

Network Working Group
Request for Comments: 4866
Category: Standards Track

J. Arkko
Ericsson Research NomadicLab
C. Vogt
Universitaet Karlsruhe (TH)
W. Haddad
Ericsson Research
May 2007

Enhanced Route Optimization for Mobile IPv6

Status of This Memo

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

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document specifies an enhanced version of Mobile IPv6 route optimization, providing lower handoff delays, increased security, and reduced signaling overhead.

Table of Contents

1. Introduction	3
2. Objectives	4
2.1. Handoff Latency	5
2.2. Security	5
2.3. Signaling Overhead	7
3. Protocol Design	7
3.1. Cryptographically Generated Home Addresses	7
3.2. Non-Cryptographic Care-of Addresses	8
3.3. Semi-Permanent Security Associations	8
3.4. Initial Home Address Tests	8
3.5. Concurrent Care-of Address Tests	9
3.6. Credit-Based Authorization	9
3.7. Parallel Home and Correspondent Registrations	10
4. Protocol Operation	10
4.1. Sending Binding Update Messages	10
4.2. Receiving Binding Update Messages	18
4.3. Sending Binding Acknowledgment Messages	22

4.4. Receiving Binding Acknowledgment Messages	23
4.5. Sending CGA Parameters	25
4.6. Receiving CGA Parameters	26
4.7. Sending Permanent Home Keygen Tokens	27
4.8. Receiving Permanent Home Keygen Tokens	28
4.9. Renewing Permanent Home Keygen Tokens	28
4.10. Handling Payload Packets	28
4.11. Credit Aging	31
4.12. Simultaneous Movements	32
5. Option Formats and Status Codes	32
5.1. CGA Parameters Option	32
5.2. Signature Option	33
5.3. Permanent Home Keygen Token Option	34
5.4. Care-of Test Init Option	35
5.5. Care-of Test Option	35
5.6. CGA Parameters Request Option	36
5.7. Status Codes	36
6. Security Considerations	38
6.1. Home Address Ownership	39
6.2. Care-of Address Ownership	41
6.3. Credit-Based Authorization	43
6.4. Time Shifting Attacks	46
6.5. Replay Attacks	47
6.6. Resource Exhaustion	47
6.7. IP Address Ownership of Correspondent Node	47
7. Protocol Constants and Configuration Variables	49
8. IANA Considerations	50
9. Acknowledgments	50
10. References	51
10.1. Normative References	51
10.2. Informative References	51

1. Introduction

Mobile IPv6 route optimization [1] enables mobile and correspondent nodes to communicate via a direct routing path despite changes in IP connectivity on the mobile node side. Both end nodes use a stable "home address" in identifying the mobile node at stack layers above IP, while payload packets are sent or received via a "care-of address" that routes to the mobile node's current network attachment. Mobile IPv6 swaps the home and care-of addresses when a payload packet traverses the IP layer. The association between a mobile node's home address and care-of address is called a "binding" for the mobile node. It is the responsibility of the mobile node to update its binding at the correspondent node through a "correspondent registration" when it changes IP connectivity. A correspondent registration further involves the mobile node's home agent, which proxies the mobile node at the home address and mainly serves as a relay for payload packets exchanged with correspondent nodes that do not support route optimization. The mobile node keeps the home agent up to date about its current care-of address by means of "home registrations".

From a security perspective, the establishment of a binding during a correspondent registration requires the correspondent node to verify the mobile node's ownership of both the home address and the care-of address. Unprecedented impersonation and flooding threats [5] would arise if correspondent nodes took liberties with respect to these obligations. A correspondent registration hence incorporates a "home address test" and a "care-of address test", collectively called the "return routability procedure". These tests allow the correspondent node to probe the mobile node's reachability at the home and care-of addresses in an ad hoc, non-cryptographic manner. Successful reachability verification at both IP addresses indicates (though it does not guarantee) the mobile node's ownership of the IP addresses, and hence that a binding between the home address and the care-of address is legitimate.

