPv6 header.

The home address for the binding MUST be determined as follows:

- o If the Home Address destination option is present, the home address is the address in that option.
- o Otherwise, the home address is the Source Address field in the packet's IPv6 header.

9.5.2. Requests to Cache a Binding

This section describes the processing of a valid Binding Update that requests a node to cache a binding, for which the Home Registration (H) bit is not set in the Binding Update.

In this case, the receiving node SHOULD create a new entry in its Binding Cache for this home address, or update its existing Binding Cache entry for this home address, if such an entry already exists. The lifetime for the Binding Cache entry is initialized from the Lifetime field specified in the Binding Update, although this lifetime MAY be reduced by the node caching the binding; the lifetime for the Binding Cache entry MUST NOT be greater than the Lifetime value specified in the Binding Update. Any Binding Cache entry MUST be deleted after the expiration of its lifetime.

Note that if the mobile node did not request a Binding Acknowledgement, then it is not aware of the selected shorter lifetime. The mobile node may thus use route optimization and send packets with the Home Address destination option. As discussed in Section 9.3.1, such packets will be dropped if there is no binding. This situation is recoverable, but can cause temporary packet loss.

The correspondent node MAY refuse to accept a new Binding Cache entry if it does not have sufficient resources. A new entry MAY also be refused if the correspondent node believes its resources are utilized more efficiently in some other purpose, such as serving another mobile node with higher amount of traffic. In both cases the correspondent node SHOULD return a Binding Acknowledgement with status value 130.

9.5.3 Requests to Delete a Binding

This section describes the processing of a valid Binding Update that requests a node to delete a binding when the Home Registration (H) bit is not set in the Binding Update.

Any existing binding for the given home address MUST be deleted. A Binding Cache entry for the home address MUST NOT be created in response to receiving the Binding Update.

If the Binding Cache entry was created by use of return routability nonces, the correspondent node MUST ensure that the same nonces are not used again with the particular home and care-of address. If both nonces are still valid, the correspondent node has to remember the particular combination of nonce indexes, addresses, and sequence number as illegal until at least one of the nonces has become too old.

9.5.4. Sending Binding Acknowledgements

A Binding Acknowledgement may be sent to indicate receipt of a Binding Update as follows:

- o If the Binding Update was discarded as described in Section 9.2 or Section 9.5.1, a Binding Acknowledgement MUST NOT be sent. Otherwise the treatment depends on the following rules.
- o If the Acknowledge (A) bit set is set in the Binding Update, a Binding Acknowledgement MUST be sent. Otherwise, the treatment depends on the below rule.
- o If the node rejects the Binding Update due to an expired nonce index, sequence number being out of window (Section 9.5.1), or insufficiency of resources (Section 9.5.2), a Binding Acknowledgement MUST be sent. If the node accepts the Binding Update, the Binding Acknowledgement SHOULD NOT be sent.

If the node accepts the Binding Update and creates or updates an entry for this binding, the Status field in the Binding Acknowledgement MUST be set to a value less than 128. Otherwise, the Status field MUST be set to a value greater than or equal to 128. Values for the Status field are described in Section 6.1.8 and in the IANA registry of assigned numbers [19].

If the Status field in the Binding Acknowledgement contains the value 136 (expired home nonce index), 137 (expired care-of nonce index), or 138 (expired nonces) then the message MUST NOT include the Binding Authorization Data mobility option. Otherwise, the Binding Authorization Data mobility option MUST be included, and MUST meet the specific authentication requirements for Binding Acknowledgements as defined in Section 5.2.

If the Source Address field of the IPv6 header that carried the Binding Update does not contain a unicast address, the Binding Acknowledgement MUST NOT be sent and the Binding Update packet MUST be silently discarded. Otherwise, the acknowledgement MUST be sent to the Source Address. Unlike the treatment of regular packets, this addressing procedure does not use information from the Binding Cache. However, a routing header is needed in some cases. If the Source Address is the home address of the mobile node, i.e., the Binding Update did not contain a Home Address destination option, then the Binding Acknowledgement MUST be sent to that address and the routing header MUST NOT be used. Otherwise, the Binding Acknowledgement MUST be sent using a type 2 routing header which contains the mobile node's home address.

9.5.5. Sending Binding Refresh Requests

If a Binding Cache entry being deleted is still in active use when sending packets to a mobile node, then the next packet sent to the mobile node will be routed normally to the mobile node's home link. Communication with the mobile node continues, but the tunneling from the home network creates additional overhead and latency in delivering packets to the mobile node.

If the sender knows that the Binding Cache entry is still in active use, it MAY send a Binding Refresh Request message to the mobile node in an attempt to avoid this overhead and latency due to deleting and recreating the Binding Cache entry. This message is always sent to the home address of the mobile node.

The correspondent node MAY retransmit Binding Refresh Request messages as long as the rate limitation is applied. The correspondent node MUST stop retransmitting when it receives a Binding Update.

9.6. Cache Replacement Policy

Conceptually, a node maintains a separate timer for each entry in its Binding Cache. When creating or updating a Binding Cache entry in response to a received and accepted Binding Update, the node sets the timer for this entry to the specified Lifetime period. Any entry in a node's Binding Cache MUST be deleted after the expiration of the Lifetime specified in the Binding Update from which the entry was created or last updated.

Each node's Binding Cache will, by necessity, have a finite size. A node MAY use any reasonable local policy for managing the space within its Binding Cache.

A node MAY choose to drop any entry already in its Binding Cache in order to make space for a new entry. For example, a "least-recently used" (LRU) strategy for cache entry replacement among entries should work well, unless the size of the Binding Cache is substantially insufficient. When entries are deleted, the correspondent node MUST follow the rules in Section 5.2.8 in order to guard the return routability procedure against replay attacks.

If the node sends a packet to a destination for which it has dropped the entry from its Binding Cache, the packet will be routed through the mobile node's home link. The mobile node can detect this and establish a new binding if necessary.

However, if the mobile node believes that the binding still exists, it may use route optimization and send packets with the Home Address destination option. This can create temporary packet loss, as discussed earlier, in the context of binding lifetime reductions performed by the correspondent node (Section 9.5.2).

10. Home Agent Operation

10.1. Conceptual Data Structures

Each home agent **MUST** maintain a Binding Cache and Home Agents List.

The rules for maintaining a Binding Cache are the same for home agents and correspondent nodes and have already been described in Section 9.1.

The Home Agents List is maintained by each home agent, recording information about each router on the same link that is acting as a home agent. This list is used by the dynamic home agent address discovery mechanism. A router is known to be acting as a home agent, if it sends a Router Advertisement in which the Home Agent (H) bit is set. When the lifetime for a list entry (defined below) expires, that entry is removed from the Home Agents List. The Home Agents List is similar to the Default Router List conceptual data structure maintained by each host for Neighbor Discovery [12]. The Home Agents List **MAY** be implemented in any manner consistent with the external behavior described in this document.

Each home agent maintains a separate Home Agents List for each link on which it is serving as a home agent. A new entry is created or an existing entry is updated in response to receipt of a valid Router Advertisement in which the Home Agent (H) bit is set. Each Home Agents List entry conceptually contains the following fields:

- o The link-local IP address of a home agent on the link. This address is learned through the Source Address of the Router Advertisements [12] received from the router.
- o One or more global IP addresses for this home agent. Global addresses are learned through Prefix Information options with the Router Address (R) bit set and received in Router Advertisements from this link-local address. Global addresses for the router in a Home Agents List entry **MUST** be deleted once the prefix associated with that address is no longer valid [12].

- o The remaining lifetime of this Home Agents List entry. If a Home Agent Information Option is present in a Router Advertisement received from a home agent, the lifetime of the Home Agents List entry representing that home agent is initialized from the Home Agent Lifetime field in the option (if present); otherwise, the lifetime is initialized from the Router Lifetime field in the received Router Advertisement. If Home Agents List entry lifetime reaches zero, the entry MUST be deleted from the Home Agents List.
- o The preference for this home agent; higher values indicate a more preferable home agent. The preference value is taken from the Home Agent Preference field in the received Router Advertisement, if the Router Advertisement contains a Home Agent Information Option and is otherwise set to the default value of 0. A home agent uses this preference in ordering the Home Agents List when it sends an ICMP Home Agent Address Discovery message.

10.2. Processing Mobility Headers

All IPv6 home agents MUST observe the rules described in Section 9.2 when processing Mobility Headers.

10.3. Processing Bindings

10.3.1. Primary Care-of Address Registration

When a node receives a Binding Update, it MUST validate it and determine the type of Binding Update according to the steps described in Section 9.5.1. Furthermore, it MUST authenticate the Binding Update as described in Section 5.1. An authorization step specific for the home agent is also needed to ensure that only the right node can control a particular home address. This is provided through the home address unequivocally identifying the security association that must be used.

This section describes the processing of a valid and authorized Binding Update when it requests the registration of the mobile node's primary care-of address.

To begin processing the Binding Update, the home agent MUST perform the following sequence of tests:

- o If the node implements only correspondent node functionality, or has not been configured to act as a home agent, then the node MUST reject the Binding Update. The node MUST also return a Binding Acknowledgement to the mobile node, in which the Status field is set to 131 (home registration not supported).

- o Else, if the home address for the binding (the Home Address field in the packet's Home Address option) is not an on-link IPv6 address with respect to the home agent's current Prefix List, then the home agent MUST reject the Binding Update and SHOULD return a Binding Acknowledgement to the mobile node, in which the Status field is set to 132 (not home subnet).
- o Else, if the home agent chooses to reject the Binding Update for any other reason (e.g., insufficient resources to serve another mobile node as a home agent), then the home agent SHOULD return a Binding Acknowledgement to the mobile node, in which the Status field is set to an appropriate value to indicate the reason for the rejection.
- o A Home Address destination option MUST be present in the message. It MUST be validated as described in Section 9.3.1 with the following additional rule. The Binding Cache entry existence test MUST NOT be done for IPsec packets when the Home Address option contains an address for which the receiving node could act as a home agent.

If home agent accepts the Binding Update, it MUST then create a new entry in its Binding Cache for this mobile node or update its existing Binding Cache entry, if such an entry already exists. The Home Address field as received in the Home Address option provides the home address of the mobile node.

The home agent MUST mark this Binding Cache entry as a home registration to indicate that the node is serving as a home agent for this binding. Binding Cache entries marked as a home registration MUST be excluded from the normal cache replacement policy used for the Binding Cache (Section 9.6) and MUST NOT be removed from the Binding Cache until the expiration of the Lifetime period.

Unless this home agent already has a binding for the given home address, the home agent MUST perform Duplicate Address Detection [13] on the mobile node's home link before returning the Binding Acknowledgement. This ensures that no other node on the home link was using the mobile node's home address when the Binding Update arrived. If this Duplicate Address Detection fails for the given home address or an associated link local address, then the home agent MUST reject the complete Binding Update and MUST return a Binding Acknowledgement to the mobile node, in which the Status field is set to 134 (Duplicate Address Detection failed). When the home agent sends a successful Binding Acknowledgement to the mobile node, the home agent assures to the mobile node that its address(es) will be kept unique by the home agent for as long as the lifetime was granted for the binding.

The specific addresses, which are to be tested before accepting the Binding Update and later to be defended by performing Duplicate Address Detection, depend on the setting of the Link-Local Address Compatibility (L) bit, as follows:

- o L=0: Defend only the given address. Do not derive a link-local address.
- o L=1: Defend both the given non link-local unicast (home) address and the derived link-local. The link-local address is derived by replacing the subnet prefix in the mobile node's home address with the link-local prefix.

The lifetime of the Binding Cache entry depends on a number of factors:

- o The lifetime for the Binding Cache entry MUST NOT be greater than the Lifetime value specified in the Binding Update.
- o The lifetime for the Binding Cache entry MUST NOT be greater than the remaining valid lifetime for the subnet prefix in the mobile node's home address specified with the Binding Update. The remaining valid lifetime for this prefix is determined by the home agent based on its own Prefix List entry [12].

The remaining preferred lifetime SHOULD NOT have any impact on the lifetime for the binding cache entry.

The home agent MUST remove a binding when the valid lifetime of the prefix associated with it expires.

- o The home agent MAY further decrease the specified lifetime for the binding, for example based on a local policy. The resulting lifetime is stored by the home agent in the Binding Cache entry, and this Binding Cache entry MUST be deleted by the home agent after the expiration of this lifetime.

Regardless of the setting of the Acknowledge (A) bit in the Binding Update, the home agent MUST return a Binding Acknowledgement to the mobile node constructed as follows:

- o The Status field MUST be set to a value indicating success. The value 1 (accepted but prefix discovery necessary) MUST be used if the subnet prefix of the specified home address is deprecated, or becomes deprecated during the lifetime of the binding, or becomes invalid at the end of the lifetime. The value 0 MUST be used

otherwise. For the purposes of comparing the binding and prefix lifetimes, the prefix lifetimes are first converted into units of four seconds by ignoring the two least significant bits.

- o The Key Management Mobility Capability (K) bit is set if the following conditions are all fulfilled, and cleared otherwise:
 - * The Key Management Mobility Capability (K) bit was set in the Binding Update.
 - * The IPsec security associations between the mobile node and the home agent have been established dynamically.
 - * The home agent has the capability to update its endpoint in the used key management protocol to the new care-of address every time it moves.

Depending on the final value of the bit in the Binding Acknowledgement, the home agent SHOULD perform the following actions:

K = 0

Discard key management connections, if any, to the old care-of address. If the mobile node did not have a binding before sending this Binding Update, discard the connections to the home address.

K = 1

Move the peer endpoint of the key management protocol connection, if any, to the new care-of address. For an IKE phase 1 connection, this means that any IKE packets sent to the peer are sent to this address, and packets from this address with the original ISAKMP cookies are accepted.

Note that RFC 2408 [8] Section 2.5.3 gives specific rules that ISAKMP cookies must satisfy: they must depend on specific parties and can only be generated by the entity itself. Then it recommends a particular way to do this, namely a hash of IP addresses. With the K bit set to 1, the recommended implementation technique does not work directly. To satisfy the two rules, the specific parties must be treated as the original IP addresses, not the ones in use at the specific moment.

- o The Sequence Number field MUST be copied from the Sequence Number given in the Binding Update.

- o The Lifetime field MUST be set to the remaining lifetime for the binding as set by the home agent in its home registration Binding Cache entry for the mobile node, as described above.
- o If the home agent stores the Binding Cache entry in nonvolatile storage, then the Binding Refresh Advice mobility option MUST be omitted. Otherwise, the home agent MAY include this option to suggest that the mobile node refreshes its binding before the actual lifetime of the binding ends.

If the Binding Refresh Advice mobility option is present, the Refresh Interval field in the option MUST be set to a value less than the Lifetime value being returned in the Binding Acknowledgement. This indicates that the mobile node SHOULD attempt to refresh its home registration at the indicated shorter interval. The home agent MUST still retain the registration for the Lifetime period, even if the mobile node does not refresh its registration within the Refresh period.

The rules for selecting the Destination IP address (and possibly routing header construction) for the Binding Acknowledgement to the mobile node are the same as in Section 9.5.4.

In addition, the home agent MUST follow the procedure defined in Section 10.4.1 to intercept packets on the mobile node's home link addressed to the mobile node, while the home agent is serving as the home agent for this mobile node. The home agent MUST also be prepared to accept reverse tunneled packets from the new care-of address of the mobile node, as described in Section 10.4.5. Finally, the home agent MUST also propagate new home network prefixes, as described in Section 10.6.

10.3.2. Primary Care-of Address De-Registration

A binding may need to be de-registered when the mobile node returns home or when the mobile node knows that it will not have any care-of addresses in the visited network.

A Binding Update is validated and authorized in the manner described in the previous section; note that when the mobile node de-registers when it is at home, it may not include the Home Address destination option, in which case the mobile node's home address is the source IP address of the de-registration Binding Update. This section describes the processing of a valid Binding Update that requests the receiving node to no longer serve as its home agent, de-registering its primary care-of address.

To begin processing the Binding Update, the home agent MUST perform the following test:

- o If the receiving node has no entry marked as a home registration in its Binding Cache for this mobile node, then this node MUST reject the Binding Update and SHOULD return a Binding Acknowledgement to the mobile node, in which the Status field is set to 133 (not home agent for this mobile node).

If the home agent does not reject the Binding Update as described above, then it MUST delete any existing entry in its Binding Cache for this mobile node. Then, the home agent MUST return a Binding Acknowledgement to the mobile node, constructed as follows:

- o The Status field MUST be set to a value 0, indicating success.
- o The Key Management Mobility Capability (K) bit is set or cleared and actions based on its value are performed as described in the previous section. The mobile node's home address is used as its new care-of address for the purposes of moving the key management connection to a new endpoint.
- o The Sequence Number field MUST be copied from the Sequence Number given in the Binding Update.
- o The Lifetime field MUST be set to zero.
- o The Binding Refresh Advice mobility option MUST be omitted.

