

MIP4 Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 28, 2011

H. Deng
China Mobile
H. Levkowitz
Netnod
V. Devarapalli
WiChorus
S. Gundavelli
Cisco Systems
B. Haley
Hewlett-Packard Company
October 25, 2010

Generic Notification Message for Mobile IPv4
draft-ietf-mip4-generic-notification-message-16

Abstract

This document specifies protocol enhancements that allow Mobile IPv4 entities to send and receive explicit notification messages using a Mobile IPv4 message type designed for this purpose.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on April 28, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	5
2. Terminology	6
3. Notification Message - Usage Scenarios	7
3.1. Notification Message - Examples	7
3.2. Notification Message - Topology	7
3.2.1. Notification Message between a Home Agent and a Mobile Node	8
3.2.2. Notification Message between a Foreign Agent and a Mobile Node	8
3.2.3. Notification Message between a Home Agent and a Foreign Agent	9
4. Generic Notification Message and Considerations	10
4.1. Generic Notification Message	10
4.2. Generic Notification Acknowledgment Message	13
4.3. Notification Retransmission	16
4.4. General Implementation Considerations	17
4.5. Mobile Node Considerations	17
4.5.1. Receiving Generic Notification Messages	17
4.5.2. Sending Generic Notification Acknowledgement Messages	19
4.5.3. Sending Generic Notification Messages	19
4.5.4. Receiving Generic Notification Acknowledgement Messages	20
4.6. Foreign Agent Consideration	21
4.6.1. Receiving Generic Notification Messages	21
4.6.2. Sending Generic Notification Acknowledgement Messages	23
4.6.3. Sending Generic Notification Messages	24
4.6.4. Receiving Generic Notification Acknowledgement Messages	24
4.7. Home Agent Consideration	25
4.7.1. Sending Generic Notification Messages	25
4.7.2. Receiving Generic Notification Acknowledgement Messages	26
4.7.3. Receiving Generic Notification Messages	26
4.7.4. Sending Generic Notification Acknowledgement Messages	28
5. Future Extensibility	29
5.1. Examples of Possible Extensions	29
5.2. Extension Specification	29
6. IANA Considerations	31
7. Security Considerations	32
7.1. Replay Protection for GNM, GNAM messages	32
7.1.1. Replay Protection using Timestamps	33
7.1.2. Replay Protection using Nonces	34
7.2. Non-authentication Extensions Handling in Foreign Agent	34

8. Acknowledgments	35
9. References	36
9.1. Normative References	36
9.2. Informative References	36
Authors' Addresses	37

1. Introduction

In some situations, there is a need for Mobile IPv4 entities, such as the home agent(HA), foreign agent(FA) and mobile node(MN) to send and receive asynchronous notification messages during a mobility session. 'Asynchronous messages' in this context is used to mean messages which are not synchronous with the Registration Request and Registration Reply messages of the base Mobile IP Specification [RFC3344]. The base Mobile IP Specification does not have a provision for this.

This document defines a generic message and a notification model that can be used by Mobile IPv4 entities to send various notifications. It also defines a corresponding acknowledgement message to allow for reliable delivery of notifications. Only the following extensions may be present in these new messages, as defined by this document:

- MN-HA Authentication Extension
- MN-FA Authentication Extension
- FA-HA Authentication Extension
- Message String Extension

The semantics of receiving a generic notification message with a Message String Extension are null; i.e., it has no effect on the state of a mobile node's existing registration. See Section 3.1 for some application examples that motivate the new messages defined in this document.

2. Terminology

It is assumed that the reader is familiar with the terminology used in [RFC4917], [RFC3344]. In addition, this document frequently uses the following terms:

Notification Message

A message from a mobility agent to a MN or other mobility agent to asynchronously notify it about an event that is relevant to the mobility service it is currently providing.

Generic Notification Message

A Notification Message in the context of Mobile IPv4 with a well-defined envelope format and extensibility, and with certain limitations on how extensions may be defined and used, but otherwise generally available for notification purposes within the Mobile IPv4 protocol. Abbreviated 'GNM' in this document.

Generic Notification Acknowledgement Message

An acknowledgement of a received Generic Notification Message. Abbreviated 'GNAM' in this document.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119, [RFC2119].

3. Notification Message - Usage Scenarios

3.1. Notification Message - Examples

The simplest usage scenario for a notification message is one where the notification has no semantic meaning within the protocol; it is only carrying a message which can be displayed to a user or an operator (depending on which is the receiving entity -- see more on this below, in Section 3.2). Examples of such usage is messages from operator to user about billing or service related events ("You have used nearly all of your prepaid quota; there is only XX MB left -- please purchase further service if you are going to need it."; or "You have now used data transfer services for the amount of \$XXX since your last bill; this is above the notification threshold for your account.") or messages about service interruptions, and more. These examples are all supported by the use of the Mobile IPv4 Generic Notification Message together with the Message String Extension, as defined in this document.

There are also other examples, which cannot be implemented solely using the messages and extensions defined in this document. Some of these are described briefly below, and covered slightly more extensively in Section 5.

One example of an application of an extended Generic Notification Message is that during handover between CDMA 2000 1x EV-DO and Wireless LAN, the PPP resource on the CDMA side has to be removed on the FA (PDSN) to avoid over-charging subscribers. To address this, the Registration Revocation Message was defined in [RFC3543], but it would have been preferable to have had it defined as a separate message (i.e., the Generic Notification Message) with a Registration Revocation extension.

Other applications are HA switch over (before HA decide to go off-line it would like to notify the MNs to register with another candidate HA), NEMO prefix changes (MN is notified by HA about NEMO prefix changes and service or billing related events, which is an operational requirement), Load balancing (HA wants to move some of the registered MNs to other HAs), Service Termination (due to end of prepaid time), and Service Interruption (due to system maintenance).

3.2. Notification Message - Topology

There are several scenarios where a mobility agent could initiate notification events. Some of these are described in the following Sections.

3.2.1. Notification Message between a Home Agent and a Mobile Node

3.2.1.1. Mobile Registered using a Foreign Agent Care-of Address

In this case, the HA cannot directly notify the MN, but must send the notification via the FA, vice versa.

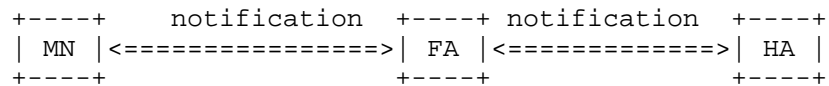


Figure 1: HA notifies MN or MN notifies HA through FA

3.2.1.2. Mobile Registered using a Co-located Care-of Address

In this case, the MN has registered with the home agent directly, so the notification message can go directly to the MN.

The notification mechanism as specified here does not support the case of Co-located CoA mode with registration through a FA (due to the 'R' bit being set in the FA's advertisement messages).

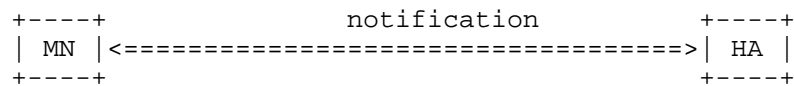


Figure 2: HA directly notifies MN or MN directly notifies HA

3.2.2. Notification Message between a Foreign Agent and a Mobile Node

There are two cases where a FA may send notification messages to a MN, one where it is relaying a message, the other where the notification is triggered by a message from another network entity, for example a AAA node(notification messages between a AAA entity and the FA could be based on RADIUS or Diameter, but this is out of scope for this document). If the notification is initiated by a FA, the FA may need to also notify the HA about the event.

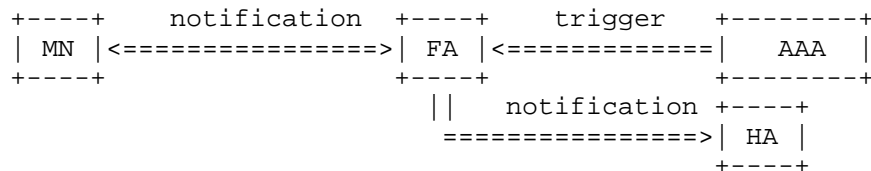


Figure 3: FA notifies MN

3.2.3. Notification Message between a Home Agent and a Foreign Agent

The HA may also need to send a notification to the FA, but not to the MN, The FA may also need to send a notification to the HA, as illustrated below:

```
+-----+ notification +-----+
|  FA  |<=====>|  HA  |
+-----+          +-----+
```

Figure 4: HA notifies FA or FA notifies HA

4. Generic Notification Message and Considerations

This section describes in detail the Generic Notification Message (GNM), Generic Notification Acknowledgement Message (GNAM), and some considerations related to the handling of these messages in the MN, FA and HA.

The MN and HA MUST maintain the following information, FA also needs to maintain both the HA's and MN's direction the below information:

- the IP source address of the Registration Request/Reply
- the IP destination address of the Registration Request/Reply
- the UDP source port of the Registration Request/Reply
- the UDP destination port of the Registration Request/Reply

The sending node always sends the GNM message following the same procedure for sending Registration Request as in Section 3.3 of [RFC3344] and the receiving node follows the same procedure for Registration Reply as in Section 3.4. of [RFC3344] when sending GNAM.

4.1. Generic Notification Message

A GNM is sent by a mobility agent to inform another mobility agent, or a MN, of MIP-related information in the form of a Message String Extension [RFC4917]. These messages MUST use the same IP and UDP headers as any previous Registration Request (RRQ) or Reply (RRP) message to the same entity. This would support NAT traversal and ensure same security association used for GNM/GNAM and RRQ/RRP. The GNM is defined as follows:

IP Fields:

Source Address	Typically copied from the destination address of the last Registration Reply/Request message that the agent received from the agent to which it is sending the GNM.
Destination Address	Copied from the source address of the last Registration Reply/Request message that the agent received from the agent to which it is sending the GNM.

UDP Fields:

Source Port Typically copied from the destination port of the last Registration Reply/Request message that the agent received from the agent to which it is sending the GNM.

Destination Port Copied from the source port of the last Registration Reply/Request message that the agent received from the agent to which it is sending the GNM.