The advantage of the return routability procedure is that it is lightweight and does not depend on a public-key infrastructure or on a preexisting relationship between the mobile node and the correspondent node. This facilitates a broad deployment. On the other hand, the procedure has an adverse impact on handoff delays since both the home address test and the care-of address test consist of an end-to-end message exchange between the mobile node and the correspondent node. The latency of the home address test may be particularly high because it routes through the home agent. The return routability procedure is also vulnerable to attackers that are in a position where they can interpose in the home or care-of address test. The value of interposing is limited in that the return

routability procedure must be repeated in intervals of at most 7 minutes, even in the absence of changes in IP connectivity on the mobile node side. But this comes at the cost of an increased signaling overhead. Much effort has therefore gone into improvements for Mobile IPv6 route optimization [6] that mitigate these disadvantages.

This document specifies Enhanced Route Optimization, an amendment to route optimization in base Mobile IPv6. Enhanced Route Optimization secures a mobile node's home address against impersonation through an interface identifier that is cryptographically and verifiably bound [2] to the public component of the mobile node's public/private-key pair. The mobile node proves ownership of the home address by providing evidence that it knows the corresponding private key. An initial home address test validates the home address prefix; subsequent home address tests are unnecessary. Enhanced Route Optimization further allows mobile and correspondent nodes to resume bidirectional communications in parallel with pursuing a care-of address test. The latency of the home and care-of address tests are therefore eliminated in most cases. The use of cryptographically generated home addresses also mitigates the threat of impersonators that can interpose on the home address test and thereby facilitate longer binding lifetimes. This leads to increased security and a reduction in signaling overhead. Cryptographically generated home addresses and concurrent care-of address tests are preferably applied together, but a mobile node may choose to use only one of these enhancements.

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 [3].

2. Objectives

The design of route optimization in base Mobile IPv6 is in many ways conservative, leaving room to optimize handoff delay, security, and signaling overhead. Enhanced Route Optimization tackles these issues and thus constitutes a more progressive variant of Mobile IPv6.

Despite any Mobile IPv6 optimizations, it is important to take into account that mobility-related activities elsewhere in the protocol stack may have their own impact. For example, attachment procedures, access control, and authentication at the link layer contribute their own handoff delays. So do IP layer tasks such as router discovery, neighbor discovery, movement detection, and IP address configuration. The handoff delays and signaling overhead of Mobile IPv6 are

typically small compared to the total delay and overhead. The improvements of Enhanced Route Optimization hence ought to be seen in view of the entire protocol stack.

2.1. Handoff Latency

The typical handoff delay in base Mobile IPv6 route optimization is one round-trip time between the mobile node and the home agent for the home registration, one round-trip time between the mobile node and the home agent plus one round-trip time between the home agent and the correspondent node for the return routability procedure, and one one-way time from the mobile node to the correspondent node for the propagation of the Binding Update message. (The assumption here is that the latency of the return routability procedure is dominated by the home address test.) The first payload packet sent to the new care-of address requires one additional one-way time to propagate from the correspondent node to the mobile node. The mobile node can resume transmissions right after it has dispatched the Binding Update message. But if it requests a Binding Acknowledgment message from the correspondent node, communications are usually delayed until this is received.

Handoff delays in base Mobile IPv6 route optimization are additive to other delays at the IP layer or link layer. They can cause perceptible quality degradations for interactive and real-time applications. TCP bulk-data transfers are likewise affected since long handoff latencies may lead to successive retransmission timeouts and degraded throughput [7]. An objective of Enhanced Route Optimization is hence a reduction of the handoff latency.