In addition, the home agent MUST stop intercepting packets on the mobile node's home link that are addressed to the mobile node (Section 10.4.1).

The rules for selecting the Destination IP address (and, if required, routing header construction) for the Binding Acknowledgement to the mobile node are the same as in the previous section. When the Status field in the Binding Acknowledgement is greater than or equal to 128 and the Source Address of the Binding Update is on the home link, the home agent MUST send it to the mobile node's link layer address (retrieved either from the Binding Update or through Neighbor Solicitation).

10.4. Packet Processing

10.4.1. Intercepting Packets for a Mobile Node

While a node is serving as the home agent for mobile node it **MUST** attempt to intercept packets on the mobile node's home link that are addressed to the mobile node.

In order to do this, when a node begins serving as the home agent it **MUST** multicast onto the home link a Neighbor Advertisement message [12] on behalf of the mobile node. For the home address specified in the Binding Update, the home agent sends a Neighbor Advertisement message [12] to the all-nodes multicast address on the home link to advertise the home agent's own link-layer address for this IP address on behalf of the mobile node. If the Link-Layer Address Compatibility (L) flag has been specified in the Binding Update, the home agent **MUST** do the same for the link-local address of the mobile node.

All fields in each Neighbor Advertisement message **SHOULD** be set in the same way they would be set by the mobile node if it was sending this Neighbor Advertisement [12] while at home, with the following exceptions:

- o The Target Address in the Neighbor Advertisement **MUST** be set to the specific IP address for the mobile node.
- o The Advertisement **MUST** include a Target Link-layer Address option specifying the home agent's link-layer address.
- o The Router (R) bit in the Advertisement **MUST** be set to zero.
- o The Solicited Flag (S) in the Advertisement **MUST NOT** be set, since it was not solicited by any Neighbor Solicitation.
- o The Override Flag (O) in the Advertisement **MUST** be set, indicating that the Advertisement **SHOULD** override any existing Neighbor Cache entry at any node receiving it.
- o The Source Address in the IPv6 header **MUST** be set to the home agent's IP address on the interface used to send the advertisement.

Any node on the home link that receives one of the Neighbor Advertisement messages (described above) will update its Neighbor Cache to associate the mobile node's address with the home agent's link layer address, causing it to transmit any future packets normally destined to the mobile node to the mobile node's home agent.

Since multicasting on the local link (such as Ethernet) is typically not guaranteed to be reliable, the home agent MAY retransmit this Neighbor Advertisement message up to MAX_NEIGHBOR_ADVERTISEMENT (see [12]) times to increase its reliability. It is still possible that some nodes on the home link will not receive any of the Neighbor Advertisements, but these nodes will eventually be able to detect the link-layer address change for the mobile node's address through use of Neighbor Unreachability Detection [12].

While a node is serving as a home agent for some mobile node, the home agent uses IPv6 Neighbor Discovery [12] to intercept unicast packets on the home link addressed to the mobile node. In order to intercept packets in this way, the home agent MUST act as a proxy for this mobile node and reply to any received Neighbor Solicitations for it. When a home agent receives a Neighbor Solicitation, it MUST check if the Target Address specified in the message matches the address of any mobile node for which it has a Binding Cache entry marked as a home registration.

If such an entry exists in the home agent's Binding Cache, the home agent MUST reply to the Neighbor Solicitation with a Neighbor Advertisement giving the home agent's own link-layer address as the link-layer address for the specified Target Address. In addition, the Router (R) bit in the Advertisement MUST be set to zero. Acting as a proxy in this way allows other nodes on the mobile node's home link to resolve the mobile node's address and for the home agent to defend these addresses on the home link for Duplicate Address Detection [12].

10.4.2. Processing Intercepted Packets

For any packet sent to a mobile node from the mobile node's home agent (in which the home agent is the original sender of the packet), the home agent is operating as a correspondent node of the mobile node for this packet and the procedures described in Section 9.3.2 apply. The home agent then uses a routing header to route the packet to the mobile node by way of the primary care-of address in the home agent's Binding Cache.

While the mobile node is away from home, the home agent intercepts any packets on the home link addressed to the mobile node's home address, as described in Section 10.4.1. In order to forward each intercepted packet to the mobile node, the home agent MUST tunnel the packet to the mobile node using IPv6 encapsulation [15]. When a home agent encapsulates an intercepted packet for forwarding to the mobile node, the home agent sets the Source Address in the new tunnel IP header to the home agent's own IP address and sets the Destination Address in the tunnel IP header to the mobile node's primary care-of

address. When received by the mobile node, normal processing of the tunnel header [15] will result in decapsulation and processing of the original packet by the mobile node.

However, packets addressed to the mobile node's link-local address MUST NOT be tunneled to the mobile node. Instead, these packets MUST be discarded and the home agent SHOULD return an ICMP Destination Unreachable, Code 3, message to the packet's Source Address (unless this Source Address is a multicast address). Packets addressed to the mobile node's site-local address SHOULD NOT be tunneled to the mobile node by default.

Interception and tunneling of the following multicast addressed packets on the home network are only done if the home agent supports multicast group membership control messages from the mobile node as described in the next section. Tunneling of multicast packets to a mobile node follows similar limitations to those defined above for unicast packets addressed to the mobile node's link-local and site-local addresses. Multicast packets addressed to a multicast address with link-local scope [3], to which the mobile node is subscribed, MUST NOT be tunneled to the mobile node. These packets SHOULD be silently discarded (after delivering to other local multicast recipients). Multicast packets addressed to a multicast address with a scope larger than link-local, but smaller than global (e.g., site-local and organization-local [3], to which the mobile node is subscribed, SHOULD NOT be tunneled to the mobile node. Multicast packets addressed with a global scope, to which the mobile node has successfully subscribed, MUST be tunneled to the mobile node.

Before tunneling a packet to the mobile node, the home agent MUST perform any IPsec processing as indicated by the security policy data base.

10.4.3. Multicast Membership Control

This section is a prerequisite for the multicast data packet forwarding, described in the previous section. If this support is not provided, multicast group membership control messages are silently ignored.

In order to forward multicast data packets from the home network to all the proper mobile nodes, the home agent SHOULD be capable of receiving tunneled multicast group membership control information from the mobile node in order to determine which groups the mobile node has subscribed to. These multicast group membership messages are Listener Report messages specified in MLD [17] or in other protocols such as [37].

The messages are issued by the mobile node, but sent through the reverse tunnel to the home agent. These messages are issued whenever the mobile node decides to enable reception of packets for a multicast group or in response to an MLD Query from the home agent. The mobile node will also issue multicast group control messages to disable reception of multicast packets when it is no longer interested in receiving multicasts for a particular group.

To obtain the mobile node's current multicast group membership the home agent must periodically transmit MLD Query messages through the tunnel to the mobile node. These MLD periodic transmissions will ensure the home agent has an accurate record of the groups in which the mobile node is interested despite packet losses of the mobile node's MLD group membership messages.

All MLD packets are sent directly between the mobile node and the home agent. Since all of these packets are destined to a link-scope multicast address and have a hop limit of 1, there is no direct forwarding of such packets between the home network and the mobile node. The MLD packets between the mobile node and the home agent are encapsulated within the same tunnel header used for other packet flows between the mobile node and home agent.

Note that at this time, even though a link-local source is used on MLD packets, no functionality depends on these addresses being unique, nor do they elicit direct responses. All MLD messages are sent to multicast destinations. To avoid ambiguity on the home agent, due to mobile nodes which may choose identical link-local source addresses for their MLD function, it is necessary for the home agent to identify which mobile node was actually the issuer of a particular MLD message. This may be accomplished by noting which tunnel such an MLD arrived by, which IPsec SA was used, or by other distinguishing means.

This specification puts no requirement on how the functions in this section and the multicast forwarding in Section 10.4.2 are to be achieved. At the time of this writing it was thought that a full IPv6 multicast router function would be necessary on the home agent, but it may be possible to achieve the same effects through a "proxy MLD" application coupled with kernel multicast forwarding. This may be the subject of future specifications.

10.4.4. Stateful Address Autoconfiguration

This section describes how home agents support the use of stateful address autoconfiguration mechanisms such as DHCPv6 [29] from the mobile nodes. If this support is not provided, then the M and O bits must remain cleared on the Mobile Prefix Advertisement Messages. Any mobile node which sends DHCPv6 messages to the home agent without this support will not receive a response.

If DHCPv6 is used, packets are sent with link-local source addresses either to a link-scope multicast address or a link-local address. Mobile nodes desiring to locate a DHCPv6 service may reverse tunnel standard DHCPv6 packets to the home agent. Since these link-scope packets cannot be forwarded onto the home network, it is necessary for the home agent to either implement a DHCPv6 relay agent or a DHCPv6 server function itself. The arriving tunnel or IPsec SA of DHCPv6 link-scope messages from the mobile node must be noted so that DHCPv6 responses may be sent back to the appropriate mobile node. DHCPv6 messages sent to the mobile node with a link-local destination must be tunneled within the same tunnel header used for other packet flows.

10.4.5. Handling Reverse Tunneled Packets

Unless a binding has been established between the mobile node and a correspondent node, traffic from the mobile node to the correspondent node goes through a reverse tunnel. Home agents MUST support reverse tunneling as follows:

- o The tunneled traffic arrives to the home agent's address using IPv6 encapsulation [15].
- o Depending on the security policies used by the home agent, reverse tunneled packets MAY be discarded unless accompanied by a valid ESP header. The support for authenticated reverse tunneling allows the home agent to protect the home network and correspondent nodes from malicious nodes masquerading as a mobile node.
- o Otherwise, when a home agent decapsulates a tunneled packet from the mobile node, the home agent MUST verify that the Source Address in the tunnel IP header is the mobile node's primary care-of address. Otherwise, any node in the Internet could send traffic through the home agent and escape ingress filtering limitations. This simple check forces the attacker to know the current location of the real mobile node and be able to defeat ingress filtering. This check is not necessary if the reverse-tunneled packet is protected by ESP in tunnel mode.

10.4.6. Protecting Return Routability Packets

The return routability procedure, described in Section 5.2.5, assumes that the confidentiality of the Home Test Init and Home Test messages is protected as they are tunneled between the home agent and the mobile node. Therefore, the home agent MUST support tunnel mode IPsec ESP for the protection of packets belonging to the return routability procedure. Support for a non-null encryption transform and authentication algorithm MUST be available. It is not necessary to distinguish between different kinds of packets during the return routability procedure.

Security associations are needed to provide this protection. When the care-of address for the mobile node changes as a result of an accepted Binding Update, special treatment is needed for the next packets sent using these security associations. The home agent MUST set the new care-of address as the destination address of these packets, as if the outer header destination address in the security association had changed [21].

The above protection SHOULD be used with all mobile nodes. The use is controlled by configuration of the IPsec security policy database both at the mobile node and at the home agent.

As described earlier, the Binding Update and Binding Acknowledgement messages require protection between the home agent and the mobile node. The Mobility Header protocol carries both these messages as well as the return routability messages. From the point of view of the security policy database these messages are indistinguishable. When IPsec is used to protect return routability signaling or payload packets, this protection MUST only be applied to the return routability packets entering the IPv6 encapsulated tunnel interface between the mobile node and the home agent. This can be achieved, for instance, by defining the security policy database entries specifically for the tunnel interface. That is, the policy entries are not generally applied on all traffic on the physical interface(s) of the nodes, but rather only on traffic that enters the tunnel. This makes use of per-interface security policy database entries [4] specific to the tunnel interface (the node's attachment to the tunnel [11]).

10.5. Dynamic Home Agent Address Discovery

This section describes how a home agent can help mobile nodes to discover the addresses of the home agents. The home agent keeps track of the other home agents on the same link and responds to queries sent by the mobile node.

10.5.1. Receiving Router Advertisement Messages

For each link on which a router provides service as a home agent, the router maintains a Home Agents List recording information about all other home agents on that link. This list is used in the dynamic home agent address discovery mechanism, described in Section 10.5. The information for the list is learned through receipt of the periodic unsolicited multicast Router Advertisements, in a manner similar to the Default Router List conceptual data structure maintained by each host for Neighbor Discovery [12]. In the construction of the Home Agents List, the Router Advertisements are from each (other) home agent on the link and the Home Agent (H) bit is set in them.

On receipt of a valid Router Advertisement, as defined in the processing algorithm specified for Neighbor Discovery [12], the home agent performs the following steps in addition to any steps already required of it by Neighbor Discovery:

- o If the Home Agent (H) bit in the Router Advertisement is not set, delete the sending node's entry in the current Home Agents List (if one exists). Skip all the following steps.
- o Otherwise, extract the Source Address from the IP header of the Router Advertisement. This is the link-local IP address on this link of the home agent sending this Advertisement [12].
- o Determine the preference for this home agent. If the Router Advertisement contains a Home Agent Information Option, then the preference is taken from the Home Agent Preference field in the option; otherwise, the default preference of 0 MUST be used.
- o Determine the lifetime for this home agent. If the Router Advertisement contains a Home Agent Information Option, then the lifetime is taken from the Home Agent Lifetime field in the option; otherwise, the lifetime specified by the Router Lifetime field in the Router Advertisement SHOULD be used.
- o If the link-local address of the home agent sending this Advertisement is already present in this home agent's Home Agents List and the received home agent lifetime value is zero, immediately delete this entry in the Home Agents List.
- o Otherwise, if the link-local address of the home agent sending this Advertisement is already present in the receiving home agent's Home Agents List, reset its lifetime and preference to the values determined above.

- o If the link-local address of the home agent sending this Advertisement is not already present in the Home Agents List maintained by the receiving home agent, and the lifetime for the sending home agent is non-zero, create a new entry in the list, and initialize its lifetime and preference to the values determined above.
- o If the Home Agents List entry for the link-local address of the home agent sending this Advertisement was not deleted as described above, determine any global address(es) of the home agent based on each Prefix Information option received in this Advertisement in which the Router Address (R) bit is set (Section 7.2). Add all such global addresses to the list of global addresses in this Home Agents List entry.

A home agent SHOULD maintain an entry in its Home Agents List for each valid home agent address until that entry's lifetime expires, after which time the entry MUST be deleted.

As described in Section 11.4.1, a mobile node attempts dynamic home agent address discovery by sending an ICMP Home Agent Address Discovery Request message to the Mobile IPv6 Home-Agents anycast address [16] for its home IP subnet prefix. A home agent receiving a Home Agent Address Discovery Request message that serves this subnet SHOULD return an ICMP Home Agent Address Discovery Reply message to the mobile node with the Source Address of the Reply packet set to one of the global unicast addresses of the home agent. The Home Agent Addresses field in the Reply message is constructed as follows:

- o The Home Agent Addresses field SHOULD contain all global IP addresses for each home agent currently listed in this home agent's own Home Agents List (Section 10.1).
- o The IP addresses in the Home Agent Addresses field SHOULD be listed in order of decreasing preference values, based either on the respective advertised preference from a Home Agent Information option or on the default preference of 0 if no preference is advertised (or on the configured home agent preference for this home agent itself).
- o Among home agents with equal preference, their IP addresses in the Home Agent Addresses field SHOULD be listed in an order randomized with respect to other home agents with equal preference every time a Home Agent Address Discovery Reply message is returned by this home agent.
- o If more than one global IP address is associated with a home agent, these addresses SHOULD be listed in a randomized order.

- o The home agent SHOULD reduce the number of home agent IP addresses so that the packet fits within the minimum IPv6 MTU [11]. The home agent addresses selected for inclusion in the packet SHOULD be those from the complete list with the highest preference. This limitation avoids the danger of the Reply message packet being fragmented (or rejected by an intermediate router with an ICMP Packet Too Big message [14]).

10.6. Sending Prefix Information to the Mobile Node

10.6.1. List of Home Network Prefixes

Mobile IPv6 arranges to propagate relevant prefix information to the mobile node when it is away from home, so that it may be used in mobile node home address configuration and in network renumbering. In this mechanism, mobile nodes away from home receive Mobile Prefix Advertisements messages. These messages include Prefix Information Options for the prefixes configured on the home subnet interface(s) of the home agent.

If there are multiple home agents, differences in the advertisements sent by different home agents can lead to an inability to use a particular home address when changing to another home agent. In order to ensure that the mobile nodes get the same information from different home agents, it is preferred that all of the home agents on the same link be configured in the same manner.

To support this, the home agent monitors prefixes advertised by itself and other home agents on the home link. In RFC 2461 [12] it is acceptable for two routers to advertise different sets of prefixes on the same link. For home agents, the differences should be detected for a given home address because the mobile node communicates only with one home agent at a time and the mobile node needs to know the full set of prefixes assigned to the home link. All other comparisons of Router Advertisements are as specified in Section 6.2.7 of RFC 2461.