The UDP header is followed by the Mobile IP fields shown below:

```

      0                               1                               2                               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   MD   |A|  Reserved  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Home Address
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Home Agent Address
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Care-of Address
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Identification
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Extensions...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type (To be assigned by IANA)

MD: Message Direction

This memo defines the semantics of the following MD field value:

- 0 -- Message sent by the HA to the MN
- 1 -- Message sent by the HA to the FA
- 2 -- Message sent by the MN to the HA
- 3 -- Message sent by the MN to the FA
- 4 -- Message sent by the FA to the MN

5 -- Message sent by the FA to the HA

A

This bit indicates whether the notification message MUST be acknowledged by the recipient. If "A" bit has been set during the message, but the sender doesn't receive any acknowledgement message, then the sender will have to re-send the notification message again.

Set to "1" to indicate that acknowledgement is REQUIRED.

Set to "0" to indicate that acknowledgement is OPTIONAL.

Reserved

MUST be sent as 0, and ignored when received.

Home Address

The home IP address of the mobile node.

Home Agent Address

The IP address of the mobile node's HA.

Care-of Address

The mobile node's care-of address, either the Co-located Care-of Address or the foreign agent care-of address.

Identification

A 64-bit number, constructed by the sender, used for matching GNM with GNAM, and for protecting against replay attacks of notification messages. See Section 7.1.1 and Section 7.1.2 for more on the use of timestamps and nonces in this field. Support for the use of timestamps is REQUIRED and support for nonces is OPTIONAL.

Extensions

The fixed portion of the GNM is followed by one or more extensions which may be used with this message, and by one or more authentication extensions as defined in Section 3.5 of [RFC3344].

Apart from the Authentication Extensions mentioned below, only one extension is defined in this document as permitted for use with

the GNM: the Message String Extension defined in [RFC4917].

This document requires the MN-HA Authentication Extension (AE) to be used when this message is sent between the MN and the HA; MN-FA AE and FA-HA AE are OPTIONAL. This document also requires the use of the MN-FA AE when this message is sent between the MN and the FA; where the MN-HA AE and FA-HA AE are not needed. This document finally require the use of the FA-HA AE when this message is sent between the FA and the HA, and the MN-HA AE and MN-FA AE are not needed. This could be determined based on the "MD" value. See Sections 3.6.1.3 and 3.7.2.2 of [RFC3344] for the rules on the order of these extensions as they appear in Mobile IPv4 RRQ and RRP messages. The same rules are applicable to GNM and GNAM.

4.2. Generic Notification Acknowledgment Message

A GNAM is sent by mobility agents or MNs to indicate the successful receipt of a GNM.

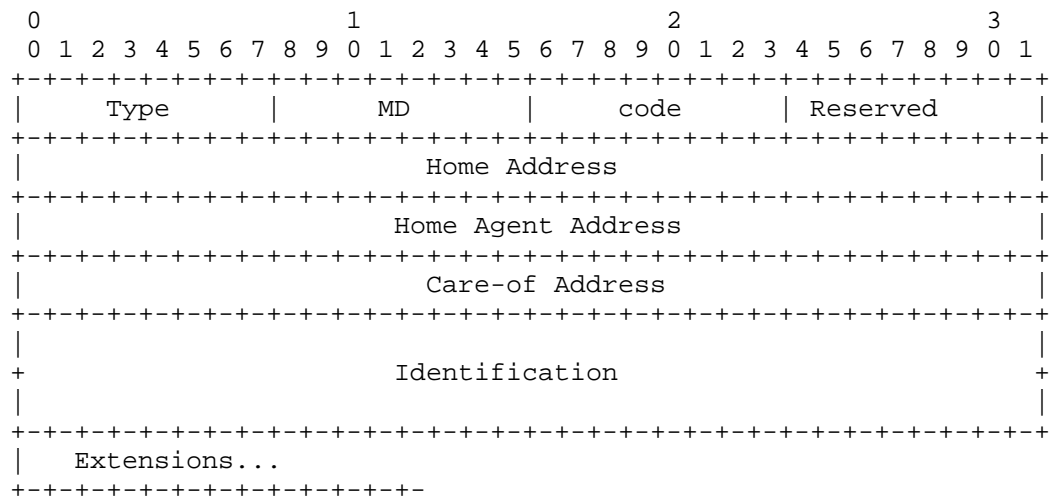
IP Fields:

Source Address	Typically copied from the destination address of the GNM to which the agent is replying.
Destination Address	Copied from the source address of the GNM to which the agent is replying.

UDP Fields:

Source Port	Copied from the destination port of the corresponding GNM.
Destination Port	Copied from the source port of the corresponding GNM.

The UDP header is followed by the Mobile IP fields shown below:



Type (To be assigned by IANA)

MD: Message Direction

This memo defines the semantics of the following MD field value:

- 0 -- Message sent by the HA to the MN
- 1 -- Message sent by the HA to the FA
- 2 -- Message sent by the MN to the HA
- 3 -- Message sent by the MN to the FA
- 4 -- Message sent by the FA to the MN
- 5 -- Message sent by the FA to the HA

code

A value indicating the result of the GNM. See below for a list of currently defined Code values.

Notification successful

- 0 -- notification accepted

Notification denied by the HA

- 128 -- reason unspecified
- 129 -- administratively prohibited
- 130 -- insufficient resources
- 131 -- mobile node failed authentication
- 132 -- foreign agent failed authentication
- 133 -- notification Identification mismatch

Notification denied by the FA

- 64 -- reason unspecified
- 65 -- administratively prohibited
- 66 -- insufficient resources
- 67 -- mobile node failed authentication
- 68 -- home agent failed authentication
- 69 -- notification Identification mismatch

Notification denied by the mobile node

- 192 -- reason unspecified
- 193 -- administratively prohibited
- 194 -- insufficient resources
- 195 -- foreign agent failed authentication
- 196 -- home agent failed authentication
- 197 -- notification Identification mismatch

Home Address

The home IP address of the mobile node.

Home Agent Address

The IP address of the sender's home agent.

Care-of Address

The mobile node's care-of address, either the Co-located Care-of Address or the foreign agent care-of address.

Identification

A 64-bit number used for matching GNM message with GNAM message and for protecting against replay attacks of registration messages. See Section 7.1.1 and Section 7.1.2 for more on the use of timestamps and nonces in this field. Support for the use of timestamps is REQUIRED and support for nonces is OPTIONAL. The value is based on the Identification field from the GNM message from the sender, and on the style of replay protection used in the security context between the sender and its receiver (defined by the mobility security association between them, and SPI value in the authorization-enabling extension).

Extensions

The fixed portion of the GNAM is followed by one or more extensions which may be used with this message, and by one or more authentication extensions as defined in Section 3.5 of [RFC3344].

This document REQUIRES the MN-HA Authentication Extension (AE) to be used when this message is sent between the MN and the HA; MN-FA AE and FA-HA AE are OPTIONAL. This document also requires the use of the MN-FA AE when this message is sent between the MN and the FA; where the MN-HA AE and FA-HA AE are not needed. This document finally requires the use of the FA-HA AE when this message is sent between the FA and the HA, and the MN-HA AE and MN-FA AE are not needed. This could be determined based on the "MD" value. See Sections 3.6.1.3 and 3.7.2.2 of [RFC3344] for the rules on the order of these extensions as they appear in Mobile IPv4 RRQ and RRP messages. The same rules are applicable to GNM and GNAM.

4.3. Notification Retransmission

If "A" flag has been set during the GNM message, but the sender doesn't receive any GNAM message within a reasonable time, then another GNM will be retransmitted. When timestamps are used, a new registration Identification is chosen for each retransmission; Thus it counts as a new GNM. When nonces are used, the unanswered GNM message is retransmitted unchanged; thus the retransmission does not count as a new GNM (Section 7.1). In this way a retransmission will not require the receiver to re-synchronize with the sender by issuing another nonce in the case in which the original GNM message (rather than its GNAM message) was lost by the network.

The maximum time until a new GNM message is sent SHOULD be no greater than the requested Lifetime of the last GNM message. The minimum value SHOULD be large enough to account for the size of the messages, twice the round trip time for transmission to the receiver, and at least an additional 100 milliseconds to allow for processing the messages before responding. The round trip time for transmission to the receiver will be at least as large as the time REQUIRED to transmit the messages at the link speed of the sender's current point of attachment. Some circuits add another 200 milliseconds of satellite delay in the total round trip time to the receiver. The minimum time between GNM MUST NOT be less than 1 second. Each successive retransmission timeout period SHOULD be at least twice the previous period, as long as that is less than the maximum as specified above.

4.4. General Implementation Considerations

Implementations of this specifications should provide support for management of the various settings related to the notification messages. In particular, it should be possible to do the following:

- * List the notification messages supported
- * Show enabled/disabled status for notification message support, overall and in detail.
- * Show the value of the maximum and minimum retransmission times.
- * Enable and disable notification support entirely.
- * Enable and disable the individual notification messages supported.
- * Set the value of the maximum and minimum retransmission times described in Section 4.3.

4.5. Mobile Node Considerations

It is possible that the MN MAY receive a GNM from a FA or HA. Both in the case of FA-CoA and Co-located CoA, the MN MAY reply with a GNAM based on the "A" flag in the GNM message.

4.5.1. Receiving Generic Notification Messages

When the MN is using FA-CoA and receives a Notification message, if the "MD" value is 0, it means that the notification message came from the HA. If the "MD" value is 4, the notification came from the FA.

If this notification message came from a FA and the MN accepts the FA's GNM, then it will process the notification extension according to the specific rules for that extension.

The MN MUST check for the presence of an authorization-enabling extension, and perform the indicated authentication. Exactly one authorization-enabling extension MUST be present in the GNM, if this message came from a FA, then MN-FA AE MUST be present. If no MN-FA AE is found, or if more than one MN-FA AE is found, or if the Authenticator is invalid, then the MN MUST reject the GNM and MAY send a GNAM to the FA with Code 195, including an Identification field computed in accordance with the rules specified in Section 7.1. The MN MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The MN MUST check that the Identification field is correct using the context selected by the SPI within mandatory authentication extension like MN-FA AE or MN-HA AE. See Section 7.1 for a description of how this is performed. If incorrect, the MN MUST reject the GNM and MAY send a GNAM to the initiator with Code 197, including an Identification field computed in accordance with the rules specified in Section 7.1. The MN MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The MN MUST also check that the extensions present in the Generic Notification Message are permitted for use with the GNM. If not, the MN MUST silently discard the message. It MUST NOT do any further processing with such a notification, though it SHOULD log the error.

After this, the MN MAY reply GNAM back to the FA. If the "A" flag is set in the GNM, then the MN MUST send the GNAM.

If this notification message came from the HA, relayed by the FA, or is a Co-located CoA, then the MN-HA AE MUST be checked and the MN MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, then the MN MUST reject the GNM and MAY send a GNAM to the initiator with Code 196, including an Identification field computed in accordance with the rules specified in Section 7.1. The MN MUST do no further processing with such a notification, though it SHOULD log the error as a security exception. If the MN accepts the HA's GNM, then it will process it according to the specific rules for that extension. After that, the MN MAY reply with a GNAM with Code 0 back to the HA based on the "A" flag in the GNM.

4.5.2. Sending Generic Notification Acknowledgement Messages

Both in the case of a Co-located CoA and FA-CoA, the MN MAY reply with a GNAM based on the "A" flag in the GNM as follows:

If the GNM was initiated from the FA to the MN ("MD" value is set to 4), then MN-FA AE MUST be the last extension in order to protect all other non-authentication extensions as defined in Section 3.5.3 of [RFC3344].

In the case of a FA-CoA, the source address is the MN's address, the destination address is the FA's address.

The Code field of the GNAM is chosen in accordance with the rules specified in Section 4.2. When replying to an accepted notification, a MN SHOULD respond with Code 0.

There are a number of reasons the MN might reject a notification such as administrative in nature returning a GNAM with a code of 193, similarly and provides the Code value 192 or 194 for the unspecified reason and insufficient resources.

If the GNM was initiated from the HA to the MN ("MD" value is set to 0) and in the case of Co-located CoA, then MN-HA AE MUST be the last extension in order to protect all other non-authentication extensions as defined in Section 3.5.2 of [RFC3344]

In the case of a FA-CoA, the source address is the MN's HoA address and the destination address is the FA's address ("MD" value is set to 2), the ordering of the extension is: any non-authentication Extensions used only by the HA, followed by the MN-HA AE defined in Section 3.5.2 of [RFC3344], followed by any non-authentication Extensions used only by the FA, followed by the MN-FA AE defined in Section 3.5.3 of [RFC3344].

4.5.3. Sending Generic Notification Messages

The MN may either send a GNM to notify the FA or HA.

If the message is sent to the FA, then the source address is the MN's address, and the destination address is the FA's address

If the FA is the target of this notification message, then the "MD" value is set to 3, MN-FA AE MUST be the last extension in order to protect all other non-authentication extensions. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344].

If the FA is working only as a relay agent, then the "MD" value is set to 2, and the ordering of the extension is: the notification extension, followed by any non-authentication extension expected to be used by HA, followed by MN-HA AE defined in Section 3.5.2 of [RFC3344], followed by any non-authentication Extensions used only by the FA, followed by The MN-FA AE defined in Section 3.5.3 of [RFC3344]. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344].

In the case of a Co-located CoA, the MN MAY send a notification message directly to the HA if it needs to be notified. The "MD" value is set to 2, and the ordering of the extension is: the notification extension, followed by any non-authentication extension expected to be used by HA, followed by MN-HA AE defined in Section 3.5.2 of [RFC3344].

The MN chooses the Identification field in accordance with the style of replay protection it uses with its HA. This is part of the mobility security association the MN shares with its HA. See Section 7.1 for the method by which the MN computes the Identification field.

4.5.4. Receiving Generic Notification Acknowledgement Messages

In the case of a FA-CoA, if the MN receives this message, and the "MD" value is set to 0, it means that the GNAM came from HA

If the "MD" value is set to 4, then the MN-FA AE MUST be checked, and the MN MUST check the Authenticator value in the Extension. If no MN-FA AE is found, or if more than one MN-FA AE is found, or if the Authenticator is invalid, then the MN MUST silently discard the GNAM.

In addition, the low-order 32 bits of the Identification field in the GNAM MUST be compared to the low-order 32 bits of the Identification field in the most recent GNM sent to the replying agent. If they do not match, then the GNAM MUST be silently discarded.

If the "MD" value is set to 0, then the MN-HA AE MUST be checked, and the MN MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, then the MN MUST silently discard the GNAM. If the MN accepted this message, then the MN MAY also process it based on the notification event.

In the case of a Co-located CoA, if the MN received this message, then the MN-HA AE MUST be checked, and the MN MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid,

then the MN MUST silently discard the Notification Acknowledgement message.

4.6. Foreign Agent Consideration

The FA may initiate a GNM to the MN or the HA. Additionally, the FA also relays GNM and GNAM messages between the MN and its HA as long as there is an active binding for the MN at the FA.

4.6.1. Receiving Generic Notification Messages

If the FA receives a GNM, and the "MD" value is set to 0, then it means that the HA is asking the FA to relay the message to the MN. If the "MD" value is set to 1, then it means that the target of the notification is the FA. If the "MD" value is set to 2, then it means that the MN is asking the FA to relay the message to the HA. If the "MD" value is set to 3, then it means that the notification came from the MN to the FA.

If the "MD" value is set to 0, then the FA MAY validate the FA-HA AE if present. If the FA-HA AE is invalid, then all extensions between the HA-MN AE and the HA-FA AE MUST be removed, FA SHOULD relay the GNM to the MN's home address as specified in the Home Address field of the GNM, MN will eventually validate the MN-HA AE to ensure that all information sent to the MN is integrity protected. If the FA-HA AE is valid, FA MUST relay the GNM to the MN's home address as specified in the Home Address field of the GNM. The FA MUST NOT modify any of the fields beginning with the fixed portion of the GNM through the MN-HA AE or other authentication extension supplied by the HA as an authorization-enabling extension for the MN.

Furthermore, the FA MUST process and remove any extensions following the MN-HA AE. If the FA shares a mobility security association with the MN, the FA MAY append any of its own non-authentication extensions which of relevance to the MN. In this case, the FA MUST append the MN-FA AE after these non-authentication extensions.

If the "MD" value is set to 1, the FA-HA AE MUST be checked, and the FA MUST check the Authenticator value in the Extension. If no FA-HA AE is found, or if more than one FA-HA AE is found, or if the Authenticator is invalid, the FA MUST reject the GNM and MAY send a GNAM to the HA with Code 68, including an Identification field computed in accordance with the rules specified in Section 7.1. The FA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The FA MUST check that the Identification field is correct using the context selected by the SPI within mandatory FA-HA AE. See

Section 7.1 for a description of how this is performed. If incorrect, the FA MUST reject the GNM and MAY send a GNAM to the initiator with Code 69, including an Identification field computed in accordance with the rules specified in Section 7.1. The FA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The FA MUST also check that the extensions present in the Generic Notification Message are permitted for use with the GNM. If not, the FA MUST silently discard the message. It MUST NOT do any further processing with such a notification, though it SHOULD log the error.

If FA accepts the HA's GNM, it will process it based on the specific rules for that extension. The FA MAY then reply with a GNAM with Code 0 back to the MN based on the "A" flag in the GNM.

In the case of a FA-CoA and if the "MD" value is set to 2, if the FA received this message, and if the MN-FA AE is present, the MN-FA AE MUST be checked, and the FA MUST check the Authenticator value in the Extension. If no MN-FA AE is found, or if more than one MN-FA AE is found, or if the Authenticator is invalid, the FA MUST silently discard the GNM message. If MN-FA is valid, FA MUST relay the GNM to the HA's address as specified in the Home Agent Address field of the GNM, HA will eventually validate the MN-HA AE to ensure that all information sent to the HA is integrity protected. The FA MUST NOT modify any of the fields beginning with the fixed portion of the GNM through the MN-HA AE or other authentication extension supplied by the MN as an authorization-enabling extension for the HA.

Furthermore, the FA MUST process and remove any Extensions following the MN-HA AE, and MAY append any of its own non-authentication Extensions of relevance to the HA if applicable, and MUST append the FA-HA AE, if the FA shares a mobility security association with the HA.

If the "MD" value is set to 3, the MN-FA AE MUST be checked, and the FA MUST check the Authenticator value in the Extension which is the same as the Section 3.7.2.1 of [RFC3344]. If no MN-FA AE is found, or if more than one MN-FA AE is found, or if the Authenticator is invalid, the FA MUST reject the GNM and MAY send a GNAM to the MN with Code 67, including an Identification field computed in accordance with the rules specified in Section 7.1. The FA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The FA MUST check that the Identification field is correct using the context selected by the SPI within mandatory MN-FA AE. See Section 7.1 for a description of how this is performed. If

incorrect, the FA MUST reject the GNM and MAY send a GNAM to the initiator with Code 69, including an Identification field computed in accordance with the rules specified in Section 7.1. The FA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

If FA accepts the MN's GNM, it will process it based on the specific rules for that extension. The FA MAY then reply with a GNAM with Code 0 back to the MN based on the "A" flag in the GNM.

4.6.2. Sending Generic Notification Acknowledgement Messages

The FA may need to either relay a GNAM message between the MN and the HA or send one as a response to a GNM message that was sent to it. In both cases, the GNAM message is defined as follows:

The source address is the FA address, the destination address is HA's or MN's home address.

The Code field of the GNAM is chosen in accordance with the rules specified in Section 4.2. When replying to an accepted notification, a FA SHOULD respond with Code 0.

There are a number of reasons the FA might reject a notification such as administrative in nature returning a GNAM with a code of 65, similarly and provides the Code value 64 or 66 for the unspecified reason and insufficient resources.

If the FA is only relaying this message to the HA, the FA MUST NOT modify any of the fields beginning with the fixed portion of the GNAM through the including the MN-HA AE or other authentication extension supplied by the MN as an authorization-enabling extension for the MN. Furthermore, the foreign agent MUST process and remove any Extensions following the MN-HA AE. If the FA shares a mobility security association with the HA, the FA MAY append any of its own non-authentication extensions which of relevance to the HA, In this case the FA MUST append the FA-HA AE after these non-authentication extensions.

If the notification message is from the HA to the FA then the "MD" value is set to 5 and the ordering of the extension is: any non-authentication Extensions used only by the FA, followed by The FA-HA AE defined in Section 3.5.4 of [RFC3344].

If the notification message is from the MN to the FA then the "MD" value is set to 4 and the ordering of the extension is: any non-authentication Extensions used only by the FA, followed by The MN-FA AE defined in Section 3.5.3 of [RFC3344].

4.6.3. Sending Generic Notification Messages

If the FA is initiating a notification to the MN using the GNM, it MAY also notify the HA as well.

In the message to the MN, the source address is the FA address, the destination address is the MN's address, the "MD" value is set to 4, and the ordering of the extension is: the notification extension, followed by any non-authentication Extensions used only by the MN, followed by The MN-FA AE defined in Section 3.5.3 of [RFC3344]. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344] except the payload is the notification other than registration.

In the message to the HA, the source address is the FA's address, the destination address is the HA's address (the "MD" value is set to 5), and the ordering of the extension is: notification extension, followed by any non-authentication Extensions used only by the HA, followed by The FA-HA AE defined in Section 3.5.4 of [RFC3344]. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344] except the payload is the notification other than registration.

4.6.4. Receiving Generic Notification Acknowledgement Messages

In the case of a FA-CoA, if the FA receives this message, and the "MD" value is set to 3, it means that the notification acknowledgement message came from the MN, otherwise it came from the HA.

If the "MD" value is set to 1, the FA-HA AE MUST be checked, and the FA MUST check the Authenticator value in the Extension. If no FA-HA AE is found, or if more than one FA-HA AE is found, or if the Authenticator is invalid, the FA MUST silently discard the Notification Acknowledgement message. If the FA accepted this message, the FA MAY also process it based on the notification event.

If the "MD" value is set to 3, if the MN-FA AE is present, it MUST be checked, and the FA MUST check the Authenticator value in the Extension. If no MN-FA AE is found, or if more than one MN-FA AE is found, or if the Authenticator is invalid, the FA MUST silently discard the GNAM message. If the FA accepted this message, the FA MAY also process it based on the notification event.

In the case of a FA-CoA and if the "MD" value is set to 2, if the FA received this message, and if the MN-FA AE is present, the MN-FA AE MUST be checked, and the FA MUST check the Authenticator value in the Extension. If no MN-FA AE is found, or if more than one MN-FA AE is

found, or if the Authenticator is invalid, the FA MUST silently discard the GNAM message. If FA accepted the MN's GNAM message, it MUST relay this message to the HA. The FA MUST NOT modify any of the fields beginning with the fixed portion of the GNAM message through the including the MN-HA AE or other authentication extension supplied by the HA as an authorization-enabling extension for the MN. Furthermore, the FA MUST process and remove any Extensions following the MN-HA AE and MAY append any of its own non-authentication Extensions of relevance to the HA, if applicable, and MUST append the FA-HA AE, if the FA shares a mobility security association with the HA.

4.7. Home Agent Consideration

The HA MAY initiate a GNM message to both the mobile node and FA, and it also MAY receive a GNAM message from both the FA and MN. The HA also MAY receive a GNM message from the FA, but only when there is a binding for a MN. If the HA receives a GNM from a FA and there is no corresponding MN registration, the HA SHOULD drop the GNM message.

4.7.1. Sending Generic Notification Messages

In the case of a FA-CoA, the HA may either send a GNM to notify the FA, or have the FA relay the GNM to the MN if the MN needs to be notified.

If the message is from the HA to the FA, the source address is the HA's address, and the destination address is the FA's address

If the FA is working only as a relay agent, the "MD" value is set to 0, and the ordering of the extension is: the notification extension, followed by any non-authentication extension expected to be used by MN, followed by MN-HA AE defined in Section 3.5.2 of [RFC3344], followed by any non-authentication Extensions used only by the FA, followed by The FA-HA AE defined in Section 3.5.4 of [RFC3344]. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344].

If the FA is the target of this notification message, then the "MD" value is set to 1, and the ordering of the extension is: the notification extension, followed by any non-authentication Extensions used only by the FA, followed by The FA-HA AE defined in Section 3.5.4 of [RFC3344]. Computing Authentication Extension Value is the same as Section 3.5.1 of [RFC3344].

In the case of a Co-located CoA, the HA MAY send a notification message directly to the MN if it needs to be notified. The "MD" value is set to 0, and the ordering of the extension is: the

notification extension, followed by any non-authentication extension expected to be used by MN, followed by MN-HA AE defined in Section 3.5.2 of [RFC3344].

4.7.2. Receiving Generic Notification Acknowledgement Messages

In the case of a FA-CoA, if the HA receives this message, and the "MD" value is set to 2, it means that the GNAM message came from MN.

If the "MD" value is set to 5, and the HA accepted this message, the HA MAY also process it based on the notification event. The FA-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no FA-HA AE is found, or if more than one FA-HA AE is found, or if the Authenticator is invalid, the HA MUST silently discard the GNAM message.

If the "MD" value is set to 2, in the case of a FA-CoA, and if FA-HA AE is present, the FA-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If more than one FA-HA AE is found, or if the Authenticator is invalid, the HA MUST silently discard the GNAM message. Anyway, MN-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, the HA MUST silently discard the GNAM. If the HA accepted this message, the HA MAY also process it based on the notification event.

If the "MD" value is set to 2, in the case of a Co-located CoA, MN-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, the HA MUST silently discard the GNAM. If the HA accepted this message, the HA MAY also process it based on the notification event.

4.7.3. Receiving Generic Notification Messages

The HA MAY receive a GNM message sent from the FA. When the HA receives this message, if the the "MD" value is set to 5, this message came from FA. FA-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no FA-HA AE is found, or if more than one FA-HA AE is found, or if the Authenticator is invalid, the HA MUST reject the GNM and MAY send a GNAM to the FA with Code 132, including an Identification field computed in accordance with the rules specified in Section 7.1. The HA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception.

The HA MUST check that the Identification field is correct using the

context selected by the SPI within mandatory authentication extension like MN-HA AE or FA-HA AE. See Section 7.1 for a description of how this is performed. If incorrect, the HA MUST reject the GNM and MAY send a GNAM to the initiator with Code 133, including an Identification field computed in accordance with the rules specified in Section 7.1. The HA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception. If HA accepts the FA's GNM message, it will process it based on the notification extension. Furthermore, the HA MAY reply with a GNAM message with Code 0 back to the FA based on the "A" flag in the GNM message.

If the the "MD" value is set to 2, this message come from MN, in the case of FA-COA, if FA-HA AE is present, it MUST be checked, and the HA MUST check the Authenticator value in the Extension. If more than one FA-HA AE Extension is found, or if the Authenticator is invalid, the HA MUST reject the GNM and MAY send a GNAM to the FA with Code 132, including an Identification field computed in accordance with the rules specified in Section 7.1. The HA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception. And MN-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, the HA MUST reject the GNM and MAY send a GNAM to the MN with Code 131, including an Identification field computed in accordance with the rules specified in Section 7.1. The HA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception. If HA accepts the MN's GNM message, it will process it based on the notification extension. Furthermore, the HA MAY reply with a GNAM message back to the MN with Code 0 based on the "A" flag in the GNM message.

If the the "MD" value is set to 2, in the case of a Co-located CoA, the MN-HA AE MUST be checked, and the HA MUST check the Authenticator value in the Extension. If no MN-HA AE is found, or if more than one MN-HA AE is found, or if the Authenticator is invalid, the HA MUST reject the GNM and MAY send a GNAM to the MN with Code 131, including an Identification field computed in accordance with the rules specified in Section 7.1. The HA MUST do no further processing with such a notification, though it SHOULD log the error as a security exception. If HA accepts the MN's GNM message, it will process it based on the notification extension. Furthermore, the HA MAY reply with a GNAM message back to the MN with Code 0 based on the "A" flag in the GNM message.

The HA MUST also check that the extensions present in the Generic Notification Message are permitted for use with the GNM. If not, the HA MUST silently discard the message. It MUST NOT do any further

processing with such a notification, though it SHOULD log the error.

4.7.4. Sending Generic Notification Acknowledgement Messages

If the GNM message came from the FA only, and if the "A" flag is set in the GNM message, then the HA MUST send a GNAM message. The message is as follows: The source address is HA's address, the destination address is the FA's address, the "MD" value is set to 1. The ordering of the extension is: any non-authentication Extensions used only by the FA, followed by The Foreign-Home Authentication extension defined in Section 3.5.4 of [RFC3344].

The Code field of the GNAM is chosen in accordance with the rules specified in Section 4.2. When replying to an accepted GNM, a MN SHOULD respond with Code 0.

If the GNM message came from the MN, and if the "A" flag is set in the GNM message, then the HA MUST send a GNAM message. The message is as follows: The source address is HA's address, the destination address is the FA's address, the "MD" value is set to 0. The ordering of the extension is: any non-authentication Extensions used only by the MN, followed by the MN-HA AE defined in Section 3.5.2 of [RFC3344], optionally followed by any non-authentication Extensions used only by the FA, optionally followed by The MN-FA AE defined in Section 3.5.3 of [RFC3344]

5. Future Extensibility

This document defines the Generic Notification Message used with the Message String Extension [RFC4917].

It is however possible to define new notification-related extensions for use with the Generic Notification Message, for cases where the notification is intended to have a semantic content and is intended for the HA, FA or MN, rather than for the user.

5.1. Examples of Possible Extensions

One example of such usage, which would have been defined in this document if it hadn't already been defined as a separate message is the Registration Revocation Message [RFC3543]. This is a message sent from the HA to FA(s) or MN to notify the receiving node that a currently active registration is being revoked. The use case for this is clearly laid out in [RFC3543].

Another example would be managed maintenanceswitch-over between HA instances, where a HA due to go down for maintenance could direct the MNs registered with it to re-register with another specified HA. Such a message could also be used for managed load balancing. There is currently no support for such forced switch-over in the Mobile IPv4 protocol.

Yet another example is when the prefix set handled by an MIPv4 NEMO [RFC5177] HA changes; to ensure proper routing, the mobile router needs to be notified about the change so that its internal routing rules may be updated.

One final example is home network changes which require host configuration changes, for instance a change of address for the DNS server or another network server; again this is a case where the HA would want to notify the MN of the change, so that service interruptions can be avoided.

5.2. Extension Specification

In order to avoid making the MIPv4 Generic Notification Message a generic protocol extension mechanism by which new protocol mechanisms could be implemented without appropriate discussion and approval, any new extensions which are to be used with the Generic Notification Message must be registered with IANA, where registration is limited by the 'RFC Required' policy defined in [RFC5226]

If additional extensions are specified for use with the Generic Notification Message, the practice exemplified in [RFC3344] and

related specification should be followed. Generally it has not been necessary so far to provide versioning support within individual extensions; in a few cases it has been necessary to define new extensions with new extension numbers where a generalizations of a pre-existing extension has been needed, and with the current rate of extension number consumption that seems to be an acceptable approach.

If at some point extensions are specified for use with the Generic Notification Message which overlap pre-existing notification messages, the authors of the specification should consider providing a method to flag which notification messages are supported, and which notification message usage is requested, in a manner similar to the way tunnelling method capabilities and usage requests are flagged in the Mobile IPv4 Base Specification [RFC3344].

Encoded in the extension number of Mobile IPv4 extensions is the notion of 'skippable' and 'not skippable' extensions; see Section 1.8 of [RFC3344]. This notion is also applicable when extensions are used with the Generic Notification Message: It is not required that a receiver understand a skippable extension, but a non-skippable extension needs to be handled according to Section 1.8 of [RFC3344] (i.e., the message must be silently discarded if the extension is not recognized). This document does not specify any change from the Mobile IPv4 Base Specification [RFC3344] in this respect.

6. IANA Considerations

This document defines two new messages, the Generic Notification Message described in Section 4.1, and the Generic Notification Acknowledgement Message, described in Section 4.2. The message numbers for these two message numbers are to be allocated from the same number space used by the Registration Request and Registration Reply messages in [RFC3344].

The Generic Notification Message may only carry extensions which are explicitly permitted for use with this message. This document defines 4 extensions which are permitted, in Section 4.1. IANA must establish a register of Mobile IPv4 extensions which are permitted for use with the Generic Notification Message. Approval of new extensions which are permitted for use with the Generic Notification Message requires that they be defined in an RFC according to the 'RFC Required' policy described in [RFC5226].

The Generic Notification Acknowledgement message, specified in Section 4.2, has a Code field. The number space for the Code field values is new, and also specified in Section 4.2. The Code number space is structured according to whether the notification was successful, or whether the HA denied the notification, or whether FA denied the notification, or whether MN denied the notification, as follows:

0	Success Code
64-69	Error Codes from the FA
128-133	Error Codes from the HA
192-197	Error Codes from the MN

Approval of new Code values require expert review.

7. Security Considerations

This specification operates with the security constraints and requirements of [RFC3344]. This means that when these message is transmitted between the MN and the HA, MN-HA AE is REQUIRED, when this message is transmitted between the MN and the FA, MN-FA AE is REQUIRED, when this message is transmitted between the FA and the HA, FA-HA AE is REQUIRED. It extends the operations of MN, HA and FA defined in [RFC3344] to notify each other about some events. The GNM message defined in the specification could carry information that modifies the mobility bindings. Therefore the message MUST be integrity protected. Replay protection MUST also be guaranteed.

RFC 3344 provides replay protection only for registration requests sent by the MN. There is no mechanism for replay protection for messages initiated by a FA or a HA. The 64-bit Identification field specified in this document (Section 4.1 and 4.2) for the GNM message is used to provide replay protection for the notification messages initiated by the FA or HA.

7.1. Replay Protection for GNM, GNAM messages

The Identification field is used to let the receiving node verify that a GNM has been freshly generated by the sending node, not replayed by an attacker from some previous registration. Two methods are described in this section: timestamps (REQUIRED) and "nonces" (OPTIONAL). All senders and receivers MUST implement timestamp-based replay protection. These nodes MAY also implement nonce-based replay protection

The style of replay protection in effect between any two peer nodes among MN, FA and HA is part of the mobile security association. A sending node and its receiving node MUST agree on which method of replay protection will be used. The interpretation of the Identification field depends on the method of replay protection as described in the subsequent subsections.

Whatever method is used, the low-order 32 bits of the Identification MUST be copied unchanged from the GNM to the GNAM. The receiver uses those bits (and the sender's source address) to match GNAM with corresponding replies. The receiver MUST verify that the low-order 32 bits of any GNAM are identical to the bits it sent in the GNM.

The Identification in a new GNM MUST NOT be the same as in an immediately preceding GNM, and SHOULD NOT repeat while the same security context is being used between the MN and the HA.

7.1.1.1. Replay Protection using Timestamps

The basic principle of timestamp replay protection is that the node generating a message inserts the current time of day, and the node receiving the message checks that this timestamp is sufficiently close to its own time of day. Unless specified differently in the security association between the nodes, a default value of 7 seconds MAY be used to limit the time difference. This value SHOULD be greater than 3 seconds. Obviously the two nodes must have adequately synchronized time-of-day clocks. As with any messages, time synchronization messages may be protected against tampering by an authentication mechanism determined by the security context between the two nodes.

In this document, the timestamps are used, the sender MUST set the Identification field to a 64-bit value formatted as specified by the Network Time Protocol (NTP) [RFC5905]. The low-order 32 bits of the NTP format represent fractional seconds. Note, however, that when using timestamps, the 64-bit Identification used in a GNM message from the sender MUST be greater than that used in any previous GNM message, as the receiver uses this field also as a sequence number. Without such a sequence number, it would be possible for a delayed duplicate of an earlier GNM message to arrive at the receiver (within the clock synchronization required by the receiver), and thus be applied out of order, mistakenly altering the sender's current status.

Upon receipt of a GNM message with an authorization-enabling extension, the receiver MUST check the Identification field for validity. In order to be valid, the timestamp contained in the Identification field MUST be close enough to the receiver's time of day clock and the timestamp MUST be greater than all previously accepted timestamps for the requesting sender. Time tolerances and re-synchronization details are specific to a particular mobility security association.

If the timestamp is valid, the receiver copies the entire Identification field into the GNAM it returns the GNAM message to the sender. If the timestamp is not valid, the receiver copies only the low-order 32 bits into the GNAM, and supplies the high-order 32 bits from its own time of day. In this latter case, the receiver MUST reject the registration by returning Code 69/133/197 (identification mismatch) in the GNAM message.

Furthermore, the receiver MUST verify that the low-order 32 bits of the Identification in the GNAM are identical to those in the rejected GNM attempt, before using the high-order bits for clock re-synchronization.

7.1.2. Replay Protection using Nonces

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

The receiver may be expected to have resources for computing pseudo-random numbers useful as nonces, according to [RFC4086]. It inserts a new nonce as the high-order 32 bits of the identification field of every GNAM message. The receiver copies the low-order 32 bits of the Identification from the GNM message into the low-order 32 bits of the Identification in the GNAM message. When the sender receives an authenticated GNAM message from the receiver, it saves the high-order 32 bits of the identification for use as the high-order 32 bits of its next GNM message.

The sender is responsible for generating the low-order 32 bits of the Identification in each GNM message. Ideally it should generate its own random nonces. However it may use any expedient method, including duplication of the random value sent by the receiver. The method chosen is of concern only to the sender, because it is the node that checks for valid values in the GNAM message. The high-order and low-order 32 bits of the identification chosen SHOULD both differ from their previous values. The receiver uses a new high-order value and the sender uses a new low-order value for each registration message.

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

7.2. Non-authentication Extensions Handling in Foreign Agent

When the FA is relaying the GNM message between the MN and the HA, and if the FA does not share a mobility security association with the MN or HA, all non-authentication extensions between MN and FA, or FA and HA are not protected; In this case, all non-authentication extensions should be silently discarded.

8. Acknowledgments

The author appreciate the efforts of Ahmad Muhanna for his detail reviewing of this document and his many contributions to the text of this document. The author also wants to thank Kent Leung, Peng Yang and Peter McCann et al. for their helping developing this document. Thanks to Alexey Melnikov, Sean Turner, Ralph Droms, Charles E. Perkins, Russ Housley, Magnus Westerlund, Lars Eggert, Dan Romascanu, Tim Polk, Amanda Baber, Sebastian Thalanany, and Joseph Salowey's discussion and comments. Thanks to Jari Arkko for each step of this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3543] Glass, S. and M. Chandra, "Registration Revocation in Mobile IPv4", RFC 3543, August 2003.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC4917] Sastry, V., Leung, K., and A. Patel, "Mobile IPv4 Message String Extension", RFC 4917, June 2007.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

9.2. Informative References

- [RFC5177] Leung, K., Dommety, G., Narayanan, V., and A. Petrescu, "Network Mobility (NEMO) Extensions for Mobile IPv4", RFC 5177, April 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

Authors' Addresses

Hui Deng
China Mobile
53A,Xibianmennei Ave.,
Xuanwu District,
Beijing 100053
China

Email: denghui02@gmail.com

Henrik Levkowetz
Netnod
Franzengatan 5
S-104 25, Stockholm
SWEDEN

Email: henrik@levkowetz.com

Vijay Devarapalli
WiChorus
3590 North First St
San Jose, CA
USA

Email: dvijay@gmail.com

Sri Gundavelli
Cisco Systems
170 W.Tasman Drive
San Jose, CA 95134
USA

Email: sgundave@cisco.com

Brian Haley
Hewlett-Packard Company
110 Spitbrook Road
Nashua, NH 03062
USA

Email: brian.haley@hp.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 18, 2011

P. Yegani
Juniper Networks
K. Leung
Cisco Systems
A. Lior
Bridgewater Systems
K. Chowdhury
J. Navali
Cisco Systems
Oct 15, 2010

GRE Key Extension for Mobile IPv4
draft-ietf-mip4-gre-key-extension-03.txt

Abstract

The GRE specification contains a Key field, which MAY contain a value that is used to identify a particular GRE data stream. This specification defines a new Mobile IP extension that is used to exchange the value to be used in the GRE Key field. This extension further allows the Mobility Agents to set up the necessary protocol interfaces prior to receiving the mobile's traffic. The new extension allows a foreign agent to request GRE tunneling without disturbing the Home Agent behavior specified for Mobile IPv4. GRE tunneling with the Key field allows the operators to have home networks that consist of multiple Virtual Private Networks (VPNs), which may have overlapping home addresses. When the tuple < Care of Address, Home Address and Home Agent Address > is the same across multiple subscriber sessions, GRE tunneling will provide a means for the FA and HA to identify data streams for the individual sessions based on the GRE key. In the absence of this key identifier, the data streams cannot be distinguished from each other, a significant drawback when using IP-in-IP tunneling.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference

material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 18, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
2. Terminology	4
3. GRE-Key Extension	4
4. Operation and Use of the GRE-Key Extension	4
4.1. Foreign Agent Requirements for GRE Tunneling Support	4
4.2. Home Agent Requirements for GRE Tunneling Support	5
4.3. Mobile Node Requirements for GRE Tunneling Support	6
5. GRE Key Extension and Tunneling Procedures	6
6. IANA Considerations	7
7. Security Considerations	7
8. Acknowledgements	8
9. Normative references	8
Authors' Addresses	8

1. Introduction

This document specifies a new extension for use by Foreign Agents operating Mobile IP for IPv4. The new extension allows a foreign agent to request GRE tunneling without disturbing the Home Agent behavior specified for Mobile IPv4 [RFC3344]. This extension contains the GRE key [RFC2890] required for establishing a GRE tunnel between the FA and the HA.

GRE tunneling with the Key field allows the operators to have home networks that consist of multiple Virtual Private Networks (VPNs), which may have overlapping home addresses. When the tuple < Care of Address, Home Address and Home Agent Address > is the same across multiple subscriber sessions, GRE tunneling will provide a means for the FA and the HA to identify data streams for the individual sessions based on the GRE key. In the absence of this key identifier, the data streams cannot be distinguished from each other, a significant drawback when using IP-in-IP tunneling.

2. 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 [RFC2119]. Other terminology is used as already defined in [RFC3344].

3. GRE-Key Extension

The format of the GRE-Key Extension conforms to the Extension format specified for Mobile IPv4 [RFC3344]. This extension option is used by the Foreign Agent to supply GRE key and other necessary information to the Home Agent to establish a GRE tunnel between the FA and the HA.

4. Operation and Use of the GRE-Key Extension

4.1. Foreign Agent Requirements for GRE Tunneling Support

The FA MUST support IP-in-IP tunneling of datagrams for Mobile IPv4 [RFC3344]. The FA may support GRE tunneling that can be used, for example, to allow for overlapping private home IP addresses [X.S0011-D]. If the FA is capable of supporting GRE encapsulation, it should set the 'G' bit in the Flags field in the Agent Advertisement message sent to the MN during the Mobile IP session establishment.

If the MN does not set the 'G' bit, the FA MAY fall back to using IP-in-IP encapsulation for the session per [RFC3344].

If the MN does not set both the 'G' bit and the 'D' bit (i.e., the mobile node is not using a co-located care-of address), and the local policy allows the FA to override the 'G' bit setting received from the MS, the FA MUST include the GRE-Key Extension as defined in this memo in the Registration Request that it propagates to the HA. The presence of this extension is a request for GRE encapsulation that takes precedence over the setting of the 'G' bit in the Registration Request. The FA MUST NOT modify the 'G' bit in the Registration Request because it is protected by the Mobile-Home Authentication Extension.

If the FA does not support GRE encapsulation, the FA MUST reset the 'G' bit in the Agent Advertisement message. In this case, if the MN sets the 'G' bit in the Registration Request message, the FA returns a Registration Reply message to the MN with code 'Requested Encapsulation Unavailable' (72) per [RFC3344].

If the FA allows GRE encapsulation, and either the MN requested GRE encapsulation or local policy dictates using GRE encapsulation for the session and the 'D' bit is not set (i.e., the mobile node is not using a co-located care-of address), the FA MUST include the GRE Key in the GRE Key Extension in all Mobile IP Registration Requests (including initial, renewal and de-registration requests) before forwarding the request to the HA. The FA may include a GRE key of value zero in the GRE Key Extension to signal that the HA assign GRE keys in both directions. The GRE key assignment in the FA and the HA is outside the scope of this memo.

The GRE Key Extension SHALL follow the format defined in [RFC3344]. This extension SHALL be added after the MN-HA and MN-FA Challenge and MN-AAA extensions (if any) and before the FA-HA Auth extension (if any).

4.2. Home Agent Requirements for GRE Tunneling Support [RFC3344]

The HA MUST follow the procedures specified in RFC 3344 in processing this extension in Registration Request messages. If the HA receives the GRE Key Extension in a Registration Request and does not recognize this non-skippable extension, it MUST silently discard the message. The HA MUST use other alternative forms of encapsulation (e.g., IP-in-IP tunneling), when requested by the mobile node per [RFC3344].

If the HA receives the GRE Key Extension in a Registration Request

and recognizes the GRE Key Extension but is not configured to support GRE encapsulation, it MUST send an RRP with code 'Requested Encapsulation Unavailable (139)' [RFC3024] .

If the HA receives a Registration Request with a GRE Key Extension but without the 'G' bit set, the HA SHOULD treat this as if 'G' bit is set in the Registration Request i.e., the presence of GRE Key Extension indicates a request for GRE encapsulation.

If the HA receives the GRE Key Extension in a Registration Request and recognizes the GRE Key Extension as well as supports GRE encapsulation, the following procedures should apply:

The HA SHOULD accept the RRQ and send a RRP with code 'Accepted (0)'. The HA MUST assign a GRE key and include the GRE Key Extension in the RRP before sending it to the FA. The HA MUST include the GRE Key Extension in all RRP in response to any RRQ that included GRE Key Extension, when a GRE key is available for the registration.

If the HA receives the GRE Key Extension in the initial Registration Request and recognizes the GRE Key Extension carrying a GRE key value of zero, it SHOULD accept the RRQ and send a RRP with code 'Accepted (0)'. The HA MUST assign GRE keys for both directions and include these keys in the GRE Key Extension in the RRP before sending it to the FA. The HA MUST include the GRE Key Extension in the RRP in response to the initial RRQ that included GRE Key Extension, when a GRE key is available for the registration.

4.3. Mobile Node Requirements for GRE Tunneling Support

If the MN is capable of supporting GRE encapsulation, it SHOULD set the 'G' bit in the Flags field in the Registration Request per [RFC3344].

5. GRE Key Extension and Tunneling Procedures

GRE tunneling support for Mobile IP will permit asymmetric GRE keying i.e., the FA assigns a GRE key for use in encapsulated traffic and the HA can assign its own GRE key. Once the GRE keys have been exchanged, the FA uses the HA-assigned key in the encapsulating GRE header for reverse tunneling and the HA uses the FA-assigned key in the encapsulating GRE header.

The format of the GRE Key Extension is as shown below.

The GRE Key extension MAY be included in Registration Requests [RFC3344]. The GRE Key extension is used to inform the recipient of

the Mobile IP request of the value to be used in GRE's Key field.




Figure 1: GRE Key Extension

Type

To be assigned by IANA. An 8-bit identifier of the GRE Key Extension type (non-skippable)

Sub-Type

0

Length

4

Key Identifier

This is a four octet value assigned during registration and inserted in every GRE packet of the user traffic.

6. IANA Considerations

The GRE Key extension defined in this memo is a Mobile IP extension as defined in [RFC3344]. IANA should assign a Type value for this Extension from the non-skippable range (0-127).

7. Security Considerations

This specification does not introduce any new security considerations, beyond those described in [RFC3344]

Despite its name, the GRE Key extension has little to do with security. The word "Key" here is not used in the cryptographic sense of a shared secret that must be protected, but rather is used in the sense of an "index" or demultiplexing value that can be used to distinguish packets belonging to a given flow within a GRE tunnel.

8. Acknowledgements

Thanks to Jun Wang, Gopal Dommety and Sri Gundavelli for their helpful comments, offline discussions and reviewing the initial draft. Also, Pete McCann and Simon Mizikovsky provided valuable review comments.

9. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", RFC 2890, September 2000.
- [RFC3024] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.

Authors' Addresses

Parviz Yegani
Juniper Netowrks
1194 North Mathilda Ave.
Sunnyvale, California 94089
U.S.A

Phone: +1 408-759-1973
Email: pyegani@juniper.net

Kent Leung
Cisco Systems Incorporated
170 West Tasman Drive
San Jose, California 95134
U.S.A

Phone: +1 408 526 5030
Email: kleung@cisco.com

Avi Lior
Bridgewater Systems Corporation
303 Terry Fox Drive
Ottawa, Ontario K2K 3J1
Canada

Phone: +1 613-591-6655
Email: avi@bridgewaterstystems.com

Kuntal Chowdhury
Cisco Systems Incorporated
170 West Tasman Drive
San Jose, California 95134
U.S.A

Email: kchowdhu@cisco.com

Jay Navali
Cisco Systems Incorporated
170 West Tasman Drive
San Jose, California 95134
U.S.A

Email: jnavali@cisco.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2011

A. Muhanna (Ed.)
Ericsson
B. Patil
Nokia
S. Chakrabarti
IP Infusion
G. Montenegro
Microsoft Corporation
Y. Wu
ZTE USA
July 12, 2010

IPv4 Mobility Extension for Multicast and Broadcast Packets
draft-ietf-mip4-mcbc-01.txt

Abstract

This document specifies a new Mobile IPv4 extension which is used to negotiate the Multicast-Broadcast Encapsulation Delivery style in the case of Mobile IPv4 Foreign Agent Care-of Address mode registration. With this extension the mobile node is able to negotiate the type of traffic that needs to be encapsulated for delivery to the foreign agent while other types of traffic use the direct delivery style. This mechanism eliminates the tunnel overhead between the mobile node and the foreign agent. Multicast and broadcast applications on a mobile IPv4 mobile node are better served with this extension.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on January 13, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions & Terminology	5
2.1. Conventions used in this document	5
2.2. Terminology	5
3. Multicast-Broadcast Encapsulating Delivery Style	5
3.1. Multicast-Broadcast Encapsulating Delivery Extension	6
3.2. Packet Header Formats for Visited Network Traffic	7
3.3. Packet Header Formats for Homebound Traffic	8
4. Multicast-Broadcast Encapsulating delivery Style Vs RFC3024 Encapsulating delivery	8
5. Link-layer Assisted Delivery Style (LLADS)	9
6. Security Considerations	10
7. IANA Considerations	10
8. Acknowledgments	10
9. References	10
9.1. Normative references	10
9.2. Informative references	11
Appendix A. Appendix-A	11
Authors' Addresses	11

1. Introduction

The IP Mobility Protocol [RFC3344] describes multicast and broadcast packet transmission between the mobile node and the home network or visited network. Reverse Tunneling for Mobile IP [RFC3024] includes support for reverse tunneling of multicast and broadcast packets to the home network using the encapsulating delivery style between the mobile nodes and the foreign agent. However, [RFC3024] says that once the encapsulated delivery style is negotiated, all packets exchanged between the mobile node and the foreign agent must be delivered encapsulated. The delivery (of packets between the MN and FA) methods specified in the base mobile IPv4 specification [RFC3344] prevents an MN from sending unicast packets to the FA. Tunneling overhead is an issue especially on wireless links with the current specification. Multicast and broadcast applications for a MN running mobile IPv4 client software also are negatively impacted. In particular, this imposition prevents direct delivery of unicast packets from the mobile node to the foreign agent. This causes a huge tunnel overhead in the (typically) wireless medium between the mobile node and the foreign agent and indirectly makes it impossible for the mobile node to use any of the multicast and broadcast services.

Additionally, [RFC3344] sections 4.3 and 4.4 discusses multicast and broadcast routing to and from the mobile node in the presence of triangular routing and with a co-located Care-of address. Reverse tunneling for Mobile IP [RFC3024] uses the optimal direct delivery style from the mobile node via the foreign agent if only unicast traffic is being reverse tunneled. If, however, multicast or broadcast packets are also meant to be reverse tunneled, it introduces the Encapsulating Delivery Style. Unfortunately, once the encapsulating delivery style is negotiated, it applies to all reverse tunneling traffics, including unicast. [RFC3344] also mandates, in the case of FA Care-of Address mode, that all multicast and broadcast packets be delivered encapsulated to mobile node. This also imposes tunnel overhead for multicast and broadcast packets. While tunneling overhead on wired links may be acceptable, it has a higher cost and throughput impact in wireless links. Even though, Mobile IP has been deployed for 3G data services, there has not been much usage of multicast or broadcast data transfer to or from the mobile node. Services like PTT (Push-To-Talk) rely on multicast. Other services such as IPTV also use multicast to distribute streaming video to mobile nodes. Hence it is essential to ensure that the mobile IPv4 clients support multicast and broadcast packet delivery in an optimal manner.

Current mobile IPv4 specifications [RFC3344] and [RFC3024] do not clearly address multicast/broadcast packet delivery for a MN with FA

care-of-address. for example, for encapsulating delivery style, the source address of the outer and inner IP header is the home address of the mobile node as described in section 5.2.2 of [RFC3024]. In addition, section 5.4 talks about local delivery of multicast/broadcast packets in the visited network but some corner cases are not completely specified. In particular, multicast messages from the mobile node to the visited network may be needed for retrieving service information. A mobile node may use all-mobility-agent multicast as the destination address and its home-address as the source-address for local service discovery. In this case, the foreign agents must consider all messages with the all-mobility-agent multicast as the destination address as special case and reply back directly to the mobile-node. However, this scenario makes foreign agent processing a bit more complex when reverse-tunnel is setup and the mobile-node sends multicast messages towards the reverse tunnel using its home-address as the source address. The all-mobility-agents multicast address is used for router solicitation by the mobile node, so foreign agent implementations must use it as a special address. This leads to complexity if in the reverse tunnel the mobile node uses its home address as the source address for other multicast messages destined to the home and visited network.

Currently different organizations [3GPP2] define their own mechanism to obtain local information such as DNS server IP address through AAA. All Mobility-agent multicast is used for router solicitation by the mobile node and the implementation can treat this address specially at the foreign agent. However, the implementation of foreign agent needs to apply multicast-address filtering and gets very complex if the mobile client uses the home address as source address for other multicast messages destined to the home and visited network, in the reverse tunnel mode. Even if multicast packets are delivered locally, the return packet which has the destination address as the home address will be routed back all the way to the home agent of the mobile node to be tunneled back to the foreign agent and then to the mobile node. [RFC3024] recommends selective reverse tunneling by delivering packets directly to the foreign agent, while encapsulating them for reverse tunnel delivery. But the specification is not clear about the source addresses of the packets from the mobile node in case of selective direct delivery. Although it clearly states that for the mobile node which uses co-located care-of address mode.

This specification aims to clarify the delivery of multicast messages when reverse tunneling is used, adds the capability to selectively negotiates which type of traffic to be delivered using encapsulating delivery, e.g., only for multicast and broadcast packets from mobile node to foreign agent, while allowing direct delivery for other type of traffic, e.g., unicast, and explores direct delivery options of

multicast messages between the mobile node and the foreign agent by using link-layer capabilities.

Section 3 describes the new delivery extension for multicast-broadcast packets in reverse tunnel mode.

2. Conventions & Terminology

2.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.2. Terminology

All the general mobility related terminology and abbreviations are to be interpreted as defined in IP Mobility Protocol [RFC3344] and Reverse tunneling for Mobile IP [RFC3024]. The following terms are used in this document.

MN

Mobile Node.

FA

Foreign Agent.

FA-CoA

Foreign Agent as the Mobile Node Care-of Address.

3. Multicast-Broadcast Encapsulating Delivery Style

The Mobile IP reverse tunneling [RFC3024] defines the Encapsulating delivery style for delivering multicast and broadcast packets from the mobile node to the foreign agent in the FA-CoA mode. It also mandates Encapsulating delivery mode for sending multicast/broadcast packets to reverse-tunnel to home agent via the foreign agent. But [RFC3024] section 2 says that all reverse-tunneled traffic is encapsulated when Encapsulating Delivery is negotiated. The "Multicast-Broadcast Encapsulating Delivery Style" (MBEDS) extension defined in this specification applies encapsulation only to the reverse-tunneled multicast and broadcast packets, leaving direct delivery for reverse-tunneled unicast packets. The main motivation

for adding this extension is to save the overhead of additional IP header for unicast packets which consequently will enable the use of Multicast and Broadcast packets when Mobile IPv4 is in use. This procedure works for both shared media like ethernet, IEEE 802.11 and links of a point-to-point nature such as those defined by 3GPP, 3GPP2 and IEEE 802.16.

3.1. Multicast-Broadcast Encapsulating Delivery Extension

The proposed extension is used in Mobile IPv4 signaling to negotiate the Multicast-Broadcast Encapsulation Delivery Style. Foreign agents SHOULD support the Multicast-Broadcast Encapsulating Delivery Style Extension. A registration request MAY include either a regular encapsulating delivery extension (see section 3.3 in [RFC3024]) or a Multicast-Broadcast Encapsulating Delivery extension, but not both. If both extensions are present, the foreign agent will consider that an error scenario and the FA MUST reject the registration request by sending a registration reply with the code field set to "Poorly Formed Request".

If a foreign agent supports MBEDS, then the foreign agent SHOULD advertise the MBEDS extension in its router advertisement to inform the mobile node about the type of delivery style it supports. This will avoid the possibility of multiple registration requests to figure out which encapsulating mode the foreign agent supports.

If the MN includes an MBEDS extension, it MUST do so after the Mobile-Home Authentication Extension, and before the Mobile-Foreign Authentication Extension, if present. The Encapsulating Delivery Style Extension MUST NOT be included if the 'T' bit is not set in the Registration Request.

If no delivery style extension is present, Direct Delivery per RFC 3024 is assumed.

The Multicast-Broadcast Encapsulation Extension format is as in Figure 1 below.




Figure 1: Multicast-Broadcast Encapsulating Extension

Type

<IANA>

Length

8-bit unsigned integer indicating the length in octets of the Bit-Field . It is set to 2.

Bit-Field Value

A 16-bit bit-field. Value specifies what type of packets are encapsulated. The following bits are defined (0 being the right-most bit, 15 the left-most bit):

0:

All packets are encapsulated between a mobile node and a foreign agent. It is same as the Encapsulating Delivery Style in RFC3024. NOTE: obsolete EDS in 3024?.

1:

Only multicast and broadcast packets are encapsulated (MBEDS).

2:

Link-layer Assisted Delivery Style (LLAS) for local network.

All other bits values are reserved.

NOTE: Only MBEDS packets are reverse tunneled after being decapsulated at the foreign agent, not those directly destined to the foreign-agent address or all mobility agent address. These are processed locally by the foreign agent.

3.2. Packet Header Formats for Visited Network Traffic

Other than Mobile IP agent solicitation packets, there might be some multicast or broadcast packets meant for consumption at the visited network. If the mobile node can acquire a local IP address, then it MUST direct deliver the multicast and broadcast traffic for local use. If the mobile node can have only one IP address, (i.e. home address) then it MUST send all the multicast and broadcast packets encapsulated. These packets will be sent to the home network through the reverse tunnel after being decapsulated at the foreign agent;

only exceptions are the multicast solicitation messages for the mobility agent.

In some cases, the mobile node may want to send multicast or broadcast packets to visited network entities other than the foreign agent. In those cases they should always be direct delivered by acquiring a local IP address or using link-layer mechanism if possible. Please see the section 'Link-layer Assisted Delivery Style' below for details.

3.3. Packet Header Formats for Homebound Traffic

The packet format and processing for encapsulated multicast and broadcast traffic is the same as defined in section 5.2 of Reverse Tunneling for Mobile IP [RFC3024]. Additionally, the packet format and processing for unicast traffic is the same as defined in section 5.1 of the same specification.

4. Multicast-Broadcast Encapsulating delivery Style Vs RFC3024 Encapsulating delivery

RFC3024 encapsulating delivery style does not require the foreign-agent to advertise an extension as well for the mobile node efficiency. MBEDS provides an option for foreign agent to advertise the extension with supported extension types, so that a mobile node can request a delivery style that the foreign agent supports.

RFC3024 encapsulating delivery style requires all multicast, broadcast and unicast traffic to be encapsulated in order to be reverse tunneled. In MBEDS unicast packets are always direct delivered to the foreign agent. Most of the the cases a node sends unicast packets for communication with a correspondent node and occasionally it may send broadcast or multicast packets to the home network. Thus this new style of delivery relieves the overhead of encapsulation for most traffic.

MBEDS introduces TLV style extension for delivery style. Therefore, this extension can be used to negotiate different delivery styles in the future. Currently, it can be backward compatible with RFC3024 encapsulating delivery style when the value field is zero. NOTE: We should make this a bit field to allow for easier advertisement and other extensions.

A mobile node SHOULD use either RFC3024 style encapsulating delivery extension or the MBEDS extension (defined in this document), but not both at the same time. If both extensions are received at the foreign-agent, the foreign agent MUST reject the registration request

by sending a registration reply with error (70) "Poorly Formed Request".

5. Link-layer Assisted Delivery Style (LLADS)

This section discusses direct-delivery of multicast and broadcast packets between the mobile node and the foreign agent by taking advantage of link-layer mechanisms. Certain link-layers allow for direct delivery from the MN to the FA (and vice-versa) without the need for encapsulation. In effect, this is assumed by RFC 3024 for Direct Delivery Style. In this mode, a unicast packet at the IP layer is carried over a unicast link-layer delivery mechanism. For example, the FA's MAC address is the link-layer destination address, or the packet is sent on a link of a point-to-point nature as in 3G networks. Broadcast and multicast packets, however are typically sent using a link-layer broadcast or multicast mechanism: a broadcast or multicast MAC address for IEEE 802.11 networks. If, however, these packets had the FA unicast MAC address while carrying an IP layer broadcast or multicast destination, then there would be no need for encapsulation to remove the ambiguity. The packet would be unequivocally directed at, and consumed by the FA. Notice that in links of a point-to-point nature, there is no ambiguity even for multicast and broadcast packets: these are unequivocally delivered to the FA. The Link-layer Assisted Delivery Style allows for direct delivery of unicast, multicast and broadcast packets over link-layers that can support it. In particular, it requires that regardless of whether the IP layer packet is unicast, broadcast or multicast, (1) when sending from MN to FA, the FA unicast address always be used, and (2) when sending from FA to MN, the MN unicast address always be used. The FA advertises such capability per the extension defined above, and the MN requests it in its registration request.

The LLADS imposes the least amount of tunneling overhead of the delivery styles as it effectively uses the equivalent of direct delivery for unicast, broadcast and multicast. It enables the MN to deliver packets to the FA for the foreign agent to reverse tunnel them back to the MN's home network.

However LLADS does not by itself allow the MN to deliver packets such that the FA know whether or not it should reverse tunnel them, or process them as local packets (e.g., perhaps forwarding them to local services). Certain networks have the capability of enabling additional context at the link-layer to effect different classification and treatment of packets otherwise indistinguishable at the IP layer, e.g., by establishing additional PDP contexts in 3GPP or additional service flows (and the corresponding CIDs) in WiMAX networks. In such networks, it is possible for the MN and the

FA to establish additional context such that packets sent by the MN to the FA are classified correctly upon arrival into either packets meant for local consumption, or packets meant to be reverse tunneled. In the absence of any IP layer differentiation (i.e., by sending packets meant for local consumption with the MN's local care-of address as source address), such link-layer mechanisms can provide the necessary means for the FA to select the correct processing for packets received from the MN. Such link-layer mechanisms, however, are out of scope of this document.

6. Security Considerations

This draft does not introduce any security threats on the top of what is defined in IP Mobility Protocol [RFC3344]. If included, the Multicast-Broadcast Encapsulating Delivery Style extension MUST be added after the MN-HA authentication extension and before the MN-FA authentication extension, if present.

7. IANA Considerations

This document defines a new IP Mobility extension, as described in Section 3.1 and uses a type <IANA-TBD>. The Multicast-Broadcast Encapsulation Delivery Extension type is assigned from the range of values associated with the skippable IP Mobility extensions.

8. Acknowledgments

The authors like to thank Charlie Perkins, Alex Bachmutsky, De Juan Huarte Federico, Parviz Yegani, Jayshree Bharatia for their comments and contribution in shaping up this document. We also thank the WiMAX Forum NWG members for their valuable input and suggestions during the initial discussion of the problem.

9. References

9.1. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3024] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344,

August 2002.

9.2. Informative references

- [3GPP2] "3GPP2 - Third Generation Partnership Project 2: X.P0028-200", Online web site <http://www.3gpp2.org>.
- [NWG] "NWG - WiMAX Network Architecture Group", Online web site <http://www.wimaxforum.org>.

Appendix A. Appendix-A

TBD.

Authors' Addresses

Ahmad Muhanna (Editor)
Ericsson Inc.
2201 Lakeside Blvd.
Richardson, TX 75082
USA

Email: ahmad.muhanna@ericsson.com

Basavaraj Patil
Nokia
6021 Connection Drive
Irving, TX 75039
USA

Email: basavaraj.patil@nokia.com

Samita Chakrabarti
IP Infusion
1188 Arquest Street
Sunnyvale, CA
USA

Email: samitac@ipinfusion.com

Gabriel Montenegro
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
USA

Email: gabriel.montenegro@microsoft.com

Yingzhe Wu
ZTE USA
10105 Pacific Heights Blvd, Suite 250
San Diego, CA 92121
USA

Email: yingzhe.wu@zteusa.com

Mobility for IPv4 Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 19, 2011

S. Gundavelli, Ed.
K. Leung
Cisco
G. Tsirtsis
Qualcomm
H. Soliman
Elevate Technologies
A. Petrescu
CEA LIST
August 18, 2010

Flow Binding Support for Mobile IPv4
draft-ietf-mip4-multiple-tunnel-support-00.txt

Abstract

This specification defines extensions to Mobile IPv4 protocol for allowing a mobile node with multiple interfaces to register a care-of address for each of its network interfaces and to simultaneously establish multiple Mobile IP tunnels with its home agent. This essentially allows the mobile node to utilize all the available network interfaces and build an higher aggregated data pipe with the home agent for its home address traffic. Furthermore, these extensions also allow the mobile node and the home agent to negotiate flow policies for binding individual traffic flows with the registered care-of addresses.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on February 19, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions & Terminology	3
2.1. Conventions	3
2.2. Terminology	3
3. Overview	4
4. Message Extensions	5
4.1. Alternate-CoA Extension	5
4.2. Flow Identification Extension	7
5. Protocol Operation	16
5.1. Mobile Node Considerations	17
5.1.1. Using the Alternate-CoA extension	17
5.1.2. Using the Flow Identification Extension	18
5.2. Home Agent Considerations	19
5.2.1. Handling Alternate-CoA extensions	19
5.2.2. Handling Flow Identification Extensions	20
6. Routing Considerations	23
7. Protocol Configuration Variables	23
8. IANA Considerations	24
9. Security Considerations	24
10. Contributors	24
11. Acknowledgements	25
12. References	25
12.1. Normative References	25
12.2. Informative References	25

1. Introduction

With the ubiquitous availability of wireless networks supporting different access technologies, mobile devices are now equipped with multiple wireless interfaces and have the ability to connect to the network over any of those interfaces and access the network. It is desirable for the mobile node to leverage all the available network connections for accessing network services.

The operation defined in the Mobile IP Protocol [RFC3344], allows a mobile node to continue to use its home address as it moves around the internet. Based on the mode of operation, there will be a tunnel that will be set up between the home agent and the mobile node, or between the home agent and the foreign agent where the mobile node is attached. In both of these modes, there will only be one interface on the mobile node that is receiving the traffic from the home agent. However, this is not efficient and requires an approach where the mobile node can use more than one interfaces for reaching the home network. The objective being efficient use of all available links to obtain higher aggregated bandwidth for the tunneled traffic between the home agent and the mobile node.

This specification defines extensions to Mobile IPv4 protocol for allowing a mobile node with multiple interfaces to register a care-of address for each of its network interfaces and to simultaneously establish multiple Mobile IP tunnels with its home agent. Furthermore, this specification also defines extensions to allow the mobile node and the home agent to optionally negotiate flow policies for binding individual traffic flows with the registered care-of addresses.

2. Conventions & Terminology

2.1. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2.2. Terminology