2.2. Security

The return routability procedure was designed with the objective to provide a level of security that compares to that of today's non-mobile Internet [5]. As such, it protects against impersonation, denial-of-service, and flooding threats that do not exist in the non-mobile Internet, but that the introduction of mobility would introduce in the absence of appropriate countermeasures. In particular, the return routability procedure satisfies the following requirements:

- o An attacker off the path from a correspondent node to a victim should not be able to trick a correspondent node into redirecting packets, which should normally be delivered to a victim, to itself, or to a third IP address. The attacker could otherwise impersonate the victim to the correspondent node or cause denial of service against the victim. The attacker may launch these

attacks from an arbitrary position, which would not necessarily have to be on the path between the victim and the correspondent node.

- o An attacker off the path from a correspondent node to a victim should not be able to trick the correspondent node into redirecting packets, which should normally be delivered to the attacker itself, to the victim. The attacker could otherwise flood the victim with unrequested packets. Such "redirection-based flooding" may be appealing to the attacker because the burden of generating the flooding packets and sending them to the victim would be on the correspondent node rather than on the attacker. The attacker could spoof multiple correspondent nodes into flooding the same victim. This would enable the attacker to impact the victim much stronger than with a direct flooding attack, where the attacker itself would generate and send the flooding packets. Comparable amplification is today only possible through an army of compromised nodes [8]. One way to cause redirection-based flooding is this: The attacker could accomplish the initial TCP handshake for a voluminous file download through its own IP address, and subsequently bind the victim's IP address (as a care-of address) to the attacker's own IP address (or home address). The correspondent node thereby redirects the download to the victim. The attacker could spoof acknowledgments on behalf of the victim based on the sequence numbers it learned during the initial handshake in order to maintain or accelerate the download. The acknowledgments would be smaller and typically less than the full-sized segments that the correspondent node generates, hence facilitating the amplification.
- o Attackers should not be able to cause denial of service against mobile or correspondent nodes through exploiting expensive computations involved in the mobility protocol.

The return routability procedure precludes impersonation, denial of service, and redirection-based flooding by attackers that are not on the path from a correspondent node to a victim, and it is sufficiently lightweight not to expose expensive operations. But the return routability procedure fails to protect against attackers that are located on the path from the correspondent node to the victim. Applications that require a higher security level are generally advised to use end-to-end protection such as IP security (IPsec) or Transport Layer Security (TLS). But even then are they vulnerable to denial of service or flooding. Furthermore, end-to-end security mechanisms generally require mobile and correspondent nodes to be preconfigured with authentication credentials, or they depend on a public-key infrastructure. Both would hinder a wide deployment of Mobile IPv6 route optimization if it was a prerequisite for the

protocol. An objective of Enhanced Route Optimization is hence to securely authenticate mobile nodes without preconfigured credentials or a public-key infrastructure, even in the presence of attackers on the path from the correspondent node to the victim.

2.3. Signaling Overhead

A complete correspondent registration involves six message transmissions at the mobile node, totaling about 376 bytes [9]. This signaling overhead may be acceptable if movements are infrequent. For example, a mobile node that moves once every 30 minutes generates an average of 1.7 bits/s of signaling traffic. Higher mobility causes more substantial overhead, however. A cell size of 100 meters and a speed of 120 km/h yields a change in IP connectivity every 3 s and about 1,000 bits/s of signaling traffic. This is significant compared to a highly compressed voice stream with a typical data rate of 10,000 to 30,000 bits/s.

Furthermore, base Mobile IPv6 requires mobile nodes to renew a correspondent registration at least every 7 minutes. The signaling overhead amounts to 7.16 bits/s if the mobile node communicates with a stationary node [9]. It doubles if both peers are mobile. This overhead may be negligible when the nodes communicate, but it can be an issue for mobile nodes that are inactive and stay at the same location for a while. These nodes typically prefer to go to standby mode to conserve battery power. Also, the periodic refreshments consume a fraction of the wireless bandwidth that one could use more efficiently. These observations lead to the objective of Enhanced Route Optimization to reduce the signaling overhead of a base Mobile IPv6 correspondent registrations as much as possible, in particular when the mobile node does not move for a while.