10.6.2. Scheduling Prefix Deliveries

A home agent serving a mobile node will schedule the delivery of the new prefix information to that mobile node when any of the following conditions occur:

MUST:

- o The state of the flags changes for the prefix of the mobile node's registered home address.

- o The valid or preferred lifetime is reconfigured or changes for any reason other than advancing real time.
- o The mobile node requests the information with a Mobile Prefix Solicitation (see Section 11.4.2).

SHOULD:

- o A new prefix is added to the home subnet interface(s) of the home agent.

MAY:

- o The valid or preferred lifetime or the state of the flags changes for a prefix which is not used in any Binding Cache entry for this mobile node.

The home agent uses the following algorithm to determine when to send prefix information to the mobile node.

- o If a mobile node sends a solicitation, answer right away.
- o If no Mobile Prefix Advertisement has been sent to the mobile node in the last MaxMobPfxAdvInterval seconds (see Section 13), then ensure that a transmission is scheduled. The actual transmission time is randomized as described below.
- o If a prefix matching the mobile node's home registration is added on the home subnet interface or if its information changes in any way that does not deprecate the mobile node's address, ensure that a transmission is scheduled. The actual transmission time is randomized as described below.
- o If a home registration expires, cancel any scheduled advertisements to the mobile node.

The list of prefixes is sent in its entirety in all cases.

If the home agent has already scheduled the transmission of a Mobile Prefix Advertisement to the mobile node, then the home agent will replace the advertisement with a new one to be sent at the scheduled time.

Otherwise, the home agent computes a fresh value for RAND_ADV_DELAY which offsets from the current time for the scheduled transmission. First calculate the maximum delay for the scheduled Advertisement:

$\text{MaxScheduleDelay} = \min(\text{MaxMobPfxAdvInterval}, \text{Preferred Lifetime}),$

where $\text{MaxMobPfxAdvInterval}$ is as defined in Section 12. Then compute the final delay for the advertisement:

$\text{RAND_ADV_DELAY} = \text{MinMobPfxAdvInterval} +$
 $(\text{rand}() \% \text{abs}(\text{MaxScheduleDelay} - \text{MinMobPfxAdvInterval}))$

Here $\text{rand}()$ returns a random integer value in the range of 0 to the maximum possible integer value. This computation is expected to alleviate bursts of advertisements when prefix information changes. In addition, a home agent MAY further reduce the rate of packet transmission by further delaying individual advertisements, when necessary to avoid overwhelming local network resources. The home agent SHOULD periodically continue to retransmit an unsolicited Advertisement to the mobile node, until it is acknowledged by the receipt of a Mobile Prefix Solicitation from the mobile node.

The home agent MUST wait $\text{PREFIX_ADV_TIMEOUT}$ (see Section 12) before the first retransmission and double the retransmission wait time for every succeeding retransmission until a maximum number of $\text{PREFIX_ADV_RETRIES}$ attempts (see Section 12) has been tried. If the mobile node's bindings expire before the matching Binding Update has been received, then the home agent MUST NOT attempt any more retransmissions, even if not all $\text{PREFIX_ADV_RETRIES}$ have been retransmitted. In the meantime, if the mobile node sends another Binding Update without returning home, then the home agent SHOULD begin transmitting the unsolicited Advertisement again.

If some condition, as described above, occurs on the home link and causes another Prefix Advertisement to be sent to the mobile node, before the mobile node acknowledges a previous transmission, the home agent SHOULD combine any Prefix Information options in the unacknowledged Mobile Prefix Advertisement into a new Advertisement. The home agent then discards the old Advertisement.

10.6.3. Sending Advertisements

When sending a Mobile Prefix Advertisement to the mobile node, the home agent MUST construct the packet as follows:

- o The Source Address in the packet's IPv6 header MUST be set to the home agent's IP address to which the mobile node addressed its current home registration or its default global home agent address if no binding exists.

- o If the advertisement was solicited, it MUST be destined to the source address of the solicitation. If it was triggered by prefix changes or renumbering, the advertisement's destination will be the mobile node's home address in the binding which triggered the rule.
- o A type 2 routing header MUST be included with the mobile node's home address.
- o IPsec headers MUST be supported and SHOULD be used.
- o The home agent MUST send the packet as it would any other unicast IPv6 packet that it originates.
- o Set the Managed Address Configuration (M) flag if the corresponding flag has been set in any of the Router Advertisements from which the prefix information has been learned (including the ones sent by this home agent).
- o Set the Other Stateful Configuration (O) flag if the corresponding flag has been set in any of the Router Advertisements from which the prefix information has been learned (including the ones sent by this home agent).

10.6.4. Lifetimes for Changed Prefixes

As described in Section 10.3.1, the lifetime returned by the home agent in a Binding Acknowledgement MUST not be greater than the remaining valid lifetime for the subnet prefix in the mobile node's home address. This limit on the binding lifetime serves to prohibit use of a mobile node's home address after it becomes invalid.

11. Mobile Node Operation

11.1. Conceptual Data Structures

Each mobile node MUST maintain a Binding Update List.

The Binding Update List records information for each Binding Update sent by this mobile node, in which the lifetime of the binding has not yet expired. The Binding Update List includes all bindings sent by the mobile node either to its home agent or correspondent nodes. It also contains Binding Updates which are waiting for the completion of the return routability procedure before they can be sent. However, for multiple Binding Updates sent to the same destination address, the Binding Update List contains only the most recent Binding Update (i.e., with the greatest Sequence Number value) sent to that destination. The Binding Update List MAY be implemented in

any manner consistent with the external behavior described in this document.

Each Binding Update List entry conceptually contains the following fields:

- o The IP address of the node to which a Binding Update was sent.
- o The home address for which that Binding Update was sent.
- o The care-of address sent in that Binding Update. This value is necessary for the mobile node to determine if it has sent a Binding Update while giving its new care-of address to this destination after changing its care-of address.
- o The initial value of the Lifetime field sent in that Binding Update.
- o The remaining lifetime of that binding. This lifetime is initialized from the Lifetime value sent in the Binding Update and is decremented until it reaches zero, at which time this entry MUST be deleted from the Binding Update List.
- o The maximum value of the Sequence Number field sent in previous Binding Updates to this destination. The Sequence Number field is 16 bits long and all comparisons between Sequence Number values MUST be performed modulo 2^{16} (see Section 9.5.1).
- o The time at which a Binding Update was last sent to this destination, as needed to implement the rate limiting restriction for sending Binding Updates.
- o The state of any retransmissions needed for this Binding Update. This state includes the time remaining until the next retransmission attempt for the Binding Update and the current state of the exponential back-off mechanism for retransmissions.
- o A flag specifying whether or not future Binding Updates should be sent to this destination. The mobile node sets this flag in the Binding Update List entry when it receives an ICMP Parameter Problem, Code 1, error message in response to a return routability message or Binding Update sent to that destination, as described in Section 11.3.5.

The Binding Update List is used to determine whether a particular packet is sent directly to the correspondent node or tunneled via the home agent (see Section 11.3.1).

The Binding Update list also conceptually contains the following data related to running the return routability procedure. This data is relevant only for Binding Updates sent to correspondent nodes.

- o The time at which a Home Test Init or Care-of Test Init message was last sent to this destination, as needed to implement the rate limiting restriction for the return routability procedure.
- o The state of any retransmissions needed for this return routability procedure. This state includes the time remaining until the next retransmission attempt and the current state of the exponential back-off mechanism for retransmissions.
- o Cookie values used in the Home Test Init and Care-of Test Init messages.
- o Home and care-of keygen tokens received from the correspondent node.
- o Home and care-of nonce indices received from the correspondent node.
- o The time at which each of the tokens and nonces were received from the correspondent node, as needed to implement reuse while moving.

11.2. Processing Mobility Headers

All IPv6 mobile nodes MUST observe the rules described in Section 9.2 when processing Mobility Headers.

11.3. Packet Processing

11.3.1. Sending Packets While Away from Home

While a mobile node is away from home, it continues to use its home address, as well as also using one or more care-of addresses. When sending a packet while away from home, a mobile node MAY choose among these in selecting the address that it will use as the source of the packet, as follows:

- o Protocols layered over IP will generally treat the mobile node's home address as its IP address for most packets. For packets sent that are part of transport-level connections established while the mobile node was at home, the mobile node MUST use its home address. Likewise, for packets sent that are part of transport-level connections that the mobile node may still be using after moving to a new location, the mobile node SHOULD use its home address in this way. If a binding exists, the mobile node SHOULD

send the packets directly to the correspondent node. Otherwise, if a binding does not exist, the mobile node MUST use reverse tunneling.

- o The mobile node MAY choose to directly use one of its care-of addresses as the source of the packet, not requiring the use of a Home Address option in the packet. This is particularly useful for short-term communication that may easily be retried if it fails. Using the mobile node's care-of address as the source for such queries will generally have a lower overhead than using the mobile node's home address, since no extra options need be used in either the query or its reply. Such packets can be routed normally, directly between their source and destination without relying on Mobile IPv6. If application running on the mobile node has no particular knowledge that the communication being sent fits within this general type of communication, however, the mobile node should not use its care-of address as the source of the packet in this way.

The choice of the most efficient communications method is application specific, and outside the scope of this specification. The APIs necessary for controlling the choice are also out of scope.

- o While not at its home link, the mobile node MUST NOT use the Home Address destination option when communicating with link-local or site-local peers, if the scope of the home address is larger than the scope of the peer's address.

Similarly, the mobile node MUST NOT use the Home Address destination option for IPv6 Neighbor Discovery [12] packets.

Detailed operation of these cases is described later in this section and also discussed in [31].

For packets sent by a mobile node while it is at home, no special Mobile IPv6 processing is required. Likewise, if the mobile node uses any address other than one of its home addresses as the source of a packet sent while away from home, no special Mobile IPv6 processing is required. In either case, the packet is simply addressed and transmitted in the same way as any normal IPv6 packet.

For packets sent by the mobile node sent while away from home using the mobile node's home address as the source, special Mobile IPv6 processing of the packet is required. This can be done in the following two ways:

Route Optimization

This manner of delivering packets does not require going through the home network, and typically will enable faster and more reliable transmission.

The mobile node needs to ensure that a Binding Cache entry exists for its home address so that the correspondent node can process the packet (Section 9.3.1 specifies the rules for Home Address Destination Option Processing at a correspondent node). The mobile node SHOULD examine its Binding Update List for an entry which fulfills the following conditions:

- * The Source Address field of the packet being sent is equal to the home address in the entry.
- * The Destination Address field of the packet being sent is equal to the address of the correspondent node in the entry.
- * One of the current care-of addresses of the mobile node appears as the care-of address in the entry.
- * The entry indicates that a binding has been successfully created.
- * The remaining lifetime of the binding is greater than zero.

When these conditions are met, the mobile node knows that the correspondent node has a suitable Binding Cache entry.

A mobile node SHOULD arrange to supply the home address in a Home Address option, and MUST set the IPv6 header's Source Address field to the care-of address which the mobile node has registered to be used with this correspondent node. The correspondent node will then use the address supplied in the Home Address option to serve the function traditionally done by the Source IP address in the IPv6 header. The mobile node's home address is then supplied to higher protocol layers and applications.

Specifically:

- * Construct the packet using the mobile node's home address as the packet's Source Address, in the same way as if the mobile node were at home. This includes the calculation of upper layer checksums using the home address as the value of the source.
- * Insert a Home Address option into the packet with the Home Address field copied from the original value of the Source Address field in the packet.

- * Change the Source Address field in the packet's IPv6 header to one of the mobile node's care-of addresses. This will typically be the mobile node's current primary care-of address, but MUST be an address assigned to the interface on the link being used.

By using the care-of address as the Source Address in the IPv6 header, with the mobile node's home address instead in the Home Address option, the packet will be able to safely pass through any router implementing ingress filtering [26].

Reverse Tunneling

This is the mechanism which tunnels the packets via the home agent. It is not as efficient as the above mechanism, but is needed if there is no binding yet with the correspondent node.

This mechanism is used for packets that have the mobile node's home address as the Source Address in the IPv6 header, or with multicast control protocol packets as described in Section 11.3.4. Specifically:

- * The packet is sent to the home agent using IPv6 encapsulation [15].
- * The Source Address in the tunnel packet is the primary care-of address as registered with the home agent.
- * The Destination Address in the tunnel packet is the home agent's address.

Then, the home agent will pass the encapsulated packet to the correspondent node.

11.3.2. Interaction with Outbound IPsec Processing

This section sketches the interaction between outbound Mobile IPv6 processing and outbound IP Security (IPsec) processing for packets sent by a mobile node while away from home. Any specific implementation MAY use algorithms and data structures other than those suggested here, but its processing MUST be consistent with the effect of the operation described here and with the relevant IPsec specifications. In the steps described below, it is assumed that IPsec is being used in transport mode [4] and that the mobile node is using its home address as the source for the packet (from the point of view of higher protocol layers or applications, as described in Section 11.3.1):

- o The packet is created by higher layer protocols and applications (e.g., by TCP) as if the mobile node were at home and Mobile IPv6 were not being used.
- o Determine the outgoing interface for the packet. (Note that the selection between reverse tunneling and route optimization may imply different interfaces, particularly if tunnels are considered interfaces as well.)
- o As part of outbound packet processing in IP, the packet is compared against the IPsec security policy database to determine what processing is required for the packet [4].
- o If IPsec processing is required, the packet is either mapped to an existing Security Association (or SA bundle), or a new SA (or SA bundle) is created for the packet, according to the procedures defined for IPsec.
- o Since the mobile node is away from home, the mobile is either using reverse tunneling or route optimization to reach the correspondent node.

If reverse tunneling is used, the packet is constructed in the normal manner and then tunneled through the home agent.

If route optimization is in use, the mobile node inserts a Home Address destination option into the packet, replacing the Source Address in the packet's IP header with the care-of address used with this correspondent node, as described in Section 11.3.1. The Destination Options header in which the Home Address destination option is inserted MUST appear in the packet after the routing header, if present, and before the IPsec (AH [5] or ESP [6]) header, so that the Home Address destination option is processed by the destination node before the IPsec header is processed.

Finally, once the packet is fully assembled, the necessary IPsec authentication (and encryption, if required) processing is performed on the packet, initializing the Authentication Data in the IPsec header.

RFC 2402 treatment of destination options is extended as follows. The AH authentication data MUST be calculated as if the following were true:

- * the IPv6 source address in the IPv6 header contains the mobile node's home address,

- * the Home Address field of the Home Address destination option (Section 6.3) contains the new care-of address.
- o This allows, but does not require, the receiver of the packet containing a Home Address destination option to exchange the two fields of the incoming packet to reach the above situation, simplifying processing for all subsequent packet headers. However, such an exchange is not required, as long as the result of the authentication calculation remains the same.

When an automated key management protocol is used to create new security associations for a peer, it is important to ensure that the peer can send the key management protocol packets to the mobile node. This may not be possible if the peer is the home agent of the mobile node and the purpose of the security associations would be to send a Binding Update to the home agent. Packets addressed to the home address of the mobile node cannot be used before the Binding Update has been processed. For the default case of using IKE [9] as the automated key management protocol, such problems can be avoided by the following requirements when communicating with its home agent:

- o When the mobile node is away from home, it MUST use its care-of address as the Source Address of all packets it sends as part of the key management protocol (without use of Mobile IPv6 for these packets, as suggested in Section 11.3.1).
- o In addition, for all security associations bound to the mobile node's home address established by IKE, the mobile node MUST include an ISAKMP Identification Payload [8] in the IKE phase 2 exchange, giving the mobile node's home address as the initiator of the Security Association [7].

The Key Management Mobility Capability (K) bit in Binding Updates and Acknowledgements can be used to avoid the need to rerun IKE upon movements.

11.3.3. Receiving Packets While Away from Home

While away from home, a mobile node will receive packets addressed to its home address, by one of two methods:

- o Packets sent by a correspondent node, that does not have a Binding Cache entry for the mobile node, will be sent to the home address, captured by the home agent and tunneled to the mobile node.
- o Packets sent by a correspondent node that has a Binding Cache entry for the mobile node that contains the mobile node's current care-of address, will be sent by the correspondent node using a

type 2 routing header. The packet will be addressed to the mobile node's care-of address, with the final hop in the routing header directing the packet to the mobile node's home address; the processing of this last hop of the routing header is entirely internal to the mobile node, since the care-of address and home address are both addresses within the mobile node.

For packets received by the first method, the mobile node MUST check that the IPv6 source address of the tunneled packet is the IP address of its home agent. In this method, the mobile node may also send a Binding Update to the original sender of the packet as described in Section 11.7.2 and subject to the rate limiting defined in Section 11.8. The mobile node MUST also process the received packet in the manner defined for IPv6 encapsulation [15], which will result in the encapsulated (inner) packet being processed normally by upper-layer protocols within the mobile node as if it had been addressed (only) to the mobile node's home address.