All the mobility related terms used in this document are to be interpreted as defined in [RFC3344] and [RFC3753]. In addition this document uses the following terms.

Binding Identifier (BID)

It is an identifier for a specific binding of a mobile node. A mobile node, when it registers multiple bindings with its home agent using different care-of addresses, each of those bindings are given a unique identifier and this identifier is called the binding identifier. The identifier is unique within all the bindings for a given mobile node.

Flow Identifier (FID)

It is an identifier for a given IP flow, uniquely identified by source address, destination address, protocol type, source port and destination port.

3. Overview

This document presents extensions to the Mobile IP protocol for allowing a mobile node to register multiple care-of addresses over which it can be reachable. Each of the registered care-of address will be identified by a unique binding identifier (BID). There will be multiple tunnels between the mobile node and the home agent, one tunnel for each of the registered bindings. These multiple tunnel paths can be used for load balancing the mobile node's home address traffic based on the negotiated traffic policies. The extensions specified in this document additionally allow the mobile node and the home agent negotiate flow policies for binding individual traffic flows to the registered care-of addresses. In the absence of any negotiated traffic policies, these multiple tunnel paths appear to the home agent and the mobile node as alternate routing paths and the default IP forwarding behavior of per-flow load balancing will leverage all the available wireless links and will result in a larger aggregated egress traffic throughput.

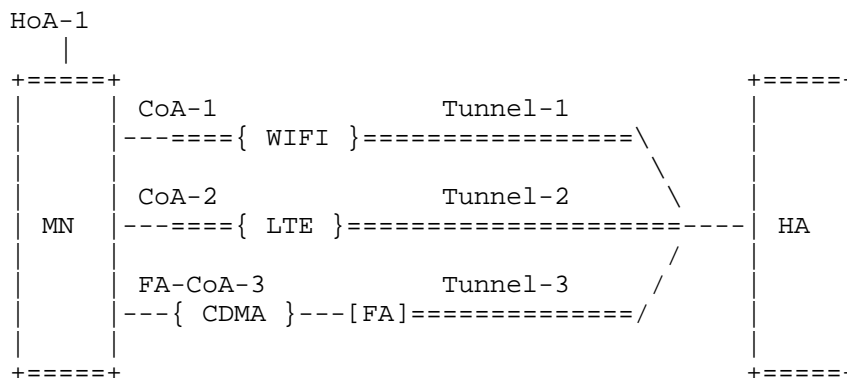


Figure 1: Mobile Node with multiple tunnels to the home agent

Figure 1, illustrates a mobile node attached to the network over three different access technologies, WiFi, LTE and CDMA. The mobile node is assigned home address, HoA-1, has care-of addresses CoA-1, CoA-2 and CoA-3 and has established tunnels Tunnel-1, Tunnel-2 and Tunnel-3 with its home agent.

Flow Id	CoA/Tunnel/BID	Negotiated Flow Policy
1.	CoA-1/Tunnel-1/BID-1	All SIP Flows over WiFi
2.	CoA-2/Tunnel-2/BID-2	All HTTP Flows over LTE value
3.	CoA-3/Tunnel-3/BID-3	All SSH Flows over CDMA

Table 1: Flow Binding Table

The above table is an example of how the individual flows are bound to different care-of addresses registered with the home agent.

4. Message Extensions

This specification defines the following new extensions.

4.1. Alternate-CoA Extension

A new skippable extension to the Mobile IPv4 header in accordance to the short extension format of [RFC3344] is defined here. This extension is for requesting the home agent to register the care-of address present in this extension as one of the alternate care-addresses through which the mobile node can be reached.

This extension MAY be added to the Registration Request only by the mobile node. This extension MUST NOT be added by the home agent or by the foreign agent either to the Registration Request or to the Registration Reply. There can be more than one instance of this extension present in the message.

This extension should be protected by Mobile Home Authentication extension [RFC3344]. As per Section 3.2 and 3.6.1.3 of [RFC3344], the mobile node MUST place this Extension before the Mobile-Home Authentication Extension in the registration messages, so that this extension is integrity protected.

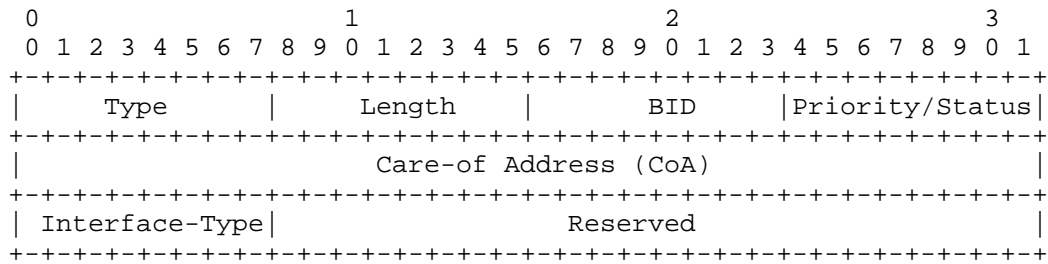


Figure 2: Alternate-CoA Extension

Type

Alternate-CoA Extension (skippable type range to be assigned by IANA)

Length

Indicates the length (in bytes) of the extension. The length does NOT include the Type and Length bytes.

BID (Binding ID)

The BID field in an 8-bit unsigned integer that identifies the binding to the CoA included in this extension and it can be used to point to an Alternate-CoA that was registered earlier.

Priority/Status

When this extension is in a Registration Request this field specifies the priority field assigned to the care-of address. The Priority field is an 8-bit unsigned integer. The receiver can utilize this priority to determine the preference of the CoA used to deliver packets. The lower the value the higher priority. A value of 255 indicates that the CoA indicated should be deregistered.

When this extension is in a Registration Reply this field indicates the status of the CoA. The Status field is an 8-bit unsigned integer. The possible status codes are listed in Table 2.

For the Status field values 0-127 indicate success and values between 128 and 255 indicate failure. The following values are defined for the Status field:

Status	Value	Comments
Accepted	0	The CoA is registered
BID Changed	1	The BID associated with an existing CoA was changed to the new value
Reject	128	The CoA is rejected
Unknown BID	129	The BID was not recognized

Table 2: Values for the Alternate-CoA Status field

Care-Of Address (CoA)

The CoA field is an 32-bit ipaddr. Set to an alternative care-of address to the one included in the Registration Request header. This field may not be included if the extension is included in a Registration Request and if the BID field is set to the BID of CoA registered earlier. In addition this field may not be included if the extension is included in a Registration Reply message.

Interface Type

Type of interface through which the mobile node is connected. The permitted values for this are from the Access Technology Type registry defined in [RFC5213].

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

4.2. Flow Identification Extension

A new skippable extension to the Mobile IPv4 header in accordance to the short extension format of [RFC3344] is defined here. This extension is included in the Registration Request and Registration Reply messages. This extension contains information that allows the home agent to identify a traffic flow and route it to a given address. There can be more than one instance of this extension present in the message.

This extension should be protected by Mobile Home Authentication extension [RFC3344]. As per Section 3.2 and 3.6.1.3 of [RFC3344], the mobile node MUST place this Extension before the Mobile-Home Authentication Extension in the registration messages, so that this extension is integrity protected.

A Flow Identification extension is designed to populate and edit a mobile node classifier in the home agent. A classifier selects packets based on the content of packet headers according to defined rules. The Flow Identification extension defines a line in such a classifier.

The Flow Identification extension has a flexible format that allows different fields to appear in the extension based on the way the mobile node chooses to represent the flow. The flags following the length field indicate which of the fields used to identify the flow are present in the extension. As a result, there is no fixed format for the flow identification extension. This may result in slight complexity in the implementation; however, this extension will minimize the total length of the extension sent, which is particularly important for bandwidth-limited wireless links.




Figure 3: Flow Identification Extension

Type

Flow Identification Extension (skippable type range. Two values to be assigned for IPv4 and IPv6 by IANA)

Length

Indicates the length (in bytes) of the extension. The length does NOT include the Type and Length bytes.

FID

The Flow Identifier field is an 8-bit unsigned integer identifying a flow. This field is used to refer to an existing flow or to identify a new flow.

Priority/Status

The Priority field is an 8-bit unsigned integer. When this extension is in a Registration Request this field specifies the priority field assigned to the filter rule defined by this extension. The receiver can utilize this priority to determine the order of application of the filter rules defined by the sender. The lower the value the higher priority (i.e., it is checked earlier against each packet). A value of 255 indicates that the filter rule indicated should be deregistered.

The Status field is an 8-bit unsigned integer. When this extension is in a registration reply this field indicates the status of the filter rule. The possible status codes are listed in Table 3.

For the Status field values 0-127 indicate success and values between 128 and 255 indicate failure. The following values are defined for the Status field:

Status	Value	Comments
Accepted	0	Flow binding successful
Reject	128	Flow binding rejected, reason unspecified.
Poorly Formed	129	Flow Identification extension poorly formed
Admin Prohibited	130	Administratively prohibited
Unknown FID	131	The FID is not recognized
Unknown BID	132	A BID included in the extension is not registered.

Table 3: Values for the Flow Identification Status field

Action

When this extension is in a Registration Request this field specifies the action that needs to be taken by the receiver. The field SHOULD be set to zero by the home agent in the registration reply and SHOULD be ignored by the mobile node. See defined values in Table 4.

The following values are reserved for the Action field.

Action	Value	Comments
Drop	0	Drop matching packets. A filter rule indicating a drop action MUST include a single BID byte, the value of which MAY be set to 255 by the sender and the value of which SHOULD be ignored by the receiver.
Forward	1	Forward matching packets to the 1st BID in the list of BIDs the filter rule is pointing to. If the 1st BID becomes invalid (i.e., the corresponding CoA is deregistered) use the next BID in the list.
X-Cast	2	Forward one copy of each matching packet to the list of BIDs this filter rule is pointing to.

Table 4: Values for the IPv4 and IPv6 Flow Descriptor Action field

F-Type

The Filter Type (F-Type) field identifies the type of Filter Descriptor included in the extension. Filter Descriptors in addition to the ones defined in this document can be defined in other documents but all Filter Descriptors MUST indicate their own length.

The following values are defined.

F-Type	Value	Comments
Do not Change	0	The already registered filter for the FID of the extension must be used
IPv4 Filter	1	An IPv4 Filter Descriptor follows, see Figure 4
IPv6 Filter	2	An IPv6 Filter Descriptor follows, see Figure 5

Table 5

Filter Descriptor

The Filter Descriptor field defines a filter. This field is further defined in Figure 4 and in Figure 5 depending on the value of the F-Type field of this extension.

BIDs

Indicates the BIDs to which the Filter Rule Descriptor points to, in order of appearance. Note that if a filter rule does not point to any valid BIDs, the filter rule itself becomes invalid.

BID	Value	Comments
Do not Change	0	The already registered filter for the FID of the extension must be used
BID	1-254	These values point to one of BIDs registered with Alternate-CoA extension, in order of appearance. Multiple BID bytes can be included to point to more than one BIDs
Default List	255	the default set of BIDs, registered with Alternate-CoA extensions MUST be used

Table 6

If the Type field of the Flow Identification extension indicates an IPv4 Flow then the Filter Rule Descriptor is as specified below. This fields in the message are identical to the format specified in Section 3.1 of [I-D.ietf-mext-binary-ts]. Please refer to that document for parameter description.

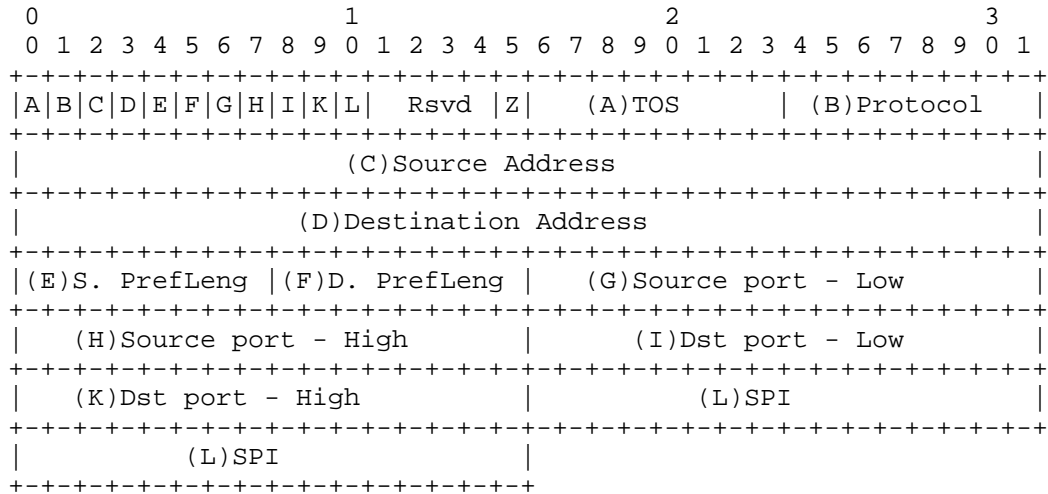


Figure 4: IPv4 Filter Rule Descriptor

Flags (A-L)

Each flag indicates whether the corresponding field is present in the message

(A)TOS - Type of Service

The TOS field in the data packet as seen by the home agent.

(B)Protocol

An 8-bit unsigned integer representing the value of the transport protocol number associated with the port numbers in data packets.

(C)Source Address

This field identifies the source address of data packets as seen by the home agent that is, the 32-bit IPv4 address of the correspondent node.

(D)Destination Address

This field identifies the destination address of data packets as seen by the home agent. When included this field must be set to one of the registered home addresses of the mobile node. It is a 32-bit IPv4 address.

(E)Source Prefix Length

This field includes the prefix length for the source address. This field can only be included if the Source Address field is included.

(F)Destination Prefix Length

This field includes the prefix length for the destination address. If The Destination Address field is included then it refers to that field; otherwise it refers to the home address field of the Registration Request header.

(G)Source Port - Low

This field identifies the lowest source port number within a range of port numbers that will be used in data packets, as seen by the home agent.

(H)Source Port - High

This field identifies the highest source port number within a range of port numbers that will be used in data packets, as seen by the home agent. If a single port is indicated then this field SHOULD NOT be included. If it is included it SHOULD be set to the value of the Source Port ? Low field.

(I)Destination Port - Low

This field identifies the lowest destination port number within a range of port numbers that will be used in data packets as seen by the home agent.

(K)Destination Port - High

This field identifies the highest destination port number within a range of port numbers that will be used in data packets as seen by the home agent. If a single port is indicated then this field SHOULD NOT be included. If it is included it SHOULD be set to the value of the Dst Port ? Low field.

(L)SPI - Security Parameter Index

The SPI field in the data packet as seen by the home agent.

If the Type field of the Flow Identification extension indicates an IPv6 Flow then the Filter Rule Descriptor is as as specified below. The fields in the message are identical to the format specified in Section 3.2 of [I-D.ietf-mext-binary-ts]. [I-D.ietf-mext-binary-ts]. The descriptor format is presented below for convenience. specified in [I-D.ietf-mext-binary-ts]. Its presented below for convenience.




Figure 5: IPv6 Filter Rule Descriptor

Flags (A-M)

Each flag indicates whether the corresponding field is present in the message

CS - Class of Service

The CS field in the data packet as seen by the home agent.

(B)Protocol

An 8-bit unsigned integer representing value of the transport protocol number associated with the port numbers in data packets.

(C)Source Address

This field identifies the source address of data packets as seen by the home agent. That is, the address of the correspondent node and it is a 128-bit IPv6 address.

(D)Destination Address

This field identifies the destination address of the data packet as seen by the home agent. When included this field must be set to one of the registered home addresses of the mobile node and it is a 128-bit IPv6 address.

(E)Source Prefix Length

This field includes the prefix for the source address. This field can only be included if the Source Address field is included .

(F)Destination Prefix Length

This field includes the prefix for the destination address. If The Destination Address field is included then it refers to that field otherwise it refers to the home address field of the registration request header.

(G)Source Port - Low

This field identifies the lowest source port number within a range of port numbers that will be used in data packets, as seen by the home agent.

(H)Source Port - High

This field identifies the highest source port number within a range of port numbers that will be used in data packets, as seen by the home agent. If a single port is indicated then this field SHOULD NOT be included. If it is included it SHOULD be set to the value of the Source Port ? Low field.

(I)Destination Port - Low

This field identifies the lowest destination port number within a range of port numbers that will be used in data packets as seen by the home agent.

(K)Destination Port - High

This field identifies the highest destination port number within a range of port numbers that will be used in data packets as seen by the home agent. If a single port is indicated then this field SHOULD NOT be included. If it is included it SHOULD be set to the value of the Dst Port ? Low field.

(L)SPI - Security Parameter Index

The SPI field in the data packet as seen by the home agent.

(M)Flow Label

The Flow Label field in the data packet as seen by the home agent.

5. Protocol Operation

This specification allows a mobile node to register multiple CoAs using the Alternate-CoA extension and associate different flows with different CoAs by using the Flow Identification extension.

When multiple CoAs are registered without any specific flow associated with them, the registered CoAs are treated as alternative paths to the mobile's current location. The CoAs are ranked by the Priority field in the Alternate-CoA extension and all traffic to the mobile's registered HoA(s) SHOULD be sent to the CoA with the lowest priority. If a CoA is deregistered, the CoA with the next lowest priority SHOULD become the default path for the mobile's traffic.

Note that, the HA MAY be configured with a local policy that takes advantage of multiple CoAs in a certain way. For example, x-casting across the registered CoAs MAY be used by the HA without any further signaling from the mobile; this is a configuration issue and outside the scope of this document.

When the Flow Identification extensions are also used, however, the mobile can indicate which flow is to be associated with which CoA. A single flow MAY be associated with more than one CoAs, while many flows MAY be associated with the same CoA. The effect of associating flows with CoA ofcourse depends on the action defined for that flow.

The Flow Identification extension is variable length and several fields might be omitted as required. When the extension is sent to deregister a filter rule (Priority set to 255) only the first line of Figure 3 needs to be sent (i.e., the first 4 bytes). If the priority and/or action values need to be changed for an existing FID then the F-Type MUST be set to 0 and one BID byte set to 0 MUST be included, indicating no changes to the filter and the BIDs associated with it. The Filter Descriptor of a given FID can be changed by sending the

extension including the new Filter Descriptor and a single BID byte set to 0. The BIDs associated with a given FID can be changed by sending the extension with F-Type set to 0 (and not including a Filter Descriptor). The F-Type (when not set to 0) indicates the type of Filter Descriptor used. In this specification we define Filter Descriptors for IPv4 and IPv6; other Filter Descriptors MAY be defined in separate documents.

5.1. Mobile Node Considerations

A mobile MAY send an Alternate-CoA extension with the CoA field matching the CoA field in the Mobile IP message header to check whether the HA supports the extensions defined in this specification. Since the extensions defined here are skippable, if the registration reply does not include the Alternate-CoA extensions sent by the mobile, the mobile knows that the HA does not support this specification. If, however, the HA returns the Alternate-CoA extensions in the reply, the HA does support this specification.

5.1.1. Using the Alternate-CoA extension

A mobile MAY include one or more Alternate-CoA extensions in each Registration Request message. If the mobile has already registered a CoA without using the Alternate-CoA extension and the mobile wants to register an additional CoA, the original and the new CoAs MUST be sent in the new registration as Alternate-CoA extensions so that they can be ranked with priorities and be associated with BIDs. In other words the new message will include an Alternate-CoA with the CoA field set to the CoA registered in the earlier message.

Unless multiple Alternate-CoA extensions are included in the same Registration Request message, the different CoAs will have different lifetimes associated with them. Each CoA MAY be refreshed individually by sending a Registration Request with that CoA in an Alternate-CoA extension. Alternatively, multiple CoAs can be refreshed at the same time by sending a Registration Request with multiple Alternate-CoA extensions.

If an earlier registered CoA is not included in a Registration Request it does not mean that the CoA is deregistered. Instead CoAs are deregistered when their lifetimes expire or when they are explicitly deregistered by the mobile node.

A mobile MAY deregister any CoA by setting its priority to 255. Note that the mobile can change the priority of a given CoA by sending an Alternate-CoA extension with the BID field set to the BID of the CoA in question, the priority field to the new value (or 255 for deregistration), and without including the CoA field.

A mobile MAY replace the CoA associated with a given BID by sending an Alternate-CoA with the BID field set to the BID of an existing CoA and the priority and CoA fields to their new values.

5.1.2. Using the Flow Identification Extension

The Flow Identification extensions allow a mobile to control a mobile specific classifier table present in the Home Agent memory. Each Flow Identification extension defines one filter rule line in that classifier, the output of which is one or more BIDs pointing to one or more of the registered CoAs.

Each filter rule in the classifier table can be referenced by the FID of the Flow Identification extension that created it. If the mobile wants to change the priority of a filter rule it can send a Flow Identification extension including the FID of the filter rule and setting the Priority field to the new value (or 255 for deregistration), and without including the Filter Rule Definition or any BIDs.

Filter rules do not need to be refreshed explicitly. A filter rule is valid as long as it points to a valid BID, i.e., a registered CoA. If a filter rule does not point to any valid BIDs it will be removed.

Any filter rule in the classifier table can be replaced by a new filter rule by sending a Flow Identification extension with the FID field set to the FID of the filter rule to be replaced and the rest of the extension defining the new filter rule, priority and the BIDs it points to.

Each Flow Identification extension is ranked according to its priority field. The lower the value of the priority field the higher its priority (i.e., it is checked earlier against each packet). As in most classifiers, filter rules with the same priority SHOULD be non-overlapping, otherwise the result is undefined. Overlapping filter rules SHOULD have different priorities.

Mobiles SHOULD define a default filter rule for traffic that does not match any other rule. The default filter rule MAY be defined with a Filter Identification extension with a high priority value (so it is checked last) and with the Filter Descriptor with all the flags set to 0 and the action field set to an appropriate value (e.g., forward). Note that such a Filter Descriptor will match all packets.

A mobile node can use the Flow Identification extension to associate a given flow with one or more of the registered CoAs. The mobile MUST register its CoAs with the Alternate-CoA extension in order to associate flows with them, using the BID as a handle. One or more

Flow Identification extensions and one or more Alternate-CoA extensions MAY be present in the same message.

If a Flow Identification extension includes a BID field set to the value 155 then the filter rule points to all the registered CoAs. The order of the CoAs for such a filter rule is dictated by the priority level of each BID, taken by the Priority field of the Alternate-CoA used to register them. If one or more BIDs are present in the Flow Identification extension then the filter rule points to the specific BIDs included in the extension. Note that the list of BIDs in the Flow Identification extension is ordered and its significance depends on the action indicated by the action field in the same extension.

5.2. Home Agent Considerations

5.2.1. Handling Alternate-CoA extensions

A Home Agent that supports this specification SHOULD ignore the "S" flag (Simultaneous Bindings) in the Registration Request message header when the same message includes Alternate-CoA extensions. In other words, the mechanisms defined in this specification override the mechanism defined by the "S" flag in [RFC3344].

If an Alternate-CoA extension is received by an HA in a Registration Request message, the HA SHOULD include a corresponding Alternate-CoA extension in the registration reply message. The BID of Alternate-CoA extension MUST be copied from the BID of the Alternate-CoA extension in the corresponding Registration Request and the Status field SHOULD be set to an appropriate value (e.g., indicating accept, reject etc).

When a valid Registration Request message includes one or more accepted Alternate-CoA extensions the HA MUST include the accepted CoAs in the mobility bindings table which binds the registered home address(es) with the registered CoAs together with their BIDs, priorities and lifetimes. The BID and priority of a CoA is indicated in the Alternate-CoA extension, while the lifetime is inherited from the lifetime of the registration reply message that accepted them as registered CoAs. Thus, different Alternate-CoAs will have different lifetimes if they are registered with different registration request messages, but they will have the same lifetime if they are included in the same Registration Request.

The CoAs are ranked according to their priority; the lowest the value of the priority field the higher their ranking. If an Alternate-CoA is rejected then the HA MUST NOT include it in the mobility bindings table. If the lifetime of an Alternate-CoA expires, the

corresponding CoA MUST be removed from the mobility bindings table.

If an Alternate-CoA extension is received with a BID that matches an existing BID then:

The HA MUST check the priority field of the extension in question. If the priority field is set to 255 (indicating deregistration) the CoA MUST be removed from the mobility bindings table and from any filter rules that point to it.

If the priority is set to any other value, the HA MUST check the CoA field of the same extension. If the CoA field is not included, the priority of the CoA, identified by the BID included in the extension, MUST be updated with the indicated value.

If the CoA field exists and matches the CoA that the BID field points to in the HA mobility bindings table, the priority of that CoA is again updated.

If the CoA field exists and is different from the CoA the BID field points to in the HA mobility bindings table, the HA SHOULD update its table with the new CoA and priority for that BID.

If an Alternate-CoA extension is received with a BID that does not match an existing BID then:

The HA MUST check the CoA field of the extension. If the CoA field is not included, the HA SHOULD include an Alternate-CoA extension in the registration reply with a BID copied from the corresponding extension in the request message and the Status field set to "Unknown BID."

If the CoA field exists, the HA MUST store the BID, CoA and priority values in the mobility bindings table for the mobile. The CoA MUST be ranked with the other registered CoAs according to the value of the priority field.

If the CoA field exists but it matches a CoA that is already registered with a different BID the HA MAY replace the old BID with the new BID and indicate a "BID changed" in the Status field of the corresponding Alternate-CoA extension included in the registration reply message.

5.2.2. Handling Flow Identification Extensions

If a Flow Identification extension is received by an HA in a Registration Request message, the HA SHOULD include a corresponding Flow Identification extension in the registration reply message. The

FID of the Flow Identification extension in the reply message MUST be copied from the FID of the Flow Identification extension in the corresponding Registration Request and the Status field SHOULD be set to an appropriate value (e.g., indicating accept, reject etc).

When a valid Registration Request message includes one or more accepted Filter Identification extensions the HA MUST include the accepted filter rules in the mobile specific classifier table which associates the order list of filter rules with the BIDs they point to. The priority of a filter rule, the description of the filter rule, the action and the BID(s) the filter rule is associated with are indicated in the Flow Identification extension.

The filter rules are ranked according to their priority. Filter rules MUST be ranked from lowest to higher priority. If a filter rule is rejected then it MUST NOT be included in the mobile specific classifier.

Each filter rule in the mobile specific classifier is valid as long as it points to a valid BID, i.e., a registered CoA. If a filter rule does not point to any valid BIDs the HA MUST remove it from the mobile specific classifier.

If the HA receives a Flow Identification extension, it SHOULD first check the FID field of that extension.

If the value of the FID field does not match any of the FIDs in the mobile specific classifier, the HA SHOULD include the filter rule described in the extension in the mobile specific classifier table. The new filter rule MUST be ranked according to the priority field indicated in the Flow Identification extension.

If one or more BIDs are included then the filter rule MUST point to the list of BIDs in the order they appear.

If any of the including BIDs do not match one of the registered BIDs in the mobile bindings table for this mobile the HA MUST disregard the Flow Identification extension and MUST return a reply message with a Flow Identification extension that includes the FID of the corresponding extension in the request message and the Status field set to an appropriate value e.g., "Unknown BID."

If a BID of value 255 is included, the filter rule MUST point to the default list of BIDs. This is the list of BIDs in the mobility bindings table for this mobile.

If a BID of value 0 is included the HA MUST disregard the Flow Identification extension and MUST return a reply message with a Flow Identification extension that includes the FID of the corresponding extension in the request message and the Status field set to an appropriate value e.g., "Unknown BID."

If the value of the FID field matches any of the FIDs in the mobile specific classifier the HA SHOULD then check the priority field of the Flow Identification extension. If the priority field is set to 255 then the filter rule associated with the FID in the Flow Identification extensions MUST be removed from the mobile specific classifier table.

If the priority field, however, is set to a value other than 255 the HA SHOULD check the Filter Description field of the Flow Identification extension.

If the Filter Description is not included (F-Type field set to 0) and the BID field is set to 0, the HA MUST adjust the ranking of the filter rule corresponding to the FID according to the priority field in the Flow Identification extension.

If any BIDs are also included in the Flow Identification extensions then the list of BIDs associated with that filter rule MUST also be replaced by the list provided in the Flow Identification extension. If a BID field set to 255 is included then the filter rules MUST be re-pointed to the default list of BIDs registered with Alternate-CoA extensions.

Note a BID field set to 0 is included the BIDs list for this filter rule in the mobility specific classifier table MUST NOT be changed.

If the priority field, however, is set to a value other than 255 and the Filter Description field is included then the HA MUST replace the corresponding filter rule in the mobile specific classifier table with the filter rule in the Flow Identification extension.

If any BIDs are also included in the Flow Identification extensions then the list of BIDs associated with that filter rule MUST also be replaced by the list provided in the Flow Identification extension. If a BID field set to 255 is included then the filter rules MUST be re-pointed to the

default list of BIDs registered with Alternate-CoA extensions.

Note that if a BID field set to 0 is included the BIDs field, the list of BIDs this filter rule points to MUST NOT be changed from its previous configuration.

6. Routing Considerations

This document allows the mobility entities to optionally exchange flow policies. In the absence of negotiated traffic flow policies, this document recommends the use of per-flow load balancing.

Most IP devices support the two alternative traffic load-balancing schemes, Per-flow and Per-packet load balancing. These load balancing schemes allow the forwarding device to evenly distribute traffic based on the criteria of per-packet or on a per-flow basis, across all the available equal-cost links through which a destination can be reached. The default forwarding behavior of Per-flow load balancing will ensure a given flow always takes the same path and will eliminate any packet re-ordering issues and that is critical for delay sensitive traffic. Whereas the per-destination load balancing scheme leverages all the paths much more affectively, but with the potential issue of packet re-ordering on the receiver end. A host can choose to enable any of these approaches.

7. Protocol Configuration Variables

The following protocol configuration variables are required for system management and these variables MUST be configurable on all the mobility entities. The configured values for these protocol variables MUST survive service restarts.

`EnableMultipleTunnelSupport.`

This flag indicates whether or not the mobility entity on which this protocol variable is configured needs to enable Multiple Tunnel support feature. This protocol variable is applicable to home agent, foreign agent and the mobile node.

The default value for this flag is set to value of 1, indicating that the multiple tunnel support SHOULD be enabled.

When the value for this flag is set to value of 0, multiple tunnel support SHOULD be disabled.

8. IANA Considerations

This document proposes two new extensions that require a type number to be assigned by IANA.

Section 4.1 defines a new Mobile IP extension, the Alternate-CoA Extension. Its a skippable extension to the Mobile IPv4 header in accordance to the short extension format of [RFC3344]. The type number for this extension needs to be assigned by IANA.

Section 4.2 defines a new Mobile IP extension, the Flow Identification Extension. Its a skippable extension to the Mobile IPv4 header in accordance to the short extension format of [RFC3344]. The type number for this extension needs to be assigned by IANA.

9. Security Considerations

This specification allows a mobile node to establish multiple tunnels with the home agent, by registering a care-of address for each of its active roaming interfaces. This essentially allows the mobile node's home address to be reachable through all of its active and registered roaming interfaces. This specification also allows the mobile node to bind traffic flows to the registered care-of addresses. This new capability has no impact on the protocol security.

The Mobile IP message extensions, defined in this document are to be carried in Mobile IP Registration messages and these messages are authenticated and integrity protected as described in [RFC3344].

Therefore, this specification does not weaken the security of Mobile IP Protocol, or, introduce any new vulnerabilities.

10. Contributors

This document reflects discussions and contributions from the following people:

Srinivasa Kanduru

skanduru@gmail.com

Vince Park

vpark@qualcomm.com

11. Acknowledgements

We like to thank Qin Wu, Shahriar Rahman, Mohana Jeyatharan, Yungui Wang, Hui Deng Behcet Sarikaya, Jouni Korhonen, Michaela Vanderveen and Antti Makela for their review and comments on this draft.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [RFC3024] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3519] Levkowetz, H. and S. Vaarala, "Mobile IP Traversal of Network Address Translation (NAT) Devices", RFC 3519, April 2003.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

12.2. Informative References

- [I-D.ietf-mext-binary-ts] Tsirtsis, G., Giaretta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", draft-ietf-mext-binary-ts-04 (work in progress), February 2010.
- [I-D.ietf-mext-flow-binding] Soliman, H., Tsirtsis, G., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and NEMO Basic Support", draft-ietf-mext-flow-binding-06 (work

in progress), March 2010.

[RFC3753]

Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.

[RFC5177]

Leung, K., Dommety, G., Narayanan, V., and A. Petrescu, "Network Mobility (NEMO) Extensions for Mobile IPv4", RFC 5177, April 2008.

[RFC5648]

Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", RFC 5648, October 2009.

Authors' Addresses

Sri Gundavelli (editor)
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: sgundave@cisco.com

Kent Leung
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: kleung@cisco.com

George Tsirtsis
Qualcomm

EMail: tsirtsis@qualcomm.com

Hesham Soliman
Elevate Technologies

EMail: hesham@elevatemobile.com

Alexandru Petrescu
CEA LIST
Communicating Systems Laboratory, Point Courrier 94
Gif-sur-Yvette F-91191
France

Phone: +33 169089223
EMail: alexandru.petrescu@cea.fr

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 21, 2011

A. Makela
Aalto University
J. Korhonen
Nokia Siemens Networks
October 18, 2010

Home Agent assisted Route Optimization between Mobile IPv4 Networks
draft-ietf-mip4-nemo-haaro-02

Abstract

This document describes a Home Agent assisted Route Optimization functionality to IPv4 Network Mobility Protocol. The function is designed to facilitate optimal routing in cases where all nodes are connected to a single Home Agent, thus the use case is Route Optimization within single organization or similar entity. The functionality adds possibility to discover eligible peer nodes based on information received from Home Agent, Network Prefixes they represent, and how to establish direct tunnel between such nodes.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on April 21, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction and motivations	4
2. Terms and definitions	6
3. Mobile IPv4 route optimization between mobile networks	7
3.1. Maintaining route optimization information	8
3.1.1. Advertising route-optimizable prefixes	8
3.1.2. Route Optimization cache	10
3.2. Return routability procedure	12
3.2.1. Router keys	13
3.2.2. Nonces	13
3.2.3. Updating Router keys and Nonces	13
3.3. Mobile-Correspondent Router operations	15
3.3.1. Triggering Route Optimization	16
3.3.2. Mobile Router routing tables	16
3.3.3. Inter-Mobile Router registration	16
3.3.4. Inter-Mobile Router tunnels	19
3.3.5. Constructing route-optimized packets	20
3.3.6. Handovers and Mobile Routers leaving network	20
3.4. Convergence and synchronization issues	21
4. Data compression schemes	22
4.1. Prefix compression	22
4.2. Realm compression	24
4.2.1. Encoding of compressed realms	24
4.2.2. Searching algorithm	25
4.2.3. Encoding example	26
5. New Mobile IPv4 messages and extensions	28
5.1. Mobile router Route optimization capability	28
5.2. Route optimization reply	29
5.3. Mobile-Correspondent authentication extension	30
5.4. Care-of address Extension	31
5.5. Route optimization prefix advertisement	31
5.6. Home-Test Init message	33
5.7. Care-of-Test Init message	33
5.8. Home Test message	34
5.9. Care-of test message	35
6. Special Considerations	35
6.1. NATs and stateful firewalls	35
6.2. Handling of concurrent handovers	37
6.3. Foreign Agents	37
6.4. Multiple Home Agents	37
6.5. Mutualness of Route Optimization	38

6.6. Extensibility	39
6.7. Load Balancing	39
7. Scalability	40
8. Example signaling scenarios	40
8.1. Registration request	40
8.2. Route optimization with return routability	41
8.3. Handovers	43
9. Protocol constants	45
10. IANA Considerations	45
11. Security Considerations	47
11.1. Return Routability	47
11.2. Trust relationships	47
12. Acknowledgements	48
13. References	48
13.1. Normative References	48
13.2. Informative References	48
Authors' Addresses	49

1. Introduction and motivations

Traditionally, there has been no method for route optimization in Mobile IPv4 [RFC3344] apart from an early attempt [I-D.ietf-mobileip-optim]. Unlike Mobile IPv6 [RFC3775], where Route Optimization has been included from the start, with Mobile IPv4 route optimization hasn't been addressed in a generalized scope.


Even though general route optimization may not be of interest in the scope of IPv4, there are still specific applications for Route Optimization in Mobile IPv4. This document proposes method to optimize routes between networks behind mobile routers, as defined by NEMO [RFC5177]. Although NAT and pending shortage of IPv4 addresses makes widespread deployment not feasible, using Route Optimization only in routers is still a practical scenario.

A particular use case concerns setting up redundant yet economical enterprise networks. Recently, a trend has emerged where customers prefer to maintain connectivity via multiple service providers. Reasons include redundancy, reliability and availability issues. These kinds of multi-homing scenarios have traditionally been solved by using such technologies as multihoming BGP. However, a more lightweight and economical solution is desirable.

From service provider perspective a common topology for enterprise customer network consists of one to several sites (typically headquarters and various branch offices). These sites are typically connected via various Layer 2 technologies (ATM or Frame relay PVCs), MPLS VPNs or Layer 3 site-to-site VPNs. With a Service Level Agreement, a customer can obtain a very reliable and well supported intranet connectivity. However, compared to the cost of "consumer-grade" broadband Internet access the SLA-guaranteed version can be considered very expensive. These consumer-grade options however, are not reliable approach for mission-critical applications.

Mobile IP, especially mobile routers, can be used to improve reliability of connectivity even when implemented over consumer-grade Internet access. The customer becomes a client for a virtual service provider, which does not take part in the actual access technology. The service provider has a backend system and an IP address pool that it distributes to customers. Access is provided by multiple, independent, possibly consumer-grade ISPs, with Mobile IP providing seamless handovers if service from a specific ISP fails. The drawback of this solution is that it creates a star topology; All Mobile IP tunnels end up at the service provider hosted home agent, causing heavy load at the backend. Route Optimization between mobile networks addresses this issue, by taking network load off the home agent and the backend.

An example network is pictured below:



Virtual service provider architecture using NEMOv4

In this example case, organization network consists of two sites, that are connected via 2 ISPs for redundancy reasons. Mobile IP allows fast handovers without problems of multi-homing and BGP peering between each individual ISP and the organization. The traffic however takes a non-optimal route through the virtual operator backend.

Route optimization addresses this issue, allowing traffic between Sites A and B to flow through ISP B's network, or in case of a link failure, via the ISP peering point (such as MAE-WEST). The backend will not suffer from heavy loads.

The primary design goal is to limit the load to the backend to minimum. Additional design goals include extensibility to a more generalized scope, beyond the need of a single, coordinating Home Agent.

2. Terms and definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Care-of Address (CoA)

RFC 3344 [RFC3344] defines Care-of Address as the termination point of a tunnel toward a mobile node, for datagrams forwarded to the mobile node while it is away from home. The protocol can use two different types of care-of address: a "foreign agent care-of address" which is an address of a foreign agent with which the mobile node is registered, and a "co-located care-of address", which is an externally obtained local address which the mobile node has associated with one of its own network interfaces. However, in the case of Network Mobility, foreign agents are not used, so no foreign care-of addresses are used either.

Correspondent Router (CR)

RFC 3344 [RFC3344] defines a Correspondent node as a peer with which a mobile node is communicating. Correspondent Router is a peer Mobile Router which MAY also represent one or more entire networks.

Home Address (HoA)

RFC 3344 [RFC3344] defines Home Address as an IP address that is assigned for an extended period of time to a mobile node. It remains unchanged regardless of where the node is attached to the Internet.

Home Agent (HA)

RFC 3344 [RFC3344] defines Home Agent as a router on a mobile node's home network which tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information for the mobile node. For this application, the "home network" sees limited usage.

Host Network Prefix

Network Prefix with the mask of /32. e.g. 192.0.2.254/32, consisting of a single host.

Mobility Binding

RFC 3344 [RFC3344] defines Mobility Binding as the association of Home Address with a Care-of address, along with the lifetime remaining for that association.

Mobile Network Prefix RFC 5177 [RFC5177] defines Mobile Network Prefix as the network prefix of the subnet delegated to a Mobile Router as the Mobile Network.

Mobile Router (MR)

Mobile Router as defined by RFC 5177 [RFC5177] and RFC 3344 [RFC3344]. They define a Mobile Router as a mobile node that can be a router that is responsible for the mobility of one or more entire networks moving together, perhaps on an airplane, a ship, a train, an automobile, a bicycle, or a kayak.

Route Optimization Cache

Data structure maintained by Mobile Routers on possible destinations for Route Optimization. Contains information (Home Addresses) on potential Correspondent Routers and their associated Mobile Networks.

Return Routability, RR

Procedure to bind a Mobile Router's Home Address to a Care-of address on a Correspondent Router with a degree of trust.

3. Mobile IPv4 route optimization between mobile networks

This section describes the changed functionality of Home Agent and Mobile Router compared to the base NEMOv4 operation defined in [RFC5177]. The basic premise is still the same; Mobile Routers, when registering to the Home Agent, either inform the Home Agent of the mobile network prefixes they are managing (explicit mode) or get prefixes assigned by Home Agent (implicit mode). However, instead of prefix <-> Mobile Router mapping information only remaining on the Home Agent and the single Mobile Router, this information will now be distributed to the other Mobile Routers as well.

The Home Agent-assisted Route Optimization is primarily intended for helping to optimize traffic patterns between multiple sites in an single organization or administrative domain; However, extranets can also be reached with optimized routes, as long as all Mobile Routers connect to the same Home Agent. The procedure aim to maintain backwards compatibility; With legacy nodes or routers full connectivity is always preserved even though optimal routing cannot be guaranteed.

The schema requires a Mobile Router to be able to receive messages from Home Agent and other Mobile Routers unsolicited - that is, without first initiating a request. This behavior is similar to the registration revocation procedure [RFC3543]. Many of the mechanisms are same - including the fact that advertising route optimization support upon registration implies capability to receive registration requests and return routability messages from other Mobile Routers.

Compared to IPv6, where Mobile Node <-> Correspondent node bindings are maintained via Mobility Routing header and Home Address options, Mobile IPv4 always requires the use of tunnels. Therefore, inter-mobile-router tunnel establishment has to be conducted.

3.1. Maintaining route optimization information

During registration, a joining Mobile Router MAY request information on route-optimizable network prefixes. The Mobile Router MAY also allow redistribution of information on its managed network prefixes regardless whether they are explicit or implicit (statically configured or assigned by Home Agent). These are indicated with Mobile Router Route Optimization capability extension, see Section 5.1. If the Home Agent accepts the request for Route Optimization, this is indicated with Route Optimization Reply extension (Section 5.2) in the registration reply.

Note that the redistribution of network prefix information from the Home Agent happens only during the registration signaling. There are no "routing updates" from Home Agent except during re-registrations triggered by handovers, registration timeouts and specific solicitation. The solicitation re-registration MAY occur if a Correspondent Router receives a registration request from a unknown Mobile Router (see Section 3.3.3).

3.1.1. Advertising route-optimizable prefixes

As noted, NEMO-supporting Home Agent already maintains, and in some cases assigns, information on which network prefixes are reachable behind specific Mobile Routers. Only change to this functionality is that this information can now be distributed to other Mobile Routers

upon request. This request is implied by including Route Optimization capability extension, Section 5.1, and setting the 'R' bit.

When a Home Agent receives a registration request, standard authentication and authorization procedures are conducted.

If registration is successful and the Route Optimization capability information extension was present in the registration request, the reply message **MUST** include Route Optimization Reply extension (Section 5.2) to indicate whether Route Optimization was accepted. Furthermore, the extension also informs Mobile Router if NAT was detected between Home Agent and the Mobile Router using the procedure in RFC 3519 [RFC3519], which is based on the discrepancy between requester's indicated Care-of address and packet's source address.

The reply message **MAY** also include one route optimization prefix advertisement extension which informs the Mobile Router of existing mobile network prefixes and the Mobile Routers that manage them, if eligible for redistribution. The networks **SHOULD** be included in order of priority, with the prefixes determined by policy as most desirable targets for Route Optimization listed first. The extension is constructed as shown in Section 5.5. The extension consists of a list where each Mobile Router, identified by Home Address, is listed with according prefix(es) and their respective realm(s).

Each network prefix can be associated with a realm, usually in the form 'organization.example.com'. Besides the routers in customer's own organization, the prefix list may also include other Mobile Routers, e.g. Default prefix (0.0.0.0/0) pointing towards Internet gateway for Internet connectivity, and possible extranets. The realm information can be used to make policy decisions on the Mobile Router, such as preferring optimization within specific realm only.

In a typical scenario where Network Prefixes are allocated to Mobile Routers connecting to a single Home Agent, the prefixes are usually either continuous or at least very close to each other. Due to these characteristics, an optional prefix compression mechanism is provided. Another, optional, compression scheme is in use for realm information, where realms often share same higher-level domains. These compression mechanisms are further explained in Section 4.

Upon receiving registration reply with a Route Optimization prefix advertisement extension, the Mobile Router **SHALL** insert the Mobile Router Home Addresses included in the extension as a host-prefixes to the local Route Optimization Cache if they do not already exist. If present, any additional prefixes information **SHALL** also be inserted to the Route Optimization Cache.

The Mobile Router MAY discard entries from a desired starting point onwards, due to memory or other policy related constraints. The intention of listing the prefixes in order of priority is to provide implicit guidance for this decision. If the capacity of the device allows, the Mobile Router SHOULD use information on all advertised prefixes.

3.1.2. Route Optimization cache

Mobile routers supporting route optimization will maintain a Route Optimization Cache.

The Route Optimization Cache contains mappings between potential Correspondent Router HoA's, network(s) associated with each HoA, network topology, and Return Routability procedure-related information. The Cache is populated based on information received from Home Agent in Route optimization prefix advertisements, and in registration messages from Correspondent Routers. Portions of the cache may also be configured statically.

The Route Optimization Cache contains the following information for all known Correspondent Routers. Note that some fields may contain multiple entries. For example, during handovers, there may be both old and new CoA's listed.

CR-HoA

Correspondent Router's Home Address. Primary key identifying each CR.

CR-CoAs

Correspondent Router's Care-of Address(es). May be empty if none known. Potential tunnel's destination address(es).

MR-CoAs

Mobile Router's Care-of Address(es) used with this Correspondent Router. Tunnel's source address.

Tunnels

Tunnel interface(s) associated with this Correspondent Router. The tunnel interface itself handles all the necessary operations to keep the tunnel operational, e.g. Sending keepalive messages required by UDP encapsulation.

NAT states

A table of booleans, set for all pairs of potential MR-CoA's and CR-CoA's which require NAT awareness and the behavior is known, populated either statically or based on discovery. If set to true, the MR can establish a UDP tunnel towards the CR, using this pair of CoA's. A received advertisement can indicate this to be set initially false for all respective CR's CoA's. Affects tunnel establishment direction, see Section 3.3.4 and the registration procedure in deciding which care-of-address to include in Care-of-address extension in registration reply. If the entry exists, mandates use of UDP encapsulation.

RRSTATes

Return routability state for each CR-HoA and MR-CoA pair. States are INACTIVE, IN PROGRESS and ACTIVE. If state is INACTIVE, return routability procedure must be completed before forwarding route-optimized traffic. If state is IN PROGRESS or ACTIVE, the information concerning this Correspondent Router MUST NOT be removed from Route Optimization Cache as long as tunnel to the Correspondent Router is established.

KRms

Registration management key for each CR-HoA - MR-CoA pair. This field is only used if configured statically - if the KRm was computed using Return Routability procedure, they are calculated in-situ based on nonces and router key. If configured statically, RRSTATE is permanently set to ACTIVE.

Care-of nonce indexes If the KRm was established with Return Routability procedure, contains the Care-of nonce index for each MR-CoA - CR-HoA pair.

Care-of keygen token If the KRm was established with Return Routability procedure, contains the Care-of keygen token for each MR-CoA - CR-HoA pair.

Home nonce index If the KRm was established with Return Routability procedure, contains the Home nonce index for each CR-HoA

Home keygen token If the KRm was established with Return Routability procedure, contains the Home keygen token for each CR-HoA.

Network Prefixes

A list of destination network prefixes reachable via this Correspondent Router. Includes network and prefix length, e.g. 192.0.2.0/25. Always contains at least a single entry, the CR-HoA host network prefix in the form of 192.0.2.1/32.

Realms

Each prefix may be associated with a realm. May also be empty, if realm is not provided by advertisement or configuration.

Prefix_Valid

Boolean field for each prefix - CR-HoA pair, which is set to true if this prefix's owner has been confirmed. The Host Network Prefix consisting of the Correspondent Router itself does need validation beyond Return Routability procedure. For other prefixes, the confirmation is done by soliciting the information from HA. Traffic for prefixes which have unconfirmed ownership should not be routed through the tunnel.

Information that is no longer valid due to expirations or topology changes MAY be removed from the Route Optimization Cache as desired by the Mobile Router.

3.2. Return routability procedure

The purpose of return routability procedure is to establish Care-of-Address <-> Home Address bindings in a trusted manner. The return routability procedure for Mobile IPv6 is described in [RFC3775]. Same principles apply to the Mobile IPv4 version: Two messages are sent to Correspondent Router's Home Address, one via Home Agent using Mobile Router's Home Address, and the other directly from the Mobile Router CoA, with two responses coming through same routes. Registration management key is derived from token information carried on these messages. This registration management key (KRm) can then be used to authenticate registration requests (comparable to Binding Updates in Mobile IPv6).

The Return Routability procedure is a method provided by Mobile IP protocol to establish the KRm in a relatively lightweight fashion.

If desired, the KRms can be configured to Mobile Routers statically, or using an desired external secure key provisioning mechanism. If KRm's are known to the Mobile Routers via some other mechanism, Return Routability procedure can be skipped. Such provisioning mechanisms are out of scope for this document.

Assumption on traffic patterns is that the Mobile Router that initiates the RR procedure can always send outbound messages, even when behind NAT or firewall. This basic assumption made for NAT Traversal in [RFC3519] is also applicable here. In case the Correspondent Router is behind such obstacles, it receives these messages via the reverse tunnel to CR's Home Address, thus any problem regarding the CR's connectivity is addressed during the registration to the Home Agent.

3.2.1. Router keys

Each Mobile Router maintains a 'correspondent router key', Kcr, which is MUST NOT be shared with any other entity. Kcr is used for authenticating peer Mobile Routers in the situation where mobile router is acting as a CR. This is analogous to node key, Kcn, in Mobile IPv6. Correspondent Router uses router key to verify that the keygen tokens sent by Mobile Router in registration request are its own. The router key MUST be a random number, 16 octets in length.

The Mobile Router MAY generate a new key at any time to avoid persistent key storage. If desired, it is RECOMMENDED to expire the keys in conjunction with nonces, see Section 3.2.3.

3.2.2. Nonces

Each Mobile Router also maintains one or more indexed nonces. Nonces should be generated periodically with a good random number generator. The Mobile Router may use same nonces with all Mobile Routers. Nonces may be of any length, with the RECOMMENDED length being 64 bits.

3.2.3. Updating Router keys and Nonces

The router keys and nonce updating guidelines are similar to ones in Mobile IPv6. Mobile Routers keep both the current nonce and small set of valid previous nonces whose lifetime have not expired yet. Nonce should be kept acceptable for at least MAX_TOKEN_LIFETIME (see Section 9) seconds after it has first been used in constructing a return routability response. However, the correspondent router MUST NOT accept nonces beyond MAX_NONCE_LIFETIME seconds (see Section 9) 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 keygen tokens that have been based on the last 8 $(\text{MAX_NONCE_LIFETIME} / 30)$ nonces. This results in keygen tokens being acceptable $\text{MAX_TOKEN_LIFETIME}$ to $\text{MAX_NONCE_LIFETIME}$ 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.

If Kcr is being updated, the update SHOULD be done at the same time as nonce is updated. This way, nonce indexes can be used to refer to both Kcr's and nonces.

The Return Routability procedure consists of four Mobile IP messages: Home Test Init, Care-of Test Init, Home Test and Care-of Test. They are constructed as shown in Section 5.6 through Section 5.9. If the Mobile Router has included the Mobile Router optimization capability extension in its Registration Request, it MUST be able to accept Return Routability messages. The messages are delivered as Mobile IP signaling packets. The destination address is set to Correspondent Router's HoA.

The return routability procedure begins with the Mobile Router sending HoTI and CoTI messages, each containing a 64-bit random value, the cookie. The cookie is used to bind specific signaling exchange together.

Upon receiving the HoTI or CoTI message the Correspondent Router MUST have a secret Kcr and nonce. If it does not have this material yet, it MUST produce it before continuing with the return routability procedure.

Correspondent Router responds to HoTI and CoTI messages by constructing HoT and CoT messages, respectively, as replies. The HoT message contains home init cookie, current home nonce index and home keygen token. The CoT message contains care-of init cookie, current care-of nonce index and care-of keygen token.

Return Routability procedure SHOULD be initiated only when the Route Optimization Cache's RRSTATE field for the desired Care-of Address with target Correspondent Router is INACTIVE. When Return Routability procedure is initiated, the state MUST be set to IN PROGRESS. In case of handover occurring, the Mobile Router SHOULD only send a CoTI message to obtain a new care-of keygen token; The home keygen token may still be valid. If the reply to a registration

indicates that one or both of the tokens has expired, the RRSTATE MUST be set to INACTIVE. The Return Routability procedure may then be restarted as needed.

Upon completion of Return Routability procedure, the Routing Optimization Cache's RRSTATE field is set to ACTIVE, allowing for registration requests to be sent. The Mobile Router will establish a registration management key K_{Rm} by default using MD5 hash algorithm:

$$K_{Rm} = \text{MD5}(\text{home keygen token} \mid \text{care-of keygen token})$$

When de-registering (by setting time to zero), care-of keygen token is not used. Instead the Registration management key is generated as follows:

$$K_{Rm} = \text{MD5}(\text{home keygen token})$$

Like in Mobile IPv6, the Correspondent Router does not maintain any state for the Mobile Router until after receiving a registration request.

3.3. Mobile-Correspondent Router operations

This section deals with the operation of Mobile and Correspondent Routers performing route optimization. Note that in the context of this document all routers work as both Mobile Router and Correspondent Router. The term "Mobile Router" applies to the router initiating the Route Optimization procedure, and "Correspondent Router" indicates the peer router.

Especially compared to Mobile IPv6 route optimization there are two issues that are different regarding IPv4. First of all, since Mobile IPv4 always uses tunnels, there must be a tunnel established between MR and CR's Care-of addresses. The Correspondent Router learns of Mobile Router's Care-of address as it is provided by the Registration Request. The Mobile Router learns Correspondent Router's Care-of address by a new extension, "Care-of Address", in registration reply. Second issue is rising from security standpoint: In a registration request, the Mobile Router claims to represent an arbitrary IPv4 network. If the CR has not yet received this information (HoA <-> Network prefix), it SHOULD perform a re-registration to Home Agent to verify the claim.

Additional aspect is that Mobile Router MAY use different Care-of-Address for different Correspondent Routers (and Home Agent). This is useful in situations where network provides only partial-mesh connectivity, and specific interfaces must be used to reach specific destinations. In addition, this allows for load balancing.

3.3.1. Triggering Route Optimization

Since each Mobile Router knows the eligible route-optimizable networks, the route optimization between all Correspondent Routers can be established at any time; However a better general practice is to conduct Route Optimization on-demand only. It is RECOMMENDED to start Route optimization only be started when receiving a packet where destination address is in a locally managed prefix (and the prefix is registered as route optimizable) and source address exists in the network prefixes of Route Optimization Cache. With small number of Mobile Routers, such on-demand behavior may not be necessary and full-mesh route-optimization may be in place constantly.

3.3.2. Mobile Router routing tables

Each Mobile Router maintains a routing table. In a typical situation, the Mobile Router has one or more interface(s) to the local networks, one or more interface(s) to wide-area networks (such as provided by ISPs), and a tunnel interface to the Home Agent. Additional tunnel interfaces become activated as Route Optimization is being performed.

The routing table SHOULD typically contain Network Prefixes managed by Correspondent Routers associated with established route-optimized tunnel interfaces. In addition, host-routes to Correspondent Routers' Care-of addresses SHOULD be associated with the assigned to the physical interfaces assigned with corresponding MR-CoA. If the tunneling method does not require such host-routes, these can be omitted. A default route MAY point to the reverse tunnel to the Home Agent if not overridden by prefix information.

The route for the Home Address of Correspondent Router SHOULD also be pointing towards the optimized tunnel.

If two prefixes overlap each other, e.g. 192.0.2.128/25 and 192.0.2.128/29, the standard longest match rule for routing is in effect. However, overlapping private address SHOULD be considered an error situation. Any aggregation for routes in private address space SHOULD be conducted only at HA.

3.3.3. Inter-Mobile Router registration

If route optimization between Mobile Router and Correspondent Router is desired, either Return Routability procedure must have been performed (See Section 3.2), or key K_{RM} must be pre-shared between the Mobile and Correspondent Router. If either condition applies, a Mobile Router MAY send a registration request to the Correspondent

Router's HoA from desired interface.

The registration request's source address and Care-of address field are set to the address of desired outgoing interface on the Mobile Router. The address MAY be same as the Care-of address used with Home Agent. The registration request MUST include Mobile-Correspondent Authentication extension defined in Section 5.3 and SHOULD include Mobile Network Request Extension defined in [RFC5177]. If present, the Mobile Network Request Extension MUST contain the network prefixes, as if registering in explicit mode. If timestamps are used, the Correspondent Router MUST check the identification field for validity. The registration request MUST include Home Address. The Authenticator field is hashed with the key K_{Rm}.

The Correspondent Router relies to the request with a Registration Reply. The registration reply MUST include Mobile-Correspondent Authentication extension defined in Section 5.3 and, if Mobile Network Request Extension was present in the request, a Mobile Network Acknowledgement extension.

The encapsulation can be set as desired, except in the case where the Route Optimization Cache Entry has NAT entries for the Correspondent Router, or the Mobile Router itself is known to be behind NAT or firewall. If either of the conditions apply, registration request MUST specify UDP encapsulation. It is RECOMMENDED to always use UDP encapsulation to facilitate detecting of path failures via keepalive mechanism.

The Correspondent Router first checks the registration request's authentication against K_{cr} and nonce indexes negotiated during Return Routability procedure. This ensures that the registration request is coming from a correct Mobile Router. If the check fails, an appropriate registration reply code is sent (see Section (Section 10). If the failure is due to nonce index expiring, the Mobile Router sets RRSTATE for the CR to INACTIVE. Return routability procedure MAY then be initiated again.

If the check passes, the Correspondent Router MUST check whether the Mobile Router already exists in it's Route Optimization Cache and is associated with the prefixes included in the request (Prefixes are present and Flag HA is true for each prefix).

If the check against the cache fails, the Correspondent Router SHOULD send a re-registration request to Home Agent with the 'S' (solicitation) bit set, thus obtaining the latest information on Network Prefixes managed by incoming Mobile Router. If, even after this update, the prefixes still don't match, the reply's Mobile Network Acknowledgement code MUST be set to "MOBNET_UNAUTHORIZED".

The registration can also be rejected completely. This verification is done to protect against Mobile Routers claiming to represent arbitrary networks; However, since Home Agent is assumed to provide trusted information, it can authorize Mobile Router's claim. If the environment itself is considered trusted, the Correspondent Router can, as a policy, accept registrations from without this check; however, this is NOT RECOMMENDED as a general practice.

If the prefixes match, the Correspondent Router MAY accept the registration. If the CR chooses to accept, the CR MUST check if a tunnel to the Mobile Router already exists. If the tunnel does NOT exist or has wrong endpoints (CoAs), a new tunnel MUST be established and Route Optimization Cache updated. The reply MUST include a list of eligible care-of-addresses for the tunnel in Section 5.4, with which the Mobile Router may establish a tunnel with. The reply MUST also include Mobile-Correspondent Authentication extensionSection 5.3.

Upon receiving the registration reply, the Mobile Router MUST check if a tunnel to the Correspondent Router already exists. If the tunnel does NOT exist, or has wrong endpoints (CoAs), a new tunnel MUST be established and Route Optimization Cache updated. This is covered in detail in Section 3.3.4.

The Correspondent Router's routing table MUST be updated to include the Mobile Router's networks are reachable via the direct tunnel to the Mobile Router.

After the tunnel is established, the Mobile Router MAY update it's routing tables to reach all Correspondent Router's Prefixes via the tunnel, although it is RECOMMENDED to wait for the Correspondent Router to perform it's own, explicit registration. This is primarily a policy decision depending on the network environment. See Section 6.5.

Due to the fact that the route optimization procedures may occur concurrently at two Mobile Routers, each working as each other's Correspondent Router, there may be a situation where two routers are attempting to establish separate tunnels between them at the same time. If a router with a smaller Home Address (meaning a normal 32-bit integer comparison treating IPv4 addresses as 32-bit unsigned integers) receives a registration request (in CR role) while its own registration request (sent in MR role) is still pending, the attempt should be rejected with reply code "concurrent registration". If receiving such an indication, the recipient should not attempt to re-register again until a grace period has passed without route optimization occurring.

3.3.4. Inter-Mobile Router tunnels

Inter-Mobile Router tunnel establishment follows establishing standard reverse tunnels to the Home Agent. The registration request to Correspondent Router includes information on the desired encapsulation. It is RECOMMENDED to use UDP encapsulation. In the cases of GRE [RFC2784], IP over IP [RFC2003] or minimal encapsulation [RFC2004] no special considerations regarding the reachability are necessary; The tunnel has no stateful information; The packets are simply encapsulated within the GRE, IP, or minimal header.

The tunnel origination point for the Correspondent Router is its Care-of Address, not the Home Address where the registration requests were sent. This is different from creation of the Reverse Tunnel to Home Agent, which reuses the channel from registration signaling.

Special considerations rise from using UDP encapsulation, especially in cases where one of the Mobile Routers is located behind NAT or firewall. A deviation from RFC 3519 [RFC3519] is that keepalives should be sent both from ends of the tunnel to detect path failures after the initial keepalive has been sent - this allows both MR and CR to detect path failures.

The initial UDP keepalive SHOULD be sent by the MR. Only after first keepalive is successfully completed, SHOULD the tunnel be considered eligible for traffic. If reply to the initial keepalive is not received, the MR may opt to attempt sending the keepalive with other Care-of addresses provided by the registration reply to check whether they provide better connectivity, or if all of these fail, perform a re-registration via alternative interface, or deregister completely. See Section 6.1. Once the initial keepalive packet has reached the CR and reply has been sent, the CR MAY start sending it's own keepalives.

The original specification for UDP encapsulation suggests a keepalive interval default of 110 seconds. However, to provide fast response time and switching to alternate paths, it is RECOMMENDED, if power and other constraints allow, to use considerably shorter periods, adapting to the perceived latency as needed. However, the maximum amount of keepalives should at no point exceed MAX_UPDATE_RATE times in second. The purpose of keepalive is not to keep NAT or firewall mappings in place, but serve as a mechanism to provide fast response in case of path failures.

If both the Mobile Router and the Correspondent Router are behind separate NATs, route optimization cannot be performed between them. Possibilities to set up mutual tunneling when both routers are behind NAT, are outside the scope of this document. However, some of these

issues are addressed in Section 6.1.

The designations "MR" and "CR" only apply to the initial tunnel-establishment phase. Once a tunnel is established between two routers, either of them can opt to either tear down the tunnel or perform a handover. Signaling messages have to be authenticated with valid Krm.

3.3.5. Constructing route-optimized packets

All packets received by the Mobile Router are forwarded using normal routing rules according to the routing table. There are no special considerations when constructing the packets, the tunnel interface's own processes will encapsulate any packet automatically.

3.3.6. Handovers and Mobile Routers leaving network

Handovers and connection breakdowns can be categorized as either ungraceful or graceful, also known as "break-before-make" (bbm) and "make-before-break" (mbb) situations.

As with establishment, the "Mobile Router" discussed here is the router wishing to change connectivity state, "Correspondent Router" being the peer.

When a Mobile Router wishes to leave network, it SHOULD, in addition to sending the registration request to the Home Agent with lifetime set to zero, also send such a request to all known Correspondent Routers. The Correspondent Router(s), upon accepting this request and sending the reply, will check if it's Route Optimization Cache contains any prefixes associated with the requesting Mobile Router. These entries should be removed and routing table updated accordingly (traffic for the prefixes will be forwarded via the Home Agent again). The tunnel MUST then be destroyed. A short grace period SHOULD be used to allow possible in-transit packets to be received correctly.

In the case of a handover, the Correspondent Router simply needs to update the tunnel's destination to the Mobile Router's new Care-of Address. Mobile Router SHOULD keep accepting packets from both old and new care-of Addresses for a short grace period, typically in the order of ten seconds. In the case of UDP encapsulation, the port numbers SHOULD be reused if possible.

If the Mobile Router was unable to send the re-registration request before handover, it MUST send it immediately after handover has been completed and tunnel with the Home Agent is established. Since Care-of Address(es) changing invalidates the Krm, at it is

RECOMMENDED to conduct partial Return Routability by sending CoTI message via the new Care-of-Address and obtaining new care-of keygen token. In all cases, necessary tokens have to also be acquired if the existing ones have expired.

If a reply is not received for a registration request to a Correspondent Router, any routes to the network prefixes managed by the Correspondent Router **MUST** be removed from the routing table, thus causing the user traffic to be forwarded via the Home Agent.

3.4. Convergence and synchronization issues

The information the Home Agent maintains on Mobile Network prefixes and the Mobile Routers' Route Optimization Caches do not need to be explicitly synchronized. This is based on the assumption is that at least some of the traffic between nodes inside mobile networks is always bidirectional. If using on-demand route optimization, this also implies that when a node in a mobile network talks to a node in another mobile network, if the initial packet does not trigger Route Optimization, the reply packet will.

Consider a situation with three mobile networks, A, B, C handled by three Mobile Routers, MR A, MR B and MR C respectively. If they register to a Home Agent in this order, the situation goes as follows:

MR A registers; Receives no information on other networks from HA, as no other MR has registered yet.

MR B registers; Receives information on mobile network A being reachable via MR A.

MR C registers; Receives information on both of the other mobile networks.

If a node in mobile network C receives traffic from mobile network A, the route optimization is straightforward; MR C already has network A in its Route Optimization Cache. Thus, packet reception triggers Route Optimization towards MR A. When MR C registers to MR A (after Return Routability procedure is completed), MR A does not have information on mobile network C; Thus it will perform a re-registration to the Home Agent on-demand. This allows MR A to verify that MR C is indeed managing network C.

If a node in mobile network B receives to traffic from mobile network C, MR B has no information on network C. No route optimization is triggered. However, when the node in network B replies and the reply reaches MR C, route optimization happens as above. Further examples

of signaling are in Section 8.

Even in the very rare case of completely unidirectional traffic from an entire network, the re-registrations to the Home Agent caused by timeouts will eventually cause convergence. However, this should be treated as a special case.

Note that all Mobile Routers are connected to same Home Agent. For possibilities concerning multiple Home Agents, see Section 6.4

4. Data compression schemes

This section defines the two compression formats used in Route Optimization Prefix Advertisement extensions.

4.1. Prefix compression

The prefix-compression is based on the idea that prefixes usually share common properties. The scheme is simple delta-compression. In the prefix information advertisement, Section 5.5, the D bit indicates whether receiving a "master" or a "delta" prefix. This, combined with the Prefix Length information, allows for compression and decompression of prefix information.

If D=0, what follows in the "Prefix" field are bits 1..n of the a new master prefix, where n is PLen. This is rounded up to nearest full octet. Thus, prefix lengths of /4 and /8 take 1 octet, /12 and /16 take 2 octets, /20 and /24 three, and larger than that full 4 octets.

If D=1, what follows in the "Prefix" field are bits m..PLen of the prefix, where m is the first changed bit of previous master prefix, with padding from master prefix filling the field to full octet. Maximum value of Plen-m is 8 (that is, delta MUST fit into one octet). If this is not possible, a new master prefix has to be declared.

Determining the order of prefix transmission should be based on saving maximum space during transmission.

Example of compression and transmitted data, where network prefixes 192.0.2.0/28, 192.0.2.64/26 and 192.0.2.128/25 are transmitted are illustrated in Figure 1. Because of the padding to full octets, redundant information is also sent. The bit-patterns being transmitted are:

+= shows the prefix mask
 --- shows the master prefix for delta coded prefixes
 192.0.2.0/28, D=0