3. Protocol Design

Enhanced Route Optimization consists of a set of optimizations that collectively afford the achievement of the objectives discussed in Section 2. These optimizations are summarized in the following.

3.1. Cryptographically Generated Home Addresses

A Mobile IPv6 binding is conceptually a packet redirection from a home address to a care-of address. The home address is the source of the redirection and the care-of address is the destination. The packets to be redirected can hence be identified based on the home address. This motivates a cryptographic ownership proof for the home address. Enhanced Route Optimization applies cryptographically generated home addresses for this purpose [10][11]. In general, a Cryptographically Generated Address (CGA) provides a strong,

cryptographic binding between its interface identifier and the CGA owner's public key. This facilitates a cryptographic home address ownership proof without a public-key infrastructure, enabling other nodes to securely and autonomously authenticate the CGA owner as such, modulo the correctness of the CGA's subnet prefix. Cryptographically generated home addresses can supersede home address tests with the exception of an initial test for validating the home address prefix. This facilitates lower handoff delays and longer binding lifetimes, as well as reduced signaling overhead for mobile nodes that temporarily do not move. Enhanced Route Optimization also optionally enables the correspondent node to prove ownership of its IP address.

3.2. Non-Cryptographic Care-of Addresses

In contrast to a home address, a care-of address does not have identifying functionality. There is hence little benefit in a cryptographic ownership proof of a care-of address. Given that the care-of address is the destination of a packet redirection, it is rather the mobile node's reachability at the care-of address that matters. Enhanced Route Optimization uses care-of address tests for this purpose, but allows correspondent nodes to send packets to a new care-of address before the mobile node has been found to be reachable there.

3.3. Semi-Permanent Security Associations

CGA-based authentication involves public-key cryptography and is hence computationally much less efficient than authentication through a shared secret key. The technique further requires a substantial amount of supplementary CGA parameters to be piggybacked onto protected messages. Enhanced Route Optimization mitigates these disadvantages in that it utilizes an initial CGA-based authentication to securely exchange a secret permanent home keygen token between a mobile node and a correspondent node. The permanent home keygen token is used to authenticate the mobile node more efficiently in subsequent correspondent registrations. Mobile and correspondent nodes renew the permanent home keygen token on an infrequent basis. The token is therefore neither constant nor short-lived, which is why the security association between the mobile node and the correspondent node is called "semi-permanent".

3.4. Initial Home Address Tests

An initial home address test is necessary despite a cryptographic proof of home address ownership to protect against spoofed subnet prefixes in home addresses. In the complete absence of home address tests, a malicious node could cryptographically generate a home

address with the subnet prefix of a victim network, and request a correspondent node to register a binding between this spoofed home address and the attacker's own care-of address. The attacker then tricks the correspondent node into sending a stream of packets to the care-of address and subsequently deregisters the binding or lets it expire. The consequence is that the correspondent node redirects the packet stream "back" to the home address, causing the victim network to be flooded with unrequested packets. To preclude such misuse, an initial home address test is required for the mobile node and the correspondent node to establish a semi-permanent security association. The home address test is, if possible, executed in proactive manner so as to save a potentially costly message exchange via the home agent during the critical handoff period. The home address test does not need to be repeated upon subsequent movements.

3.5. Concurrent Care-of Address Tests

Enhanced Route Optimization allows a correspondent node to send payload packets to a mobile node's new care-of address before the mobile node has been found to be reachable at the care-of address. When the mobile node changes IP connectivity, it first updates its binding at the correspondent node to the new care-of address without providing a proof of reachability. The correspondent node registers the new care-of address on a tentative basis and sets it to UNVERIFIED state. Payload packets can then be exchanged bidirectionally via the new care-of address, while the mobile node's reachability at the new care-of address is verified concurrently. The correspondent node moves the care-of address to VERIFIED state once reachability verification completes.