For packets received by the second method, the following rules will result in the packet being processed normally by upper-layer protocols within the mobile node as if it had been addressed to the mobile node's home address.

A node receiving a packet addressed to itself (i.e., one of the node's addresses is in the IPv6 destination field) follows the next header chain of headers and processes them. When it encounters a type 2 routing header during this processing, it performs the following checks. If any of these checks fail, the node MUST silently discard the packet.

- o The length field in the routing header is exactly 2.
- o The segments left field in the routing header is 1 on the wire. (But implementations may process the routing header so that the value may become 0 after the routing header has been processed, but before the rest of the packet is processed.)
- o The Home Address field in the routing header is one of the node's home addresses, if the segments left field was 1. Thus, in particular the address field is required to be a unicast routable address.

Once the above checks have been performed, the node swaps the IPv6 destination field with the Home Address field in the routing header, decrements segments left by one from the value it had on the wire, and resubmits the packet to IP for processing the next header.

Conceptually, this follows the same model as in RFC 2460. However, in the case of type 2 routing header this can be simplified since it is known that the packet will not be forwarded to a different node.

The definition of AH requires the sender to calculate the AH integrity check value of a routing header in the same way it appears in the receiver after it has processed the header. Since IPsec headers follow the routing header, any IPsec processing will operate on the packet with the home address in the IP destination field and segments left being zero. Thus, the AH calculations at the sender and receiver will have an identical view of the packet.

11.3.4. Routing Multicast Packets

A mobile node that is connected to its home link functions in the same way as any other (stationary) node. Thus, when it is at home, a mobile node functions identically to other multicast senders and receivers. Therefore, this section describes the behavior of a mobile node that is not on its home link.

In order to receive packets sent to some multicast group, a mobile node must join that multicast group. One method, in which a mobile node MAY join the group, is via a (local) multicast router on the foreign link being visited. In this case, the mobile node MUST use its care-of address and MUST NOT use the Home Address destination option when sending MLD packets [17].

Alternatively, a mobile node MAY join multicast groups via a bi-directional tunnel to its home agent. The mobile node tunnels its multicast group membership control packets (such as those defined in [17] or in [37]) to its home agent, and the home agent forwards multicast packets down the tunnel to the mobile node. A mobile node MUST NOT tunnel multicast group membership control packets until (1) the mobile node has a binding in place at the home agent, and (2) the latter sends at least one multicast group membership control packet via the tunnel. Once this condition is true, the mobile node SHOULD assume it does not change as long as the binding does not expire.

A mobile node that wishes to send packets to a multicast group also has two options:

1. Send directly on the foreign link being visited.

The application is aware of the care-of address and uses it as a source address for multicast traffic, just like it would use a stationary address. The mobile node MUST NOT use Home Address destination option in such traffic.

2. Send via a tunnel to its home agent.

Because multicast routing in general depends upon the Source Address used in the IPv6 header of the multicast packet, a mobile node that tunnels a multicast packet to its home agent MUST use its home address as the IPv6 Source Address of the inner multicast packet.

Note that direct sending from the foreign link is only applicable while the mobile node is at that foreign link. This is because the associated multicast tree is specific to that source location and any change of location and source address will invalidate the source specific tree or branch and the application context of the other multicast group members.

This specification does not provide mechanisms to enable such local multicast session to survive hand-off and to seamlessly continue from a new care-of address on each new foreign link. Any such mechanism, developed as an extension to this specification, needs to take into account the impact of fast moving mobile nodes on the Internet multicast routing protocols and their ability to maintain the integrity of source specific multicast trees and branches.

While the use of bidirectional tunneling can ensure that multicast trees are independent of the mobile nodes movement, in some case such tunneling can have adverse affects. The latency of specific types of multicast applications (such as multicast based discovery protocols) will be affected when the round-trip time between the foreign subnet and the home agent is significant compared to that of the topology to be discovered. In addition, the delivery tree from the home agent in such circumstances relies on unicast encapsulation from the agent to the mobile node. Therefore, bandwidth usage is inefficient compared to the native multicast forwarding in the foreign multicast system.

11.3.5. Receiving ICMP Error Messages

Any node that does not recognize the Mobility header will return an ICMP Parameter Problem, Code 1, message to the sender of the packet. If the mobile node receives such an ICMP error message in response to a return routability procedure or Binding Update, it SHOULD record in its Binding Update List that future Binding Updates SHOULD NOT be sent to this destination. Such Binding Update List entries SHOULD be removed after a period of time in order to allow for retrying route optimization.

New Binding Update List entries MUST NOT be created as a result of receiving ICMP error messages.

Correspondent nodes that have participated in the return routability procedure MUST implement the ability to correctly process received packets containing a Home Address destination option. Therefore, correctly implemented correspondent nodes should always be able to recognize Home Address options. If a mobile node receives an ICMP Parameter Problem, Code 2, message from some node indicating that it does not support the Home Address option, the mobile node SHOULD log the error and then discard the ICMP message.

11.3.6. Receiving Binding Error Messages

When a mobile node receives a packet containing a Binding Error message, it should first check if the mobile node has a Binding Update List entry for the source of the Binding Error message. If the mobile node does not have such an entry, it MUST ignore the message. This is necessary to prevent a waste of resources on, e.g., return routability procedure due to spoofed Binding Error messages.

Otherwise, if the message Status field was 1 (unknown binding for Home Address destination option), the mobile node should perform one of the following two actions:

- o If the mobile node has recent upper layer progress information, which indicates that communications with the correspondent node are progressing, it MAY ignore the message. This can be done in order to limit the damage that spoofed Binding Error messages can cause to ongoing communications.
- o If the mobile node has no upper layer progress information, it MUST remove the entry and route further communications through the home agent. It MAY also optionally start a return routability procedure (see Section 5.2).

If the message Status field was 2 (unrecognized MH Type value), the mobile node should perform one of the following two actions:

- o If the mobile node is not expecting an acknowledgement or response from the correspondent node, the mobile node SHOULD ignore this message.
- o Otherwise, the mobile node SHOULD cease the use of any extensions to this specification. If no extensions had been used, the mobile node should cease the attempt to use route optimization.

11.4. Home Agent and Prefix Management

11.4.1. Dynamic Home Agent Address Discovery

Sometimes when the mobile node needs to send a Binding Update to its home agent to register its new primary care-of address, as described in Section 11.7.1, the mobile node may not know the address of any router on its home link that can serve as a home agent for it. For example, some nodes on its home link may have been reconfigured while the mobile node has been away from home, such that the router that was operating as the mobile node's home agent has been replaced by a different router serving this role.

In this case, the mobile node MAY attempt to discover the address of a suitable home agent on its home link. To do so, the mobile node sends an ICMP Home Agent Address Discovery Request message to the Mobile IPv6 Home-Agents anycast address [16] for its home subnet prefix. As described in Section 10.5, the home agent on its home link that receives this Request message will return an ICMP Home Agent Address Discovery Reply message. This message gives the addresses for the home agents operating on the home link.

The mobile node, upon receiving this Home Agent Address Discovery Reply message, MAY then send its home registration Binding Update to any of the unicast IP addresses listed in the Home Agent Addresses field in the Reply. For example, the mobile node MAY attempt its home registration to each of these addresses, in turn, until its registration is accepted. The mobile node sends a Binding Update to an address and waits for the matching Binding Acknowledgement, moving on to the next address if there is no response. The mobile node MUST, however, wait at least InitialBindackTimeoutFirstReg seconds (see Section 13) before sending a Binding Update to the next home agent. In trying each of the returned home agent addresses, the mobile node SHOULD try each of them in the order they appear in the Home Agent Addresses field in the received Home Agent Address Discovery Reply message.

If the mobile node has a current registration with some home agent (the Lifetime for that registration has not yet expired), then the mobile node MUST attempt any new registration first with that home agent. If that registration attempt fails (e.g., timed out or rejected), the mobile node SHOULD then reattempt this registration with another home agent. If the mobile node knows of no other suitable home agent, then it MAY attempt the dynamic home agent address discovery mechanism described above.

If, after a mobile node transmits a Home Agent Address Discovery Request message to the Home Agents Anycast address, it does not receive a corresponding Home Agent Address Discovery Reply message within INITIAL_DHAAD_TIMEOUT (see Section 12) seconds, the mobile node MAY retransmit the same Request message to the same anycast address. This retransmission MAY be repeated up to a maximum of DHAAD_RETRIES (see Section 12) attempts. Each retransmission MUST be delayed by twice the time interval of the previous retransmission.

11.4.2. Sending Mobile Prefix Solicitations

When a mobile node has a home address that is about to become invalid, it SHOULD send a Mobile Prefix Solicitation to its home agent in an attempt to acquire fresh routing prefix information. The new information also enables the mobile node to participate in renumbering operations affecting the home network, as described in Section 10.6.

The mobile node MUST use the Home Address destination option to carry its home address. The mobile node MUST support and SHOULD use IPsec to protect the solicitation. The mobile node MUST set the Identifier field in the ICMP header to a random value.

As described in Section 11.7.2, Binding Updates sent by the mobile node to other nodes MUST use a lifetime no greater than the remaining lifetime of its home registration of its primary care-of address. The mobile node SHOULD further limit the lifetimes that it sends on any Binding Updates to be within the remaining valid lifetime (see Section 10.6.2) for the prefix in its home address.

When the lifetime for a changed prefix decreases, and the change would cause cached bindings at correspondent nodes in the Binding Update List to be stored past the newly shortened lifetime, the mobile node MUST issue a Binding Update to all such correspondent nodes.

These limits on the binding lifetime serve to prohibit use of a mobile node's home address after it becomes invalid.

11.4.3. Receiving Mobile Prefix Advertisements

Section 10.6 describes the operation of a home agent to support boot time configuration and renumbering a mobile node's home subnet while the mobile node is away from home. The home agent sends Mobile Prefix Advertisements to the mobile node while away from home, giving "important" Prefix Information options that describe changes in the prefixes in use on the mobile node's home link.

The Mobile Prefix Solicitation is similar to the Router Solicitation used in Neighbor Discovery [12], except it is routed from the mobile node on the visited network to the home agent on the home network by usual unicast routing rules.

When a mobile node receives a Mobile Prefix Advertisement, it MUST validate it according to the following test:

- o The Source Address of the IP packet carrying the Mobile Prefix Advertisement is the same as the home agent address to which the mobile node last sent an accepted home registration Binding Update to register its primary care-of address. Otherwise, if no such registrations have been made, it SHOULD be the mobile node's stored home agent address, if one exists. Otherwise, if the mobile node has not yet discovered its home agent's address, it MUST NOT accept Mobile Prefix Advertisements.
- o The packet MUST have a type 2 routing header and SHOULD be protected by an IPsec header as described in Section 5.4 and Section 6.8.
- o If the ICMP Identifier value matches the ICMP Identifier value of the most recently sent Mobile Prefix Solicitation and no other advertisement has yet been received for this value, then the advertisement is considered to be solicited and will be processed further.

Otherwise, the advertisement is unsolicited, and MUST be discarded. In this case the mobile node SHOULD send a Mobile Prefix Solicitation.

Any received Mobile Prefix Advertisement not meeting these tests MUST be silently discarded.

For an accepted Mobile Prefix Advertisement, the mobile node MUST process Managed Address Configuration (M), Other Stateful Configuration (O), and the Prefix Information Options as if they arrived in a Router Advertisement [12] on the mobile node's home link. (This specification does not, however, describe how to acquire home addresses through stateful protocols.) Such processing may result in the mobile node configuring a new home address, although due to separation between preferred lifetime and valid lifetime, such changes should not affect most communications by the mobile node, in the same way as for nodes that are at home.

This specification assumes that any security associations and security policy entries that may be needed for new prefixes have been pre-configured in the mobile node. Note that while dynamic key

management avoids the need to create new security associations, it is still necessary to add policy entries to protect the communications involving the home address(es). Mechanisms for automatic set-up of these entries are outside the scope of this specification.

11.5. Movement

11.5.1. Movement Detection

The primary goal of movement detection is to detect L3 handovers. This section does not attempt to specify a fast movement detection algorithm which will function optimally for all types of applications, link-layers and deployment scenarios; instead, it describes a generic method that uses the facilities of IPv6 Neighbor Discovery, including Router Discovery and Neighbor Unreachability Detection. At the time of this writing, this method is considered well enough understood to recommend for standardization, however it is expected that future versions of this specification or other specifications may contain updated versions of the movement detection algorithm that have better performance.

Generic movement detection uses Neighbor Unreachability Detection to detect when the default router is no longer bi-directionally reachable, in which case the mobile node must discover a new default router (usually on a new link). However, this detection only occurs when the mobile node has packets to send, and in the absence of frequent Router Advertisements or indications from the link-layer, the mobile node might become unaware of an L3 handover that occurred. Therefore, the mobile node should supplement this method with other information whenever it is available to the mobile node (e.g., from lower protocol layers).

When the mobile node detects an L3 handover, it performs Duplicate Address Detection [13] on its link-local address, selects a new default router as a consequence of Router Discovery, and then performs Prefix Discovery with that new router to form new care-of address(es) as described in Section 11.5.2. It then registers its new primary care-of address with its home agent as described in Section 11.7.1. After updating its home registration, the mobile node then updates associated mobility bindings in correspondent nodes that it is performing route optimization with as specified in Section 11.7.2.

Due to the temporary packet flow disruption and signaling overhead involved in updating mobility bindings, the mobile node should avoid performing an L3 handover until it is strictly necessary. Specifically, when the mobile node receives a Router Advertisement from a new router that contains a different set of on-link prefixes,

if the mobile node detects that the currently selected default router on the old link is still bi-directionally reachable, it should generally continue to use the old router on the old link rather than switch away from it to use a new default router.

Mobile nodes can use the information in received Router Advertisements to detect L3 handovers. In doing so the mobile node needs to consider the following issues:

- o There might be multiple routers on the same link, thus hearing a new router does not necessarily constitute an L3 handover.
- o When there are multiple routers on the same link they might advertise different prefixes. Thus even hearing a new router with a new prefix might not be a reliable indication of an L3 handover.
- o The link-local addresses of routers are not globally unique, hence after completing an L3 handover the mobile node might continue to receive Router Advertisements with the same link-local source address. This might be common if routers use the same link-local address on multiple interfaces. This issue can be avoided when routers use the Router Address (R) bit, since that provides a global address of the router.

In addition, the mobile node should consider the following events as indications that an L3 handover may have occurred. Upon receiving such indications, the mobile node needs to perform Router Discovery to discover routers and prefixes on the new link, as described in Section 6.3.7 of RFC 2461 [12].

- o If Router Advertisements that the mobile node receives include an Advertisement Interval option, the mobile node may use its Advertisement Interval field as an indication of the frequency with which it should expect to continue to receive future Advertisements from that router. This field specifies the minimum rate (the maximum amount of time between successive Advertisements) that the mobile node should expect. If this amount of time elapses without the mobile node receiving any Advertisement from this router, the mobile node can be sure that at least one Advertisement sent by the router has been lost. The mobile node can then implement its own policy to determine how many lost Advertisements from its current default router constitute an L3 handover indication.
- o Neighbor Unreachability Detection determines that the default router is no longer reachable.

- o With some types of networks, notification that an L2 handover has occurred might be obtained from lower layer protocols or device driver software within the mobile node. While further details around handling L2 indications as movement hints is an item for further study, at the time of writing this specification the following is considered reasonable:

An L2 handover indication may or may not imply L2 movement and L2 movement may or may not imply L3 movement; the correlations might be a function of the type of L2 but might also be a function of actual deployment of the wireless topology.

Unless it is well-known that an L2 handover indication is likely to imply L3 movement, instead of immediately multicasting a router solicitation it may be better to attempt to verify whether the default router is still bi-directionally reachable. This can be accomplished by sending a unicast Neighbor Solicitation and waiting for a Neighbor Advertisement with the solicited flag set. Note that this is similar to Neighbor Unreachability detection but it does not have the same state machine, such as the STALE state.

If the default router does not respond to the Neighbor Solicitation it makes sense to proceed to multicasting a Router Solicitation.

11.5.2. Forming New Care-of Addresses

After detecting that it has moved a mobile node SHOULD generate a new primary care-of address using normal IPv6 mechanisms. This SHOULD also be done when the current primary care-of address becomes deprecated. A mobile node MAY form a new primary care-of address at any time, but a mobile node MUST NOT send a Binding Update about a new care-of address to its home agent more than MAX_UPDATE_RATE times within a second.

In addition, a mobile node MAY form new non-primary care-of addresses even when it has not switched to a new default router. A mobile node can have only one primary care-of address at a time (which is registered with its home agent), but it MAY have an additional care-of address for any or all of the prefixes on its current link. Furthermore, since a wireless network interface may actually allow a mobile node to be reachable on more than one link at a time (i.e., within wireless transmitter range of routers on more than one separate link), a mobile node MAY have care-of addresses on more than one link at a time. The use of more than one care-of address at a time is described in Section 11.5.3.