```

      0              1              2              3
    1 2 3 4 5 6 7 8  9 0 1 2 3 4 5 6  7 8 9 0 1 2 3 4  5 6 7 8 9 0 1 2
+-----+-----+-----+-----+-----+-----+-----+-----+
|1|1|0|0|0|0|0|0|. |0|0|0|0|0|0|0|0|. |0|0|0|0|0|0|1|0|. |0|0|0|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
      ^                                         ^
+----- encoded -----+
                                         ^
                                         ^
                                         +-pad-+
  
```

192.0.2.64/26, D=1

```

      0              1              2              3
    1 2 3 4 5 6 7 8  9 0 1 2 3 4 5 6  7 8 9 0 1 2 3 4  5 6 7 8 9 0 1 2
+-----+-----+-----+-----+-----+-----+-----+-----+
|1|1|0|0|0|0|0|0|. |0|0|0|0|0|0|0|0|. |0|0|0|0|0|0|1|0|. |0|1|0|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
                                         ^         ^
+--- encoded ---+
      ^         ^
      +- padding -+
  
```

192.0.2.128/25, D=1

```

      0              1              2              3
    1 2 3 4 5 6 7 8  9 0 1 2 3 4 5 6  7 8 9 0 1 2 3 4  5 6 7 8 9 0 1 2
+-----+-----+-----+-----+-----+-----+-----+-----+
|1|1|0|0|0|0|0|0|. |0|0|0|0|0|0|0|0|. |0|0|0|0|0|0|1|0|. |1|0|0|0|0|0|0|0|
+-----+-----+-----+-----+-----+-----+-----+-----+
                                         ^         ^
+--- encoded ---+
      ^         ^
      +- padding -+
  
```

Figure 1: Prefix Compression Example

First prefix, 192.0.2.0/28, is considered a master prefix and is transmitted in full. The PLen of 28 bits determines that all four octets must be transmitted. If the prefix would have been e.g. 192.0.2.0/24, three octets would have sufficed since 24 bits fit into 3 octets.

For the following prefixes, the D=1. Thus, they are deltas of the previous prefix where D was zero.

192.0.2.64/26 includes bits 19-26 (full octet). Bits 19-25 are copied from master prefix, but bit 26 is changed to 1. The final notation in binary is "1001", or 0x09.

192.0.2.128/25 includes bits 18-25 (full octet). Bits 18-24 are copied from master prefix, but bit 25 is changed to 1. The final notation in binary is "101", or 0x05.

The final encoding thus becomes:

Prefix	Plen	D	Transmitted Prefix
192.0.2.0/28	28	0	0xc0 0x00 0x02 0x00
192.0.2.64/26	26	1	0x09
192.0.2.128/25	25	1	0x05

It should be noted that in this case the order of prefix transmission would not affect compression efficiency. If prefix 192.0.2.128/25 would have been considered the master prefix and the others as deltas instead, the resulting encoding still fits into one octet for the subsequent prefixes. There would be no need to declare a new master prefix.

4.2. Realm compression

4.2.1. Encoding of compressed realms

In order to reduce the size of messages, the system introduces a realm compression scheme, which reduces the size of realms in a message. The compression scheme is a simple dynamically updated dictionary based algorithm, which is designed to compress arbitrary length text strings. In this scheme, an entire realm, a single label or a list of labels may be replaced with an index to a previous occurrence of the same string stored in the dictionary. The realm compression defined in this specification was inspired by the RFC 1035 [RFC1035] DNS domain name label compression. Our algorithm is, however, improved to gain more compression.

When compressing realms, the dictionary is first reset and does not contain a single string. The realms are processed one by one so the algorithm does not expect to see them all or the whole message at once. The state of the compressor is the current content of the dictionary. The realms are compressed label by label or as a list of labels. The dictionary can hold maximum 128 strings. Thus, when adding the 129th string into the dictionary, the dictionary MUST first be reset to the initial state (i.e. Emptied) and the index of

the string will become 0.

The encoding of an index to the dictionary or an uncompressed run of octets representing a single label has purposely been made simple and the whole encoding works on an octet granularity. The encoding of an uncompressed label takes the form of a one octet:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+---+---+---+
|0|   LENGTH   | 'length' octets long string.. |
+---+---+---+---+---+---+---+---+---+---+

```

This encoding allows label lengths from 1 to 127 octets. A label length of zero (0) is not allowed. The "label length" tag octet is then followed by up to 127 octets of the actual encoded label string.

The index to the dictionary (the "label index" tag octet) takes the form of a one octet:

```

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

```

The above encodings do not allow generating an output octet value of zero (0). The encapsulating Mobile IPv4 extension makes use of this property and uses the value of zero (0) to mark the end of compressed realm or to indicate an empty realm. It is also possible to encode the complete realm using only "label length" tags. In this case no compression takes place. This allows the sender to skip compression, for example to reduce computation requirements when generating messages. However, the receiver MUST always be prepared to receive compressed realms.

4.2.2. Searching algorithm

When compressing the input realm, the dictionary is searched for a matching string. If no match could be found, the last label is removed from the right-hand side of the used input realm. The search is repeated until the whole input realm has been processed. If no match was found at all, then the first label of the original input realm is encoded using the "label length" tag and the label is inserted into the dictionary. The previously described search is repeated with the remaining part of the input realm, if any. If nothing remains, the realm encoding is complete.

When a matching string is found in the dictionary the matching part of the input realm is encoded using the "label index" tag. The matching part of the input realm is removed and the search is repeated with the remaining part of the input realm, if any. If nothing remains, the octet value of zero (0) is inserted to mark the end of encoded realm.

The search algorithm also maintains the "longest non-matching string" for each input realm. Each time the search in dictionary fails and a new label gets encoded using the "label length" tag and inserted into the dictionary, the "longest non-matching string" is concatenated by this label including the separating "." (dot, i.e. Hexadecimal 0x2e). When a match is found in the dictionary the "longest non-matching string" is reset (i.e. Emptied). Once the whole input realm has been processed and encoded, all possible suffixes longer than one label are taken from the string and inserted into the dictionary.

4.2.3. Encoding example

This section shows an example how to encode a set of realms using the specified realm compression algorithm. For example, a message might need to compress the realms "foo.example.com", "bar.foo.example.com", "buz.foo.example.org", "example.com" and "bar.example.com.org". The following example shows the processing of input realms on the left side and the contents of the dictionary on the right hand side. The example uses hexadecimal representation of numbers.

COMPRESSOR:	DICTIONARY:
1) Input "foo.example.com"	
Search("foo.example.com")	
Search("foo.example")	
Search("foo")	
Encode(0x03,'f','o','o')	0x00 "foo"
+--> "longest non-matching string" = "foo"	
Search("example.com")	
Search("example")	
Encode(0x07,'e','x','a','m','p','l','e')	0x01 "example"
+--> "longest non-matching string" = "foo.example"	
Search("com")	
Encode(0x03,'c','o','m')	0x02 "com"
+--> "longest non-matching string" = "foo.example.com"	
	0x03 "foo.example.com"
	0x04 "example.com"
Encode(0x00)	
2) Input "bar.foo.example.com"	
Search("bar.foo.example.com")	
Search("bar.foo.example")	

```
Search("bar.foo")
Search("bar")
Encode(0x03,'b','a','r')           0x05 "bar"
+--> "longest non-matching string" = "bar"
Search("foo.example.com") -> match to 0x03
Encode(0x83)
+--> "longest non-matching string" = NUL
Encode(0x00)
3) Input "buz.foo.example.org"
Search("buz.foo.example.org")
Search("buz.foo.example")
Search("buz.foo")
Search("buz")
Encode(0x03,'b','u','z')           0x06 "buz"
+--> "longest non-matching string" = "buz"
Search("foo.example.org")
Search("foo.example")
Search("foo") -> match to 0x00
Encode(0x80)
+--> "longest non-matching string" = NUL
Search("example.org")
Search("example") -> match to 0x01
Encode(0x81)
+--> "longest non-matching string" = NUL
Search("org")
Encode(0x03,'o','r','g')           0x07 "org"
+--> "longest non-matching string" = "org"
Encode(0x00)
4) Input "example.com"
Search("example.com") -> match to 0x04
Encode(0x84)
Encode(0x00)
5) Input "bar.example.com.org"
Search("bar.example.com.org")
Search("bar.example.com")
Search("bar.example")
Search("bar") -> match to 0x05
Encode(0x85)
Search("example.com.org")
Search("example.com") -> match to 0x04
Encode(0x84)
Search("org") -> match to 0x07
Encode(0x87)
Encode(0x00)
```

As can be seen from the example, due the greedy approach of encoding matches, the search algorithm and the dictionary update function is not the most optimal one. However, we do not claim the algorithm

would be the most efficient. It functions efficiently enough for most inputs. In this example, the original input realm data was 79 octets and the compressed output excluding the end mark is 35 octets.

5. New Mobile IPv4 messages and extensions

This section describes the construction of all new information elements.

5.1. Mobile router Route optimization capability

This skippable extension MAY be sent by a Mobile Router to a Home Agent in the registration request message.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      Sub-type      |A|R|S|O| Rsvd |
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               Optional Mobile Router HoA                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type TBA_T1. Skippable; If Home Agent does not support route optimization advertisements, it can ignore this request and simply not include any information in the reply. "Short" extension format.

Sub_Type TBA_ST1_1

Reserved Set to zero, MUST be ignored on reception.

A Advertise my networks. If 'A' bit is set, the Home Agent is allowed to advertise the networks managed by this Mobile Router to other Mobile Routers. This also indicates that the Mobile Router is capable of receiving route optimization binding updates. In effect, this allows the Mobile Router to work in Correspondent Router role.

R Request mobile network information. If 'R' bit is set, the Home Agent MAY respond with information about mobile networks in the same domain.

S Soliciting prefixes managed by specific Mobile Router. The Mobile Router is specified in the Optional Mobile Router HoA field.

- O Explicitly specifying the requesting Router is only able to initiate outgoing connections, not accept any incoming ones, due to NAT device, stateful firewall, or similar issue on any interface. This is reflected by the Home Agent in the reply, and distributed in Prefix Advertisements to outer Mobile Routers.

Optional Mobile Router HoA

Solicited Mobile Router's Home Address.

5.2. Route optimization reply

This non-skippable extension MUST be sent by a Home Agent to a Mobile Router in the registration reply message, if Mobile Router indicated support for Route Optimization in registration message and Home Agent supports Route Optimization.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      Sub-Type      | O | N | S |      Code      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type TBA_T2 (Non-skippable), "short" extension format

Sub-Type TBA_ST2_1.

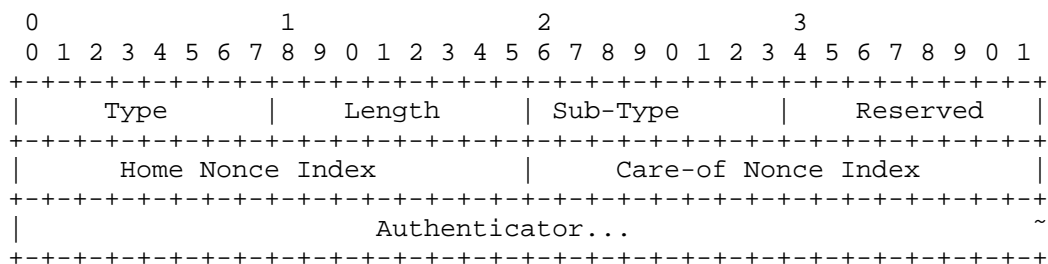
- O The 'O' flag in Mobile Router Optimization capability extension was set during registration.
- N Presence of NAT was detected by Home Agent. This informs the Mobile Router that it is located behind NAT. The detection procedure is specified in RFC 3519 [RFC3519], and is based on the discrepancy between registration packet's source address and indicated Care-of Address. The Mobile Router can use this information to make decisions about Route Optimization strategy.
- S Responding to a solicitation. If 'S' bit was present in Mobile router Route optimization capability extension (Section 5.1), this is set, otherwise unset.

The Reply code indicates whether Route Optimization has been accepted. Values of 0..15 indicate assent and values 16..63 indicate Route Optimization is not done.

- 0 Will do Route Optimization
- 16 Route Optimization declined, reason unspecified.

5.3. Mobile-Correspondent authentication extension

Mobile-Correspondent authentication extension is included in registration requests sent from Mobile Router to Correspondent Router. The existence of this extension indicates that the message is not destined to a Home Agent, but another Mobile Router. The format is similar to the other Authentication Extensions defined in [RFC3344], with SPIs replaced by Nonce Indexes.



The Home Nonce Index field tells the Correspondent Router which nonce value to use when producing the home keygen token. The Care-of Nonce Index field is ignored in requests to remove a binding. Otherwise, it tells the Correspondent Router which nonce value to use when producing the Care-of Keygen Token.

Type TBA_T2 (non-skippable). "Short" extension format.

Sub-Type TBA_ST_2_2

Reserved Set to zero, MUST be ignored on reception.

Home Nonce Index

Home Nonce Index in use.

Care-of Nonce Index

Care-of Index in use.

Authenticator

Authenticator field, by default constructed with HMAC_MD5 (KRm2, Protected Data)

The protected data, just like on other cases where Authenticator is used, consists of

- o the UDP payload (i.e., the Registration Request or Registration Reply data),
- o all prior Extensions in their entirety, and
- o the Type, Length, and Nonce Indexes of this Extension.

5.4. Care-of address Extension

The Care-of Address extension is added to a registration reply sent by the Correspondent Router to inform the Mobile Router of the upcoming tunnel endpoint.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      Sub-type      |      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
      1-n Times the following information structure
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Care-of Address                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type TBA_T2 (Non-skippable), "short" extension format

Sub-Type TBA_ST_2_3

Care-of Address

Care-of address(es) which may be used for tunnel with Mobile Router, in order of priority. Multiple CoA's MAY be listed to facilitate faster NAT traversal.