3.6. Credit-Based Authorization

Concurrent care-of address tests without additional protection would enable an attacker to trick a correspondent node into temporarily redirecting payload packets, which would otherwise be addressed to the attacker itself, to the IP address of a victim. Such "redirection-based flooding" [5] may be appealing to the attacker because the correspondent node (not the attacker) generates the flooding packets and sends them to the victim. This enables the attacker to amplify the strength of the attack to a significant degree compared to a direct flooding attack where the attacker itself would generate the flooding packets.

Enhanced Route Optimization protects against redirection-based flooding attacks through the use of Credit-Based Authorization. Credit-Based Authorization manages the effort that a correspondent node expends in sending payload packets to a care-of address in UNVERIFIED state so as to ensure that a redirection-based flooding

attack cannot be more effective than direct flooding. The ability to send unrequested packets is an inherent property of packet-oriented networks, and direct flooding is a threat that results from this. Since direct flooding exists with and without mobility support, and redirection-based flooding attacks cannot be any more efficient than this, Credit-Based Authorization increases the security level provided by Enhanced Route Optimization with respect to flooding to that of the non-mobile Internet. Enhanced Route Optimization therefore satisfies the objective to provide a security level comparable to that of the non-mobile Internet.

The measuring and limiting of effort are technically realized through the concept of "credit", which a correspondent node maintains to put its own effort in relation to the effort that a mobile node expends during regular communications with the correspondent node. The correspondent node increases the credit for payload packets it receives from a care-of address of the mobile node in VERIFIED state, and it reduces the credit in proportion to its own effort for sending payload packets to a care-of address of the mobile node in UNVERIFIED state.

3.7. Parallel Home and Correspondent Registrations

Enhanced Route Optimization enables mobile nodes to pursue a correspondent registration in parallel with the respective home registration. This reduces handoff delays compared to base Mobile IPv6, which requires mobile nodes to wait for a Binding Acknowledgment message indicating a successful home registration before they initiate a correspondent registration.

4. Protocol Operation

Enhanced Route Optimization allows a mobile node to securely authenticate to a correspondent node based on the CGA property of its home address, and to request a concurrent care-of address test for increased handoff efficiency. Depending on whether the mobile node wishes to take advantage of either or both of these enhancements, the messages exchanged during a correspondent registration are different. This is described in the following.

4.1. Sending Binding Update Messages

A mobile node may initiate a correspondent registration for any of the following reasons:

- o To establish a new binding at a correspondent node while away from its home link so that subsequent packets will be route-optimized and no longer be routed through the mobile node's home agent.

- o To update an existing binding at the correspondent node while moving from one point of IP attachment to another.
- o To follow up an early Binding Update message with a complete Binding Update message after receiving a Binding Acknowledgment message with a Care-of Test option.
- o To refresh an existing binding at the correspondent node without changing the current point of IP attachment.
- o To request the correspondent node to renew an existing permanent home keygen token shared between the mobile node and the correspondent node (see Section 4.5).
- o To request the correspondent node to deregister an existing binding.