As described in Section 4, in order to form a new care-of address, a mobile node MAY use either stateless [13] or stateful (e.g., DHCPv6 [29]) Address Autoconfiguration. If a mobile node needs to use a source address (other than the unspecified address) in packets sent as a part of address autoconfiguration, it MUST use an IPv6 link-local address rather than its own IPv6 home address.

RFC 2462 [13] specifies that in normal processing for Duplicate Address Detection, the node SHOULD delay sending the initial Neighbor Solicitation message by a random delay between 0 and `MAX_RTR_SOLICITATION_DELAY`. Since delaying DAD can result in significant delays in configuring a new care-of address when the Mobile Node moves to a new link, the Mobile Node preferably SHOULD NOT delay DAD when configuring a new care-of address. The Mobile Node SHOULD delay according to the mechanisms specified in RFC 2462 unless the implementation has a behavior that desynchronizes the steps that happen before the DAD in the case that multiple nodes experience handover at the same time. Such desynchronizing behaviors might be due to random delays in the L2 protocols or device drivers, or due to the movement detection mechanism that is used.

11.5.3. Using Multiple Care-of Addresses

As described in Section 11.5.2, a mobile node MAY use more than one care-of address at a time. Particularly in the case of many wireless networks, a mobile node effectively might be reachable through multiple links at the same time (e.g., with overlapping wireless cells), on which different on-link subnet prefixes may exist. The mobile node MUST ensure that its primary care-of address always has a prefix that is advertised by its current default router. After selecting a new primary care-of address, the mobile node MUST send a Binding Update containing that care-of address to its home agent. The Binding Update MUST have the Home Registration (H) and Acknowledge (A) bits set its home agent, as described on Section 11.7.1.

To assist with smooth handovers, a mobile node SHOULD retain its previous primary care-of address as a (non-primary) care-of address, and SHOULD still accept packets at this address, even after registering its new primary care-of address with its home agent. This is reasonable, since the mobile node could only receive packets at its previous primary care-of address if it were indeed still connected to that link. If the previous primary care-of address was allocated using stateful Address Autoconfiguration [29], the mobile node may not wish to release the address immediately upon switching to a new primary care-of address.

Whenever a mobile node determines that it is no longer reachable through a given link, it SHOULD invalidate all care-of addresses associated with address prefixes that it discovered from routers on the unreachable link which are not in the current set of address prefixes advertised by the (possibly new) current default router.

11.5.4. Returning Home

A mobile node detects that it has returned to its home link through the movement detection algorithm in use (Section 11.5.1), when the mobile node detects that its home subnet prefix is again on-link. The mobile node SHOULD then send a Binding Update to its home agent, to instruct its home agent to no longer intercept or tunnel packets for it. In this home registration, the mobile node MUST set the Acknowledge (A) and Home Registration (H) bits, set the Lifetime field to zero, and set the care-of address for the binding to the mobile node's own home address. The mobile node MUST use its home address as the source address in the Binding Update.

When sending this Binding Update to its home agent, the mobile node must be careful in how it uses Neighbor Solicitation [12] (if needed) to learn the home agent's link-layer address, since the home agent will be currently configured to intercept packets to the mobile node's home address using Duplicate Address Detection (DAD). In particular, the mobile node is unable to use its home address as the Source Address in the Neighbor Solicitation until the home agent stops defending the home address.

Neighbor Solicitation by the mobile node for the home agent's address will normally not be necessary, since the mobile node has already learned the home agent's link-layer address from a Source Link-Layer Address option in a Router Advertisement. However, if there are multiple home agents it may still be necessary to send a solicitation. In this special case of the mobile node returning home, the mobile node MUST multicast the packet, and in addition set the Source Address of this Neighbor Solicitation to the unspecified address (0:0:0:0:0:0:0:0). The target of the Neighbor Solicitation MUST be set to the mobile node's home address. The destination IP address MUST be set to the Solicited-Node multicast address [3]. The home agent will send a multicast Neighbor Advertisement back to the mobile node with the Solicited flag (S) set to zero. In any case, the mobile node SHOULD record the information from the Source Link-Layer Address option or from the advertisement, and set the state of the Neighbor Cache entry for the home agent to REACHABLE.

The mobile node then sends its Binding Update to the home agent's link-layer address, instructing its home agent to no longer serve as a home agent for it. By processing this Binding Update, the home

agent will cease defending the mobile node's home address for Duplicate Address Detection and will no longer respond to Neighbor Solicitations for the mobile node's home address. The mobile node is then the only node on the link receiving packets at the mobile node's home address. In addition, when returning home prior to the expiration of a current binding for its home address, and configuring its home address on its network interface on its home link, the mobile node MUST NOT perform Duplicate Address Detection on its own home address, in order to avoid confusion or conflict with its home agent's use of the same address. This rule also applies to the derived link-local address of the mobile node, if the Link Local Address Compatibility (L) bit was set when the binding was created. If the mobile node returns home after the bindings for all of its care-of addresses have expired, then it SHOULD perform DAD.

After the Mobile Node sends the Binding Update, it MUST be prepared to reply to Neighbor Solicitations for its home address. Such replies MUST be sent using a unicast Neighbor Advertisement to the sender's link-layer address. It is necessary to reply, since sending the Binding Acknowledgement from the home agent may require performing Neighbor Discovery, and the mobile node may not be able to distinguish Neighbor Solicitations coming from the home agent from other Neighbor Solicitations. Note that a race condition exists where both the mobile node and the home agent respond to the same solicitations sent by other nodes; this will be only temporary, however, until the Binding Update is accepted.

After receiving the Binding Acknowledgement for its Binding Update to its home agent, the mobile node MUST multicast onto the home link (to the all-nodes multicast address) a Neighbor Advertisement [12], to advertise the mobile node's own link-layer address for its own home address. The Target Address in this Neighbor Advertisement MUST be set to the mobile node's home address, and the Advertisement MUST include a Target Link-layer Address option specifying the mobile node's link-layer address. The mobile node MUST multicast such a Neighbor Advertisement for each of its home addresses, as defined by the current on-link prefixes, including its link-local address and site-local address. The Solicited Flag (S) in these Advertisements MUST NOT be set, since they were not solicited by any Neighbor Solicitation. The Override Flag (O) in these Advertisements MUST be set, indicating that the Advertisements SHOULD override any existing Neighbor Cache entries at any node receiving them.

Since multicasting on the local link (such as Ethernet) is typically not guaranteed to be reliable, the mobile node MAY retransmit these Neighbor Advertisements [12] up to MAX_NEIGHBOR_ADVERTISEMENT times to increase their reliability. It is still possible that some nodes

on the home link will not receive any of these Neighbor Advertisements, but these nodes will eventually be able to recover through use of Neighbor Unreachability Detection [12].

Note that the tunnel via the home agent typically stops operating at the same time that the home registration is deleted.

11.6. Return Routability Procedure

This section defines the rules that the mobile node must follow when performing the return routability procedure. Section 11.7.2 describes the rules when the return routability procedure needs to be initiated.

11.6.1. Sending Test Init Messages

A mobile node that initiates a return routability procedure MUST send (in parallel) a Home Test Init message and a Care-of Test Init messages. However, if the mobile node has recently received (see Section 5.2.7) one or both home or care-of keygen tokens, and associated nonce indices for the desired addresses, it MAY reuse them. Therefore, the return routability procedure may in some cases be completed with only one message pair. It may even be completed without any messages at all, if the mobile node has a recent home keygen token and has previously visited the same care-of address so that it also has a recent care-of keygen token. If the mobile node intends to send a Binding Update with the Lifetime set to zero and the care-of address equal to its home address - such as when returning home - sending a Home Test Init message is sufficient. In this case, generation of the binding management key depends exclusively on the home keygen token (Section 5.2.5).

A Home Test Init message MUST be created as described in Section 6.1.3.

A Care-of Test Init message MUST be created as described in Section 6.1.4. When sending a Home Test Init or Care-of Test Init message the mobile node MUST record in its Binding Update List the following fields from the messages:

- o The IP address of the node to which the message was sent.
- o The home address of the mobile node. This value will appear in the Source Address field of the Home Test Init message. When sending the Care-of Test Init message, this address does not appear in the message, but represents the home address for which the binding is desired.

- o The time at which each of these messages was sent.
- o The cookies used in the messages.

Note that a single Care-of Test Init message may be sufficient even when there are multiple home addresses. In this case the mobile node MAY record the same information in multiple Binding Update List entries.

11.6.2. Receiving Test Messages

Upon receiving a packet carrying a Home Test message, a mobile node MUST validate the packet according to the following tests:

- o The Source Address of the packet belongs to a correspondent node for which the mobile node has a Binding Update List entry with a state indicating that return routability procedure is in progress. Note that there may be multiple such entries.
- o The Binding Update List indicates that no home keygen token has been received yet.
- o The Destination Address of the packet has the home address of the mobile node, and the packet has been received in a tunnel from the home agent.
- o The Home Init Cookie field in the message matches the value stored in the Binding Update List.

Any Home Test message not satisfying all of these tests MUST be silently ignored. Otherwise, the mobile node MUST record the Home Nonce Index and home keygen token in the Binding Update List. If the Binding Update List entry does not have a care-of keygen token, the mobile node SHOULD continue waiting for the Care-of Test message.

Upon receiving a packet carrying a Care-of Test message, a mobile node MUST validate the packet according to the following tests:

- o The Source Address of the packet belongs to a correspondent node for which the mobile node has a Binding Update List entry with a state indicating that return routability procedure is in progress. Note that there may be multiple such entries.
- o The Binding Update List indicates that no care-of keygen token has been received yet.
- o The Destination Address of the packet is the current care-of address of the mobile node.

- o The Care-of Init Cookie field in the message matches the value stored in the Binding Update List.

Any Care-of Test message not satisfying all of these tests MUST be silently ignored. Otherwise, the mobile node MUST record the Care-of Nonce Index and care-of keygen token in the Binding Update List. If the Binding Update List entry does not have a home keygen token, the mobile node SHOULD continue waiting for the Home Test message.

If after receiving either the Home Test or the Care-of Test message and performing the above actions, the Binding Update List entry has both the home and the care-of keygen tokens, the return routability procedure is complete. The mobile node SHOULD then proceed with sending a Binding Update as described in Section 11.7.2.

Correspondent nodes from the time before this specification was published may not support the Mobility Header protocol. These nodes will respond to Home Test Init and Care-of Test Init messages with an ICMP Parameter Problem code 1. The mobile node SHOULD take such messages as an indication that the correspondent node cannot provide route optimization, and revert back to the use of bidirectional tunneling.

11.6.3. Protecting Return Routability Packets

The mobile node MUST support the protection of Home Test and Home Test Init messages as described in Section 10.4.6.

When IPsec is used to protect return routability signaling or payload packets, the mobile node MUST set the source address it uses for the outgoing tunnel packets to the current primary care-of address. The mobile node starts to use a new primary care-of address immediately after sending a Binding Update to the home agent to register this new address.

11.7. Processing Bindings

11.7.1. Sending Binding Updates to the Home Agent

After deciding to change its primary care-of address as described in Section 11.5.1 and Section 11.5.2, a mobile node MUST register this care-of address with its home agent in order to make this its primary care-of address.

Also, if the mobile node wants the services of the home agent beyond the current registration period, the mobile node should send a new Binding Update to it well before the expiration of this period, even if it is not changing its primary care-of address. However, if the

home agent returned a Binding Acknowledgement for the current registration with Status field set to 1 (accepted but prefix discovery necessary), the mobile node should not try to register again before it has learned the validity of its home prefixes through mobile prefix discovery. This is typically necessary every time this Status value is received, because information learned earlier may have changed.

To register a care-of address or to extend the lifetime of an existing registration, the mobile node sends a packet to its home agent containing a Binding Update, with the packet constructed as follows:

- o The Home Registration (H) bit MUST be set in the Binding Update.
- o The Acknowledge (A) bit MUST be set in the Binding Update.
- o The packet MUST contain a Home Address destination option, giving the mobile node's home address for the binding.
- o The care-of address for the binding MUST be used as the Source Address in the packet's IPv6 header, unless an Alternate Care-of Address mobility option is included in the Binding Update. This option MUST be included in all home registrations, as the ESP protocol will not be able to protect care-of addresses in the IPv6 header. (Mobile IPv6 implementations that know they are using IPsec AH to protect a particular message might avoid this option. For brevity the usage of AH is not discussed in this document.)
- o If the mobile node's link-local address has the same interface identifier as the home address for which it is supplying a new care-of address, then the mobile node SHOULD set the Link-Local Address Compatibility (L) bit.
- o If the home address was generated using RFC 3041 [18], then the link local address is unlikely to have a compatible interface identifier. In this case, the mobile node MUST clear the Link-Local Address Compatibility (L) bit.
- o If the IPsec security associations between the mobile node and the home agent have been established dynamically, and the mobile node has the capability to update its endpoint in the used key management protocol to the new care-of address every time it moves, the mobile node SHOULD set the Key Management Mobility Capability (K) bit in the Binding Update. Otherwise, the mobile node MUST clear the bit.

- o The value specified in the Lifetime field MUST be non-zero and SHOULD be less than or equal to the remaining valid lifetime of the home address and the care-of address specified for the binding.

Mobile nodes that use dynamic home agent address discovery should be careful with long lifetimes. If the mobile node loses the knowledge of its binding with a specific home agent, registering a new binding with another home agent may be impossible as the previous home agent is still defending the existing binding. Therefore, to ensure that mobile nodes using home agent address discovery do not lose information about their binding, they SHOULD de-register before losing this information, or use small lifetimes.

The Acknowledge (A) bit in the Binding Update requests the home agent to return a Binding Acknowledgement in response to this Binding Update. As described in Section 6.1.8, the mobile node SHOULD retransmit this Binding Update to its home agent until it receives a matching Binding Acknowledgement. Once reaching a retransmission timeout period of MAX_BINDACK_TIMEOUT, the mobile node SHOULD restart the process of delivering the Binding Update, but trying instead the next home agent returned during dynamic home agent address discovery (see Section 11.4.1). If there was only one home agent, the mobile node instead SHOULD continue to periodically retransmit the Binding Update at this rate until acknowledged (or until it begins attempting to register a different primary care-of address). See Section 11.8 for information about retransmitting Binding Updates.

With the Binding Update, the mobile node requests the home agent to serve as the home agent for the given home address. Until the lifetime of this registration expires, the home agent considers itself the home agent for this home address.

Each Binding Update MUST be authenticated as coming from the right mobile node, as defined in Section 5.1. The mobile node MUST use its home address - either in the Home Address destination option or in the Source Address field of the IPv6 header - in Binding Updates sent to the home agent. This is necessary in order to allow the IPsec policies to be matched with the correct home address.

When sending a Binding Update to its home agent, the mobile node MUST also create or update the corresponding Binding Update List entry, as specified in Section 11.7.2.

The last Sequence Number value sent to the home agent in a Binding Update is stored by the mobile node. If the sending mobile node has no knowledge of the correct Sequence Number value, it may start at

any value. If the home agent rejects the value, it sends back a Binding Acknowledgement with a status code 135, and the last accepted sequence number in the Sequence Number field of the Binding Acknowledgement. The mobile node MUST store this information and use the next Sequence Number value for the next Binding Update it sends.

If the mobile node has additional home addresses, then the mobile node SHOULD send an additional packet containing a Binding Update to its home agent to register the care-of address for each such other home address.

The home agent will only perform DAD for the mobile node's home address when the mobile node has supplied a valid binding between its home address and a care-of address. If some time elapses during which the mobile node has no binding at the home agent, it might be possible for another node to autoconfigure the mobile node's home address. Therefore, the mobile node MUST treat the creation of a new binding with the home agent using an existing home address, the same as creation of a new home address. In the unlikely event that the mobile node's home address is autoconfigured as the IPv6 address of another network node on the home network, the home agent will reply to the mobile node's subsequent Binding Update with a Binding Acknowledgement containing a Status of 134 (Duplicate Address Detection failed). In this case, the mobile node MUST NOT attempt to re-use the same home address. It SHOULD continue to register the care-of addresses for its other home addresses, if any. (Mechanisms outlined in Appendix B.5 may in the future allow mobile nodes to acquire new home addresses to replace the one for which Status 134 was received.)

11.7.2. Correspondent Registration

When the mobile node is assured that its home address is valid, it can initiate a correspondent registration with the purpose of allowing the correspondent node to cache the mobile node's current care-of address. This procedure consists of the return routability procedure followed by a registration.

This section defines when the correspondent registration is to be initiated and the rules to follow while it is being performed.

After the mobile node has sent a Binding Update to its home agent, registering a new primary care-of address (as described in Section 11.7.1), the mobile node SHOULD initiate a correspondent registration for each node that already appears in the mobile node's Binding Update List. The initiated procedures can be used to either update or delete binding information in the correspondent node.