5.5. Route optimization prefix advertisement

This non-skippable extension MAY be sent by a Home Agent to a Mobile Router in the registration reply message. The extension is only included when explicitly requested by the Mobile Router in the registration request message. Implicit prioritization of prefixes is caused by the order of extensions.


```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   Sub-type   |   Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
1-n times the following information structure
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|D|M| PLen/Info | Optional Mobile Router HoA, 4 octets      ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~               | Optional Prefix, 1,2,3 or 4 octets        ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~               Realm (1-n characters)                      ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type TBA_T3 (Non-skippable), "long" extension format

Sub-Type TBA_ST3_1

D Delta. If D=1, the prefix is a delta from last Prefix where D=0. MUST be zero on first information structure, MAY be zero or one on subsequent information structures. If D=1, the Prefix field is one octet in length. See Section 4.1 for details.

M Mobile Router HoA bit. If M=1, the next field is Mobile Router HoA, and Prefix and Realm are omitted. If M=0, the next field is Prefix followed by Realm, and Mobile Router HoA is omitted. For the first information structure, M MUST be set to 1. If M=1, the D bit is set to zero and ignored upon reception.

PLen/Info This field is interpreted differently depending on whether M is set or not. If M=0, this indicates the length of the prefix advertised. 6 bits, allows for values from 0 to 63, of which 33-63 are illegal. If M=1, the Information field can be set to zero to indicate no specific information, or to 1 to indicate "outbound connections only". This indicates that the target Mobile Router can only initiate, not receive, connections on any of it's interfaces (apart from the reverse tunnel to HA). This is set if the Mobile Router has explicitly requested it by the 'O' flag in Mobile router Route optimization capability extension (Section 5.1).

Mobile router HoA


Mobile Router's Home address. All prefixes in the following information structures where M=0 are maintained

by this Mobile Router. This field is present only when M = 1.

Prefix The IPv4 prefix advertised. If D=0, the field length is Plen bits, rounded up to nearest full octet. Least-significant bits starting off Plen (and are zeros) are omitted. If D=1, field length is one octet. This field is present only when M = 0.

Realm The Realm that is associated with the advertised Mobile Router HoA and prefix. If empty, MUST be set to '\0'. For realm encoding and optional compression scheme, refer to Section 4.2. This field is present only when M = 0.

5.6. Home-Test Init message



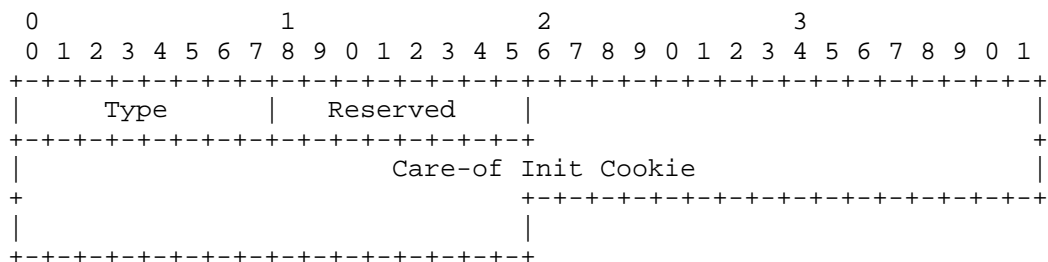
Type TBA_MIP1

Reserved Set to zero, MUST be ignored on reception.

Home Init Cookie

64-bit field which contains a random value, the Home Init Cookie.

5.7. Care-of-Test Init message



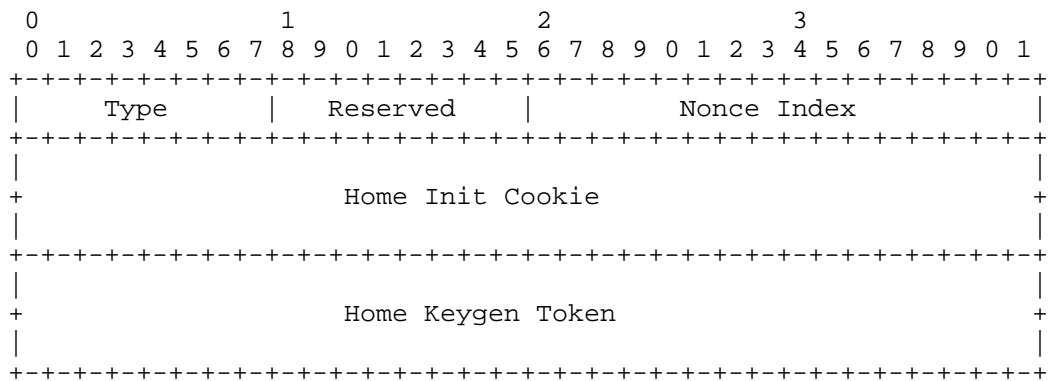
Type TBA_MIP2

Reserved Set to zero, MUST be ignored on reception.

Care-of Init Cookie

64-bit field which contains a random value, the Care-of Init Cookie.

5.8. Home Test message



Type TBA_MIP3

Reserved Set to zero, MUST be ignored on reception.

Nonce Index

This field will be echoed back by the Mobile Router to the Correspondent Router in a subsequent registration request's authentication extension.


Home Init Cookie

64-bit field which contains a random value, the Home Init Cookie.

Home Keygen Token

This field contains the 64 bit home keygen token used in the Return Routability procedure. Generated from cookie + nonce.

5.9. Care-of test message



Type TBA_MIP4

Reserved Set to zero, MUST be ignored on reception.

Care-of Nonce Index

This field will be echoed back by the Mobile Router to the Correspondent Router in a subsequent registration requests' authentication extension.

Care-of Init Cookie

64-bit field which contains a random value, the Home Init Cookie.

Care-of Keygen Token

This field contains the 64 bit home keygen token used in the Return Routability procedure.

6. Special Considerations

6.1. NATs and stateful firewalls

Mechanisms described in MIP NAT traversal [RFC3519] allow the Home Agent to work with Mobile Routers situated behind a NAT device or a stateful firewall. Furthermore, the Home Agent may also detect whether NAT device is located between the Mobile Node and the HA. Mobile Router may also explicitly state it is behind a NAT or firewall on all interfaces, and this information is passed on to the other Mobile Routers with the information field in Route optimization

prefix advertisement extension (Section 5.5). Home Agent may also detect presence of NAT and informs the registering Mobile Router with the 'N' flag in Route Optimization Reply extension (Section 5.2). In the case of one or both of the routers is known to be behind NAT or similarly impaired (not being able to accept incoming connections), the tunnel establishment procedure SHOULD take this into account.

In the case where Mobile Router is behind NAT (or firewall) and Correspondent Router is not, the Mobile Router will, when tunnel has been established, send keepalive messages (ICMP echo requests) through the tunnel. Until a reply has been received, the tunnel SHOULD NOT be considered active. Once reply has been received, NAT mapping is in place and traffic can be sent.

Source address may change due to NAT in CoTI and Registration Request messages. This does not affect the process - the hash values are calculated by the translated address, and the Registration Request will also appear from the same translated address.

Unlike in communication with the Home Agent, in the case of Route Optimization the path used for signaling is not used for tunneled packets, as signaling always uses Home Addresses, and MR <-> CR tunnel is from CoA to CoA. It is assumed that even though port numbers may change, NAT processing rarely allocates more than one external IP address to a single internal address, thus the IP address seen in the Registration Request and Tunnel packets remains the same. However, the UDP source port number may be different in Registration Request and incoming tunnel packets due to port translation. This must not cause an error situation - the Correspondent Router MUST be able to accept tunneling packets from a different UDP source port than what was used in the Registration Request.

Since Mobile Routers may have multiple interfaces connecting to several different networks, it might be possible that specific Mobile Routers may only be able to perform Route Optimization using specific Care-of-address pairs, obtained from specific networks, for example in a case where two Mobile Routers have an interface behind same NAT. Similar case may be applicable to nested NATs. In such cases, Mobile Router MAY attempt to detect eligible Care-of-Address pairs by performing a registration and attempting to establish a tunnel (sending keepalives) with each Care-of-Address listed in the Registration Reply's Care-of-Address extension. The eligible pairs should be recorded in Route Optimization cache. If tunnel cannot be established with any CoA's, the Mobile Router MAY attempt to repeat the procedure with alternative interfaces. The above information on network topology can also be configured to the Mobile Routers either statically or via some external feedback mechanism.

If both the Mobile Router and the Correspondent Router are behind two separate NATs, some sort of proxy or hole-punching technique may be applicable. This is out of scope of this document.

6.2. Handling of concurrent handovers

If both Mobile Router and Correspondent Router move at the same time, this causes no issues from signaling perspective, as all requests are always sent from a Care-of-Address to Home Addresses. Thus, the recipient will always receive the request and can send the reply. This applies even in break-before-make situations where both MR and CR get disconnected at same time - once the connectivity is restored, one end-point of the signaling messages is always the Home Address of respective router, and it is up to the Home Agent to provide reachability.

6.3. Foreign Agents

Since Foreign Agents have been dropped from Network Mobility for Mobile IPv4 work, they are not considered here.

6.4. Multiple Home Agents

Mobile Routers can negotiate and perform route optimization without the assistance of Home Agent - if they can discover each others existence and thus know where to send registration messages. This document only addresses a logically single Home Agent that distributes network prefix information to the Mobile Routers. Problems arise from possible trust relationships; In this document the Home Agent serves as a way to provide verification that a specific network is managed by a specific router.

If Route Optimization is desired between nodes attached to separate Home Agents, there are several possibilities. Note that standard high availability redundancy protocols, such as VRRP, can be utilized; However, in such case the Home Agent is still a single logical entity even if consisting of more than a single node.

Several possibilities exist for achieving Route Optimization between Mobile Routers attached to separate Home Agents, such as a new discovery/probing protocol, routing protocol between Home Agents or DNS SRV records, or a common AAA architecture. There already is a framework for HA to retrieve information from AAA so it can be considered as the most viable possibility. See Section 6.6 for information on possibility to generalize the method.

Any discovery/probing protocols are out of scope for this document.

6.5. Mutualness of Route Optimization

The procedure as specified is asymmetric; That is, if bidirectional route optimization is desired while maintaining consistency, the route optimization (RR check and registration) has to be performed in both directions, but this is not strictly necessary. This is primarily a policy decision depending on how often the mobile prefixes are reconfigured.

Consider the case where two networks, A and B, are handled by Mobile Routers A and B respectively. If the routers are set up in such a fashion that Route Optimization is triggered when a packet is received from a Network Prefix in Route Optimization Cache, the following occurs if a node in network A starts sending ICMP echo requests (pinging) a node in network B.

MR B sees the incoming ICMP echo request packet, which is travelling inside the reverse tunnel to the Home Agent. MR B sees that the destination is in network B, and furthermore, source is in network A which exists in the cache. This triggers Route Optimization processing. Until RO is active, the ping packets (echo requests and replies) are routed via the reverse tunnel.

MR B completes RR procedure and registration with MR A, which thus becomes a Correspondent Router for MR B. A tunnel is created between the routers. MR A updates its routing tables so that network B is reachable via MR A <-> MR B tunnel.

The traffic pattern is now that packets from network B to network A are sent over the direct tunnel, but the packets from A to B are transmitted via the Home Agent and reverse tunnels. MR A now performs its own registration towards MR B. Upon completion, MR A notices that a tunnel to MR B already exists, but updates its routing table so that network B is now reachable via the MR A <-> MR B tunnel. From this point onward, traffic is bidirectional.

In this scenario, if MR A does NOT perform a separate route optimization (RR check and registration), but instead simply updates its routing table to reach network B via the tunnel, problems may arise if MR B has started to manage another network B' before the information has propagated to MR A. The end result is that MR B starts to receive packets for network B' via the Home Agent and for network B via direct tunnel. If Reverse Path checking or similar mechanism is in use on MR B, packets from network A could be black holed.

Whether to perform this mutual registration or not thus depends on the situation, and whether Mobile Routers are going to start managing

additional Network Prefixes during operation.

6.6. Extensibility

The design considerations include several mechanisms which might not be strictly necessary if Route Optimization would only be desired between individual customer sites in a managed network. The registration procedure (with the optional Return Routability part), which allows for Correspondent Routers to learn Mobile Router's Care-of Addresses is not strictly necessary; The CoA's could have been provided by HA directly.

However, this approach allows the method to be extended to a more generic route optimization. The primary driver for having Home Agent to work as a centralized information distributor is to provide Mobile Routers with the knowledge of not only the other routers, but to provide information on which networks are managed by which routers.

The Home Agent provides the information on all feasible nodes with which it is possible to establish Route Optimization. If representing a whole Mobile Network is not necessary, in effect the typical Mobile Node <-> Correspondent Node situation, the mechanisms in this document work just as well - only problem is discovering if the target Correspondent Node can provide Route Optimization capability. This can be performed by not including any prefixes in the information extension, just the HoA address of Mobile Router.

In addition, with Route Optimization for single node, checks on whether a Mobile Router is allowed to represent specific networks are unnecessary since there are none.

Correspondent node/router discovery protocols (whether they are based on probing or a centralized directory beyond the single Home Agent) are outside the scope of this document.

6.7. Load Balancing

The design simply provides possibility to create optimal paths between Mobile Routers; It doesn't dictate what should be the user traffic using these paths. One possible approach in helping facilitate load balancing and utilizing all available paths is presented in [I-D.ietf-mip4-multiple-tunnel-support], which effectively allows for multiple Care-of addresses for a single Home Address. In addition, per-tunnel load balancing is possible by using separate Care-of-Addresses for separate tunnels.

7. Scalability

Home Agent assisted Route Optimization scalability issues stem from the general Mobile IPv4 architecture which is based on tunnels. Creating, maintaining and destroying tunnel interfaces can cause load on the Mobile Routers. However, the MRs can always fall back to normal, reverse tunnelled routing if resource constraints are apparent.

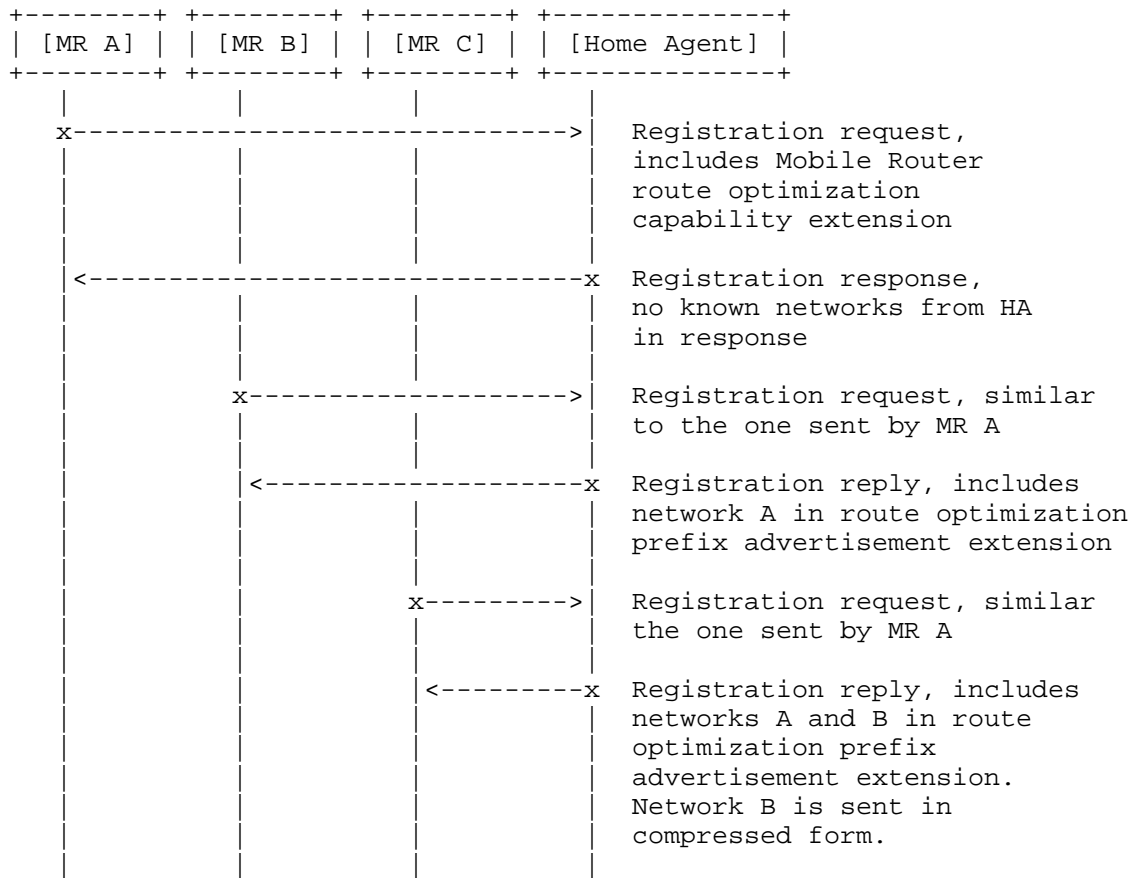
If there is a large number of optimization-capable prefixes, maintaining state for all of these may be an issue also, due to limits on routing table sizes.

Registration responses from Home Agent to Mobile Router may provide information on large number of network prefixes. If thousands of networks are involved, the registration reply messages are bound to grow very large. The prefix- and realm compression mechanisms defined in Section 4 mitigates this problem to an extent. There will, however, be some practical upper limit after which point some other delivery mechanism for the prefix information will be needed.

8. Example signaling scenarios

8.1. Registration request

The following example signaling assumes that there are three Mobile Routers, MR A, B, C, each managing network prefixes A, B, and C. At the beginning, no networks are registered to the Home Agent. Any AAA processing at the Home Agent is omitted from the diagram.



8.2. Route optimization with return routability

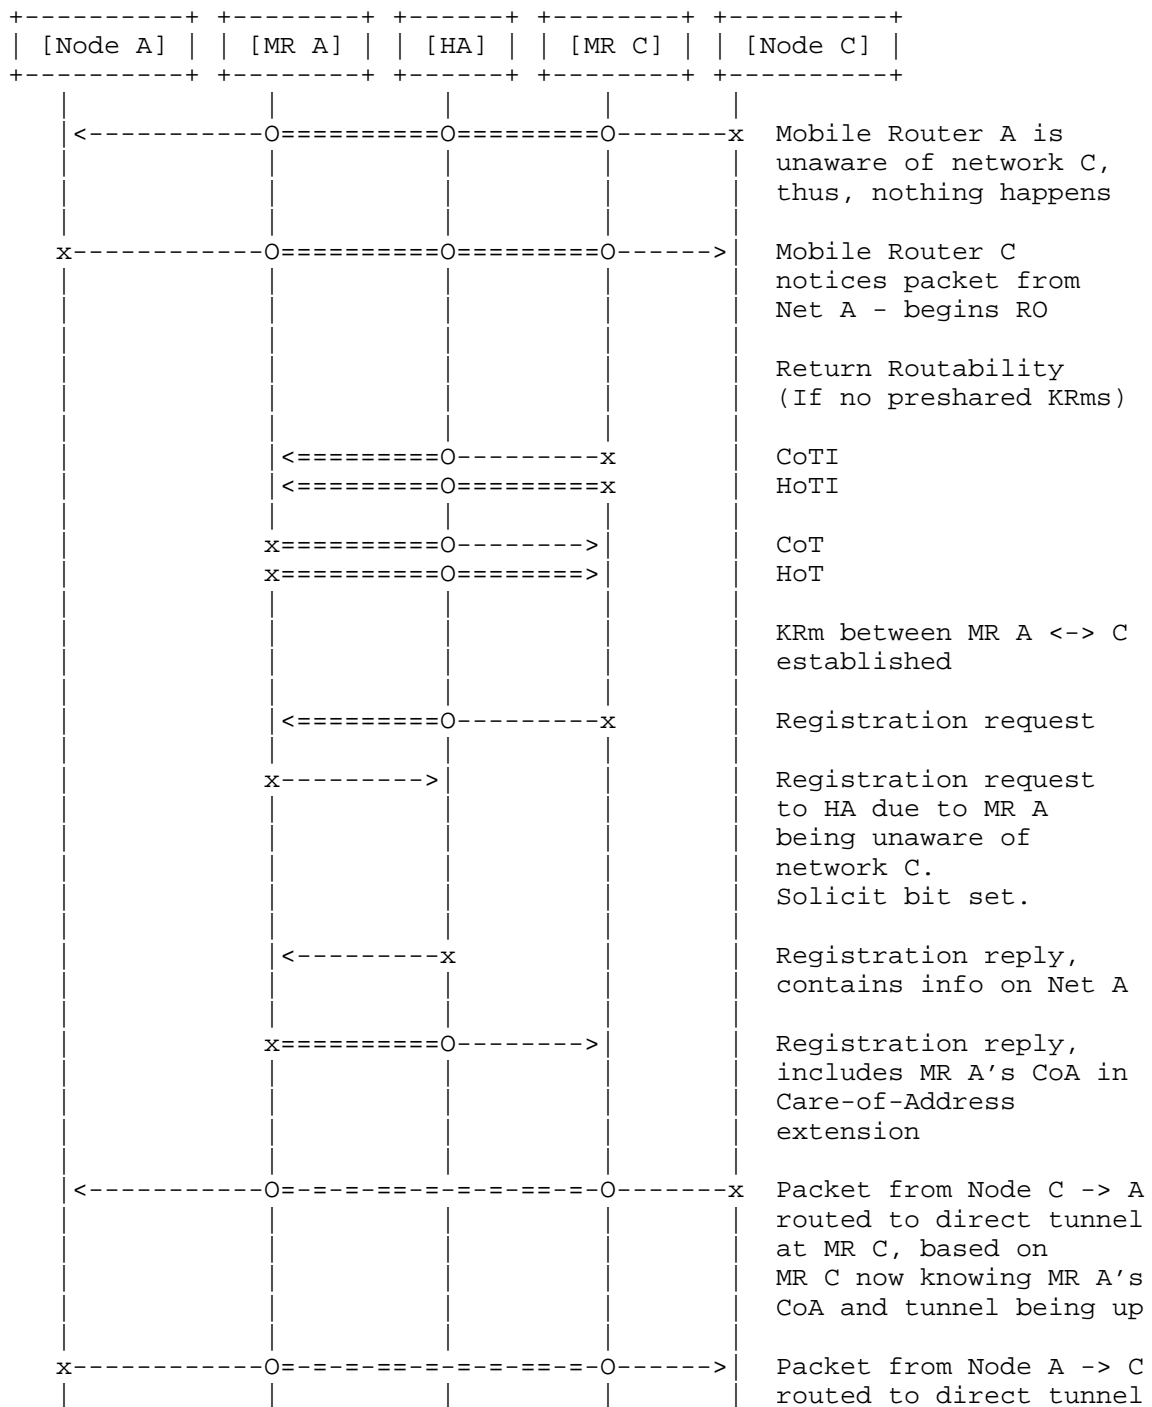
The following example signaling has same network setup as in Section 8.1 - Three mobile routers, each corresponding to their respective network. Node A is in network A and Node C is in network C.

At the beginning, no mobile routers know KRM's of each other. If the KRM's would be pre-shared or provisioned with some other method, the Return Routability messages can be omitted. Signaling in Section 8.1 has occurred, thus MR A is not aware of the other networks, and MR C is aware of networks A and B.