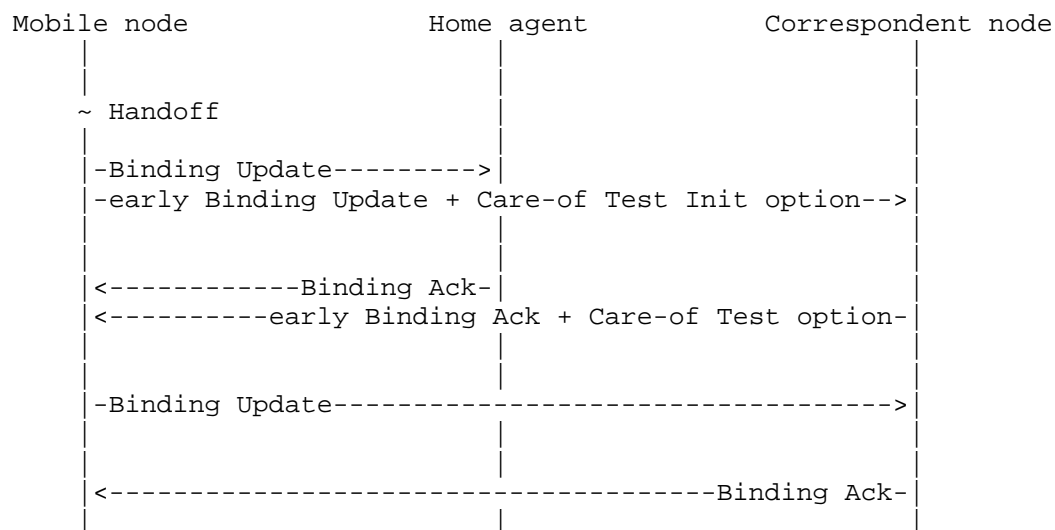


Figure 1: Correspondent registration with authentication by a proof of the mobile node's knowledge of a permanent home keygen token; concurrent care-of address test

In any of these cases, the mobile node sends a Binding Update message to the correspondent node. The Binding Update message is authenticated by one of the following three authentication methods:

- o If the mobile node's home address is a CGA, but the mobile node does not have a permanent home keygen token in its Binding Update List entry for the correspondent node, the mobile node SHOULD

authenticate the Binding Update message based on the CGA property of its home address. This requires the mobile node to send its CGA parameters and signature to the correspondent node and to pass a check of reachability at the home address.

- o If the mobile node's home address is a CGA, and the mobile node has a permanent home keygen token in its Binding Update List entry for the correspondent node, the mobile node MUST authenticate the Binding Update message by a proof of its knowledge of the permanent home keygen token.
- o If the mobile node's home address is not a CGA, the mobile node MUST authenticate the Binding Update message through a proof of reachability at its home address.

The lifetime requested by the mobile node in the Lifetime field of the Binding Update message MUST NOT exceed `MAX_CGA_BINDING_LIFETIME` (see Section 7) if the Binding Update message is to be authenticated based on the CGA property of the mobile node's home address or by a proof of the mobile node's knowledge of a permanent home keygen token. If the selected authentication method is a proof of the mobile node's reachability at the home address, the lifetime MUST NOT exceed `MAX_RR_BINDING_LIFETIME` [1]. It is RECOMMENDED in all cases that the mobile node requests the maximum permitted lifetime in order to avoid unnecessary binding refreshes and thus reduce signaling overhead. The Lifetime field of a Binding Update message that requests the deletion of an existing binding at the correspondent node MUST be set to zero.

If the selected authentication method is by way of the CGA property of the mobile node's home address, the mobile node includes its CGA parameters and signature in the Binding Update message by adding one or more CGA Parameters options (see Section 5.1) directly followed by a Signature option (see Section 5.2). This is described in Section 4.5. Once a permanent home keygen token has been obtained from the correspondent node, the mobile node MUST authenticate all subsequent Binding Update messages by a proof of its knowledge of this permanent home keygen token until either the binding lifetime expires, the permanent home keygen token is renewed, or the mobile node explicitly deregisters the binding at the correspondent node. This ensures that an attacker on the path from the correspondent node to the mobile node's home address cannot downgrade the mobile node's chosen authentication method to a proof of reachability at the home address. The mobile node MAY choose to ignore the CGA property of its home address and authenticate Binding Update messages through a proof of reachability at the home address. However, this behavior increases the vulnerability to on-path attackers and is therefore NOT RECOMMENDED.