For nodes that do not appear in the mobile node's Binding Update List, the mobile node MAY initiate a correspondent registration at any time after sending the Binding Update to its home agent.

Considerations regarding when (and if) to initiate the procedure depend on the specific movement and traffic patterns of the mobile node and are outside the scope of this document.

In addition, the mobile node MAY initiate the correspondent registration in response to receiving a packet that meets all of the following tests:

- o The packet was tunneled using IPv6 encapsulation.
- o The Destination Address in the tunnel (outer) IPv6 header is equal to any of the mobile node's care-of addresses.
- o The Destination Address in the original (inner) IPv6 header is equal to one of the mobile node's home addresses.
- o The Source Address in the tunnel (outer) IPv6 header differs from the Source Address in the original (inner) IPv6 header.
- o The packet does not contain a Home Test, Home Test Init, Care-of Test, or Care-of Test Init message.

If a mobile node has multiple home addresses, it becomes important to select the right home address to use in the correspondent registration. The used home address MUST be the Destination Address of the original (inner) packet.

The peer address used in the procedure MUST be determined as follows:

- o If a Home Address destination option is present in the original (inner) packet, the address from this option is used.
- o Otherwise, the Source Address in the original (inner) IPv6 header of the packet is used.

Note that the validity of the original packet is checked before attempting to initiate a correspondent registration. For instance, if a Home Address destination option appeared in the original packet, then rules in Section 9.3.1 are followed.

A mobile node MAY also choose to keep its topological location private from certain correspondent nodes, and thus need not initiate the correspondent registration.

Upon successfully completing the return routability procedure, and after receiving a successful Binding Acknowledgement from the Home Agent, a Binding Update MAY be sent to the correspondent node.

In any Binding Update sent by a mobile node, the care-of address (either the Source Address in the packet's IPv6 header or the Care-of Address in the Alternate Care-of Address mobility option of the Binding Update) MUST be set to one of the care-of addresses currently in use by the mobile node or to the mobile node's home address. A mobile node MAY set the care-of address differently for sending Binding Updates to different correspondent nodes.

A mobile node MAY also send a Binding Update to such a correspondent node, instructing it to delete any existing binding for the mobile node from its Binding Cache, as described in Section 6.1.7. Even in this case a successful completion of the return routability procedure is required first.

If the care-of address is not set to the mobile node's home address, the Binding Update requests that the correspondent node create or update an entry for the mobile node in the correspondent node's Binding Cache. This is done in order to record a care-of address for use in sending future packets to the mobile node. In this case, the value specified in the Lifetime field sent in the Binding Update SHOULD be less than or equal to the remaining lifetime of the home registration and the care-of address specified for the binding. The care-of address given in the Binding Update MAY differ from the mobile node's primary care-of address.

If the Binding Update is sent to the correspondent node, requesting the deletion of any existing Binding Cache entry it has for the mobile node, the care-of address is set to the mobile node's home address and the Lifetime field set to zero. In this case, generation of the binding management key depends exclusively on the home keygen token (Section 5.2.5). The care-of nonce index SHOULD be set to zero in this case. In keeping with the Binding Update creation rules below, the care-of address MUST be set to the home address if the mobile node is at home, or to the current care-of address if it is away from home.

If the mobile node wants to ensure that its new care-of address has been entered into a correspondent node's Binding Cache, the mobile node needs to request an acknowledgement by setting the Acknowledge (A) bit in the Binding Update.

A Binding Update is created as follows:

- o The current care-of address of the mobile node MUST be sent either in the Source Address of the IPv6 header, or in the Alternate Care-of Address mobility option.
- o The Destination Address of the IPv6 header MUST contain the address of the correspondent node.
- o The Mobility Header is constructed according to rules in Section 6.1.7 and Section 5.2.6, including the Binding Authorization Data (calculated as defined in Section 6.2.7) and possibly the Nonce Indices mobility options.
- o The home address of the mobile node MUST be added to the packet in a Home Address destination option, unless the Source Address is the home address.

Each Binding Update MUST have a Sequence Number greater than the Sequence Number value sent in the previous Binding Update to the same destination address (if any). The sequence numbers are compared modulo 2^{16} , as described in Section 9.5.1. There is no requirement, however, that the Sequence Number value strictly increase by 1 with each new Binding Update sent or received, as long as the value stays within the window. The last Sequence Number value sent to a destination in a Binding Update is stored by the mobile node in its Binding Update List entry for that destination. If the sending mobile node has no Binding Update List entry, the Sequence Number SHOULD start at a random value. The mobile node MUST NOT use the same Sequence Number in two different Binding Updates to the same correspondent node, even if the Binding Updates provide different care-of addresses.

The mobile node is responsible for the completion of the correspondent registration, as well as any retransmissions that may be needed (subject to the rate limitation defined in Section 11.8).

11.7.3. Receiving Binding Acknowledgements

Upon receiving a packet carrying a Binding Acknowledgement, a mobile node MUST validate the packet according to the following tests:

- o The packet meets the authentication requirements for Binding Acknowledgements defined in Section 6.1.8 and Section 5. That is, if the Binding Update was sent to the home agent, underlying IPsec protection is used. If the Binding Update was sent to the correspondent node, the Binding Authorization Data mobility option MUST be present and have a valid value.

- o The Binding Authorization Data mobility option, if present, MUST be the last option and MUST not have trailing padding.
- o The Sequence Number field matches the Sequence Number sent by the mobile node to this destination address in an outstanding Binding Update.

Any Binding Acknowledgement not satisfying all of these tests MUST be silently ignored.

When a mobile node receives a packet carrying a valid Binding Acknowledgement, the mobile node MUST examine the Status field as follows:

- o If the Status field indicates that the Binding Update was accepted (the Status field is less than 128), then the mobile node MUST update the corresponding entry in its Binding Update List to indicate that the Binding Update has been acknowledged; the mobile node MUST then stop retransmitting the Binding Update. In addition, if the value specified in the Lifetime field in the Binding Acknowledgement is less than the Lifetime value sent in the Binding Update being acknowledged, the mobile node MUST subtract the difference between these two Lifetime values from the remaining lifetime for the binding as maintained in the corresponding Binding Update List entry (with a minimum value for the Binding Update List entry lifetime of 0). That is, if the Lifetime value sent in the Binding Update was L_{update} , the Lifetime value received in the Binding Acknowledgement was L_{ack} , and the current remaining lifetime of the Binding Update List entry is L_{remain} , then the new value for the remaining lifetime of the Binding Update List entry should be

$$\max((L_{remain} - (L_{update} - L_{ack})), 0)$$

where $\max(X, Y)$ is the maximum of X and Y . The effect of this step is to correctly manage the mobile node's view of the binding's remaining lifetime (as maintained in the corresponding Binding Update List entry) so that it correctly counts down from the Lifetime value given in the Binding Acknowledgement, but with the timer countdown beginning at the time that the Binding Update was sent.

Mobile nodes SHOULD send a new Binding Update well before the expiration of this period in order to extend the lifetime. This helps to avoid disruptions in communications which might otherwise be caused by network delays or clock drift.

- o Additionally, if the Status field value is 1 (accepted but prefix discovery necessary), the mobile node SHOULD send a Mobile Prefix Solicitation message to update its information about the available prefixes.
- o If the Status field indicates that the Binding Update was rejected (the Status field is greater than or equal to 128), then the mobile node can take steps to correct the cause of the error and retransmit the Binding Update (with a new Sequence Number value), subject to the rate limiting restriction specified in Section 11.8. If this is not done or it fails, then the mobile node SHOULD record in its Binding Update List that future Binding Updates SHOULD NOT be sent to this destination.

The treatment of a Binding Refresh Advice mobility option within the Binding Acknowledgement depends on where the acknowledgement came from. This option MUST be ignored if the acknowledgement came from a correspondent node. If it came from the home agent, the mobile node uses the Refresh Interval field in the option as a suggestion that it SHOULD attempt to refresh its home registration at the indicated shorter interval.

If the acknowledgement came from the home agent, the mobile node examines the value of the Key Management Mobility Capability (K) bit. If this bit is not set, the mobile node SHOULD discard key management protocol connections, if any, to the home agent. The mobile node MAY also initiate a new key management connection.

If this bit is set, the mobile node SHOULD move its own endpoint in the key management protocol connections to the home agent, if any. The mobile node's new endpoint should be the new care-of address. For an IKE phase 1 connection, this means that packets sent to this address with the original ISAKMP cookies are accepted.

11.7.4. Receiving Binding Refresh Requests

When a mobile node receives a packet containing a Binding Refresh Request message, the mobile node has a Binding Update List entry for the source of the Binding Refresh Request, and the mobile node wants to retain its binding cache entry at the correspondent node, then the mobile node should start a return routability procedure. If the mobile node wants to have its binding cache entry removed, it can either ignore the Binding Refresh Request and wait for the binding to time out, or at any time, it can delete its binding from a correspondent node with an explicit binding update with a zero lifetime and the care-of address set to the home address. If the

mobile node does not know if it needs the binding cache entry, it can make the decision in an implementation dependent manner, such as based on available resources.

Note that the mobile node should be careful to not respond to Binding Refresh Requests for addresses not in the Binding Update List to avoid being subjected to a denial of service attack.

If the return routability procedure completes successfully, a Binding Update message SHOULD be sent, as described in Section 11.7.2. The Lifetime field in this Binding Update SHOULD be set to a new lifetime, extending any current lifetime remaining from a previous Binding Update sent to this node (as indicated in any existing Binding Update List entry for this node), and the lifetime SHOULD again be less than or equal to the remaining lifetime of the home registration and the care-of address specified for the binding. When sending this Binding Update, the mobile node MUST update its Binding Update List in the same way as for any other Binding Update sent by the mobile node.

11.8. Retransmissions and Rate Limiting

The mobile node is responsible for retransmissions and rate limiting in the return routability procedure, registrations, and in solicited prefix discovery.

When the mobile node sends a Mobile Prefix Solicitation, Home Test Init, Care-of Test Init or Binding Update for which it expects a response, the mobile node has to determine a value for the initial retransmission timer:

- o If the mobile node is sending a Mobile Prefix Solicitation, it SHOULD use an initial retransmission interval of INITIAL_SOLICIT_TIMER (see Section 12).
- o If the mobile node is sending a Binding Update and does not have an existing binding at the home agent, it SHOULD use InitialBindackTimeoutFirstReg (see Section 13) as a value for the initial retransmission timer. This long retransmission interval will allow the home agent to complete the Duplicate Address Detection procedure mandated in this case, as detailed in Section 11.7.1.
- o Otherwise, the mobile node should use the specified value of INITIAL_BINDACK_TIMEOUT for the initial retransmission timer.

If the mobile node fails to receive a valid matching response within the selected initial retransmission interval, the mobile node SHOULD retransmit the message until a response is received.

The retransmissions by the mobile node MUST use an exponential back-off process in which the timeout period is doubled upon each retransmission, until either the node receives a response or the timeout period reaches the value MAX_BINDACK_TIMEOUT. The mobile node MAY continue to send these messages at this slower rate indefinitely.

The mobile node SHOULD start a separate back-off process for different message types, different home addresses and different care-of addresses. However, in addition an overall rate limitation applies for messages sent to a particular correspondent node. This ensures that the correspondent node has a sufficient amount of time to respond when bindings for multiple home addresses are registered, for instance. The mobile node MUST NOT send Mobility Header messages of a particular type to a particular correspondent node more than MAX_UPDATE_RATE times within a second.

Retransmitted Binding Updates MUST use a Sequence Number value greater than that used for the previous transmission of this Binding Update. Retransmitted Home Test Init and Care-of Test Init messages MUST use new cookie values.

12. Protocol Constants

DHAAD_RETRIES	4 retransmissions
INITIAL_BINDACK_TIMEOUT	1 second
INITIAL_DHAAD_TIMEOUT	3 seconds
INITIAL_SOLICIT_TIMER	3 seconds
MAX_BINDACK_TIMEOUT	32 seconds
MAX_NONCE_LIFETIME	240 seconds
MAX_TOKEN_LIFETIME	210 seconds
MAX_RR_BINDING_LIFETIME	420 seconds
MAX_UPDATE_RATE	3 times
PREFIX_ADV_RETRIES	3 retransmissions
PREFIX_ADV_TIMEOUT	3 seconds

13. Protocol Configuration Variables

MaxMobPfxAdvInterval	Default: 86,400 seconds
MinDelayBetweenRAs	Default: 3 seconds, Min: 0.03 seconds
MinMobPfxAdvInterval	Default: 600 seconds
InitialBindackTimeoutFirstReg	Default: 1.5 seconds

Home agents MUST allow the first three variables to be configured by system management, and mobile nodes MUST allow the last variable to be configured by system management.

The default value for InitialBindackTimeoutFirstReg has been calculated as 1.5 times the default value of RetransTimer [12] times the default value of DupAddrDetectTransmits [13].

The value MinDelayBetweenRAs overrides the value of the protocol constant MIN_DELAY_BETWEEN_RAS, as specified in RFC 2461 [12]. This variable SHOULD be set to MinRtrAdvInterval, if MinRtrAdvInterval is less than 3 seconds.

14. IANA Considerations

This document defines a new IPv6 protocol, the Mobility Header, described in Section 6.1. This protocol has been assigned protocol number 135.

This document also creates a new name space "Mobility Header Type", for the MH Type field in the Mobility Header. The current message types are described starting from Section 6.1.2, and are the following:

- 0 Binding Refresh Request
- 1 Home Test Init
- 2 Care-of Test Init
- 3 Home Test
- 4 Care-of Test
- 5 Binding Update
- 6 Binding Acknowledgement
- 7 Binding Error

Future values of the MH Type can be allocated using Standards Action or IESG Approval [10].

Furthermore, each mobility message may contain mobility options as described in Section 6.2. This document defines a new name space "Mobility Option" to identify these options. The current mobility options are defined starting from Section 6.2.2 and are the following:

- 0 Pad1
- 1 PadN
- 2 Binding Refresh Advice
- 3 Alternate Care-of Address
- 4 Nonce Indices
- 5 Authorization Data

Future values of the Option Type can be allocated using Standards Action or IESG Approval [10].

Finally, this document creates a third new name space "Status Code" for the Status field in the Binding Acknowledgement message. The current values are described in Section 6.1.8, and are the following:

- 0 Binding Update accepted
- 1 Accepted but prefix discovery necessary
- 128 Reason unspecified
- 129 Administratively prohibited
- 130 Insufficient resources
- 131 Home registration not supported
- 132 Not home subnet
- 133 Not home agent for this mobile node
- 134 Duplicate Address Detection failed
- 135 Sequence number out of window
- 136 Expired home nonce index
- 137 Expired care-of nonce index
- 138 Expired nonces
- 139 Registration type change disallowed

Future values of the Status field can be allocated using Standards Action or IESG Approval [10].

All fields labeled "Reserved" are only to be assigned through Standards Action or IESG Approval.

This document also defines a new IPv6 destination option, the Home Address option, described in Section 6.3. This option has been assigned the Option Type value 0xC9.

This document also defines a new IPv6 type 2 routing header, described in Section 6.4. The value 2 has been allocated by IANA.

In addition, this document defines four ICMP message types, two used as part of the dynamic home agent address discovery mechanism, and two used in lieu of Router Solicitations and Advertisements when the mobile node is away from the home link. These messages have been assigned ICMPv6 type numbers from the informational message range:

- o The Home Agent Address Discovery Request message, described in Section 6.5;
- o The Home Agent Address Discovery Reply message, described in Section 6.6;
- o The Mobile Prefix Solicitation, described in Section 6.7; and
- o The Mobile Prefix Advertisement, described in Section 6.8.

This document also defines two new Neighbor Discovery [12] options, which have been assigned Option Type values within the option numbering space for Neighbor Discovery messages:

- o The Advertisement Interval option, described in Section 7.3; and
- o The Home Agent Information option, described in Section 7.4.

15. Security Considerations

15.1. Threats

Any mobility solution must protect itself against misuses of the mobility features and mechanisms. In Mobile IPv6, most of the potential threats are concerned with false Bindings, usually resulting in Denial-of-Service attacks. Some of the threats also pose potential for Man-in-the-Middle, Hijacking, Confidentiality, and Impersonation attacks. The main threats this protocol protects against are the following:

- o Threats involving Binding Updates sent to home agents and correspondent nodes. For instance, an attacker might claim that a certain mobile node is currently at a different location than it really is. If a home agent accepts such spoofed information sent to it, the mobile node might not get traffic destined to it. Similarly, a malicious (mobile) node might use the home address of a victim node in a forged Binding Update sent to a correspondent node.

These pose threats against confidentiality, integrity, and availability. That is, an attacker might learn the contents of packets destined to another node by redirecting the traffic to itself. Furthermore, an attacker might use the redirected packets in an attempt to set itself as a Man-in-the-Middle between a mobile and a correspondent node. This would allow the attacker to impersonate the mobile node, leading to integrity and availability problems.