```

===== Traffic inside Mobile IP tunnel to/from HA
----- Traffic inside Mobile IP tunnel between MRs
----- Traffic outside Mobile IP tunnel

```

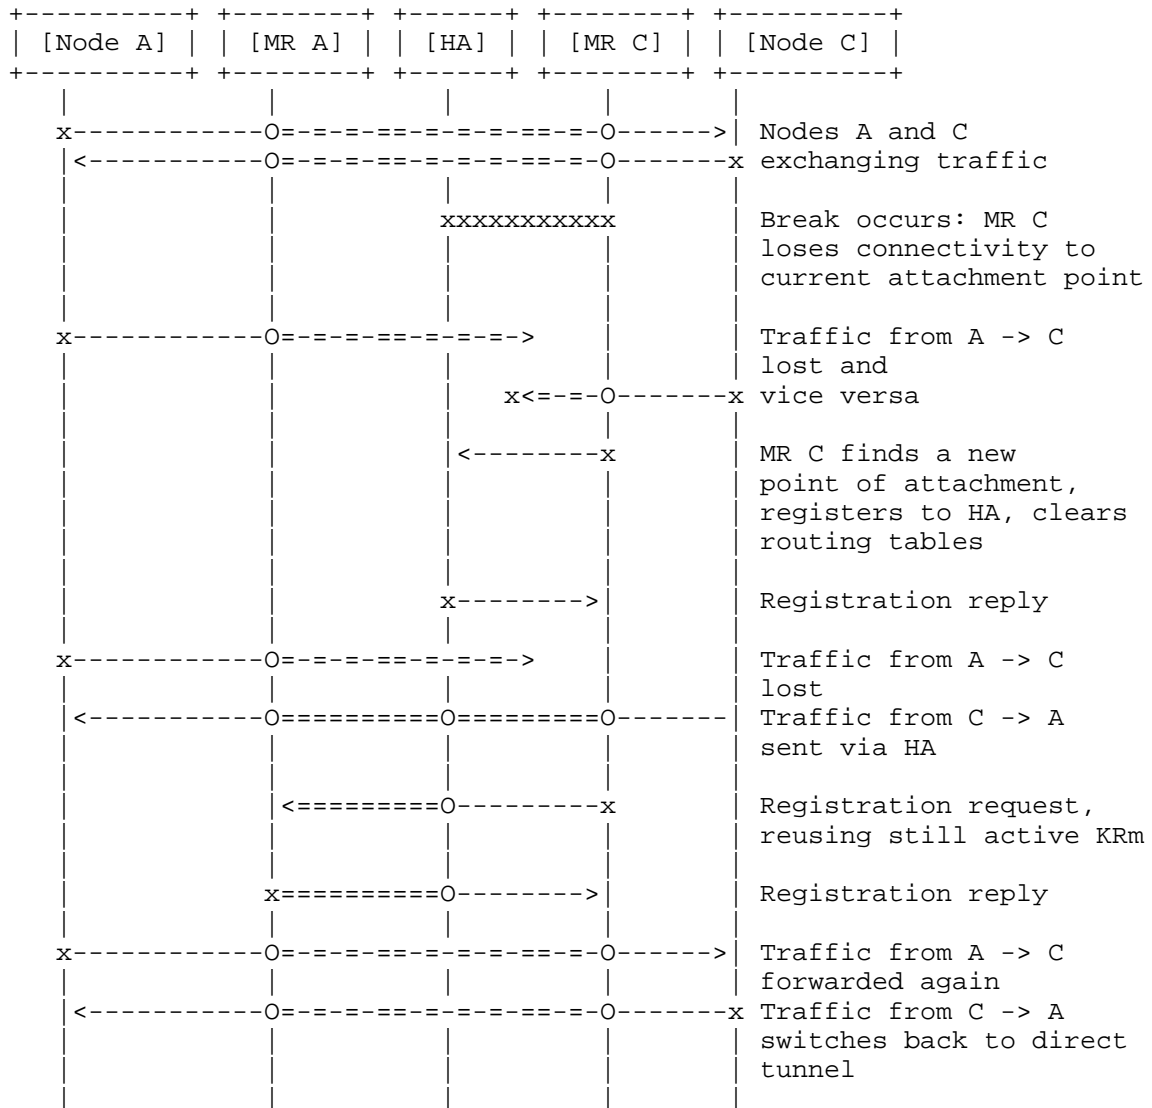


					at MR A, based on MR A
					now knowing MR C's CoA
					and tunnel being up

8.3. Handovers

In this example signaling, MR C changes care-of address while Route Optimization between MR A is operating and data is being transferred. Both cases where the handover is graceful ("make before break") and ungraceful ("break before make") occur in similar fashion, except in the graceful version no packets get lost.

===== Traffic inside Mobile IP tunnel to/from HA
 ==----- Traffic inside Mobile IP tunnel between MRs
 ----- Traffic outside Mobile IP tunnel



9. Protocol constants

MAX_NONCE_LIFETIME	240 seconds
MAX_TOKEN_LIFETIME	210 seconds
MAX_RR_BINDING_LIFETIME	420 seconds
MAX_UPDATE_RATE	5 times

10. IANA Considerations

IANA has assigned rules for the existing registries "Mobile IP Messages" and "Mobile IPv4 numbers" in RFC 3344 [RFC3344]. Numbering spaces for Mobile IP messages and for Extensions that may appear in Mobile IP control messages (those sent to and from UDP port number 434) should be modified.

New Mobile IP control message extension and message type values are needed for the messages and extensions listed in Section 5. The Route Optimization authentication processing requires four new message type numbers. In addition, there is a skippable extension which requires it's own type number. The rest of the new extensions are non-skippable, and grouped under two new types as subtypes. Other type is for extensions in "short" format and other for single extension in "long" extension format.

New Mobile IP registration reply code values are needed for responses from Correspondent Routers. The Route Optimization requires three new reply codes. In addition, a new allocation guideline for "Correspondent Router reply codes" are needed.

The new MIP message types are listed below:

Value	Name
TBA_MIP1	Home-Test Init message
TBA_MIP2	Care-of-Test Init message
TBA_MIP3	Home Test message
TBA_MIP4	Care-of Test message

Table 1: New Values for Mobile IP Message types

The new MIP control message extension types are listed below:

Value	Name
TBA_T1, 128-255	Mobile router Route optimization indication
TBA_T2, 0-127	Route Optimization Extensions
TBA_T3, 0-127	Route Optimization data

Table 2: New Values and Names for Extensions in Mobile IP Control messages

Three new number spaces have been created for the Values and Names for the Sub-Type for Route Optimization-related Extensions. This number spaces are initially defined to hold the following entries, allocated by this document:

Value	Name
TBA_ST1_1	Mobile router Route optimization capability
TBA_ST2_1	Route optimization reply
TBA_ST2_2	Mobile-Correspondent authentication extension
TBA_ST2_3	Care-of address Extension
TBA_ST3_1	Route optimization prefix advertisement

Table 3: New Values and Names for the Sub-type Route Optimization Extension

Note to RFC Editor: this section may be removed on publication as an RFC.

Three new registration reply codes have been created for Code Values for Mobile IP Registration Reply Messages. Following values are added:

Value	Name
TBA_C1	Expired Home nonce Index
TBA_C2	Expired Care-of nonce Index
TBA_C3	Expired nonces
TBA_C4	Concurrent registration

11. Security Considerations

There are two primary security issues: Other relates to return routability check, which establishes that a specific Care-of address is, indeed, managed by a specific Home Address. Other issue is trust relationships and arbitrary router claiming to represent arbitrary network.

The end-user traffic can be protected using normal IPSec mechanisms.

11.1. Return Routability

The Return Routability check's security has been vetted with Mobile IPv6. There are no large differences apart from requiring a separate ICMP message for connectivity check, and replay attack protection, which in this case uses Mobile IPv4 timestamps in registration request's identification field instead of sequence numbers.

The Return Routability procedure does not establish any kind of state information on the Correspondent router, mitigating Denial of Service attacks. State information is only maintained after a Registration request has been accepted.

11.2. Trust relationships

The network of trust relationships in Home Agent assisted Route Optimization solve the issues where arbitrary Correspondent Router can trust an arbitrary Mobile Router that it is indeed the proper route to reach an arbitrary mobile network.

It is assumed that all Mobile Routers have a trust relationship with the Home Agent. Thus, they trust information provided by Home Agent.

The Home Agent provides information matching Home Addresses and network prefixes. Each Mobile Router trusts this information.

Mobile Routers may perform Return Routability procedure between each other. This creates a trusted association between Mobile Router Home Address and Care-of Address. The Mobile Router also claims to represent a specific network. This information is not trustworthy as such.

The claim can be verified by checking the Home Address <-> network prefix information received, either earlier, or due to on-demand request, from the Home Agent. If they match, the Mobile Router's claim is authentic. If the network is considered trusted, a policy decision can be made to skip this check. Exact definitions on situations where such decision can be made are out of scope of this

document. The RECOMMENDED general practice is to perform the check.

12. Acknowledgements

Thanks to Jyrki Soini and Kari Laihonen for initial reviews. This work was supported by TEKES as part of the Future Internet program of TIVIT (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT).

13. References

13.1. Normative References

- [RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [RFC2004] Perkins, C., "Minimal Encapsulation within IP", RFC 2004, October 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3519] Levkowetz, H. and S. Vaarala, "Mobile IP Traversal of Network Address Translation (NAT) Devices", RFC 3519, April 2003.
- [RFC5177] Leung, K., Dommety, G., Narayanan, V., and A. Petrescu, "Network Mobility (NEMO) Extensions for Mobile IPv4", RFC 5177, April 2008.

13.2. Informative References

- [I-D.ietf-mip4-multiple-tunnel-support]
Gundavelli, S., Leung, K., Tsirtsis, G., Soliman, H., and A. Petrescu, "Flow Binding Support for Mobile IPv4", draft-ietf-mip4-multiple-tunnel-support-00 (work in progress), August 2010.
- [I-D.ietf-mobileip-optim]

Perkins, C. and D. Johnson, "Route Optimization in Mobile IP", September 2001.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC3543] Glass, S. and M. Chandra, "Registration Revocation in Mobile IPv4", RFC 3543, August 2003.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

Authors' Addresses

Antti Makela
Aalto University
P.O. BOX 13000
FIN-00076 Aalto
FINLAND

Phone: +358 9 451 5590
Email: antti.makela@tkk.fi

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
FINLAND

Email: jouni.nospam@gmail.com

MIP4 Working Group
Internet-Draft
Obsoletes: 3344 (if approved)
Intended status: Standards Track
Expires: October 10, 2010

C. Perkins, Ed.
WiChorus Inc.
April 8, 2010

IP Mobility Support for IPv4, revised
draft-ietf-mip4-rfc3344bis-10

Abstract

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

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on October 10, 2010.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
1.1. Protocol Requirements	5
1.2. Goals	6
1.3. Assumptions	6
1.4. Applicability	6
1.5. New Architectural Entities	7
1.6. Terminology	7
1.7. Protocol Overview	10
1.8. Message Format and Protocol Extensibility	14
1.9. Type-Length-Value Extension Format for Mobile IP Extensions	16
1.10. Long Extension Format	17
1.11. Short Extension Format	18
2. Agent Discovery	19
2.1. Agent Advertisement	19
2.1.1. Mobility Agent Advertisement Extension	21
2.1.2. Prefix-Lengths Extension	24
2.1.3. One-byte Padding Extension	25
2.2. Agent Solicitation	25
2.3. Foreign Agent and Home Agent Considerations	25
2.3.1. Advertised Router Addresses	26
2.3.2. Sequence Numbers and Rollover Handling	27
2.4. Mobile Node Considerations	27
2.4.1. Registration Required	28
2.4.2. Move Detection	28
2.4.3. Returning Home	30
2.4.4. Sequence Numbers and Rollover Handling	30
3. Registration	31
3.1. Registration Overview	31
3.2. Authentication	33
3.3. Registration Request	33
3.4. Registration Reply	36
3.5. Registration Extensions	39
3.5.1. Computing Authentication Extension Values	39
3.5.2. Mobile-Home Authentication Extension	40
3.5.3. Mobile-Foreign Authentication Extension	41
3.5.4. Foreign-Home Authentication Extension	42

3.6. Mobile Node Considerations	43
3.6.1. Sending Registration Requests	44
3.6.2. Receiving Registration Replies	48
3.6.3. Registration Retransmission	51
3.7. Foreign Agent Considerations	52
3.7.1. Configuration and Registration Tables	52
3.7.2. Receiving Registration Requests	53
3.7.3. Receiving Registration Replies	56
3.8. Home Agent Considerations	58
3.8.1. Configuration and Registration Tables	59
3.8.2. Receiving Registration Requests	60
3.8.3. Sending Registration Replies	64
4. Routing Considerations	68
4.1. Encapsulation Types	68
4.2. Unicast Datagram Routing	68
4.2.1. Mobile Node Considerations	68
4.2.2. Foreign Agent Considerations	69
4.2.3. Home Agent Considerations	70
4.3. Broadcast Datagrams	71
4.4. Multicast Datagram Routing	72
4.5. Mobile Routers	73
4.6. ARP, Proxy ARP, and Gratuitous ARP	75
5. Security Considerations	79
5.1. Message Authentication Codes	79
5.2. Areas of Security Concern in this Protocol	79
5.3. Key Management	79
5.4. Picking Good Random Numbers	80
5.5. Privacy	80
5.6. Ingress Filtering	80
5.7. Replay Protection for Registration Requests	81
5.7.1. Replay Protection using Timestamps	81
5.7.2. Replay Protection using Nonces	82
6. IANA Considerations	84
6.1. Mobile IP Message Types	84
6.2. Extensions to RFC 1256 Router Advertisement	85
6.3. Extensions to Mobile IP Registration Messages	85
6.4. Code Values for Mobile IP Registration Reply Messages	85
7. Acknowledgments	87
8. References	89
8.1. Normative References	89
8.2. Informative References	90
Appendix A. Pre-RFC5378 Disclaimer	93
Appendix B. Link-Layer Considerations	94
Appendix C. TCP Considerations	95
C.1. TCP Timers	95
C.2. TCP Congestion Management	95
Appendix D. Example Scenarios	96
D.1. Registering with a Foreign Agent Care-of Address	96

D.2. Registering with a Co-Located Care-of Address	96
D.3. Deregistration	97
Appendix E. Applicability of Prefix-Lengths Extension	98
Appendix F. Interoperability Considerations	99
Appendix G. Changes since RFC 3344	100
Appendix H. Example Messages	102
H.1. Example ICMP Agent Advertisement Message Format	102
H.2. Example Registration Request Message Format	102
H.3. Example Registration Reply Message Format	103
Author's Address	105

1. Introduction

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

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

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

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

Changes between this revised specification for Mobile IP and the original specifications (see [44],[14],[15],[20],[4]) are detailed in Appendix G.

1.1. Protocol Requirements

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

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

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

1.2. Goals

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

1.3. Assumptions

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

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

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

1.4. Applicability

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

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

1.5. New Architectural Entities

Mobile IP introduces the following new functional entities:

Mobile Node

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

Home Agent

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

Foreign Agent

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

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

1.6. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1].

In addition, this document frequently uses the following terms:

Authorization-enabling extension

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

In this document, all uses of authorization-enabling extension refer to authentication extensions that enable the Registration Request message to be acceptable to the home agent. Using additional protocol structures specified outside of this document, it may be possible for the mobile node to provide authentication of its registration to the home agent, by way of another authenticating entity within the network that is acceptable to the home agent (for example, see RFC 2794 [2]).

Agent Advertisement

An advertisement message constructed by attaching a special Extension to a router advertisement [5] message.

Authentication

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

Care-of Address

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

Correspondent Node

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

Foreign Network

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

Gratuitous ARP

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

Home Address

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

Home Network

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

Link

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

Link-Layer Address

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

Mobility Agent

Either a home agent or a foreign agent.

Mobility Binding

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

Mobility Security Association

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

Node

A host or a router.

Nonce

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

Security Parameter Index (SPI)

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

Tunnel

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

Virtual Network

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

Visited Network

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

Visitor List

The list of mobile nodes visiting a foreign agent.

1.7. Protocol Overview

The following support services are defined for Mobile IP:

Agent Discovery

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

prospective agents are present.

Registration

When the mobile node is away from home, it registers its care-of address with its home agent. Depending on its method of attachment, the mobile node will register either directly with its home agent, or through a foreign agent which forwards the registration to the home agent.

silently discard

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

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

- o Mobility agents (i.e., foreign agents and home agents) advertise their presence via Agent Advertisement messages (Section 2). A mobile node may optionally solicit an Agent Advertisement message from any locally attached mobility agents through an Agent Solicitation message.
- o A mobile node receives these Agent Advertisements and determines whether it is on its home network or a foreign network.
- o When the mobile node detects that it is located on its home network, it operates without mobility services. If returning to its home network from being registered elsewhere, the mobile node deregisters with its home agent, through exchange of a Registration Request and Registration Reply message with it.
- o When a mobile node detects that it has moved to a foreign network, it obtains a care-of address on the foreign network. The care-of address can either be determined from a foreign agent's advertisements (a foreign agent care-of address), or by some external assignment mechanism such as DHCP [34] (a co-located care-of address).
- o The mobile node operating away from home then registers its new care-of address with its home agent through exchange of a Registration Request and Registration Reply message with it, possibly via a foreign agent (Section 3).

- o Datagrams sent to the mobile node's home address are intercepted by its home agent, tunneled by the home agent to the mobile node's care-of address, received at the tunnel endpoint (either at a foreign agent or at the mobile node itself), and finally delivered to the mobile node (Section 4.2.3).
- o In the reverse direction, datagrams sent by the mobile node are generally delivered to their destination using standard IP routing mechanisms, not necessarily passing through the home agent.

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

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

- a. A "foreign agent care-of address" is a care-of address provided by a foreign agent through its Agent Advertisement messages. In this case, the care-of address is an IP address of the foreign agent. In this mode, the foreign agent is the endpoint of the tunnel and, upon receiving tunneled datagrams, decapsulates them and delivers the inner datagram to the mobile node. This mode of acquisition is preferred because it allows many mobile nodes to share the same care-of address and therefore does not place unnecessary demands on the already limited IPv4 address space.
- b. A "co-located care-of address" is a care-of address acquired by the mobile node as a local IP address through some external means, which the mobile node then associates with one of its own network interfaces. The address may be dynamically acquired as a temporary address by the mobile node such as through DHCP [34], or may be owned by the mobile node as a long-term address for its use only while visiting some foreign network. Specific external methods of acquiring a local IP address for use as a co-located care-of address are beyond the scope of this document. When using a co-located care-of address, the mobile node serves as the endpoint of the tunnel and itself performs decapsulation of the datagrams tunneled to it.

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

because it requires a pool of addresses within the foreign network to be made available to visiting mobile nodes. It is difficult to efficiently maintain pools of addresses for each subnet that may permit mobile nodes to visit.

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

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

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

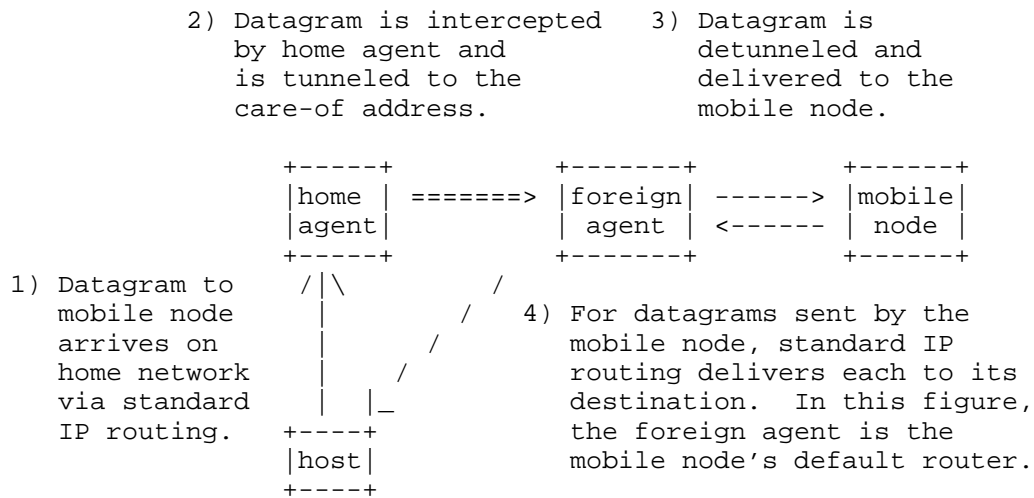


Figure 1: Operation of Mobile IPv4

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

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

1.8. Message Format and Protocol Extensibility

Mobile IP defines a set of new control messages, sent with UDP [17] using well-known port number 434. The following two message types are defined in this document:

1 Registration Request

3 Registration Reply

Up-to-date values for the message types for Mobile IP control messages are specified in the IANA online database [48].

In addition, for Agent Discovery, Mobile IP makes use of the existing Router Advertisement and Router Solicitation messages defined for ICMP Router Discovery [5].

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

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

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

- o The first set consists of those Extensions which may appear in Mobile IP control messages (those sent to and from UDP port number 434). In this document, the following Types are defined for Extensions appearing in Mobile IP control messages:

- 0 One-byte Padding (encoded with no Length nor Data field)
 - 32 Mobile-Home Authentication
 - 33 Mobile-Foreign Authentication
 - 34 Foreign-Home Authentication

- o The second set consists of those extensions which may appear in ICMP Router Discovery messages [5]. In this document, the following Types are defined for Extensions appearing in ICMP Router Discovery messages:

- 0 One-byte Padding (encoded with no Length nor Data field)
 - 16 Mobility Agent Advertisement
 - 19 Prefix-Lengths

Each individual Extension is described in detail in a separate section later in this document. Up-to-date values for these Extension Type numbers are specified in the IANA online database [48].

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

The type field in the Mobile IP extension structure can support up to 255 (skippable and not skippable) uniquely identifiable extensions. When an Extension numbered in either of these sets within the range 0 through 127 is encountered but not recognized, the message containing that Extension MUST be silently discarded. When an Extension

numbered in the range 128 through 255 is encountered which is not recognized, that particular Extension is ignored, but the rest of the Extensions and message data MUST still be processed. The Length field of the Extension is used to skip the Data field in searching for the next Extension.

Unless additional structure is utilized for the extension types, new developments or additions to Mobile IP might require so many new extensions that the available space for extension types might run out. Two new extension structures are proposed to solve this problem. Certain types of extensions can be aggregated, using subtypes to identify the precise extension, for example as has been done with the Generic Authentication Keys extensions [46]. In many cases, this may reduce the rate of allocation for new values of the type field.

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

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

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

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

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




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

Type

Indicates the particular type of Extension.

Length


Indicates the length (in bytes) of the data field within this Extension. The length does NOT include the Type and Length bytes.

Data

The particular data associated with this Extension. This field may be zero or more bytes in length. The format and length of the data field is determined by the type and length fields.

1.10. Long Extension Format

This format is applicable for non-skippable extensions which carry information more than 256 bytes. Skippable extensions can never use the long format, because the receiver is not required to include parsing code and is likely to treat the 8 bits immediately following the Type as the Length field.



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

Type

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

Sub-Type

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

Length

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

Data

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

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

1.11. Short Extension Format

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

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |   Length   |   Sub-Type   |   Data ...   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

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

Type

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

Sub-Type

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

Length

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

Data

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

2. Agent Discovery

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

Mobile IP extends ICMP Router Discovery [5] as its primary mechanism for Agent Discovery. An Agent Advertisement is formed by including a Mobility Agent Advertisement Extension in an ICMP Router Advertisement message (Section 2.1). An Agent Solicitation message is identical to an ICMP Router Solicitation, except that its IP TTL MUST be set to 1 (Section 2.2). This section describes the message formats and procedures by which mobile nodes, foreign agents, and home agents cooperate to realize Agent Discovery.

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

No authentication is required for Agent Advertisement and Agent Solicitation messages. They MAY be authenticated using the IP Authentication Header [9], which is unrelated to the messages described in this document. Further specification of the way in which Advertisement and Solicitation messages may be authenticated is outside of the scope of this document.

2.1. Agent Advertisement

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

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

Link-Layer Fields

Destination Address

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

IP Fields

TTL

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

Destination Address

As specified for ICMP Router Discovery [5], the IP destination address of an multicast Agent Advertisement MUST be either the "all systems on this link" multicast address (224.0.0.1) [6] or the "limited broadcast" address (255.255.255.255). The subnet-directed broadcast address of the form <prefix>.<-1> cannot be used since mobile nodes will not generally know the prefix of the foreign network. When the Agent Advertisement is unicast to a mobile node, the IP home address of the mobile node SHOULD be used as the Destination Address.

ICMP Fields

Code

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

0 The mobility agent handles common traffic -- that is, it acts as a router for IP datagrams not necessarily related to mobile nodes.

16 The mobility agent does not route common traffic. However, all foreign agents MUST (minimally) forward to a default router any datagrams received from a registered mobile node (Section 4.2.2).

Lifetime

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

Router Address(es)

See Section 2.3.1 for a discussion of the addresses that may appear in this portion of the Agent Advertisement.


Num Addrs

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

If sent periodically, the nominal interval at which Agent Advertisements are sent SHOULD be no longer than 1/3 of the advertisement Lifetime given in the ICMP header. This interval MAY be shorter than 1/3 the advertised Lifetime. This allows a mobile node to miss three successive advertisements before deleting the agent from its list of valid agents. The actual transmission time for each advertisement SHOULD be slightly randomized [5] in order to avoid synchronization and subsequent collisions with other Agent Advertisements that may be sent by other agents (or with other Router Advertisements sent by other routers). Note that this field has no relation to the "Registration Lifetime" field within the Mobility Agent Advertisement Extension defined below.

2.1.1. Mobility Agent Advertisement Extension

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



Type

16

Length

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

Sequence Number

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

Registration Lifetime

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

R

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

B

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

H

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

F

Foreign agent. This agent offers service as a foreign agent on the link on which this Agent Advertisement message is sent.

M

Minimal encapsulation. This agent implements receiving tunneled datagrams that use minimal encapsulation [15].

G

GRE encapsulation. This agent implements receiving tunneled datagrams that use GRE encapsulation [13].

r

Sent as zero; ignored on reception. SHOULD NOT be allocated for any other uses.

T

Foreign agent supports reverse tunneling as specified in [12].

U

Mobility agent supports UDP Tunnelling as specified in [27].

X

Mobility agent supports Registration Revocation as specified in [28].

I

Foreign agent supports Regional Registration as specified in [29].

reserved

Sent as zero; ignored on reception.

Care-of Address(es)

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

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

'F' bit is not also set. Furthermore, at least one of the 'F' bit and the 'H' bit MUST be set in any Agent Advertisement message sent.

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

2.1.2. Prefix-Lengths Extension

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

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      | Prefix Length |      ....
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type

19 (Prefix-Lengths Extension)

Length

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

Prefix Length(s)

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

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

2.1.3. One-byte Padding Extension

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

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

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

Type 0 (One-byte Padding Extension)

2.2. Agent Solicitation

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

2.3. Foreign Agent and Home Agent Considerations

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

The same procedures, defaults, and constants are used in Agent Advertisement messages and Agent Solicitation messages as specified for ICMP Router Discovery [5], except that:

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

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

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

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

2.3.1. Advertised Router Addresses

The ICMP Router Advertisement portion of the Agent Advertisement MAY contain one or more router addresses. An agent SHOULD only put its own addresses, if any, in the advertisement. Whether or not its own address appears in the Router Addresses, a foreign agent MUST route datagrams it receives from registered mobile nodes (Section 3.7).

2.3.2. Sequence Numbers and Rollover Handling

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

2.4. Mobile Node Considerations

Every mobile node MUST implement Agent Solicitation. Solicitations SHOULD only be sent in the absence of Agent Advertisements and when a care-of address has not been determined through a link-layer protocol or other means. The mobile node uses the same procedures, defaults, and constants for Agent Solicitation as specified for ICMP Router Solicitation messages [5], except that the mobile node MAY solicit more often than once every three seconds, and that a mobile node that is currently not connected to any foreign agent MAY solicit more times than MAX_SOLICITATIONS.

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

While still searching for an agent, the mobile node MUST NOT increase the rate at which it sends Solicitations unless it has received a positive indication that it has moved to a new link. After successfully registering with an agent, the mobile node SHOULD also increase the rate at which it will send Solicitations when it next begins searching for a new agent with which to register. The increased solicitation rate MAY revert to the maximum rate, but then MUST be limited in the manner described above. In all cases, the recommended solicitation intervals are nominal values. Mobile nodes MUST randomize their solicitation times around these nominal values as specified for ICMP Router Discovery [5].

Mobile nodes MUST process received Agent Advertisements. A mobile

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

If there is more than one advertised address, the mobile node SHOULD pick the first address for its initial registration attempt. If the registration attempt fails with a status Code indicating rejection by the foreign agent, the mobile node MAY retry the attempt with each subsequent advertised address in turn.

When multiple methods of agent discovery are in use, the mobile node SHOULD first attempt registration with agents including Mobility Agent Advertisement Extensions in their advertisements, in preference to those discovered by other means. This preference maximizes the likelihood that the registration will be recognized, thereby minimizing the number of registration attempts.

A mobile node MUST ignore reserved bits in Agent Advertisements, as opposed to discarding such advertisements. In this way, new bits can be defined later, without affecting the ability for mobile nodes to use the advertisements even when the newly defined bits are not understood.

2.4.1. Registration Required

When the mobile node receives an Agent Advertisement with the 'R' bit set, the mobile node SHOULD register through the foreign agent, even when the mobile node might be able to acquire its own co-located care-of address. This feature is intended to allow sites to enforce visiting policies (such as accounting) which require exchanges of authorization.

If formerly reserved bits require some kind of monitoring/enforcement at the foreign link, foreign agents implementing the new specification for the formerly reserved bits can set the 'R' bit. This has the effect of forcing the mobile node to register through the foreign agent, so the foreign agent could then monitor/enforce the policy.

2.4.2. Move Detection

Two primary mechanisms are provided for mobile nodes to detect when they have moved from one subnet to another. Other mechanisms MAY

also be used. When the mobile node detects that it has moved, it SHOULD register (Section 3) with a suitable care-of address on the new foreign network. However, the mobile node MUST NOT register more frequently than once per second on average, as specified in Section 3.6.3.

2.4.2.1. Algorithm 1

The first method of move detection is based upon the Lifetime field within the main body of the ICMP Router Advertisement portion of the Agent Advertisement. A mobile node SHOULD record the Lifetime received in any Agent Advertisements, until that Lifetime expires. If the mobile node fails to receive another advertisement from the same agent within the specified Lifetime, it SHOULD assume that it has lost contact with that agent. If the mobile node has previously received an Agent Advertisement from another agent for which the Lifetime field has not yet expired, the mobile node MAY immediately attempt registration with that other agent. Otherwise, the mobile node SHOULD attempt to discover a new agent with which to register.

2.4.2.2. Algorithm 2

The second method uses network prefixes. The Prefix-Lengths Extension MAY be used in some cases by a mobile node to determine whether or not a newly received Agent Advertisement was received on the same subnet as the mobile node's current care-of address. If the prefixes differ, the mobile node MAY assume that it has moved. If a mobile node is currently using a foreign agent care-of address, the mobile node SHOULD NOT use this method of move detection unless both the current agent and the new agent include the Prefix-Lengths Extension in their respective Agent Advertisements; if this Extension is missing from one or both of the advertisements, this method of move detection SHOULD NOT be used. Similarly, if a mobile node is using a co-located care-of address, it SHOULD NOT use this method of move detection unless the new agent includes the Prefix-Lengths Extension in its Advertisement and the mobile node knows the network prefix of its current co-located care-of address. On the expiration of its current registration, if this method indicates that the mobile node has moved, rather than re-registering with its current care-of address, a mobile node MAY choose instead to register with a the foreign agent sending the new Advertisement with the different network prefix. The Agent Advertisement on which the new registration is based MUST NOT have expired according to its Lifetime field.

2.4.3. Returning Home

A mobile node can detect that it has returned to its home network when it receives an Agent Advertisement from its own home agent. If so, it SHOULD deregister with its home agent (Section 3). Before attempting to deregister, the mobile node SHOULD configure its routing table appropriately for its home network (Section 4.2.1). In addition, if the home network is using ARP [16], the mobile node MUST follow the procedures described in Section 4.6 with regard to ARP, proxy ARP, and gratuitous ARP.

2.4.4. Sequence Numbers and Rollover Handling

If a mobile node detects two successive values of the sequence number in the Agent Advertisements from the foreign agent with which it is registered, the second of which is less than the first and inside the range 0 to 255, the mobile node SHOULD register again. If the second value is less than the first but is greater than or equal to 256, the mobile node SHOULD assume that the sequence number has rolled over past its maximum value (0xffff), and that reregistration is not necessary (Section 2.3).

3. Registration

Mobile IP registration provides a flexible mechanism for mobile nodes to communicate their current reachability information to their home agent. It is the method by which mobile nodes:

- o request forwarding services when visiting a foreign network,
- o inform their home agent of their current care-of address,
- o renew a registration which is due to expire, and/or
- o deregister when they return home.

Registration messages exchange information between a mobile node, (optionally) a foreign agent, and the home agent. Registration creates or modifies a mobility binding at the home agent, associating the mobile node's home address with its care-of address for the specified Lifetime.

Several other (optional) capabilities are available through the registration procedure, which enable a mobile node to:

- o discover its home address, if the mobile node is not configured with this information.
- o maintain multiple simultaneous registrations, so that a copy of each datagram will be tunneled to each active care-of address
- o deregister specific care-of addresses while retaining other mobility bindings, and
- o discover the address of a home agent if the mobile node is not configured with this information.

3.1. Registration Overview

Mobile IP defines two different registration procedures, one via a foreign agent that relays the registration to the mobile node's home agent, and one directly with the mobile node's home agent. The following rules determine which of these two registration procedures to use in any particular circumstance:

- o If a mobile node is registering a foreign agent care-of address, the mobile node **MUST** register via that foreign agent.
- o If a mobile node is using a co-located care-of address, and receives an Agent Advertisement from a foreign agent on the link

on which it is using this care-of address, the mobile node SHOULD register via that foreign agent (or via another foreign agent on this link) if the 'R' bit is set in the received Agent Advertisement message.

- o If a mobile node is otherwise using a co-located care-of address, the mobile node MUST register directly with its home agent.
- o If a mobile node has returned to its home network and is (de)registering with its home agent, the mobile node MUST register directly with its home agent.

Both registration procedures involve the exchange of Registration Request and Registration Reply messages (Section 3.3 and Section 3.4). When registering via a foreign agent, the registration procedure requires the following four messages:

- a. The mobile node sends a Registration Request to the prospective foreign agent to begin the registration process.
- b. The foreign agent processes the Registration Request and then relays it to the home agent.
- c. The home agent sends a Registration Reply to the foreign agent to grant or deny the Request.
- d. The foreign agent processes the Registration Reply and then relays it to the mobile node to inform it of the disposition of its Request.

When the mobile node instead registers directly with its home agent, the registration procedure requires only the following two messages:

- a. The mobile node sends a Registration Request to the home agent.
- b. The home agent sends a Registration Reply to the mobile node, granting or denying the Request.

The registration messages defined in Section 3.3 and Section 3.4 use the User Datagram Protocol (UDP) [17]. A nonzero UDP checksum SHOULD be included in the header, and MUST be checked by the recipient. A zero UDP checksum SHOULD be accepted by the recipient. The behavior of the mobile node and the home agent with respect to their mutual acceptance of packets with zero UDP checksums SHOULD be defined as part of the mobility security association which exists between them.

3.2. Authentication

Each mobile node, foreign agent, and home agent MUST be able to support a mobility security association for mobile entities, indexed by their SPI and IP address. In the case of the mobile node, this must be its Home Address. See Section 5.1 for requirements for support of authentication algorithms. Registration messages between a mobile node and its home agent MUST be authenticated with an authorization-enabling extension, e.g. the Mobile-Home Authentication Extension (Section 3.5.2). This extension MUST be the first authentication extension; other foreign agent-specific extensions MAY be added to the message after the mobile node computes the authentication.

3.3. Registration Request

A mobile node registers with its home agent using a Registration Request message so that its home agent can create or modify a mobility binding for that mobile node (e.g., with a new lifetime). The Request may be relayed to the home agent by the foreign agent through which the mobile node is registering, or it may be sent directly to the home agent in the case in which the mobile node is registering a co-located care-of address.

IP fields:

Source Address

Typically the interface address from which the message is sent.

Destination Address

Typically that of the foreign agent or the home agent.

See Section 3.6.1.1 and Section 3.7.2.2 for details.

UDP fields:


Source Port

variable

Destination Port

434

The UDP header is followed by the Mobile IP fields shown below:



Type

1 (Registration Request)

S

Simultaneous bindings. If the 'S' bit is set, the mobile node is requesting that the home agent retain its prior mobility bindings, as described in Section 3.6.1.2.

B

Broadcast datagrams. If the 'B' bit is set, the mobile node requests that the home agent tunnel to it any broadcast datagrams that it receives on the home network, as described in Section 4.3.

D

Decapsulation by mobile node. If the 'D' bit is set, the mobile node will itself decapsulate datagrams which are sent to the care-of address. That is, the mobile node is using a co-located care-of address.

M

Minimal encapsulation. If the 'M' bit is set, the mobile node requests that its home agent use minimal encapsulation [16] for datagrams tunneled to the mobile node.

G

GRE encapsulation. If the 'G' bit is set, the mobile node requests that its home agent use GRE encapsulation [13] for datagrams tunneled to the mobile node.

r

Sent as zero; ignored on reception. SHOULD NOT be allocated for any other uses.

T

Reverse Tunneling requested; see [12].

x

Sent as zero; ignored on reception.

Lifetime

The number of seconds remaining before the registration is considered expired. A value of zero indicates a request for deregistration. A value of 0xffff indicates infinity.

Home Address

The IP address of the mobile node.

Home Agent

The IP address of the mobile node's home agent.

Care-of Address

The IP address for the end of the tunnel.

Identification

A 64-bit number, constructed by the mobile node, used for matching Registration Requests with Registration Replies, and for protecting against replay attacks of registration messages. See Section 5.4 and Section 5.7.

Extensions

The fixed portion of the Registration Request is followed by one or more of the Extensions listed in Section 3.5. An

authorization-enabling extension MUST be included in all Registration Requests. See Section 3.6.1.3 and Section 3.7.2.2 for information on the relative order in which different extensions, when present, MUST be placed in a Registration Request message.

3.4. Registration Reply

A mobility agent typically returns a Registration Reply message to a mobile node which has sent a Registration Request message. If the mobile node is requesting service from a foreign agent, that foreign agent will typically receive the Reply from the home agent and subsequently relay it to the mobile node. Reply messages contain the necessary codes to inform the mobile node about the status of its Request, along with the lifetime granted by the home agent, which MAY be smaller than the original Request.

The foreign agent MUST NOT increase the Lifetime selected by the mobile node in the Registration Request, since the Lifetime is covered by an authentication extension which enables authorization by the home agent. Such an extension contains authentication data which cannot be correctly (re)computed by the foreign agent. The home agent MUST NOT increase the Lifetime selected by the mobile node in the Registration Request, since doing so could increase it beyond the maximum Registration Lifetime allowed by the foreign agent. If the Lifetime received in the Registration Reply is greater than that in the Registration Request, the Lifetime in the Request MUST be used. When the Lifetime received in the Registration Reply is less than that in the Registration Request, the Lifetime in the Reply MUST be used.

IP fields:

Source Address

Typically copied from the destination address of the Registration Request to which the agent is replying. See Section 3.7.2.3 and Section 3.8.3.2 for complete details.

Destination Address

Copied from the source address of the Registration Request to which the agent is replying

UDP fields:


Source Port

Copied from the UDP destination port of the corresponding Registration Request.

Destination Port

Copied from the source port of the corresponding Registration Request (Section 3.7.1).

The UDP header is followed by the Mobile IP fields shown below:



Type

3 (Registration Reply)

Code

A value indicating the result of the Registration Request. See below for a list of currently defined Code values.

Lifetime

If the Code field indicates that the registration was accepted, the Lifetime field is set to the number of seconds remaining before the registration is considered expired. A value of zero indicates that the mobile node has been deregistered. A value of 0xffff indicates infinity. If the Code field indicates that the registration was denied, the contents of the Lifetime field are unspecified and MUST be ignored on reception.

Home Address

The IP address of the mobile node.

Home Agent

The IP address of the mobile node's home agent.

Identification

A 64-bit number used for matching Registration Requests with Registration Replies, and for protecting against replay attacks of registration messages. The value is based on the Identification field from the Registration Request message from the mobile node, and on the style of replay protection used in the security context between the mobile node and its home agent (defined by the mobility security association between them, and SPI value in the authorization-enabling extension). See Section 5.4 and Section 5.7.

Extensions

The fixed portion of the Registration Reply is followed by one or more of the Extensions listed in Section 3.5. An authorization-enabling extension MUST be included in all Registration Replies returned by the home agent. See Section 3.7.2.2 and Section 3.8.3.3 for rules on placement of extensions to Reply messages.

The following values are defined for use within the Code field.
Registration successful:

- 0 registration accepted
- 1 registration accepted, but simultaneous mobility bindings unsupported

Registration denied by the foreign agent:

- 64 reason unspecified
- 65 administratively prohibited
- 66 insufficient resources
- 67 mobile node failed authentication
- 68 home agent failed authentication
- 69 requested Lifetime too long
- 70 poorly formed Request
- 71 poorly formed Reply

72 requested encapsulation unavailable
73 reserved and unavailable
TBD-IANA Invalid Home Agent address
77 invalid care-of address
78 registration timeout
80 home network unreachable (ICMP error received)
81 home agent host unreachable (ICMP error received)
82 home agent port unreachable (ICMP error received)
88 home agent unreachable (other ICMP error received)

Registration denied by the home agent:

128 reason unspecified
129 administratively prohibited
130 insufficient resources
131 mobile node failed authentication
132 foreign agent failed authentication
133 registration Identification mismatch
134 poorly formed Request
135 too many simultaneous mobility bindings
136 unknown home agent address

Up-to-date values of the Code field are specified in the IANA online database [48].

3.5. Registration Extensions

3.5.1. Computing Authentication Extension Values

The Authenticator value computed for each authentication Extension MUST protect the following fields from the registration message:

- o the UDP payload (that is, the Registration Request or Registration Reply data),
- o all prior Extensions in their entirety, and
- o the Type, Length, and SPI of this Extension.

The default authentication algorithm uses HMAC-MD5 [10] to compute a 128-bit "message digest" of the registration message. The data over which the HMAC is computed is defined as:

- o the UDP payload (that is, the Registration Request or Registration Reply data),
- o all prior Extensions in their entirety, and

- o the Type, Length, and SPI of this Extension.

Note that the Authenticator field itself and the UDP header are NOT included in the computation of the default Authenticator value. See Section 5.1 for information about support requirements for message authentication codes, which are to be used with the various authentication Extensions.

The Security Parameter Index (SPI) within any of the authentication Extensions defines the security context which is used to compute the Authenticator value and which MUST be used by the receiver to check that value. In particular, the SPI selects the authentication algorithm and mode (Section 5.1) and secret (a shared key, or appropriate public/private key pair) used in computing the Authenticator. In order to ensure interoperability between different implementations of the Mobile IP protocol, an implementation MUST be able to associate any SPI value with any authentication algorithm and mode which it implements. In addition, all implementations of Mobile IP MUST implement the default authentication algorithm (HMAC-MD5) specified above.

3.5.2. Mobile-Home Authentication Extension

At least one authorization-enabling extension MUST be present in all Registration Requests, and also in all Registration Replies generated by the Home Agent. The Mobile-Home Authentication Extension is always an authorization-enabling for registration messages specified in this document. This requirement is intended to eliminate problems [30] which result from the uncontrolled propagation of remote redirects in the Internet. The location of the authorization-enabling extension marks the end of the data to be authenticated by the authorizing agent interpreting that authorization-enabling extension.

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      SPI      ....
+-----+-----+-----+-----+-----+-----+-----+
|      ... SPI (cont.)      |      Authenticator ...
+-----+-----+-----+-----+-----+-----+

```

Type

Length

4 plus the number of bytes in the Authenticator.

SPI

Security Parameter Index (4 bytes). An opaque identifier (see Section 1.6).

Authenticator

(variable length) (See Section 3.5.1)

3.5.3. Mobile-Foreign Authentication Extension

This Extension MAY be included in Registration Requests and Replies in cases in which a mobility security association exists between the mobile node and the foreign agent. See Section 5.1 for information about support requirements for message authentication codes.

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      SPI      ....
+-----+-----+-----+-----+-----+-----+-----+-----+
|      ... SPI (cont.)      |      Authenticator ...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type

33

Length

4 plus the number of bytes in the Authenticator.

SPI

Security Parameter Index (4 bytes). An opaque identifier (see Section 1.6).

Authenticator

(variable length) (See Section 3.5.1)

3.5.4. Foreign-Home Authentication Extension

This Extension MAY be included in Registration Requests and Replies in cases in which a mobility security association exists between the foreign agent and the home agent, as long as the Registration Request is not a deregistration (i.e., the mobile node requested a nonzero lifetime and the home address is different than the care-of address). The Foreign-Home Authentication extension MUST NOT be applied to deregistration messages. See Section 5.1 for information about support requirements for message authentication codes.

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      |      SPI      ....
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      ... SPI (cont.)      |      Authenticator ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type

34

Length

4 plus the number of bytes in the Authenticator.

SPI

Security Parameter Index (4 bytes). An opaque identifier (see Section 1.6).

Authenticator

(variable length) (See Section 3.5.1)

In order to perform the authentication, the Home Agent and the Foreign Agent are configured with a mobility security association that is indexed by the SPI (in the appended Foreign-Home Authentication Extension) and the IP Source Address of the Registration Request. When the extension is used with a Registration Reply message, the foreign agent address MUST be used as the Destination IP address in the IP header.

When this extension is applied to a Registration Request message, the mobility security association for verifying the correctness of the authentication data is selected by the Home Agent based on the value of the Source IP Address field of the Registration Request and the

SPI of the Authentication extension. The Source IP Address will be the same as the Care-of Address field of the Registration Request (see Section 3.7.2.2)

When this extension is applied to a Registration Reply message, the mobility security association for verifying the correctness of the authentication data is selected by the foreign agent based on the value of the Home Agent Address field of the Registration Reply.

If the Care-of Address in the Registration Request is not in the Agent Advertisement, then the foreign agent MUST NOT append the Foreign-Home Authentication Extension when relaying the message to the home agent. Moreover, for a deregistration message (i.e., lifetime = 0), the foreign agent MUST NOT append the Foreign-Home Authentication Extension when relaying the message to the home agent. Consequently, when the HA receives a deregistration request that does not contain a Foreign-Home Authentication Extension it MUST NOT for this reason discard the request as part of security association processing.

3.6. Mobile Node Considerations

A mobile node MUST be configured (statically or dynamically) with a netmask and a mobility security association for each of its home agents. In addition, a mobile node MAY be configured with its home address, and the IP address of one or more of its home agents; otherwise, the mobile node MAY discover a home agent using the procedures described in Section 3.6.1.2.

If the mobile node is not configured with a home address, it MAY use the Mobile Node NAI extension [2] to identify itself, and set the Home Address field of the Registration Request to 0.0.0.0. In this case, the mobile node MUST be able to assign its home address after extracting this information from the Registration Reply from the home agent.

For each pending registration, the mobile node maintains the following information:

- o the link-layer address of the foreign agent to which the Registration Request was sent, if applicable,
- o the IP destination address of the Registration Request,
- o the care-of address used in the registration,
- o the Identification value sent in the registration,
- o the originally requested Lifetime, and
- o the remaining Lifetime of the pending registration.

A mobile node SHOULD initiate a registration whenever it detects a

change in its network connectivity. See Section 2.4.2 for methods by which mobile nodes MAY make such a determination. When it is away from home, the mobile node's Registration Request allows its home agent to create or modify a mobility binding for it. When it is at home, the mobile node's (de)Registration Request allows its home agent to delete any previous mobility binding(s) for it. A mobile node operates without the support of mobility functions when it is at home.

There are other conditions under which the mobile node SHOULD (re)register with its foreign agent, such as when the mobile node detects that the foreign agent has rebooted (as specified in Section 2.4.4) and when the current registration's Lifetime is near expiration.

In the absence of link-layer indications of changes in point of attachment, Agent Advertisements from new agents SHOULD NOT cause a mobile node to attempt a new registration, if its current registration has not expired and it is still also receiving Agent Advertisements from the foreign agent with which it is currently registered. In the absence of link-layer indications, a mobile node MUST NOT attempt to register more often than once per second.

A mobile node MAY register with a different agent when transport-layer protocols indicate excessive retransmissions. A mobile node MUST NOT consider reception of an ICMP Redirect from a foreign agent that is currently providing service to it as reason to register with a new foreign agent. Within these constraints, the mobile node MAY register again at any time.

Appendix D shows some examples of how the fields in registration messages would be set up in some typical registration scenarios.

3.6.1. Sending Registration Requests

The following sections specify details for the values the mobile node MUST supply in the fields of Registration Request messages.

3.6.1.1. IP Fields

This section provides the specific rules by which mobile nodes pick values for the IP header fields of a Registration Request.

IP Source Address:

- o When registering on a foreign network with a co-located care-of address, the IP source address MUST be the care-of address.

- o Otherwise, if the mobile node does not have a home address, the IP source address MUST be 0.0.0.0.
- o In all other circumstances, the IP source address MUST be the mobile node's home address.

IP Destination Address:

- o When the mobile node has discovered the agent with which it is registering, through some means (e.g., link-layer) that does not provide the IP address of the agent (the IP address of the agent is unknown to the mobile node), then the "All Mobility Agents" multicast address (224.0.0.11) MUST be used. In this case, the mobile node MUST use the agent's link-layer unicast address in order to deliver the datagram to the correct agent.
- o When registering with a foreign agent, the address of the agent as learned from the IP source address of the corresponding Agent Advertisement MUST be used. This MAY be an address which does not appear as an advertised care-of address in the Agent Advertisement. In addition, when transmitting this Registration Request message, the mobile node MUST use a link-layer destination address copied from the link-layer source address of the Agent Advertisement message in which it learned this foreign agent's IP address.
- o When the mobile node is registering directly with its home agent and knows the (unicast) IP address of its home agent, the destination address MUST be set to this address.
- o If the mobile node is registering directly with its home agent, but does not know the IP address of its home agent, the mobile node may use dynamic home agent address resolution to automatically determine the IP address of its home agent (Section 3.6.1.2). In this case, the IP destination address is set to the subnet-directed broadcast address of the mobile node's home network. This address MUST NOT be used as the destination IP address if the mobile node is registering via a foreign agent, although it MAY be used as the Home Agent address in the body of the Registration Request when registering via a foreign agent.

IP Time to Live:

- o The IP TTL field MUST be set to 1 if the IP destination address is set to the "All Mobility Agents" multicast address as described above. Otherwise a suitable value should be chosen in accordance with standard IP practice [18].

3.6.1.2. Registration Request Fields

This section provides specific rules by which mobile nodes pick values for the fields within the fixed portion of a Registration Request.

A mobile node MAY set the 'S' bit in order to request that the home agent maintain prior mobility binding(s). Otherwise, the home agent deletes any previous binding(s) and replaces them with the new binding specified in the Registration Request. Multiple simultaneous mobility bindings are likely to be useful when a mobile node using at least one wireless network interface moves within wireless transmission range of more than one foreign agent. IP explicitly allows duplication of datagrams. When the home agent allows simultaneous bindings, it will tunnel a separate copy of each arriving datagram to each care-of address, and the mobile node will receive multiple copies of datagrams destined to it.

The mobile node SHOULD set the 'D' bit if it is registering with a co-located care-of address. Otherwise, the 'D' bit MUST NOT be set.

A mobile node MAY set the 'B' bit to request its home agent to forward to it, a copy of broadcast datagrams received by its home agent from the home network. The method used by the home agent to forward broadcast datagrams depends on the type of care-of address registered by the mobile node, as determined by the 'D' bit in the mobile node's Registration Request:

- o If the 'D' bit is set, then the mobile node has indicated that it will decapsulate any datagrams tunneled to this care-of address itself (the mobile node is using a co-located care-of address). In this case, to forward such a received broadcast datagram to the mobile node, the home agent MUST tunnel it to this care-of address. The mobile node de-tunnels the received datagram in the same way as any other datagram tunneled directly to it.
- o If the 'D' bit is NOT set, then the mobile node has indicated that it is using a foreign agent care-of address, and that the foreign agent will thus decapsulate arriving datagrams before forwarding them to the mobile node. In this case, to forward such a received broadcast datagram to the mobile node, the home agent MUST first encapsulate the broadcast datagram in a unicast datagram addressed to the mobile node's home address, and then MUST tunnel this resulting datagram to the mobile node's care-of address.

When decapsulated by the foreign agent, the inner datagram will thus be a unicast IP datagram addressed to the mobile node, identifying to the foreign agent the intended destination of the

encapsulated broadcast datagram, and will be delivered to the mobile node in the same way as any tunneled datagram arriving for the mobile node. The foreign agent **MUST NOT** decapsulate the encapsulated broadcast datagram and **MUST NOT** use a local network broadcast to transmit it to the mobile node. The mobile node thus **MUST** decapsulate the encapsulated broadcast datagram itself, and thus **MUST NOT** set the 'B' bit in its Registration Request in this case unless it is capable of decapsulating datagrams.

The mobile node **MAY** request alternative forms of encapsulation by setting the 'M' bit and/or the 'G' bit, but only if the mobile node is decapsulating its own datagrams (the mobile node is using a co-located care-of address) or if its foreign agent has indicated support for these forms of encapsulation by setting the corresponding bits in the Mobility Agent Advertisement Extension of an Agent Advertisement received by the mobile node. Otherwise, the mobile node **MUST NOT** set these bits.

The Lifetime field is chosen as follows:

- o If the mobile node is registering with a foreign agent, the Lifetime **SHOULD NOT** exceed the value in the Registration Lifetime field of the Agent Advertisement message received from the foreign agent. When the method by which the care-of address is learned does not include a Lifetime, the default ICMP Router Advertisement Lifetime (1800 seconds) **MAY** be used.
- o The mobile node **MAY** ask a home agent to delete a particular mobility binding, by sending a Registration Request with the care-of address for this binding, with the Lifetime field set to zero (Section 3.8.2).
- o Similarly, a Lifetime of zero is used when the mobile node deregisters all care-of addresses, such as upon returning home.

The Home Address field **MUST** be set to the mobile node's home address, if this information is known. Otherwise, the Home Address **MUST** be set to zeroes.

The Home Agent field **MUST** be set to the address of the mobile node's home agent, if the mobile node knows this address. Otherwise, the mobile node **MAY** use dynamic home agent address resolution to learn the address of its home agent. In this case, the mobile node **MUST** set the Home Agent field to the subnet-directed broadcast address of the mobile node's home network. Each home agent receiving such a Registration Request with a broadcast destination address **MUST** reject the mobile node's registration and **SHOULD** return a rejection Registration Reply indicating its unicast IP address for use by the

mobile node in a future registration attempt.

The Care-of Address field MUST be set to the value of the particular care-of address that the mobile node wishes to (de)register. In the special case in which a mobile node wishes to deregister all care-of addresses, it MUST set this field to its home address.

The mobile node chooses the Identification field in accordance with the style of replay protection it uses with its home agent. This is part of the mobility security association the mobile node shares with its home agent. See Section 5.7 for the method by which the mobile node computes the Identification field.

3.6.1.3. Extensions

This section describes the ordering of any mandatory and any optional Extensions that a mobile node appends to a Registration Request. This ordering is REQUIRED:

- a. The IP header, followed by the UDP header, followed by the fixed-length portion of the Registration Request, followed by
- b. If present, any non-authentication Extensions expected to be used by the home agent or other authorizing agent (which may or may not also be useful to the foreign agent), followed by
- c. All authorization-enabling extensions (see Section 1.6), followed by
- d. If present, any non-authentication Extensions used only by the foreign agent, followed by
- e. The Mobile-Foreign Authentication Extension, if present.

Note that items (a) and (c) MUST appear in every Registration Request sent by the mobile node. Items (b), (d), and (e) are optional. However, item (e) MUST be included when the mobile node and the foreign agent share a mobility security association.

3.6.2. Receiving Registration Replies

Registration Replies will be received by the mobile node in response to its Registration Requests. Registration Replies generally fall into three categories:

- o the registration was accepted,

- o the registration was denied by the foreign agent, or
- o the registration was denied by the home agent.

The remainder of this section describes the Registration Reply handling by a mobile node in each of these three categories.

3.6.2.1. Validity Checks

Registration Replies with an invalid, non-zero UDP checksum MUST be silently discarded.