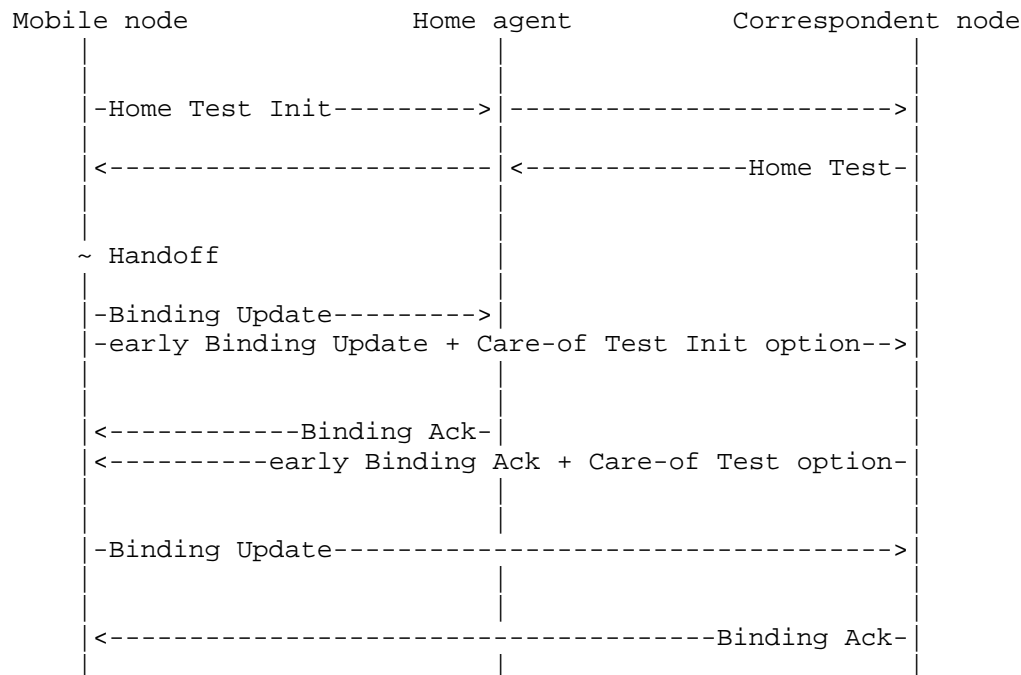


Figure 2: Correspondent registration with authentication based on reachability verification at the home address; concurrent care-of address test

The mobile node also includes its CGA parameters in the Binding Update message when it intends to renew an existing permanent home keygen token shared with the correspondent node. This is accomplished, as before, by adding to the message one or more CGA Parameters options and a Signature option.

The authenticator for the Binding Update message is calculated based on a permanent or temporary home keygen token. Which type of home keygen token the mobile node uses in calculating the authenticator depends on the authentication method:

- o If the Binding Update message is to be authenticated based on the CGA property of the mobile node's home address, the mobile node MUST use a temporary home keygen token from the correspondent node. The mobile node may already have a valid temporary home keygen token in its Binding Update List entry for the correspondent node, or it may retrieve one through the exchange of a Home Test Init message and a Home Test message.

- o If the Binding Update message is to be authenticated by a proof of the mobile node's knowledge of a permanent home keygen token, the mobile node MUST use the permanent home keygen token that is has in its Binding Update List entry for the correspondent node.
- o If the Binding Update message is to be authenticated through a proof of reachability at the home address, the mobile node MUST use a temporary home keygen token from the correspondent node. As before, the mobile node may already have a valid temporary home keygen token in its Binding Update List entry for the correspondent node, or it may retrieve one through the exchange of a Home Test Init message and a Home Test message.

Unless the purpose of the Binding Update message is to delete an existing binding at the correspondent node, the authenticator is also calculated based on a care-of keygen token. The mobile node selects this as follows:

- o If the mobile node has a valid care-of keygen token for the to-be-registered care-of address in its Binding Update List entry for the correspondent node, the mobile node MUST use this in calculating the authenticator for the Binding Update message. The Binding Update message is in this case "complete".
- o If the mobile node does not have a valid care-of keygen token in its Binding Update List entry for the correspondent node, the mobile node SHOULD define the care-of keygen token to be zero and use this in calculating the authenticator for the Binding Update message. The Binding Update message is in this case "early".
- o If the mobile node does not have a valid care-of keygen token in its Binding Update List entry for the correspondent node, the mobile node MAY choose to retrieve a care-of keygen token through the exchange of a Care-of Test Init message and a Care-of Test message, as defined in [1], without sending an early Binding Update message. In this case, the mobile node waits for receipt of the Care-of Test message and uses the care-of keygen token contained therein in calculating the authenticator for a complete Binding Update message. This approach increases the handoff latency, however, and is therefore NOT RECOMMENDED.