A malicious (mobile) node might also send Binding Updates in which the care-of address is set to the address of a victim node. If such Binding Updates were accepted, the malicious node could lure the correspondent node into sending potentially large amounts of data to the victim; the correspondent node's replies to messages sent by the malicious mobile node will be sent to the victim host or network. This could be used to cause a Distributed Denial-of-Service attack. For example, the correspondent node might be a site that will send a high-bandwidth stream of video to anyone who asks for it. Note that the use of flow-control protocols such as TCP does not necessarily defend against this type of attack, because the attacker can fake the acknowledgements. Even keeping TCP initial sequence numbers secret does not help, because the attacker can receive the first few segments (including the ISN) at its own address, and only then redirect the stream to the victim's address. These types of attacks may also be directed to networks instead of nodes. Further variations of this threat are described elsewhere [27, 34].

An attacker might also attempt to disrupt a mobile node's communications by replaying a Binding Update that the node had sent earlier. If the old Binding Update was accepted, packets destined for the mobile node would be sent to its old location as opposed to its current location.

In conclusion, there are Denial-of-Service, Man-in-the-Middle, Confidentiality, and Impersonation threats against the parties involved in sending legitimate Binding Updates, and Denial-of-Service threats against any other party.

- o Threats associated with payload packets: Payload packets exchanged with mobile nodes are exposed to similar threats as that of regular IPv6 traffic. However, Mobile IPv6 introduces the Home Address destination option, a new routing header type (type 2), and uses tunneling headers in the payload packets. The protocol must protect against potential new threats involving the use of these mechanisms.

Third parties become exposed to a reflection threat via the Home Address destination option, unless appropriate security precautions are followed. The Home Address destination option could be used to direct response traffic toward a node whose IP address appears in the option. In this case, ingress filtering would not catch the forged "return address" [36, 32].

A similar threat exists with the tunnels between the mobile node and the home agent. An attacker might forge tunnel packets between the mobile node and the home agent, making it appear that the traffic is coming from the mobile node when it is not. Note that an attacker who is able to forge tunnel packets would typically also be able to forge packets that appear to come directly from the mobile node. This is not a new threat as such. However, it may make it easier for attackers to escape detection by avoiding ingress filtering and packet tracing mechanisms. Furthermore, spoofed tunnel packets might be used to gain access to the home network.

Finally, a routing header could also be used in reflection attacks, and in attacks designed to bypass firewalls. The generality of the regular routing header would allow circumvention of IP-address based rules in firewalls. It would also allow reflection of traffic to other nodes. These threats exist with routing headers in general, even if the usage that Mobile IPv6 requires is safe.

- o Threats associated with dynamic home agent and mobile prefix discovery.

- o Threats against the Mobile IPv6 security mechanisms themselves: An attacker might, for instance, lure the participants into executing expensive cryptographic operations or allocating memory for the purpose of keeping state. The victim node would have no resources left to handle other tasks.

As a fundamental service in an IPv6 stack, Mobile IPv6 is expected to be deployed in most nodes of the IPv6 Internet. The above threats should therefore be considered as being applicable to the whole Internet.

It should also be noted that some additional threats result from movements as such, even without the involvement of mobility protocols. Mobile nodes must be capable to defend themselves in the networks that they visit, as typical perimeter defenses applied in the home network no longer protect them.

15.2. Features

This specification provides a series of features designed to mitigate the risk introduced by the threats listed above. The main security features are the following:

- o Reverse Tunneling as a mandatory feature.
- o Protection of Binding Updates sent to home agents.
- o Protection of Binding Updates sent to correspondent nodes.
- o Protection against reflection attacks that use the Home Address destination option.
- o Protection of tunnels between the mobile node and the home agent.
- o Closing routing header vulnerabilities.
- o Mitigating Denial-of-Service threats to the Mobile IPv6 security mechanisms themselves.

The support for encrypted reverse tunneling (see Section 11.3.1) allows mobile nodes to defeat certain kinds of traffic analysis.

Protecting those Binding Updates that are sent to home agents and those that are sent to arbitrary correspondent nodes requires very different security solutions due to the different situations. Mobile nodes and home agents are naturally expected to be subject to the network administration of the home domain.

Thus, they can and are supposed to have a security association that can be used to reliably authenticate the exchanged messages. See Section 5.1 for the description of the protocol mechanisms, and Section 15.3 below for a discussion of the resulting level of security.

It is expected that Mobile IPv6 route optimization will be used on a global basis between nodes belonging to different administrative domains. It would be a very demanding task to build an authentication infrastructure on this scale. Furthermore, a traditional authentication infrastructure cannot be easily used to authenticate IP addresses because IP addresses can change often. It is not sufficient to just authenticate the mobile nodes; Authorization to claim the right to use an address is needed as well. Thus, an "infrastructureless" approach is necessary. The chosen infrastructureless method is described in Section 5.2, and Section 15.4 discusses the resulting security level and the design rationale of this approach.

Specific rules guide the use of the Home Address destination option, the routing header, and the tunneling headers in the payload packets. These rules are necessary to remove the vulnerabilities associated with their unrestricted use. The effect of the rules is discussed in Section 15.7, Section 15.8, and Section 15.9.

Denial-of-Service threats against Mobile IPv6 security mechanisms themselves concern mainly the Binding Update procedures with correspondent nodes. The protocol has been designed to limit the effects of such attacks, as will be described in Section 15.4.5.

15.3. Binding Updates to Home Agent

Signaling between the mobile node and the home agent requires message integrity. This is necessary to assure the home agent that a Binding Update is from a legitimate mobile node. In addition, correct ordering and anti-replay protection are optionally needed.

IPsec ESP protects the integrity of the Binding Updates and Binding Acknowledgements by securing mobility messages between the mobile node and the home agent.

IPsec can provide anti-replay protection only if dynamic keying is used (which may not always be the case). IPsec does not guarantee correct ordering of packets, only that they have not been replayed. Because of this, sequence numbers within the Mobile IPv6 messages are used to ensure correct ordering (see Section 5.1). However, if the 16 bit Mobile IPv6 sequence number space is cycled through, or the home agent reboots and loses its state regarding the sequence

numbers, replay and reordering attacks become possible. The use of dynamic keying, IPsec anti-replay protection, and the Mobile IPv6 sequence numbers can together prevent such attacks. It is also recommended that use of non-volatile storage be considered for home agents, to avoid losing their state.

A sliding window scheme is used for the sequence numbers. The protection against replays and reordering attacks without a key management mechanism works when the attacker remembers up to a maximum of 2^{15} Binding Updates.

The above mechanisms do not show that the care-of address given in the Binding Update is correct. This opens the possibility for Denial-of-Service attacks against third parties. However, since the mobile node and home agent have a security association, the home agent can always identify an ill-behaving mobile node. This allows the home agent operator to discontinue the mobile node's service, and possibly take further actions based on the business relationship with the mobile node's owner.

Note that the use of a single pair of manually keyed security associations conflicts with the generation of a new home address [18] for the mobile node, or with the adoption of a new home subnet prefix. This is because IPsec security associations are bound to the used addresses. While certificate-based automatic keying alleviates this problem to an extent, it is still necessary to ensure that a given mobile node cannot send Binding Updates for the address of another mobile node. In general, this leads to the inclusion of home addresses in certificates in the Subject AltName field. This again limits the introduction of new addresses without either manual or automatic procedures to establish new certificates. Therefore, this specification restricts the generation of new home addresses (for any reason) to those situations where a security association or certificate for the new address already exists. (Appendix B.4 lists the improvement of security for new addresses as one of the future developments for Mobile IPv6.)

Support for IKE has been specified as optional. The following should be observed about the use of manual keying:

- o As discussed above, with manually keyed IPsec, only a limited form of protection exists against replay and reordering attacks. A vulnerability exists if either the sequence number space is cycled through, or if the home agent reboots and forgets its sequence numbers (and uses volatile memory to store the sequence numbers). Assuming the mobile node moves continuously every 10 minutes, it

takes roughly 455 days before the sequence number space has been cycled through. Typical movement patterns rarely reach this high frequency today.

- o A mobile node and its home agent belong to the same domain. If this were not the case, manual keying would not be possible [28], but in Mobile IPv6 only these two parties need to know the manually configured keys. Similarly, we note that Mobile IPv6 employs standard block ciphers in IPsec, and is not vulnerable to problems associated with stream ciphers and manual keying.
- o It is expected that the owner of the mobile node and the administrator of the home agent agree on the used keys and other parameters with some off-line mechanism.

The use of IKEv1 with Mobile IPv6 is documented in more detail in [21]. The following should be observed from the use of IKEv1:

- o It is necessary to prevent a mobile node from claiming another mobile node's home address. The home agent must verify that the mobile node trying to negotiate the SA for a particular home address is authorized for that home address. This implies that even with the use of IKE, a policy entry needs to be configured for each home address served by the home agent.

It may be possible to include home addresses in the Subject AltName field of certificate to avoid this. However, implementations are not guaranteed to support the use of a particular IP address (care-of address) while another address (home address) appears in the certificate. In any case, even this approach would require user-specific tasks in the certificate authority.

- o If preshared secret authentication is used, IKEv1 main mode cannot be used. Aggressive mode or group preshared secrets need to be used with corresponding security implications instead.

Note that, like many other issues, this is a general IKEv1 issue related to the ability to use different IP addresses, and not specifically related to Mobile IPv6. For further information, see Section 4.4 in [21].

- o Due to the problems outlined in Section 11.3.2, IKE phase 1 between the mobile node and its home agent is established using the mobile node's current care-of address. This implies that when the mobile node moves to a new location, it may have to re-establish phase 1. A Key Management Mobility Capability (K) flag

is provided for implementations that can update the IKE phase 1 endpoints without re-establishing phase 1, but the support for this behavior is optional.

- o When certificates are used, IKE fragmentation can occur as discussed in Section 7 in [21].
- o Nevertheless, even if per-mobile node configuration is required with IKE, an important benefit of IKE is that it automates the negotiation of cryptographic parameters, including the SPIs, cryptographic algorithms, and so on. Thus, less configuration information is needed.
- o The frequency of movements in some link layers or deployment scenarios may be high enough to make replay and reordering attacks possible, if only manual keying is used. IKE SHOULD be used in such cases. Potentially vulnerable scenarios involve continuous movement through small cells, or uncontrolled alternation between available network attachment points.
- o Similarly, in some deployment scenarios the number of mobile nodes may be very large. In these cases, it can be necessary to use automatic mechanisms to reduce the management effort in the administration of cryptographic parameters, even if some per-mobile node configuration is always needed. IKE SHOULD also be used in such cases.
- o Other automatic key management mechanisms exist beyond IKEv1, but this document does not address the issues related to them. We note, however, that most of the above discussion applies to IKEv2 [30] as well, at least as it is currently specified.

15.4. Binding Updates to Correspondent Nodes

The motivation for designing the return routability procedure was to have sufficient support for Mobile IPv6, without creating significant new security problems. The goal for this procedure was not to protect against attacks that were already possible before the introduction of Mobile IPv6.

The next sections will describe the security properties of the used method, both from the point of view of possible on-path attackers who can see those cryptographic values that have been sent in the clear (Section 15.4.2 and Section 15.4.3) and from the point of view of other attackers (Section 15.4.6).

15.4.1. Overview

The chosen infrastructureless method verifies that the mobile node is "live" (that is, it responds to probes) at its home and care-of addresses. Section 5.2 describes the return routability procedure in detail. The procedure uses the following principles:

- o A message exchange verifies that the mobile node is reachable at its addresses, i.e., is at least able to transmit and receive traffic at both the home and care-of addresses.
- o The eventual Binding Update is cryptographically bound to the tokens supplied in the exchanged messages.
- o Symmetric exchanges are employed to avoid the use of this protocol in reflection attacks. In a symmetric exchange, the responses are always sent to the same address the request was sent from.
- o The correspondent node operates in a stateless manner until it receives a fully authorized Binding Update.
- o Some additional protection is provided by encrypting the tunnels between the mobile node and home agent with IPsec ESP. As the tunnel also transports the nonce exchanges, the ability of attackers to see these nonces is limited. For instance, this prevents attacks from being launched from the mobile node's current foreign link, even when no link-layer confidentiality is available.

The resulting level of security is in theory the same even without this additional protection: the return routability tokens are still exposed only to one path within the whole Internet. However, the mobile nodes are often found on an insecure link, such as a public access Wireless LAN. Thus, in many cases, this addition makes a practical difference.

For further information about the design rationale of the return routability procedure, see [27, 34, 33, 32]. The mechanisms used have been adopted from these documents.

15.4.2. Achieved Security Properties

The return routability procedure protects Binding Updates against all attackers who are unable to monitor the path between the home agent and the correspondent node. The procedure does not defend against attackers who can monitor this path. Note that such attackers are in any case able to mount an active attack against the mobile node when

it is at its home location. The possibility of such attacks is not an impediment to the deployment of Mobile IPv6 because these attacks are possible regardless of whether or not Mobile IPv6 is in use.

This procedure also protects against Denial-of-Service attacks in which the attacker pretends to be mobile, but uses the victim's address as the care-of address. This would cause the correspondent node to send the victim some unexpected traffic. This procedure defends against these attacks by requiring at least the passive presence of the attacker at the care-of address or on the path from the correspondent to the care-of address. Normally, this will be the mobile node.

15.4.3. Comparison to Regular IPv6 Communications

This section discusses the protection offered by the return routability method by comparing it to the security of regular IPv6 communications. We will divide vulnerabilities into three classes: (1) those related to attackers on the local network of the mobile node, home agent, or the correspondent node, (2) those related to attackers on the path between the home network and the correspondent node, and (3) off-path attackers, i.e., the rest of the Internet.

We will now discuss the vulnerabilities of regular IPv6 communications. The on-link vulnerabilities of IPv6 communications include Denial-of-Service, Masquerading, Man-in-the-Middle, Eavesdropping, and other attacks. These attacks can be launched through spoofing Router Discovery, Neighbor Discovery and other IPv6 mechanisms. Some of these attacks can be prevented with the use of cryptographic protection in the packets.

A similar situation exists with on-path attackers. That is, without cryptographic protection, the traffic is completely vulnerable.

Assuming that attackers have not penetrated the security of the Internet routing protocols, attacks are much harder to launch from off-path locations. Attacks that can be launched from these locations are mainly Denial-of-Service attacks, such as flooding and/or reflection attacks. It is not possible for an off-path attacker to become a Man-in-the-Middle.

Next, we will consider the vulnerabilities that exist when IPv6 is used together with Mobile IPv6 and the return routability procedure. On the local link, the vulnerabilities are the same as those in IPv6, but Masquerade and Man-in-the-Middle attacks can now also be launched against future communications, and not just against current communications. If a Binding Update was sent while the attacker was present on the link, its effects remain for the lifetime of the

binding. This happens even if the attacker moves away from the link. In contrast, an attacker who uses only plain IPv6 generally has to stay on the link in order to continue the attack. Note that in order to launch these new attacks, the IP address of the victim must be known. This makes this attack feasible, mainly in the context of well-known interface IDs, such as those already appearing in the traffic on the link or registered in the DNS.

On-path attackers can exploit similar vulnerabilities as in regular IPv6. There are some minor differences, however. Masquerade, Man-in-the-Middle, and Denial-of-Service attacks can be launched with just the interception of a few packets, whereas in regular IPv6 it is necessary to intercept every packet. The effect of the attacks is the same regardless of the method, however. In any case, the most difficult task an attacker faces in these attacks is getting on the right path.

The vulnerabilities for off-path attackers are the same as in regular IPv6. Those nodes that are not on the path between the home agent and the correspondent node will not be able to receive the home address probe messages.

In conclusion, we can state the following main results from this comparison:

- o Return routability prevents any off-path attacks beyond those that are already possible in regular IPv6. This is the most important result, preventing attackers on the Internet from exploiting any vulnerabilities.
- o Vulnerabilities to attackers on the home agent link, the correspondent node link, and the path between them are roughly the same as in regular IPv6.
- o However, one difference is that in basic IPv6 an on-path attacker must be constantly present on the link or the path, whereas with Mobile IPv6, an attacker can leave a binding behind after moving away.

For this reason, this specification limits the creation of bindings to at most `MAX_TOKEN_LIFETIME` seconds after the last routability check has been performed, and limits the duration of a binding to at most `MAX_RR_BINDING_LIFETIME` seconds. With these limitations, attackers cannot take any practical advantages of this vulnerability.

- o There are some other minor differences, such as an effect to the Denial-of-Service vulnerabilities. These can be considered to be insignificant.
- o The path between the home agent and a correspondent node is typically easiest to attack on the links at either end, in particular if these links are publicly accessible wireless LANs.