In addition, the low-order 32 bits of the Identification field in the Registration Reply MUST be compared to the low-order 32 bits of the Identification field in the most recent Registration Request sent to the replying agent. If they do not match, the Reply MUST be silently discarded.

Also, the Registration Reply MUST be checked for presence of an authorization-enabling extension. For all Registration Reply messages containing a Status Code indicating status from the Home Agent, the mobile node MUST check for the presence of an authorization-enabling extension, acting in accordance with the Code field in the Reply. The rules are as follows:

- a. If the mobile node and the foreign agent share a mobility security association, exactly one Mobile-Foreign Authentication Extension MUST be present in the Registration Reply, and the mobile node MUST check the Authenticator value in the Extension. If no Mobile-Foreign Authentication Extension is found, or if more than one Mobile-Foreign Authentication Extension is found, or if the Authenticator is invalid, the mobile node MUST silently discard the Reply and SHOULD log the event as a security exception.
- b. If the Code field indicates that service is denied by the home agent, or if the Code field indicates that the registration was accepted by the home agent, exactly one Mobile-Home Authentication Extension MUST be present in the Registration Reply, and the mobile node MUST check the Authenticator value in the Extension. If the Registration Reply was generated by the home agent but no Mobile-Home Authentication Extension is found, or if more than one Mobile-Home Authentication Extension is found, or if the Authenticator is invalid, the mobile node MUST silently discard the Reply and SHOULD log the event as a security exception.

If the Code field indicates an authentication failure, either at the foreign agent or the home agent, then it is quite possible that any

authenticators in the Registration Reply will also be in error. This could happen, for example, if the shared secret between the mobile node and home agent was erroneously configured. The mobile node SHOULD log such errors as security exceptions.

3.6.2.2. Registration Request Accepted

If the Code field indicates that the request has been accepted, the mobile node SHOULD configure its routing table appropriately for its current point of attachment (Section 4.2.1).

If the mobile node is returning to its home network and that network is one which implements ARP, the mobile node MUST follow the procedures described in Section 4.6 with regard to ARP, proxy ARP, and gratuitous ARP.

If the mobile node has registered on a foreign network, it SHOULD re-register before the expiration of the Lifetime of its registration. As described in Section 3.6, for each pending Registration Request, the mobile node MUST maintain the remaining lifetime of this pending registration, as well as the original Lifetime from the Registration Request. When the mobile node receives a valid Registration Reply, the mobile node MUST decrease its view of the remaining lifetime of the registration by the amount by which the home agent decreased the originally requested Lifetime. This procedure is equivalent to the mobile node starting a timer for the granted Lifetime at the time it sent the Registration Request, even though the granted Lifetime is not known to the mobile node until the Registration Reply is received. Since the Registration Request is certainly sent before the home agent begins timing the registration Lifetime (also based on the granted Lifetime), this procedure ensures that the mobile node will re-register before the home agent expires and deletes the registration, in spite of possibly non-negligible transmission delays for the original Registration Request and Reply that started the timing of the Lifetime at the mobile node and its home agent.

3.6.2.3. Registration Request Denied

If the Code field indicates that service is being denied, the mobile node SHOULD log the error. In certain cases the mobile node may be able to "repair" the error. These include:

Code 69: (Denied by foreign agent, Lifetime too long)

In this case, the Lifetime field in the Registration Reply will contain the maximum Lifetime value which that foreign agent is willing to accept in any Registration Request. The mobile node MAY attempt to register with this same agent, using a Lifetime in

the Registration Request that MUST be less than or equal to the value specified in the Reply.

Code 133: (Denied by home agent, Identification mismatch)

In this case, the Identification field in the Registration Reply will contain a value that allows the mobile node to synchronize with the home agent, based upon the style of replay protection in effect (Section 5.7). The mobile node MUST adjust the parameters it uses to compute the Identification field based upon the information in the Registration Reply, before issuing any future Registration Requests.

Code 136: (Denied by home agent, Unknown home agent address)

This code is returned by a home agent when the mobile node is performing dynamic home agent address resolution as described in Section 3.6.1.1 and Section 3.6.1.2. In this case, the Home Agent field within the Reply will contain the unicast IP address of the home agent returning the Reply. The mobile node MAY then attempt to register with this home agent in future Registration Requests. In addition, the mobile node SHOULD adjust the parameters it uses to compute the Identification field based upon the corresponding field in the Registration Reply, before issuing any future Registration Requests.

3.6.3. Registration Retransmission

When no Registration Reply has been received within a reasonable time, another Registration Request MAY be transmitted. When timestamps are used, a new registration Identification is chosen for each retransmission; thus it counts as a new registration. When nonces are used, the unanswered Request is retransmitted unchanged; thus the retransmission does not count as a new registration (Section 5.7). In this way a retransmission will not require the home agent to resynchronize with the mobile node by issuing another nonce in the case in which the original Registration Request (rather than its Registration Reply) was lost by the network.

The maximum time until a new Registration Request is sent SHOULD be no greater than the requested Lifetime of the Registration Request. The minimum value SHOULD be large enough to account for the size of the messages, twice the round trip time for transmission to the home agent, and at least an additional 100 milliseconds to allow for processing the messages before responding. The round trip time for transmission to the home agent will be at least as large as the time required to transmit the messages at the link speed of the mobile node's current point of attachment. Some circuits add another 200

milliseconds of satellite delay in the total round trip time to the home agent. The minimum time between Registration Requests MUST NOT be less than 1 second. Each successive retransmission timeout period SHOULD be at least twice the previous period, as long as that is less than the maximum as specified above.

3.7. Foreign Agent Considerations

The foreign agent plays a mostly passive role in Mobile IP registration. It relays Registration Requests between mobile nodes and home agents, and, when it provides the care-of address, decapsulates datagrams for delivery to the mobile node. It SHOULD also send periodic Agent Advertisement messages to advertise its presence as described in Section 2.3, if not detectable by link-layer means.

A foreign agent MUST NOT transmit a Registration Request except when relaying a Registration Request received from a mobile node, to the mobile node's home agent. A foreign agent MUST NOT transmit a Registration Reply except when relaying a Registration Reply received from a mobile node's home agent, or when replying to a Registration Request received from a mobile node in the case in which the foreign agent is denying service to the mobile node. In particular, a foreign agent MUST NOT generate a Registration Request or Reply because a mobile node's registration Lifetime has expired. A foreign agent also MUST NOT originate a Registration Request message that asks for deregistration of a mobile node; however, it MUST relay well-formed (de)Registration Requests originated by a mobile node.

3.7.1. Configuration and Registration Tables

Each foreign agent MUST be configured with a care-of address. In addition, for each pending or current registration the foreign agent MUST maintain a visitor list entry containing the following information obtained from the mobile node's Registration Request:

- o the link-layer source address of the mobile node
- o the IP Source Address (the mobile node's Home Address) or its co-located care-of address (see description of the 'R' bit in Section 2.1.1)
- o the IP Destination Address (as specified in Section 3.6.1.1)
- o the UDP Source Port
- o the Home Agent address
- o the Identification field
- o the requested registration Lifetime, and
- o the remaining Lifetime of the pending or current registration.

If there is an NAI extension in the Registration Request message

(often, for example, when the mobile node's Home Address is zero), then the foreign agent MUST follow the procedures specified in RFC 2794 [2]. In particular, if the foreign agent cannot manage pending registration request records with such a zero Home Address for the mobile node, the foreign agent MUST return a Registration Reply with Code indicating NONZERO_HOMEADDR_REQD (see [2]).

The foreign agent MAY configure a maximum number of pending registrations that it is willing to maintain (typically 5). Additional registrations SHOULD then be rejected by the foreign agent with code 66. The foreign agent MAY delete any pending Registration Request after the request has been pending for more than 7 seconds; in this case, the foreign agent SHOULD reject the Request with code 78 (registration timeout).

As with any node on the Internet, a foreign agent MAY also share mobility security associations with any other nodes. When relaying a Registration Request from a mobile node to its home agent, if the foreign agent shares a mobility security association with the home agent, it MUST add a Foreign-Home Authentication Extension to the Request. In this case, when the Registration Reply has nonzero lifetime, the foreign agent MUST check the required Foreign-Home Authentication Extension in the Registration Reply from the home agent (Section 3.3 and Section 3.4). Similarly, when receiving a Registration Request from a mobile node, if the foreign agent shares a mobility security association with the mobile node, it MUST check the required Mobile-Foreign Authentication Extension in the Request and MUST add a Mobile-Foreign Authentication Extension to the Registration Reply to the mobile node.

3.7.2. Receiving Registration Requests

If the foreign agent accepts a Registration Request from a mobile node, it checks to make sure that the indicated home agent address does not belong to any network interface of the foreign agent. If not, the foreign agent then MUST relay the Request to the indicated home agent. Otherwise, if the foreign agent denies the Request, it MUST send a Registration Reply to the mobile node with an appropriate denial Code, except in cases where the foreign agent would be required to send out more than one such denial per second to the same mobile node. The following sections describe this behavior in more detail.

If the foreign agent has configured one of its network interfaces with the IP address specified by the mobile node as its home agent address, the foreign agent MUST NOT forward the request again. If the foreign agent serves the mobile node as a home agent, the foreign agent follows the procedures specified in Section 3.8.2. Otherwise,

if the foreign agent does not serve the mobile node as a home agent, the foreign agent rejects the Registration Request with code TBD-IANA (Invalid Home Agent Address).

If a foreign agent receives a Registration Request from a mobile node in its visitor list, the existing visitor list entry for the mobile node SHOULD NOT be deleted or modified until the foreign agent receives a valid Registration Reply from the home agent with a Code indicating success. The foreign agent MUST record the new pending Request as a separate part of the existing visitor list entry for the mobile node. If the Registration Request requests deregistration, the existing visitor list entry for the mobile node SHOULD NOT be deleted until the foreign agent has received a successful Registration Reply. If the Registration Reply indicates that the Request (for registration or deregistration) was denied by the home agent, the existing visitor list entry for the mobile node MUST NOT be modified as a result of receiving the Registration Reply.

3.7.2.1. Validity Checks

Registration Requests with an invalid, non-zero UDP checksum MUST be silently discarded. Requests with non-zero bits in reserved fields MUST be rejected with code 70 (poorly formed request). Requests with the 'D' bit set to 0, nonzero lifetime, and specifying a care-of address not offered by the foreign agent, MUST be rejected with code 77 (invalid care-of address).

Also, the authentication in the Registration Request MUST be checked. If the foreign agent and the mobile node share a mobility security association, exactly one Mobile-Foreign Authentication Extension MUST be present in the Registration Request, and the foreign agent MUST check the Authenticator value in the Extension. If no Mobile-Foreign Authentication Extension is found, or if more than one Mobile-Foreign Authentication Extension is found, or if the Authenticator is invalid, the foreign agent MUST silently discard the Request and SHOULD log the event as a security exception. The foreign agent also SHOULD send a Registration Reply to the mobile node with Code 67.

3.7.2.2. Forwarding a Valid Request to the Home Agent

If the foreign agent accepts the mobile node's Registration Request, it MUST relay the Request to the mobile node's home agent as specified in the Home Agent field of the Registration Request. The foreign agent MUST NOT modify any of the fields beginning with the fixed portion of the Registration Request up through and including the Mobile-Home Authentication Extension or other authentication extension supplied by the mobile node as an authorization-enabling extension for the home agent. Otherwise, an authentication failure

is very likely to occur at the home agent. In addition, the foreign agent proceeds as follows:

- o It MUST process and remove any extensions which do not precede any authorization-enabling extension.
- o It MAY append any of its own non-authentication Extensions of relevance to the home agent, if applicable, and
- o If the foreign agent shares a mobility security association with the home agent, and the Request has lifetime != 0, then it MUST append the Foreign-Home Authentication Extension,

Specific fields within the IP header and the UDP header of the relayed Registration Request MUST be set as follows:

IP Source Address

The care-of address offered by the foreign agent for the mobile node sending the Registration Request.

IP Destination Address

Copied from the Home Agent field within the Registration Request.

UDP Source Port

variable

UDP Destination Port

434

After forwarding a valid Registration Request to the home agent, the foreign agent MUST begin timing the remaining lifetime of the pending registration based on the Lifetime in the Registration Request. If this lifetime expires before receiving a valid Registration Reply, the foreign agent MUST delete its visitor list entry for this pending registration.

3.7.2.3. Denying Invalid Requests

If the foreign agent denies the mobile node's Registration Request for any reason, it SHOULD send the mobile node a Registration Reply with a suitable denial Code. In such a case, the Home Address, Home Agent, and Identification fields within the Registration Reply are copied from the corresponding fields of the Registration Request.

If the Reserved field is nonzero, the foreign agent MUST deny the Request and SHOULD return a Registration Reply with status code 70 to

the mobile node. If the Request is being denied because the requested Lifetime is too long, the foreign agent sets the Lifetime in the Reply to the maximum Lifetime value it is willing to accept in any Registration Request, and sets the Code field to 69. Otherwise, the Lifetime SHOULD be copied from the Lifetime field in the Request.

Specific fields within the IP header and the UDP header of the Registration Reply MUST be set as follows:

IP Source Address

Copied from the IP Destination Address of Registration Request, unless the "All Agents Multicast" address was used. In this case, the foreign agent's address (on the interface from which the message will be sent) MUST be used.

IP Destination Address

If the Registration Reply is generated by the Foreign Agent in order to reject a mobile node's Registration Request, and the Registration Request contains a Home Address which is not 0.0.0.0, then the IP Destination Address is copied from the Home Address field of the Registration Request. Otherwise, if the Registration Reply is received from the Home Agent, and contains a Home Address which is not 0.0.0.0, then the IP Destination Address is copied from the Home Address field of the Registration Reply. Otherwise, the IP Destination Address of the Registration Reply is set to be 255.255.255.255.

UDP Source Port

434

UDP Destination Port

Copied from the UDP Source Port of the Registration Request.

3.7.3. Receiving Registration Replies

The foreign agent updates its visitor list when it receives a valid Registration Reply from a home agent. It then relays the Registration Reply to the mobile node. The following sections describe this behavior in more detail.

If upon relaying a Registration Request to a home agent, the foreign agent receives an ICMP error message instead of a Registration Reply, then the foreign agent SHOULD send to the mobile node a Registration Reply with an appropriate "Home Agent Unreachable" failure Code

(within the range 80-95, inclusive). See Section 3.7.2.3 for details on building the Registration Reply.

3.7.3.1. Validity Checks

Registration Replies with an invalid, non-zero UDP checksum MUST be silently discarded.

When a foreign agent receives a Registration Reply message, it MUST search its visitor list for a pending Registration Request with the same mobile node home address as indicated in the Reply. If there are multiple entries with the same home address, and if the Registration Reply has the Mobile Node NAI extension [2], the foreign agent MUST use the NAI to disambiguate the pending Registration Requests with the same home address. If no matching pending Request is found, and if the Registration Reply does not correspond with any pending Registration Request with a zero mobile node home address (see Section 3.7.1), the foreign agent MUST silently discard the Reply. The foreign agent MUST also silently discard the Reply if the low-order 32 bits of the Identification field in the Reply do not match those in the Request.

Also, the authentication in the Registration Reply MUST be checked. If the foreign agent and the home agent share a mobility security association, exactly one Foreign-Home Authentication Extension MUST be present in the Registration Reply, and the foreign agent MUST check the Authenticator value in the Extension. If no Foreign-Home Authentication Extension is found, or if more than one Foreign-Home Authentication Extension is found, or if the Authenticator is invalid, the foreign agent MUST silently discard the Reply and SHOULD log the event as a security exception. The foreign agent also MUST reject the mobile node's registration and SHOULD send a Registration Reply to the mobile node with Code 68.

3.7.3.2. Forwarding Replies to the Mobile Node

A Registration Reply which satisfies the validity checks of Section 3.8.2.1 is relayed to the mobile node. The foreign agent MUST also update its visitor list entry for the mobile node to reflect the results of the Registration Request, as indicated by the Code field in the Reply. If the Code indicates that the home agent has accepted the registration and the Lifetime field is nonzero, the foreign agent SHOULD set the Lifetime in the visitor list entry to the minimum of the following two values:

- o the value specified in the Lifetime field of the Registration Reply, and

- o the foreign agent's own maximum value for allowable registration lifetime.

If, instead, the Code indicates that the Lifetime field is zero, the foreign agent MUST delete its visitor list entry for the mobile node. Finally, if the Code indicates that the registration was denied by the home agent, the foreign agent MUST delete its pending registration list entry, but not its visitor list entry, for the mobile node.

The foreign agent MUST NOT modify any of the fields beginning with the fixed portion of the Registration Reply up through and including the Mobile-Home Authentication Extension. Otherwise, an authentication failure is very likely to occur at the mobile node. In addition, the foreign agent SHOULD perform the following additional procedures:

- o It MUST process and remove any Extensions which are not covered by any authorization-enabling extension.
- o It MAY append its own non-authentication Extensions that supply information to the mobile node, if applicable, and
- o It MUST append the Mobile-Foreign Authentication Extension, if the foreign agent shares a mobility security association with the mobile node.

Specific fields within the IP header and the UDP header of the relayed Registration Reply are set according to the same rules specified in Section 3.7.2.3.

After forwarding a valid Registration Reply to the mobile node, the foreign agent MUST update its visitor list entry for this registration as follows. If the Registration Reply indicates that the registration was accepted by the home agent, the foreign agent resets its timer of the lifetime of the registration to the Lifetime granted in the Registration Reply; unlike the mobile node's timing of the registration lifetime as described in Section 3.6.2.2, the foreign agent considers this lifetime to begin when it forwards the Registration Reply message, ensuring that the foreign agent will not expire the registration before the mobile node does. On the other hand, if the Registration Reply indicates that the registration was rejected by the home agent, the foreign agent deletes its visitor list entry for this attempted registration.

3.8. Home Agent Considerations

Home agents play a reactive role in the registration process. The home agent receives Registration Requests from the mobile node (perhaps relayed by a foreign agent), updates its record of the

mobility bindings for this mobile node, and issues a suitable Registration Reply in response to each.

A home agent MUST NOT transmit a Registration Reply except when replying to a Registration Request received from a mobile node. In particular, the home agent MUST NOT generate a Registration Reply to indicate that the Lifetime has expired.

3.8.1. Configuration and Registration Tables

Each home agent MUST be configured with an IP address and with the prefix size for the home network. The home agent MUST be configured with the mobility security association of each authorized mobile node that it is serving as a home agent.

When the home agent accepts a valid Registration Request from a mobile node that it serves as a home agent, the home agent MUST create or modify the entry for this mobile node in its mobility binding list containing:

- o the mobile node's home address
- o the mobile node's care-of address
- o the Identification field from the Registration Reply
- o the remaining Lifetime of the registration

The home agent MAY optionally offer the capability to dynamically associate a home address to a mobile node upon receiving a Registration Request from that mobile node. The method by which a home address is allocated to the mobile node is beyond the scope of this document, but see [2]. After the home agent makes the association of the home address to the mobile node, the home agent MUST put that home address into the Home Address field of the Registration Reply.

The home agent MAY also maintain mobility security associations with various foreign agents. When receiving a Registration Request from a foreign agent, if the home agent shares a mobility security association with the foreign agent, the home agent MUST check the Authenticator in the required Foreign-Home Authentication Extension in the message, based on this mobility security association, unless the Lifetime field equals 0. When processing a Registration Request with Lifetime=0, the HA MAY skip checking for the presence and validity of a Foreign-Home Authentication Extension. Similarly, when sending a Registration Reply to a foreign agent, if the home agent shares a mobility security association with the foreign agent, the home agent MUST include a Foreign-Home Authentication Extension in the message, based on this mobility security association.

3.8.2. Receiving Registration Requests

If the home agent accepts an incoming Registration Request, it MUST update its record of the the mobile node's mobility binding(s) and SHOULD send a Registration Reply with a suitable Code. Otherwise (the home agent has denied the Request), it SHOULD in most cases send a Registration Reply with an appropriate Code specifying the reason the Request was denied. The following sections describe this behavior in more detail. If the home agent does not support broadcasts (see Section 4.3), it MUST ignore the 'B' bit (as opposed to rejecting the Registration Request).

3.8.2.1. Validity Checks

Registration Requests with an invalid, non-zero UDP checksum MUST be silently discarded by the home agent.

The authentication in the Registration Request MUST be checked. This involves the following operations:

- a. The home agent MUST check for the presence of at least one authorization-enabling extension, and ensure that all indicated authentications are carried out. At least one authorization-enabling extension MUST be present in the Registration Request; and the home agent MUST either check the Authenticator value in the extension or verify that the authenticator value has been checked by another agent with which it has a security association.

If the home agent receives a Registration Request from a Mobile Node with which it does not have any security association, the home agent MUST silently discard the Registration Request.

If the home agent receives a Registration Request without any authorization-enabling extension, the home agent MUST silently discard the Registration Request.

If the Authenticator is invalid, the home agent MUST reject the mobile node's registration. Further action to be taken in this case depends upon whether the Request has a valid Foreign-Home authentication extension (as follows):

- * If there is a valid Foreign-Home authentication extension, the home agent MUST send a Registration Reply with Code 131.
- * Otherwise, if there is no Foreign-Home security association, the home agent MAY send a Registration Reply with Code 131. If the home agent sends a Registration Reply, it MUST contain

a valid Mobile-Home Authentication Extension. In constructing the Reply, the home agent SHOULD choose a security association that is likely to exist in the mobile node; for example, this may be an older security association or one with a longer lifetime than the one that was attempted to be used by the mobile node in its Request. Deployments should take care when updating security associations to ensure that there is at least one common security association shared between the mobile node and home agent. In any case of a failed Authenticator, the home agent MUST then discard the Request without further processing and SHOULD log the error as a security exception.

- b. The home agent MUST check that the registration Identification field is correct using the context selected by the SPI within the authorization-enabling extension that the home agent used to authenticate the Mobile Node's Registration Request. See Section 5.7 for a description of how this is performed. If incorrect, the home agent MUST reject the Request and SHOULD send a Registration Reply to the mobile node with Code 133, including an Identification field computed in accordance with the rules specified in Section 5.7. The home agent MUST do no further processing with such a Request, though it SHOULD log the error as a security exception.
- c. If the home agent shares a mobility security association with the foreign agent, and this is a registration request (has non-zero lifetime), the home agent MUST check for the presence of a valid Foreign-Home Authentication Extension. Exactly one Foreign-Home Authentication Extension MUST be present in the Registration Request in this case, and the home agent MUST check the Authenticator value in the Extension. If no Foreign-Home Authentication Extension is found, or if more than one Foreign-Home Authentication Extension is found, or if the Authenticator is invalid, the home agent MUST reject the mobile node's registration and SHOULD send a Registration Reply to the mobile node with Code 132. The home agent MUST then discard the Request and SHOULD log the error as a security exception.
- d. If the home agent and the foreign agent do not share a mobility security association, and the Registration contains a Foreign-Home Authentication Extension, the home agent MUST discard the Request and SHOULD log the error as a security exception.

In addition to checking the authentication in the Registration Request, home agents MUST deny Registration Requests that are sent to the subnet-directed broadcast address of the home network (as opposed to being unicast to the home agent). The home agent MUST discard the

Request and SHOULD returning a Registration Reply with a Code of 136. In this case, the Registration Reply will contain the home agent's unicast address, so that the mobile node can re-issue the Registration Request with the correct home agent address.

Note that some routers change the IP destination address of a datagram from a subnet-directed broadcast address to 255.255.255.255 before injecting it into the destination subnet. In this case, home agents that attempt to pick up dynamic home agent discovery requests by binding a socket explicitly to the subnet-directed broadcast address will not see such packets. Home agent implementors should be prepared for both the subnet-directed broadcast address and 255.255.255.255 if they wish to support dynamic home agent discovery.

3.8.2.2. Accepting a Valid Request

If the Registration Request satisfies the validity checks in Section 3.8.2.1, and the home agent is able to accommodate the Request, the home agent MUST update its mobility binding list for the requesting mobile node and MUST return a Registration Reply to the mobile node. In this case, the Reply Code will be either 0 if the home agent supports simultaneous mobility bindings, or 1 if it does not. See Section 3.8.3 for details on building the Registration Reply message.

The home agent updates its record of the mobile node's mobility bindings as follows, based on the fields in the Registration Request:

- o If the Lifetime is zero and the Care-of Address equals the mobile node's home address, the home agent deletes all of the entries in the mobility binding list for the requesting mobile node. This is how a mobile node requests that its home agent cease providing mobility services.
- o If the Lifetime is zero and the Care-of Address does not equal the mobile node's home address, the home agent deletes only the entry containing the specified Care-of Address from the mobility binding list for the requesting mobile node. Any other active entries containing other care-of addresses will remain active.
- o If the Lifetime is nonzero, the home agent adds an entry containing the requested Care-of Address to the mobility binding list for the mobile node. If the 'S' bit is set and the home agent supports simultaneous mobility bindings, the previous mobility binding entries are retained. Otherwise, the home agent removes all previous entries in the mobility binding list for the mobile node.

In all cases, the home agent MUST send a Registration Reply to the source of the Registration Request, which might indeed be a different foreign agent than that whose care-of address is being (de)registered. If the home agent shares a mobility security association with the foreign agent whose care-of address is being deregistered, and that foreign agent is different from the one which relayed the Registration Request, the home agent MAY additionally send a Registration Reply to the foreign agent whose care-of address is being deregistered. The home agent MUST NOT send such a Reply if it does not share a mobility security association with the foreign agent. If no Reply is sent, the foreign agent's visitor list will expire naturally when the original Lifetime expires.

When a foreign agent relays a deregistration message containing a care-of address that it does not own, it MUST NOT add a Foreign-Home Authentication Extension to that deregistration. See Section 3.5.4 for more details.

The home agent MUST NOT increase the Lifetime above that specified by the mobile node in the Registration Request. However, it is not an error for the mobile node to request a Lifetime longer than the home agent is willing to accept. In this case, the home agent simply reduces the Lifetime to a permissible value and returns this value in the Registration Reply. The Lifetime value in the Registration Reply informs the mobile node of the granted lifetime of the registration, indicating when it SHOULD re-register in order to maintain continued service. After the expiration of this registration lifetime, the home agent MUST delete its entry for this registration in its mobility binding list.

If the Registration Request duplicates an accepted current Registration Request, the new Lifetime MUST NOT extend beyond the Lifetime originally granted. A Registration Request is a duplicate if the home address, care-of address, and Identification fields all equal those of an accepted current registration.

In addition, if the home network implements ARP [16], and the Registration Request asks the home agent to create a mobility binding for a mobile node which previously had no binding (the mobile node was previously assumed to be at home), then the home agent MUST follow the procedures described in Section 4.6 with regard to ARP, proxy ARP, and gratuitous ARP. If the mobile node already had a previous mobility binding, the home agent MUST continue to follow the rules for proxy ARP described in Section 4.6.

3.8.2.3. Denying an Invalid Request

If the Registration Request does not satisfy all of the validity checks in Section 3.8.2.1, or the home agent is unable to accommodate the Request, the home agent SHOULD return a Registration Reply to the mobile node with a Code that indicates the reason for the error. If a foreign agent was involved in relaying the Request, this allows the foreign agent to delete its pending visitor list entry. Also, this informs the mobile node of the reason for the error such that it may attempt to fix the error and issue another Request.

This section lists a number of reasons the home agent might reject a Request, and provides the Code value it should use in each instance. See Section 3.8.3 for additional details on building the Registration Reply message.

Many reasons for rejecting a registration are administrative in nature. For example, a home agent can limit the number of simultaneous registrations for a mobile node, by rejecting any registrations that would cause its limit to be exceeded, and returning a Registration Reply with error code 135. Similarly, a home agent may refuse to grant service to mobile nodes which have entered unauthorized service areas by returning a Registration Reply with a Code of 129.

Requests with non-zero bits in reserved fields MUST be rejected with code 134 (poorly formed request).

3.8.3. Sending Registration Replies

If the home agent accepts a Registration Request, it then MUST update its record of the mobile node's mobility binding(s) and SHOULD send a Registration Reply with a suitable Code. Otherwise (the home agent has denied the Request), it SHOULD in most cases send a Registration Reply with an appropriate Code specifying the reason the Request was denied. The following sections provide additional detail for the values the home agent MUST supply in the fields of Registration Reply messages.

3.8.3.1. IP/UDP Fields

This section provides the specific rules by which home agents pick values for the IP and UDP header fields of a Registration Reply.

IP Source Address

Copied from the IP Destination Address of Registration Request, unless a multicast or broadcast address was used. If the IP Destination Address of the Registration Request was a broadcast or multicast address, the IP Source Address of the Registration Reply MUST be set to the home agent's (unicast) IP address.

IP Destination Address

Copied from the IP Source Address of the Registration Request.

UDP Source Port

Copied from the UDP Destination Port of the Registration Request.

UDP Destination Port

Copied from the UDP Source Port of the Registration Request.

When sending a Registration Reply in response to a Registration Request that requested deregistration of the mobile node (the Lifetime is zero and the Care-of Address equals the mobile node's home address) and in which the IP Source Address was also set to the mobile node's home address (this is the normal method used by a mobile node to deregister when it returns to its home network), the IP Destination Address in the Registration Reply will be set to the mobile node's home address, as copied from the IP Source Address of the Request.

In this case, when transmitting the Registration Reply, the home agent MUST transmit the Reply directly onto the home network as if the mobile node were at home, bypassing any mobility binding list entry that may still exist at the home agent for the destination mobile node. In particular, for a mobile node returning home after being registered with a care-of address, if the mobile node's new Registration Request is not accepted by the home agent, the mobility binding list entry for the mobile node will still indicate that datagrams addressed to the mobile node should be tunneled to the mobile node's registered care-of address; when sending the Registration Reply indicating the rejection of this Request, this existing binding list entry MUST be ignored, and the home agent MUST transmit this Reply as if the mobile node were at home.

3.8.3.2. Registration Reply Fields

This section provides the specific rules by which home agents pick values for the fields within the fixed portion of a Registration

Reply.

The Code field of the Registration Reply is chosen in accordance with the rules specified in the previous sections. When replying to an accepted registration, a home agent SHOULD respond with Code 1 if it does not support simultaneous registrations.

The Lifetime field MUST be copied from the corresponding field in the Registration Request, unless the requested value is greater than the maximum length of time the home agent is willing to provide the requested service. In such a case, the Lifetime MUST be set to the length of time that service will actually be provided by the home agent. This reduced Lifetime SHOULD be the maximum Lifetime allowed by the home agent (for this mobile node and care-of address).

If the Home Address field of the Registration Request is nonzero, it MUST be copied into the Home Address field of the Registration Reply message. If the Home Agent cannot support the specified nonzero unicast address in the Home Address field of the Registration Request, then the Home Agent MUST reject the Registration Request with an error code of 129.

Otherwise, if the Home Address field of the Registration Request is zero as specified in Section 3.6, the home agent SHOULD arrange for the selection of a home address for the mobile node, and insert the selected address into the Home Address field of the Registration Reply message. See [2] for further relevant details in the case where mobile nodes identify themselves using an NAI instead of their IP home address.

If the Home Agent field in the Registration Request contains a unicast address of this home agent, then that field MUST be copied into the Home Agent field of the Registration Reply. Otherwise, the home agent MUST set the Home Agent field in the Registration Reply to its unicast address. In this latter case, the home agent MUST reject the registration with a suitable code (e.g., Code 136) to prevent the mobile node from possibly being simultaneously registered with two or more home agents.

3.8.3.3. Extensions

This section describes the ordering of any required and any optional Mobile IP Extensions that a home agent appends to a Registration Reply. The following ordering MUST be followed:

- a. The IP header, followed by the UDP header, followed by the fixed-length portion of the Registration Reply,

- b. If present, any non-authentication Extensions used by the mobile node (which may or may not also be used by the foreign agent),
- c. The Mobile-Home Authentication Extension,
- d. If present, any non-authentication Extensions used only by the foreign agent, and
- e. The Foreign-Home Authentication Extension, if present.

Note that items (a) and (c) MUST appear in every Registration Reply sent by the home agent. Items (b), (d), and (e) are optional. However, item (e) MUST be included when the home agent and the foreign agent share a mobility security association.

4. Routing Considerations

This section describes how mobile nodes, home agents, and (possibly) foreign agents cooperate to route datagrams to/from mobile nodes that are connected to a foreign network. The mobile node informs its home agent of its current location using the registration procedure described in Section 3. See the protocol overview in Section 1.7 for the relative locations of the mobile node's home address with respect to its home agent, and the mobile node itself with respect to any foreign agent with which it might attempt to register.

4.1. Encapsulation Types

Home agents and foreign agents **MUST** support tunneling datagrams using IP in IP encapsulation [14]. Any mobile node that uses a co-located care-of address **MUST** support receiving datagrams tunneled using IP in IP encapsulation. Minimal encapsulation [15] and GRE encapsulation [13] are alternate encapsulation methods which **MAY** optionally be supported by mobility agents and mobile nodes. The use of these alternative forms of encapsulation, when requested by the mobile node, is otherwise at the discretion of the home agent.

4.2. Unicast Datagram Routing

4.2.1. Mobile Node Considerations

When connected to its home network, a mobile node operates without the support of mobility services. That is, it operates in the same way as any other (fixed) host or router. The method by which a mobile node selects a default router when connected to its home network, or when away from home and using a co-located care-of address, is outside the scope of this document. ICMP Router Advertisement [5] is one such method.

When registered on a foreign network, the mobile node chooses a default router by the following rules:

- o If the mobile node is registered using a foreign agent care-of address, it **MAY** use its foreign agent as a first-hop router. The foreign agent's MAC address can be learned from Agent Advertisement. Otherwise, the mobile node **MUST** choose its default router from among the Router Addresses advertised in the ICMP Router Advertisement portion of that Agent Advertisement message.
- o If the mobile node is registered directly with its home agent using a co-located care-of address, then the mobile node **SHOULD** choose its default router from among those advertised in any ICMP Router Advertisement message that it receives for which its

externally obtained care-of address and the Router Address match under the network prefix. If the mobile node's externally obtained care-of address matches the IP source address of the Agent Advertisement under the network prefix, the mobile node MAY also consider that IP source address as another possible choice for the IP address of a default router. The network prefix MAY be obtained from the Prefix-Lengths Extension in the Router Advertisement, if present. The prefix MAY also be obtained through other mechanisms beyond the scope of this document.

While they are away from the home network, mobile nodes MUST NOT broadcast ARP packets to find the MAC address of another Internet node. Thus, the (possibly empty) list of Router Addresses from the ICMP Router Advertisement portion of the message is not useful for selecting a default router, unless the mobile node has some means not involving broadcast ARP and not specified within this document for obtaining the MAC address of one of the routers in the list. Similarly, in the absence of unspecified mechanisms for obtaining MAC addresses on foreign networks, the mobile node MUST ignore redirects to other routers on foreign networks.

4.2.2. Foreign Agent Considerations

Upon receipt of an encapsulated datagram sent to its advertised care-of address, a foreign agent MUST compare the inner destination address to those entries in its visitor list. When the destination does not match the address of any mobile node currently in the visitor list, the foreign agent MUST NOT forward the datagram without modifications to the original IP header, because otherwise a routing loop is likely to result. The datagram SHOULD be silently discarded. ICMP Destination Unreachable MUST NOT be sent when a foreign agent is unable to forward an incoming tunneled datagram. Otherwise, the foreign agent forwards the decapsulated datagram to the mobile node.

The foreign agent MUST NOT advertise to other routers in its routing domain, nor to any other mobile node, the presence of a mobile router (Section 4.5) or mobile node in its visitor list.

The foreign agent MUST route datagrams it receives from registered mobile nodes. At a minimum, this means that the foreign agent must verify the IP Header Checksum, decrement the IP Time To Live, recompute the IP Header Checksum, and forward such datagrams to a default router.

A foreign agent MUST NOT use broadcast ARP for a mobile node's MAC address on a foreign network. It may obtain the MAC address by copying the information from an Agent Solicitation or a Registration Request transmitted from a mobile node. A foreign agent's ARP cache

for the mobile node's IP address MUST NOT be allowed to expire before the mobile node's visitor list entry expires, unless the foreign agent has some way other than broadcast ARP to refresh its MAC address associated with the mobile node's IP address.

Each foreign agent SHOULD support the mandatory features for reverse tunneling [12].

4.2.3. Home Agent Considerations

The home agent MUST be able to intercept any datagrams on the home network addressed to the mobile node while the mobile node is registered away from home. Proxy and gratuitous ARP MAY be used in enabling this interception, as specified in Section 4.6.

The home agent must examine the IP Destination Address of all arriving datagrams to see if it is equal to the home address of any of its mobile nodes registered away from home. If so, the home agent tunnels the datagram to the mobile node's currently registered care-of address or addresses. If the home agent supports the optional capability of multiple simultaneous mobility bindings, it tunnels a copy to each care-of address in the mobile node's mobility binding list. If the mobile node has no current mobility bindings, the home agent MUST NOT attempt to intercept datagrams destined for the mobile node, and thus will not in general receive such datagrams. However, if the home agent is also a router handling common IP traffic, it is possible that it will receive such datagrams for forwarding onto the home network. In this case, the home agent MUST assume the mobile node is at home and simply forward the datagram directly onto the home network.

For multihomed home agents, the source address in the outer IP header of the encapsulated datagram MUST be the address sent to the mobile node in the home agent field of the registration reply. That is, the home agent cannot use the address of some other network interface as the source address.

See Section 4.1 regarding methods of encapsulation that may be used for tunneling. Nodes implementing tunneling SHOULD also implement the "tunnel soft state" mechanism [14], which allows ICMP error messages returned from the tunnel to correctly be reflected back to the original senders of the tunneled datagrams.

Home agents MUST decapsulate packets addressed to themselves, sent by a mobile node for the purpose of maintaining location privacy, as described in Section 5.5. This feature is also required for support of reverse tunneling [12].

If the Lifetime for a given mobility binding expires before the home agent has received another valid Registration Request for that mobile node, then that binding is deleted from the mobility binding list. The home agent **MUST NOT** send any Registration Reply message simply because the mobile node's binding has expired. The entry in the visitor list of the mobile node's current foreign agent will expire naturally, probably at the same time as the binding expired at the home agent. When a mobility binding's lifetime expires, the home agent **MUST** delete the binding, but it **MUST** retain any other (non-expired) simultaneous mobility bindings that it holds for the mobile node.

When a home agent receives a datagram, intercepted for one of its mobile nodes registered away from home, the home agent **MUST** examine the datagram to check if it is already encapsulated. If so, special rules apply in the forwarding of that datagram to the mobile node:

- o If the inner (encapsulated) Destination Address is the same as the outer Destination Address (the mobile node), then the home agent **MUST** also examine the outer Source Address of the encapsulated datagram (the source address of the tunnel). If this outer Source Address is the same as the mobile node's current care-of address, the home agent **MUST** silently discard that datagram in order to prevent a likely routing loop. If, instead, the outer Source Address is **NOT** the same as the mobile node's current care-of address, then the home agent **SHOULD** forward the datagram to the mobile node. In order to forward the datagram in this case, the home agent **MAY** simply alter the outer Destination Address to the care-of address, rather than re-encapsulating the datagram.
- o Otherwise (the inner Destination Address is **NOT** the same as the outer Destination Address), the home agent **SHOULD** encapsulate the datagram again (nested encapsulation), with the new outer Destination Address set equal to the mobile node's care-of address. That is, the home agent forwards the entire datagram to the mobile node in the same way as any other datagram (encapsulated already or not).

4.3. Broadcast Datagrams

When a home agent receives a broadcast datagram, it **MUST NOT** forward the datagram to any mobile nodes in its mobility binding list other than those that have requested forwarding of broadcast datagrams. A mobile node **MAY** request forwarding of broadcast datagrams by setting the 'B' bit in its Registration Request message (Section 3.3). For each such registered mobile node, the home agent **SHOULD** forward received broadcast datagrams to the mobile node, although it is a matter of configuration at the home agent as to which specific

categories of broadcast datagrams will be forwarded to such mobile nodes.