Attacks against the routers or switches on the path are typically harder to accomplish. The security on layer 2 of the links plays then a major role in the resulting overall network security. Similarly, security of IPv6 Neighbor and Router Discovery on these links has a large impact. If these were secured using some new technology in the future, this could change the situation regarding the easiest point of attack.

For a more in-depth discussion of these issues, see [32].

15.4.4. Replay Attacks

The return routability procedure also protects the participants against replayed Binding Updates. The attacker is unable to replay the same message due to the sequence number which is a part of the Binding Update. It is also unable to modify the Binding Update since the MAC verification would fail after such a modification.

Care must be taken when removing bindings at the correspondent node, however. If a binding is removed while the nonce used in its creation is still valid, an attacker could replay the old Binding Update. Rules outlined in Section 5.2.8 ensure that this cannot happen.

15.4.5. Denial-of-Service Attacks

The return routability procedure has protection against resource exhaustion Denial-of-Service attacks. The correspondent nodes do not retain any state about individual mobile nodes until an authentic Binding Update arrives. This is achieved through the construct of keygen tokens from the nonces and node keys that are not specific to individual mobile nodes. The keygen tokens can be reconstructed by the correspondent node, based on the home and care-of address information that arrives with the Binding Update. This means that the correspondent nodes are safe against memory exhaustion attacks except where on-path attackers are concerned. Due to the use of symmetric cryptography, the correspondent nodes are relatively safe against CPU resource exhaustion attacks as well.

Nevertheless, as [27] describes, there are situations in which it is impossible for the mobile and correspondent nodes to determine if they actually need a binding or whether they just have been fooled into believing so by an attacker. Therefore, it is necessary to consider situations where such attacks are being made.

Even if route optimization is a very important optimization, it is still only an optimization. A mobile node can communicate with a correspondent node even if the correspondent refuses to accept any Binding Updates. However, performance will suffer because packets from the correspondent node to the mobile node will be routed via the mobile's home agent rather than a more direct route. A correspondent node can protect itself against some of these resource exhaustion attacks as follows. If the correspondent node is flooded with a large number of Binding Updates that fail the cryptographic integrity checks, it can stop processing Binding Updates. If a correspondent node finds that it is spending more resources on checking bogus Binding Updates than it is likely to save by accepting genuine Binding Updates, then it may silently discard some or all Binding Updates without performing any cryptographic operations.

Layers above IP can usually provide additional information to help decide if there is a need to establish a binding with a specific peer. For example, TCP knows if the node has a queue of data that it is trying to send to a peer. An implementation of this specification is not required to make use of information from higher protocol layers, but some implementations are likely to be able to manage resources more effectively by making use of such information.

We also require that all implementations be capable of administratively disabling route optimization.

15.4.6. Key Lengths

Attackers can try to break the return routability procedure in many ways. Section 15.4.2 discusses the situation where the attacker can see the cryptographic values sent in the clear, and Section 15.4.3 discusses the impact this has on IPv6 communications. This section discusses whether attackers can guess the correct values without seeing them.

While the return routability procedure is in progress, 64 bit cookies are used to protect spoofed responses. This is believed to be sufficient, given that to blindly spoof a response a very large number of messages would have to be sent before success would be probable.

The tokens used in the return routability procedure provide together 128 bits of information. This information is used internally as input to a hash function to produce a 160 bit quantity suitable for producing the keyed hash in the Binding Update using the HMAC_SHA1 algorithm. The final keyed hash length is 96 bits. The limiting factors in this case are the input token lengths and the final keyed hash length. The internal hash function application does not reduce the entropy.

The 96 bit final keyed hash is of typical size and is believed to be secure. The 128 bit input from the tokens is broken in two pieces, the home keygen token and the care-of keygen token. An attacker can try to guess the correct cookie value, but again this would require a large number of messages (on the average 2^{63} messages for one or 2^{127} for two). Furthermore, given that the cookies are valid only for a short period of time, the attack has to keep a high constant message rate to achieve a lasting effect. This does not appear practical.

When the mobile node is returning home, it is allowed to use just the home keygen token of 64 bits. This is less than 128 bits, but attacking it blindly would still require a large number of messages to be sent. If the attacker is on the path and capable of seeing the Binding Update, it could conceivably break the keyed hash with brute force. However, in this case the attacker has to be on the path, which appears to offer easier ways for denial-of-service than preventing route optimization.

15.5. Dynamic Home Agent Address Discovery

The dynamic home agent address discovery function could be used to learn the addresses of home agents in the home network.

The ability to learn addresses of nodes may be useful to attackers because brute-force scanning of the address space is not practical with IPv6. Thus, they could benefit from any means which make mapping the networks easier. For example, if a security threat targeted at routers or even home agents is discovered, having a simple ICMP mechanism to easily find out possible targets may prove to be an additional (though minor) security risk.

Apart from discovering the address(es) of home agents, attackers will not be able to learn much from this information, and mobile nodes cannot be tricked into using wrong home agents, as all other communication with the home agents is secure.

15.6. Mobile Prefix Discovery

The mobile prefix discovery function may leak interesting information about network topology and prefix lifetimes to eavesdroppers; for this reason, requests for this information has to be authenticated. Responses and unsolicited prefix information needs to be authenticated to prevent the mobile nodes from being tricked into believing false information about the prefixes and possibly preventing communications with the existing addresses. Optionally, encryption may be applied to prevent leakage of the prefix information.

15.7. Tunneling via the Home Agent

Tunnels between the mobile node and the home agent can be protected by ensuring proper use of source addresses, and optional cryptographic protection. These procedures are discussed in Section 5.5.

Binding Updates to the home agents are secure. When receiving tunneled traffic, the home agent verifies that the outer IP address corresponds to the current location of the mobile node. This acts as a weak form of protection against spoofing packets that appear to come from the mobile node. This is particularly useful, if no end-to-end security is being applied between the mobile and correspondent nodes. The outer IP address check prevents attacks where the attacker is controlled by ingress filtering. It also prevents attacks when the attacker does not know the current care-of address of the mobile node. Attackers who know the care-of address and are not controlled by ingress filtering could still send traffic through the home agent. This includes attackers on the same local link as the mobile node is currently on. But such attackers could send packets that appear to come from the mobile node without attacking the tunnel; the attacker could simply send packets with the source address set to the mobile node's home address. However, this attack does not work if the final destination of the packet is in the home network, and some form of perimeter defense is being applied for packets sent to those destinations. In such cases it is recommended that either end-to-end security or additional tunnel protection be applied, as is usual in remote access situations.

Home agents and mobile nodes may use IPsec ESP to protect payload packets tunneled between themselves. This is useful for protecting communications against attackers on the path of the tunnel.

When site local home addresses are used, reverse tunneling can be used to send site local traffic from another location. Administrators should be aware of this when allowing such home

addresses. In particular, the outer IP address check described above is not sufficient against all attackers. The use of encrypted tunnels is particularly useful for these kinds of home addresses.

15.8. Home Address Option

When the mobile node sends packets directly to the correspondent node, the Source Address field of the packet's IPv6 header is the care-of address. Therefore, ingress filtering [26] works in the usual manner even for mobile nodes, as the Source Address is topologically correct. The Home Address option is used to inform the correspondent node of the mobile node's home address.

However, the care-of address in the Source Address field does not survive in replies sent by the correspondent node unless it has a binding for this mobile node. Also, not all attacker tracing mechanisms work when packets are being reflected through correspondent nodes using the Home Address option. For these reasons, this specification restricts the use of the Home Address option. It may only be used when a binding has already been established with the participation of the node at the home address, as described in Section 5.5 and Section 6.3. This prevents reflection attacks through the use of the Home Address option. It also ensures that the correspondent nodes reply to the same address that the mobile node sends traffic from.

No special authentication of the Home Address option is required beyond the above, but note that if the IPv6 header of a packet is covered by IPsec Authentication Header, then that authentication covers the Home Address option as well. Thus, even when authentication is used in the IPv6 header, the security of the Source Address field in the IPv6 header is not compromised by the presence of a Home Address option. Without authentication of the packet, any field in the IPv6 header, including the Source Address field or any other part of the packet and the Home Address option can be forged or modified in transit. In this case, the contents of the Home Address option is no more suspect than any other part of the packet.

15.9. Type 2 Routing Header

The definition of the type 2 routing header is described in Section 6.4. This definition and the associated processing rules have been chosen so that the header cannot be used for what is traditionally viewed as source routing. In particular, the Home Address in the routing header will always have to be assigned to the home address of the receiving node; otherwise the packet will be dropped.

Generally, source routing has a number of security concerns. These include the automatic reversal of unauthenticated source routes (which is an issue for IPv4, but not for IPv6). Another concern is the ability to use source routing to "jump" between nodes inside, as well as outside a firewall. These security concerns are not issues in Mobile IPv6, due to the rules mentioned above.

In essence the semantics of the type 2 routing header is the same as a special form of IP-in-IP tunneling where the inner and outer source addresses are the same.

This implies that a device which implements the filtering of packets should be able to distinguish between a type 2 routing header and other routing headers, as required in Section 8.3. This is necessary in order to allow Mobile IPv6 traffic while still having the option of filtering out other uses of routing headers.

16. Contributors

Tuomas Aura, Mike Roe, Greg O'Shea, Pekka Nikander, Erik Nordmark, and Michael Thomas worked on the return routability protocols eventually led to the procedures used in this protocol. The procedures described in [34] were adopted in the protocol.

Significant contributions were made by members of the Mobile IPv6 Security Design Team, including (in alphabetical order) Gabriel Montenegro, Erik Nordmark and Pekka Nikander.

17. Acknowledgements

We would like to thank the members of the Mobile IP and IPng Working Groups for their comments and suggestions on this work. We would particularly like to thank (in alphabetical order) Fred Baker, Josh Broch, Samita Chakrabarti, Robert Chalmers, Noel Chiappa, Greg Daley, Vijay Devarapalli, Rich Draves, Francis Dupont, Thomas Eklund, Jun-ichiro Itojun Hagino, Brian Haley, Marc Hasson, John Ioannidis, James Kempf, Rajeev Koodli, Krishna Kumar, T.J. Kniveton, Joe Lau, Jiwoong Lee, Aime Le Rouzic, Vesa-Matti Mantyla, Kevin Miles, Glenn Morrow, Thomas Narten, Karen Nielsen, Simon Nybroe, David Oran, Brett Pentland, Lars Henrik Petander, Basavaraj Patil, Mohan Parthasarathy, Alexandru Petrescu, Mattias Petterson, Ken Powell, Phil Roberts, Ed Remmell, Patrice Romand, Luis A. Sanchez, Jeff Schiller, Pekka Savola, Arvind Sevalkar, Keiichi Shima, Tom Soderlund, Hesham Soliman, Jim Solomon, Tapio Suihko, Dave Thaler, Benny Van Houdt, Jon-Olov Vatn, Carl E. Williams, Vladislav Yasevich, Alper Yegin, and

Xinhua Zhao, for their detailed reviews of earlier versions of this document. Their suggestions have helped to improve both the design and presentation of the protocol.

We would also like to thank the participants of the Mobile IPv6 testing event (1999), implementors who participated in Mobile IPv6 interoperability testing at Connectathons (2000, 2001, 2002, and 2003), and the participants at the ETSI interoperability testing (2000, 2002). Finally, we would like to thank the TAHI project who has provided test suites for Mobile IPv6.

18. References

18.1. Normative References

- [1] Eastlake 3rd., D., Crocker, S. and J. Schiller, "Randomness Recommendations for Security", RFC 1750, December 1994.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [4] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [5] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [6] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.
- [7] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.
- [8] Maughan, D., Schertler, M., Schneider, M. and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998.
- [9] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.
- [10] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

- [11] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [12] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [13] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [14] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.
- [15] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.
- [16] Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast Addresses", RFC 2526, March 1999.
- [17] Deering, S., Fenner, W. and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [18] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [19] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, January 2002.
- [20] National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-1, April 1995, <<http://www.itl.nist.gov/fipspubs/fip180-1.htm>>.
- [21] Arkko, J., Devarapalli, V. and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", RFC 3776, June 2004.

18.2. Informative References

- [22] Perkins, C., Ed., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [23] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [24] Perkins, C., "Minimal Encapsulation within IP", RFC 2004, October 1996.

- [25] Krawczyk, H., Bellare, M. and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [26] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [27] Aura, T. and J. Arkko, "MIPv6 BU Attacks and Defenses", Work in Progress, March 2002.
- [28] Bellovin, S., "Guidelines for Mandating Automated Key Management", Work in Progress, August 2003.
- [29] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C. and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [30] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", Work in Progress, April 2003.
- [31] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [32] Nikander, P., Aura, T., Arkko, J., Montenegro, G. and E. Nordmark, "Mobile IP version 6 Route Optimization Security Design Background", Work in Progress, April 2003.
- [33] Nordmark, E., "Securing MIPv6 BUs using return routability (BU3WAY)", Work in Progress, November 2001.
- [34] Roe, M., Aura, T., O'Shea, G. and J. Arkko, "Authentication of Mobile IPv6 Binding Updates and Acknowledgments", Work in Progress, March 2002.
- [35] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, September 2003.
- [36] Savola, P., "Security of IPv6 Routing Header and Home Address Options", Work in Progress, December 2002.
- [37] Vida, R. and L. Costa, Eds., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

Appendix A. Future Extensions

A.1. Piggybacking

This document does not specify how to piggyback payload packets on the binding related messages. However, it is envisioned that this can be specified in a separate document when issues such as the interaction between piggybacking and IPsec are fully resolved (see also Appendix A.3). The return routability messages can indicate support for piggybacking with a new mobility option.

A.2. Triangular Routing

Due to the concerns about opening reflection attacks with the Home Address destination option, this specification requires that this option be verified against the Binding Cache, i.e., there must be a Binding Cache entry for the Home Address and Care-of Address.

Future extensions may be specified that allow the use of unverified Home Address destination options in ways that do not introduce security issues.

A.3. New Authorization Methods

While the return routability procedure provides a good level of security, there exist methods that have even higher levels of security. Secondly, as discussed in Section 15.4, future enhancements of IPv6 security may cause a need to also improve the security of the return routability procedure. Using IPsec as the sole method for authorizing Binding Updates to correspondent nodes is also possible. The protection of the Mobility Header for this purpose is easy, though one must ensure that the IPsec SA was created with appropriate authorization to use the home address referenced in the Binding Update. For instance, a certificate used by IKE to create the security association might contain the home address. A future specification may specify how this is done.

A.4. Dynamically Generated Home Addresses

A future version of this specification may include functionality that allows the generation of new home addresses without requiring pre-arranged security associations or certificates even for the new addresses.

A.5. Remote Home Address Configuration

The method for initializing a mobile node's home address upon power-up or after an extended period of being disconnected from the network is beyond the scope of this specification. Whatever procedure is used should result in the mobile node having the same stateless or stateful (e.g., DHCPv6) home address autoconfiguration information it would have if it were attached to the home network. Due to the possibility that the home network could be renumbered while the mobile node is disconnected, a robust mobile node would not rely solely on storing these addresses locally.

Such a mobile node could be initialized by using the following procedure:

1. Generate a care-of address.
2. Query DNS for an anycast address associated with the FQDN of the home agent(s).
3. Perform home agent address discovery, and select a home agent.
4. Configure one home address based on the selected home agent's subnet prefix and the interface identifier of the mobile node.
5. Create security associations and security policy database entries for protecting the traffic between the selected home address and home agent.
6. Perform a home registration on the selected home agent.
7. Perform mobile prefix discovery.
8. Make a decision if further home addresses need to be configured.

This procedure is restricted to those situations where the home prefix is 64 bits and the mobile node knows its own interface identifier, which is also 64 bits.

A.6. Neighbor Discovery Extensions

Future specifications may improve the efficiency of Neighbor Discovery tasks, which could be helpful for fast movements. One factor is currently being looked at: the delays caused by the Duplicate Address Detection mechanism. Currently, Duplicate Address Detection needs to be performed for every new care-of address as the mobile node moves, and for the mobile node's link-local address on every new link. In particular, the need and the trade-offs of re-performing Duplicate Address Detection for the link-local address every time the mobile node moves on to new links will need to be examined. Improvements in this area are, however, generally applicable and progress independently from the Mobile IPv6 specification.

Future functional improvements may also be relevant for Mobile IPv6 and other applications. For instance, mechanisms that would allow recovery from a Duplicate Address Detection collision would be useful for link-local, care-of, and home addresses.

Authors' Addresses

David B. Johnson
Rice University
Dept. of Computer Science, MS 132
6100 Main Street
Houston TX 77005-1892
USA

EEmail: dbj@cs.rice.edu

Charles E. Perkins
Nokia Research Center
313 Fairchild Drive
Mountain View CA 94043
USA

EEmail: charliep@iprg.nokia.com

Jari Arkko
Ericsson
02420 Jorvas
Finland

EEmail: jari.arkko@ericsson.com

Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