If the 'D' bit was set in the mobile node's Registration Request message, indicating that the mobile node is using a co-located care-of address, the home agent simply tunnels appropriate broadcast IP datagrams to the mobile node's care-of address. Otherwise (the 'D' bit was NOT set), the home agent first encapsulates the broadcast datagram in a unicast datagram addressed to the mobile node's home address, and then tunnels this encapsulated datagram to the foreign agent. This extra level of encapsulation is required so that the foreign agent can determine which mobile node should receive the datagram after it is decapsulated. When received by the foreign agent, the unicast encapsulated datagram is detunneled and delivered to the mobile node in the same way as any other datagram. In either case, the mobile node must decapsulate the datagram it receives in order to recover the original broadcast datagram.

4.4. Multicast Datagram Routing

As mentioned previously, a mobile node that is connected to its home network functions in the same way as any other (fixed) host or router. Thus, when it is at home, a mobile node functions identically to other multicast senders and receivers. This section therefore describes the behavior of a mobile node that is visiting a foreign network.

In order to receive multicasts, a mobile node MUST join the multicast group in one of two ways. First, a mobile node MAY join the group via a (local) multicast router on the visited subnet. This option assumes that there is a multicast router present on the visited subnet. If the mobile node is using a co-located care-of address, it SHOULD use this address as the source IP address of its IGMP [6] messages. Otherwise, it MAY use its home address.

Alternatively, a mobile node which wishes to receive multicasts MAY join groups via a bi-directional tunnel to its home agent, assuming that its home agent is a multicast router. The mobile node tunnels IGMP messages to its home agent and the home agent forwards multicast datagrams down the tunnel to the mobile node. For packets tunneled to the home agent, the source address in the IP header SHOULD be the mobile node's home address.

The rules for multicast datagram delivery to mobile nodes in this case are identical to those for broadcast datagrams (Section 4.3). Namely, if the mobile node is using a co-located care-of address (the 'D' bit was set in the mobile node's Registration Request), then the home agent SHOULD tunnel the datagram to this care-of address;

otherwise, the home agent MUST first encapsulate the datagram in a unicast datagram addressed to the mobile node's home address and then MUST tunnel the resulting datagram (nested tunneling) to the mobile node's care-of address. For this reason, the mobile node MUST be capable of decapsulating packets sent to its home address in order to receive multicast datagrams using this method.

A mobile node that wishes to send datagrams to a multicast group also has two options: (1) send directly on the visited network; or (2) send via a tunnel to its home agent. Because multicast routing in general depends upon the IP source address, a mobile node which sends multicast datagrams directly on the visited network MUST use a co-located care-of address as the IP source address. Similarly, a mobile node which tunnels a multicast datagram to its home agent MUST use its home address as the IP source address of both the (inner) multicast datagram and the (outer) encapsulating datagram. This second option assumes that the home agent is a multicast router.

4.5. Mobile Routers

A mobile node can be a router that is responsible for the mobility of one or more entire networks moving together, perhaps on an airplane, a ship, a train, an automobile, a bicycle, or a kayak. The nodes connected to a network served by the mobile router may themselves be fixed nodes or mobile nodes or routers. In this document, such networks are called "mobile networks".

A mobile router MAY act as a foreign agent and provide a foreign agent care-of address to mobile nodes connected to the mobile network. Typical routing to a mobile node via a mobile router in this case is illustrated by the following example:

- a. A laptop computer is disconnected from its home network and later attached to a network port in the seat back of an aircraft. The laptop computer uses Mobile IP to register on this foreign network, using a foreign agent care-of address discovered through an Agent Advertisement from the aircraft's foreign agent.
- b. The aircraft network is itself mobile. Suppose the node serving as the foreign agent on the aircraft also serves as the default router that connects the aircraft network to the rest of the Internet. When the aircraft is at home, this router is attached to some fixed network at the airline's headquarters, which is the router's home network. While the aircraft is in flight, this router registers from time to time over its radio link with a series of foreign agents below it on the ground. This router's home agent is a node on the fixed network at the airline's headquarters.

- c. Some correspondent node sends a datagram to the laptop computer, addressing the datagram to the laptop's home address. This datagram is initially routed to the laptop's home network.
- d. The laptop's home agent intercepts the datagram on the home network and tunnels it to the laptop's care-of address, which in this example is an address of the node serving as router and foreign agent on the aircraft. Normal IP routing will route the datagram to the fixed network at the airline's headquarters.
- e. The aircraft router and foreign agent's home agent there intercepts the datagram and tunnels it to its current care-of address, which in this example is some foreign agent on the ground below the aircraft. The original datagram from the correspondent node has now been encapsulated twice: once by the laptop's home agent and again by the aircraft's home agent.
- f. The foreign agent on the ground decapsulates the datagram, yielding a datagram still encapsulated by the laptop's home agent, with a destination address of the laptop's care-of address. The ground foreign agent sends the resulting datagram over its radio link to the aircraft.
- g. The foreign agent on the aircraft decapsulates the datagram, yielding the original datagram from the correspondent node, with a destination address of the laptop's home address. The aircraft foreign agent delivers the datagram over the aircraft network to the laptop's link-layer address.

This example illustrates the case in which a mobile node is attached to a mobile network. That is, the mobile node is mobile with respect to the network, which itself is also mobile (here with respect to the ground). If, instead, the node is fixed with respect to the mobile network (the mobile network is the fixed node's home network), then either of two methods may be used to cause datagrams from correspondent nodes to be routed to the fixed node.

A home agent MAY be configured to have a permanent registration for the fixed node, that indicates the mobile router's address as the fixed host's care-of address. The mobile router's home agent will usually be used for this purpose. The home agent is then responsible for advertising connectivity using normal routing protocols to the fixed node. Any datagrams sent to the fixed node will thus use nested tunneling as described above.

Alternatively, the mobile router MAY advertise connectivity to the entire mobile network using normal IP routing protocols through a bi-directional tunnel to its own home agent. This method avoids the

need for nested tunneling of datagrams.

4.6. ARP, Proxy ARP, and Gratuitous ARP

The use of ARP [16] requires special rules for correct operation when wireless or mobile nodes are involved. The requirements specified in this section apply to all home networks in which ARP is used for address resolution.

In addition to the normal use of ARP for resolving a target node's link-layer address from its IP address, this document distinguishes two special uses of ARP:

- o A Proxy ARP [49] is an ARP Reply sent by one node on behalf of another node which is either unable or unwilling to answer its own ARP Requests. The sender of a Proxy ARP reverses the Sender and Target Protocol Address fields as described in [16], but supplies some configured link-layer address (generally, its own) in the Sender Hardware Address field. The node receiving the Reply will then associate this link-layer address with the IP address of the original target node, causing it to transmit future datagrams for this target node to the node with that link-layer address.
- o A Gratuitous ARP [45] is an ARP packet sent by a node in order to spontaneously cause other nodes to update an entry in their ARP cache. A gratuitous ARP MAY use either an ARP Request or an ARP Reply packet. In either case, the ARP Sender Protocol Address and ARP Target Protocol Address are both set to the IP address of the cache entry to be updated, and the ARP Sender Hardware Address is set to the link-layer address to which this cache entry should be updated. When using an ARP Reply packet, the Target Hardware Address is also set to the link-layer address to which this cache entry should be updated (this field is not used in an ARP Request packet).

In either case, for a gratuitous ARP, the ARP packet MUST be transmitted as a local broadcast packet on the local link. As specified in [16], any node receiving any ARP packet (Request or Reply) MUST update its local ARP cache with the Sender Protocol and Hardware Addresses in the ARP packet, if the receiving node has an entry for that IP address already in its ARP cache. This requirement in the ARP protocol applies even for ARP Request packets, and for ARP Reply packets that do not match any ARP Request transmitted by the receiving node [16].

While a mobile node is registered on a foreign network, its home agent uses proxy ARP [49] to reply to ARP Requests it receives that seek the mobile node's link-layer address. When receiving an ARP

Request, the home agent MUST examine the target IP address of the Request, and if this IP address matches the home address of any mobile node for which it has a registered mobility binding, the home agent MUST transmit an ARP Reply on behalf of the mobile node. After exchanging the sender and target addresses in the packet [49], the home agent MUST set the sender link-layer address in the packet to the link-layer address of its own interface over which the Reply will be sent.

When a mobile node leaves its home network and registers a binding on a foreign network, its home agent uses gratuitous ARP to update the ARP caches of nodes on the home network. This causes such nodes to associate the link-layer address of the home agent with the mobile node's home (IP) address. When registering a binding for a mobile node for which the home agent previously had no binding (the mobile node was assumed to be at home), the home agent MUST transmit a gratuitous ARP on behalf of the mobile node. This gratuitous ARP packet MUST be transmitted as a broadcast packet on the link on which the mobile node's home address is located. Since broadcasts on the local link (such as Ethernet) are typically not guaranteed to be reliable, the gratuitous ARP packet SHOULD be retransmitted a small number of times to increase its reliability.

When a mobile node returns to its home network, the mobile node and its home agent use gratuitous ARP to cause all nodes on the mobile node's home network to update their ARP caches to once again associate the mobile node's own link-layer address with the mobile node's home (IP) address. Before transmitting the (de)Registration Request message to its home agent, the mobile node MUST transmit this gratuitous ARP on its home network as a local broadcast on this link. The gratuitous ARP packet SHOULD be retransmitted a small number of times to increase its reliability, but these retransmissions SHOULD proceed in parallel with the transmission and processing of its (de)Registration Request.

When the mobile node's home agent receives and accepts this (de)Registration Request, the home agent MUST also transmit a gratuitous ARP on the mobile node's home network. This gratuitous ARP also is used to associate the mobile node's home address with the mobile node's own link-layer address. A gratuitous ARP is transmitted by both the mobile node and its home agent, since in the case of wireless network interfaces, the area within transmission range of the mobile node will likely differ from that within range of its home agent. The ARP packet from the home agent MUST be transmitted as a local broadcast on the mobile node's home link, and SHOULD be retransmitted a small number of times to increase its reliability; these retransmissions, however, SHOULD proceed in parallel with the transmission and processing of its (de)Registration

Reply.

While the mobile node is away from home, it MUST NOT transmit any broadcast ARP Request or ARP Reply messages. Finally, while the mobile node is away from home, it MUST NOT reply to ARP Requests in which the target IP address is its own home address unless the ARP Request is unicast by a foreign agent with which the mobile node is attempting to register or a foreign agent with which the mobile node has an unexpired registration. In the latter case, the mobile node MUST use a unicast ARP Reply to respond to the foreign agent. Note that if the mobile node is using a co-located care-of address and receives an ARP Request in which the target IP address is this care-of address, then the mobile node SHOULD reply to this ARP Request. Note also that, when transmitting a Registration Request on a foreign network, a mobile node may discover the link-layer address of a foreign agent by storing the address as it is received from the Agent Advertisement from that foreign agent, but not by transmitting a broadcast ARP Request message.

The specific order in which each of the above requirements for the use of ARP, proxy ARP, and gratuitous ARP are applied, relative to the transmission and processing of the mobile node's Registration Request and Registration Reply messages when leaving home or returning home, are important to the correct operation of the protocol.

To summarize the above requirements, when a mobile node leaves its home network, the following steps, in this order, MUST be performed:

- o The mobile node decides to register away from home, perhaps because it has received an Agent Advertisement from a foreign agent and has not recently received one from its home agent.
- o Before transmitting the Registration Request, the mobile node disables its own future processing of any ARP Requests it may subsequently receive requesting the link-layer address corresponding to its home address, except insofar as necessary to communicate with foreign agents on visited networks.
- o The mobile node transmits its Registration Request.
- o When the mobile node's home agent receives and accepts the Registration Request, it performs a gratuitous ARP on behalf of the mobile node, and begins using proxy ARP to reply to ARP Requests that it receives requesting the mobile node's link-layer address. In the gratuitous ARP, the ARP Sender Hardware Address is set to the link-layer address of the home agent. If, instead, the home agent rejects the Registration Request, no ARP processing

(gratuitous nor proxy) is performed by the home agent.

When a mobile node later returns to its home network, the following steps, in this order, MUST be performed:

- o The mobile node decides to register at home, perhaps because it has received an Agent Advertisement from its home agent.
- o Before transmitting the Registration Request, the mobile node re-enables its own future processing of any ARP Requests it may subsequently receive requesting its link-layer address.
- o The mobile node performs a gratuitous ARP for itself. In this gratuitous ARP, the ARP Sender Hardware Address is set to the link-layer address of the mobile node.
- o The mobile node transmits its Registration Request.
- o When the mobile node's home agent receives and accepts the Registration Request, it stops using proxy ARP to reply to ARP Requests that it receives requesting the mobile node's link-layer address, and then performs a gratuitous ARP on behalf of the mobile node. In this gratuitous ARP, the ARP Sender Hardware Address is set to the link-layer address of the mobile node. If, instead, the home agent rejects the Registration Request, the home agent MUST NOT make any change to the way it performs ARP processing (gratuitous nor proxy) for the mobile node. In this latter case, the home agent should operate as if the mobile node has not returned home, and continue to perform proxy ARP on behalf of the mobile node.

5. Security Considerations

The mobile computing environment is potentially very different from the ordinary computing environment. In many cases, mobile computers will be connected to the network via wireless links. Such links are particularly vulnerable to passive eavesdropping, active replay attacks, and other active attacks.

5.1. Message Authentication Codes

Home agents and mobile nodes **MUST** be able to perform authentication. The default algorithm is HMAC-MD5 [10], with a key size of 128 bits. The foreign agent **MUST** also support authentication using HMAC-MD5 and key sizes of 128 bits or greater, with manual key distribution. Keys with arbitrary binary values **MUST** be supported.

The "prefix+suffix" use of MD5 to protect data and a shared secret is considered vulnerable to attack by the cryptographic community. Where backward compatibility with existing Mobile IP implementations that use this mode is needed, new implementations **SHOULD** include keyed MD5 [19] as one of the additional authentication algorithms for use when producing and verifying the authentication data that is supplied with Mobile IP registration messages, for instance in the extensions specified in Section 3.5.2, Section 3.5.3, and Section 3.5.4.

More authentication algorithms, algorithm modes, key distribution methods, and key sizes **MAY** also be supported for all of these extensions.

5.2. Areas of Security Concern in this Protocol

The registration protocol described in this document will result in a mobile node's traffic being tunneled to its care-of address. This tunneling feature could be a significant vulnerability if the registration were not authenticated. Such remote redirection, for instance as performed by the mobile registration protocol, is widely understood to be a security problem in the current Internet if not authenticated [30]. Moreover, the Address Resolution Protocol (ARP) is not authenticated, and can potentially be used to steal another host's traffic. The use of "Gratuitous ARP" (Section 4.6) brings with it all of the risks associated with the use of ARP.

5.3. Key Management

This specification requires a strong authentication mechanism (keyed MD5) which precludes many potential attacks based on the Mobile IP registration protocol. However, because key distribution is

difficult in the absence of a network key management protocol, messages with the foreign agent are not all required to be authenticated. In a commercial environment it might be important to authenticate all messages between the foreign agent and the home agent, so that billing is possible, and service providers do not provide service to users that are not legitimate customers of that service provider.

5.4. Picking Good Random Numbers

The strength of any authentication mechanism depends on several factors, including the innate strength of the authentication algorithm, the secrecy of the key used, the strength of the key used, and the quality of the particular implementation. This specification requires implementation of keyed MD5 for authentication, but does not preclude the use of other authentication algorithms and modes. For keyed MD5 authentication to be useful, the 128-bit key must be both secret (that is, known only to authorized parties) and pseudo-random. If nonces are used in connection with replay protection, they must also be selected carefully. Eastlake, et al. [8] provides more information on generating pseudo-random numbers.

5.5. Privacy

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

5.6. Ingress Filtering

Many routers implement security policies such as "ingress filtering" [35] that do not allow forwarding of packets that have a Source Address which appears topologically incorrect. In environments where this is a problem, mobile nodes may use reverse tunneling [12] with the foreign agent supplied care-of address as the Source Address. Reverse tunneled packets will be able to pass normally through such routers, while ingress filtering rules will still be able to locate the true topological source of the packet in the same way as packets from non-mobile nodes.

5.7. Replay Protection for Registration Requests

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

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

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

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

5.7.1. Replay Protection using Timestamps

The basic principle of timestamp replay protection is that the node generating a message inserts the current time of day, and the node receiving the message checks that this timestamp is sufficiently close to its own time of day. Unless specified differently in the security association between the nodes, a default value of 7 seconds **MAY** be used to limit the time difference. This value **SHOULD** be greater than 3 seconds. Obviously the two nodes must have adequately synchronized time-of-day clocks. As with any messages, time synchronization messages may be protected against tampering by an authentication mechanism determined by the security context between the two nodes.

If timestamps are used, the mobile node **MUST** set the Identification field to a 64-bit value formatted as specified by the Network Time Protocol [11]. The low-order 32 bits of the NTP format represent fractional seconds, and those bits which are not available from a

time source SHOULD be generated from a good source of randomness. Note, however, that when using timestamps, the 64-bit Identification used in a Registration Request from the mobile node MUST be greater than that used in any previous Registration Request, as the home agent uses this field also as a sequence number. Without such a sequence number, it would be possible for a delayed duplicate of an earlier Registration Request to arrive at the home agent (within the clock synchronization required by the home agent), and thus be applied out of order, mistakenly altering the mobile node's current registered care-of address.

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

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

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

5.7.2. Replay Protection using Nonces

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

The home agent may be expected to have resources for computing pseudo-random numbers useful as nonces [8]. It inserts a new nonce as the high-order 32 bits of the identification field of every Registration Reply. The home agent copies the low-order 32 bits of the Identification from the Registration Request message into the

low-order 32 bits of the Identification in the Registration Reply. When the mobile node receives an authenticated Registration Reply from the home agent, it saves the high-order 32 bits of the identification for use as the high-order 32 bits of its next Registration Request.

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

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

6. IANA Considerations

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

- o Mobile IP message types sent to UDP port 434, as defined in Section 1.8.
- o types of extensions to Registration Request and Registration Reply messages (see Section 3.3 and Section 3.4, and also consult ([12],[43],[2],[3],[7])).
- o values for the Code in the Registration Reply message (see Section 3.4, and also consult ([12],[43],[2],[3],[7])).
- o Mobile IP defines so-called Agent Solicitation and Agent Advertisement messages. These messages are in fact Router Discovery messages [5] augmented with mobile-IP specific extensions. Thus, they do not define a new name space, but do define additional Router Discovery extensions as described below in Section 6.2. Also see Section 2.1 and consult ([3], [7])).

There are additional Mobile IP numbering spaces specified in [3].

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

In this revised specification, a new Code value (for the field in the Registration Reply message) is needed within the range typically used for Foreign Agent messages. This error code is needed to indicate the status "Invalid Home Agent Address". See Section 3.7.2 for details.

6.1. Mobile IP Message Types

Mobile IP messages are defined to be those that are sent to a message recipient at port 434 (UDP or TCP). The number space for Mobile IP messages is specified in Section 1.8. Approval of new extension numbers is subject to Expert Review, and a specification is required [22]. The currently standardized message types have the following numbers, and are specified in the indicated sections.

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

6.2. Extensions to RFC 1256 Router Advertisement

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

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

Approval of new extension numbers for use with Mobile IP is subject to Expert Review, and a specification is required [22].

6.3. Extensions to Mobile IP Registration Messages

The Mobile IP messages, specified within this document, and listed in Section 1.8 and Section 6.1, may have extensions. Mobile IP message extensions all share the same number space, even if they are to be applied to different Mobile IP messages. The number space for Mobile IP message extensions is specified within this document. Approval of new extension numbers is subject to Expert Review, and a specification is required [22].

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

6.4. Code Values for Mobile IP Registration Reply Messages

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

Code #s	Guideline
0-8	Success Codes
9-63	Allocation guidelines not specified in this document
64-127	Error Codes from the Foreign Agent
128-192	Error Codes from the Home Agent
193-200	Error Codes from the Gateway Foreign Agent [29]
201-255	Allocation guidelines not specified in this document

Approval of new Code values requires Expert Review [22].

Table 1: Guidelines for Allocation of Code Values

7. Acknowledgments

Special thanks to Steve Deering (Xerox PARC), along with Dan Duchamp and John Ioannidis (JI) (Columbia University), for forming the working group, chairing it, and putting so much effort into its early development. Columbia's early Mobile IP work can be found in [37],[38],[39].

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

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

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

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

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

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

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

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

Thanks to Vijay Devarapalli, who put in many hours to convert the
source for this text document into XML format.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Calhoun, P. and C. Perkins, "Mobile IP Network Access Identifier Extension for IPv4", RFC 2794, March 2000.
- [3] Perkins, C., Calhoun, P., and J. Bharatia, "Mobile IPv4 Challenge/Response Extensions (Revised)", RFC 4721, January 2007.
- [4] Cong, D., Hamlen, M., and C. Perkins, "The Definitions of Managed Objects for IP Mobility Support using SMIPv2", RFC 2006, October 1996.
- [5] Deering, S., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [6] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [7] Dommety, G. and K. Leung, "Mobile IP Vendor/Organization-Specific Extensions", RFC 3115, April 2001.
- [8] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [9] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [10] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [11] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.
- [12] Montenegro, G., "Reverse Tunneling for Mobile IP, revised", RFC 3024, January 2001.
- [13] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [14] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [15] Perkins, C., "Minimal Encapsulation within IP", RFC 2004,

October 1996.

- [16] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.
- [17] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.
- [18] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [19] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.
- [20] Solomon, J., "Applicability Statement for IP Mobility Support", RFC 2005, October 1996.
- [21] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [22] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

8.2. Informative References

- [23] Solomon, J. and S. Glass, "Mobile-IPv4 Configuration Option for PPP IPCP", RFC 2290, February 1998.
- [24] Montenegro, G., Dawkins, S., Kojo, M., Magret, V., and N. Vaidya, "Long Thin Networks", RFC 2757, January 2000.
- [25] Allman, M., Glover, D., and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", BCP 28, RFC 2488, January 1999.
- [26] Paxson, V. and M. Allman, "Computing TCP's Retransmission Timer", RFC 2988, November 2000.
- [27] Levkowetz, H. and S. Vaarala, "Mobile IP Traversal of Network Address Translation (NAT) Devices", RFC 3519, April 2003.
- [28] Glass, S. and M. Chandra, "Registration Revocation in Mobile IPv4", RFC 3543, August 2003.
- [29] Fogelstroem, E., Jonsson, A., and C. Perkins, "Mobile IPv4 Regional Registration", RFC 4857, June 2007.

- [30] Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, 19(2), March 1989.
- [31] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, June 2001.
- [32] Caceres, R. and L. Iftode, "Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments", IEEE Journal on Selected Areas in Communication, 13(5):850--857, June 1995.
- [33] Dawkins, S., Montenegro, G., Kojo, M., Magret, V., and N. Vaidya, "End-to-end Performance Implications of Links with Errors", BCP 50, RFC 3155, August 2001.
- [34] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [35] 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.
- [36] Jacobson, V., "Compressing TCP/IP headers for low-speed serial links", RFC 1144, February 1990.
- [37] Ioannidis, J., Duchamp, D., and G. Maguire, "IP-Based Protocols for Mobile Interworking", In Proceedings of the SIGCOMM '01 Conference: Communications Architectures and Protocols, Pages 235--245, September 1991.
- [38] Ioannidis, J. and G. Maguire, "The Design and Implementation of a Mobile Internetworking Architecture", In Proceedings of the Winter USENIX Technical Conference, Pages 489--500, January 1993.
- [39] Ioannidis, J., "Protocols for Mobile Interworking", PhD Dissertation - Columbia University in the City of New York , July 1993.
- [40] Jacobson, V., "Congestion Avoidance and Control", In Proceedings of the SIGCOMM '88 Workshop, ACM SIGCOMM, ACM Press, Pages 314--329, August 1998.
- [41] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.
- [42] McGregor, G., "The PPP Internet Protocol Control Protocol

(IPCP)", RFC 1332, May 1992.

- [43] Montenegro, G. and V. Gupta, "Sun's SKIP Firewall Traversal for Mobile IP", RFC 2356, June 1998.
- [44] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [45] Stevens, R., "TCP/IP Illustrated, Volume 1: The Protocols", Addison-Wesley, Reading, Massachusetts, 1994.
- [46] Perkins, C. and P. Calhoun, "Authentication, Authorization, and Accounting (AAA) Registration Keys for Mobile IPv4", RFC 3957, March 2005.
- [47] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [48] IANA Assigned Numbers Online Database, "Mobile IPv4 Numbers", <http://www.iana.org/assignments/mobileip-numbers> .
- [49] Postel, J., "Multi-LAN address resolution", RFC 925, October 1984.

Appendix A. Pre-RFC5378 Disclaimer

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Appendix B. Link-Layer Considerations

The mobile node MAY use link-layer mechanisms to decide that its point of attachment has changed. Such indications include the Down/Testing/Up interface status [41], and changes in cell or administration. The mechanisms will be specific to the particular link-layer technology, and are outside the scope of this document.

The Point-to-Point-Protocol (PPP) [47] and its Internet Protocol Control Protocol (IPCP) [42], negotiates the use of IP addresses.

The mobile node SHOULD first attempt to specify its home address, so that if the mobile node is attaching to its home network, the unrouted link will function correctly. When the home address is not accepted by the peer, but a transient IP address is dynamically assigned to the mobile node, and the mobile node is capable of supporting a co-located care-of address, the mobile node MAY register that address as a co-located care-of address. When the peer specifies its own IP address, that address MUST NOT be assumed to be a foreign agent care-of address or the IP address of a home agent. PPP extensions for Mobile IP have been specified in RFC 2290 [23]. Please consult that document for additional details for how to handle care-of address assignment from PPP in a more efficient manner.

Appendix C. TCP Considerations

C.1. TCP Timers

When high-delay (e.g. SATCOM) or low-bandwidth (e.g. High-Frequency Radio) links are in use, some TCP stacks may have insufficiently adaptive (non-standard) retransmission timeouts. There may be spurious retransmission timeouts, even when the link and network are actually operating properly, but just with a high delay because of the medium in use. This can cause an inability to create or maintain TCP connections over such links, and can also cause unneeded retransmissions which consume already scarce bandwidth. Vendors are encouraged to follow the algorithms in RFC 2988 [26] when implementing TCP retransmission timers. Vendors of systems designed for low-bandwidth, high-delay links should consult RFCs 2757 and 2488 [24], [25]. Designers of applications targeted to operate on mobile nodes should be sensitive to the possibility of timer-related difficulties.

C.2. TCP Congestion Management

Mobile nodes often use media which are more likely to introduce errors, effectively causing more packets to be dropped. This introduces a conflict with the mechanisms for congestion management found in modern versions of TCP [40]. Now, when a packet is dropped, the correspondent node's TCP implementation is likely to react as if there were a source of network congestion, and initiate the slow-start mechanisms [40] designed for controlling that problem. However, those mechanisms are inappropriate for overcoming errors introduced by the links themselves, and have the effect of magnifying the discontinuity introduced by the dropped packet. This problem has been analyzed by Caceres, et al. [32]. TCP approaches to the problem of handling errors that might interfere with congestion management are discussed in documents from the [pilc] working group [31] [33]. While such approaches are beyond the scope of this document, they illustrate that providing performance transparency to mobile nodes involves understanding mechanisms outside the network layer. Problems introduced by higher media error rates also indicate the need to avoid designs which systematically drop packets; such designs might otherwise be considered favorably when making engineering tradeoffs.

Appendix D. Example Scenarios

This section shows example Registration Requests for several common scenarios.

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

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

IP fields:

Source Address = mobile node's home address

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

Time to Live = 1

UDP fields:

Source Port = <any>

Destination Port = 434

Registration Request fields:

Type = 1

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

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

Home Address = the mobile node's home address

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

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

Identification = Network Time Protocol timestamp or Nonce

Extensions:

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

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

The mobile node enters a foreign network that contains no foreign agents. The mobile node obtains an address from a DHCP server [34] for use as a co-located care-of address. The mobile node supports all forms of encapsulation (IP-in-IP, minimal encapsulation, and GRE), desires a copy of broadcast datagrams on the home network, and does not want simultaneous mobility bindings:

IP fields:

Source Address = care-of address obtained from DHCP server

Destination Address = IP address of home agent

Time to Live = 64

UDP fields:

Source Port = <any>

Destination Port = 434

Registration Request fields:

Type = 1

S=0,B=1,D=1,M=1,G=1

Lifetime = 1800 (seconds)

Home Address = the mobile node's home address

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

Care-of Address = care-of address obtained from DHCP server

Identification = Network Time Protocol timestamp or Nonce

Extensions:

The Mobile-Home Authentication Extension

D.3. Deregistration

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

IP fields:

Source Address = mobile node's home address

Destination Address = IP address of home agent

Time to Live = 1

UDP fields:

Source Port = <any>

Destination Port = 434

Registration Request fields:

Type = 1

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

Lifetime = 0

Home Address = the mobile node's home address

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

Care-of Address = the mobile node's home address

Identification = Network Time Protocol timestamp or Nonce

Extensions:

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

Appendix E. Applicability of Prefix-Lengths Extension

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

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

Appendix F. Interoperability Considerations

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

However, this specification does have new features that, if used, would cause interoperability problems with older implementations. All features specified in RFC 2002 will work with the new implementations, except for V-J compression [36]. The following list details some of the possible areas of compatibility problems that may be experienced by nodes conforming to this revised specification, when attempting to interoperate with nodes obeying RFC 2002.

- o A client that expects some of the newly mandatory features (like reverse tunneling) from a foreign agent would still be interoperable as long as it pays attention to the 'T' bit.
- o Mobile nodes that use the NAI extension to identify themselves would not work with old mobility agents.
- o Mobile nodes that use a zero home address and expect to receive their home address in the Registration Reply would not work with old mobility agents.
- o Mobile nodes that attempt to authenticate themselves without using the Mobile-Home authentication extension will be unable to successfully register with their home agent.

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

Appendix G. Changes since RFC 3344

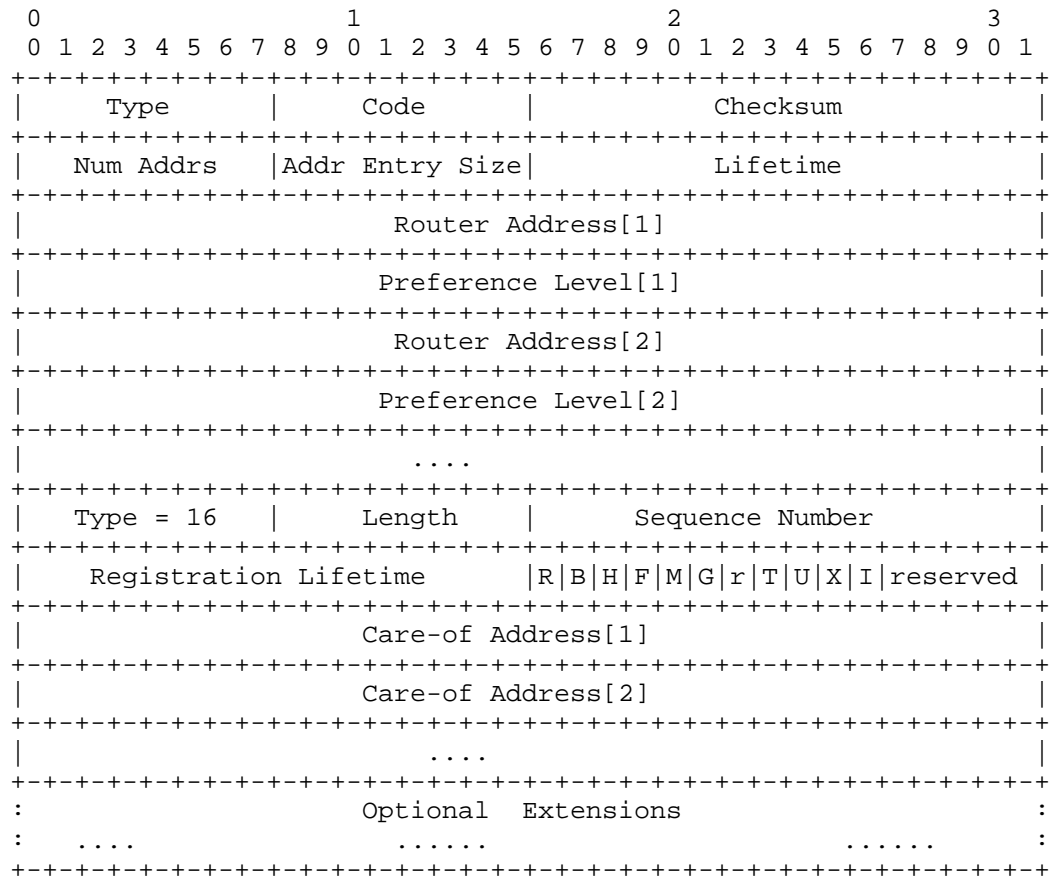
The following revisions to details of the specification in this document were made after RFC 3344 was published. A list of changes from RFC 2002 made during the development of RFC 3344 [21] may be found in the latter document. For items marked with issue numbers, more information is available by consulting the MIP4 mailing list archives.

- o Showed more bit definitions in the Agent Advertisement message structure (see Section 2.1.1). New advertisement bits have been defined by other specification documents, but not reflected in previous publications of this specification; this has led to confusion. Citations for the other specification documents have also been included.
- o (Issue 6) The behavior of the home agent was changed to avoid mandating error replies to Registration Requests that were invalidated because the foreign agent failed authentication. The intention is to make the home agent more robust against Denial of Service attacks in which the malicious device has no intention of providing a valid registration request but only wants to congest traffic on the home network. See section Section 3.8.2.1.
- o Due to non-unique assignment of IPv4 addresses in many domains, it is possible for different mobile nodes to have the same home address. If they use the NAI, the foreign agent can still distinguish them. Language was added to Section 3.7.1 and Section 3.7.3.1 to specify that the foreign agent MUST use the NAI to distinguish mobile nodes with the same home address.
- o (Issue 45) Specified that a foreign agent MUST NOT apply a Foreign-Home Authentication extension to a mobile node's deregistration request. Also, the foreign agent MUST NOT apply a Foreign-Home Authentication extension unless Care-of Address in the Registration Request matches an address advertised by the foreign agent.
- o Specified that the mobility security association to be used by the Foreign Agent and Home Agent depends upon values contained in the message data, not the IP headers.
- o (Issues 9, 18) Created a new error code for use by the foreign agent, for the case when the foreign agent does not serve the mobile node as a home agent. Formerly, the foreign agent could use error code 136 for this case.

- o (Issue 17) Specified that, if the Home Agent cannot support the requested nonzero unicast address in the Home Address field of the Registration Request, then the it MUST reject the registration with an error code of 129. See section Section 3.8.3.2.
- o (Issue 19) Specified that multiple authorization-enabling extensions may be present in the Registration Request message, but that the home agent has to (somehow) ensure that all have been checked (see section Section 3.8.3.1).
- o (Issue 20) Specified that the foreign agent SHOULD NOT modify any of the fields of the Registration Reply message that are covered by the Mobile-Home Authentication Extension, when it relays the packet to the mobile node.
- o (Issue 21) Clarified that the foreign agent removes extensions that do not precede any authorization-enabling extension, not just the Mobile-Home Authentication extension (section 3.7.3.2).
- o (Issue 44) Specified that the address advertised by the foreign agent in Agent Advertisements is the care-of address offered on that network interface, not necessarily the address of the network interface (section 3.7.2.2).
- o (Issue 45) Clarification in section 3.7.2.1 that code 77 can only apply to a Registration Request with nonzero lifetime.
- o Created a new error code for use when a Foreign Agent can detect that the Home Agent address field is incorrect.
- o Prohibited the use of the Foreign-Home Authorization Extension on deregistration messages.
- o Cleaned up some more wording having to do with authorization-enabling extensions.
- o For consistency, changed some wording about copying UDP ports.
- o Added wording to clearly not disallow dynamically configuring netmask and security information at the mobile node.
- o Revamped Changes section.
- o Updated citations.

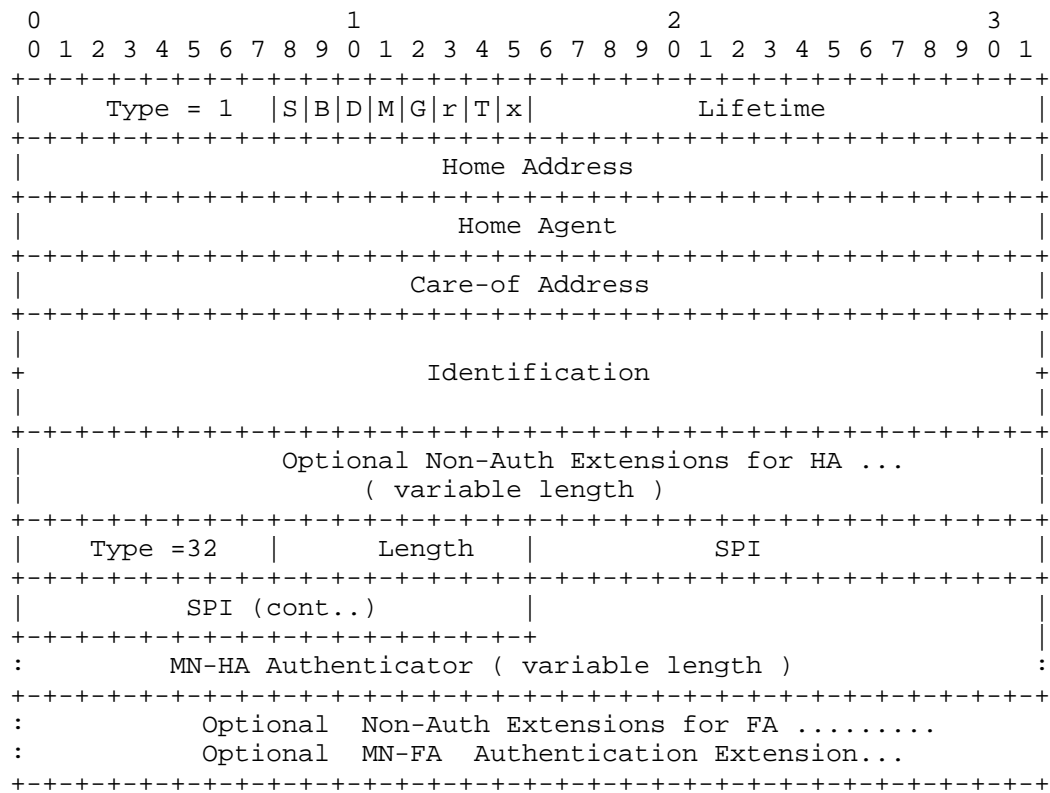
Appendix H. Example Messages

H.1. Example ICMP Agent Advertisement Message Format



H.2. Example Registration Request Message Format

The UDP header is followed by the Mobile IP fields shown below:



H.3. Example Registration Reply Message Format

The UDP header is followed by the Mobile IP fields shown below:

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Type = 3   |      Code      |      Lifetime      |
+-----+-----+-----+-----+-----+-----+-----+
|                                     Home Address          |
+-----+-----+-----+-----+-----+-----+-----+
|                                     Home Agent             |
+-----+-----+-----+-----+-----+-----+-----+
|                                     Identification          |
+-----+-----+-----+-----+-----+-----+-----+
|      Optional HA Non-Auth Extensions ...                |
|      ( variable length )                                |
+-----+-----+-----+-----+-----+-----+-----+
|  Type =32   |      Length   |      SPI           |
+-----+-----+-----+-----+-----+-----+-----+
|      SPI (cont...)                                     |
+-----+-----+-----+-----+-----+-----+-----+
:      MN-HA Authenticator ( variable length )           :
+-----+-----+-----+-----+-----+-----+-----+
:      Optional Extensions used by FA.....              :
:      Optional MN-FA Authentication Extension...          :
+-----+-----+-----+-----+-----+-----+-----+

```

Author's Address

Charles E. Perkins
WiChorus Inc.
3590 N. 1st Street, Suite 300
San Jose, CA 95134
USA

Email: charliep@computer.org