For reduced handoff delays, the mobile node SHOULD simultaneously initiate home and correspondent registrations for a particular care-of address. The mobile node SHOULD also pursue home and correspondent deregistrations in parallel if it wishes to discontinue Mobile IPv6 service while away from its home link. However, when the mobile node commits home and correspondent deregistrations after returning back to the home link after a period of roaming, the mobile

node MUST initiate the home deregistration first, and it MUST wait for a Binding Acknowledgment message indicating a successful home deregistration before it initiates the correspondent deregistration. This behavior ensures that the home agent does not proxy the mobile node's home address while the mobile node is on the home link, hence preventing interference between the mobile node and the home agent during Duplicate Address Detection. Since a home deregistration consumes only a link-local round-trip time when the mobile node pursues it from the home link, the cost of not parallelizing it with a correspondent deregistration, in terms of increased handoff delay, is typically negligible.

Moreover, when the Binding Update message for the correspondent registration is to be authenticated based on the CGA property of the mobile node's home address or through a proof of reachability at the home address, the mobile node SHOULD initiate the exchange of Home Test Init and Home Test messages prior to handoff in order to proactively elicit a fresh home keygen token from the correspondent node. This reduces handoff delays further. A Home Test Init message may be sent periodically whenever the home keygen token previously acquired from the correspondent node is about to expire. Tokens are valid for 3.5 minutes [1], so the interval between successive Home Test Init messages should be a little less. Alternatively, the mobile node may be able to send the Home Test Init message right in time if its link layer provides a trigger announcing imminent handoff. Proactive home address tests are technically feasible because a home address does not change across handoffs.

If the mobile node initiates the home address test from the home link, it MUST address the Home Test Init message directly to the correspondent node. The Home Test message will then be received directly from the correspondent node. If the home address test is initiated from a visited link, the mobile node MUST tunnel the Home Test Init message to the home agent. The Home Test message will then be tunneled back to the mobile node by the home agent. A home address test SHOULD NOT overlap with a home registration or home deregistration since this could result in the loss of the Home Test Init or Home Test message.

If the Binding Update message is early, the mobile node MUST add a Care-of Test Init option (see Section 5.4) to the message, requesting the correspondent node to return a new care-of keygen token. The Care-of Test Init option MUST follow the CGA Parameters and Signature options, if those exist in the Binding Update message. Once a responding Binding Acknowledgment message with a Care-of Test option (see Section 5.5) is received, the mobile node MUST use the care-of

keygen token contained therein in calculating the authenticator for a complete Binding Update message and send this message to the correspondent node.

If the Binding Update message is authenticated based on the CGA property of the mobile node's home address, the mobile node MAY add a CGA Parameters Request option (see Section 5.6) to the Binding Update message so as to request the correspondent node to prove ownership of its IP address within the Binding Acknowledgment message. This ownership proof enables the mobile node to verify that the permanent home keygen token returned in the Binding Acknowledgment message was generated by the right correspondent node.

The mobile node includes the nonce indices associated with the selected home and care-of keygen tokens in the Binding Update message using a Nonce Indices option [1]. The home nonce index is thereby determined as follows:

- o If the Binding Update message is to be authenticated based on the CGA property of the mobile node's home address, the mobile node uses a temporary home keygen token to calculate the authenticator for the Binding Update message, and the associated home nonce index MUST be taken from the Home Test message with which the home keygen token was obtained.
- o If the Binding Update message is to be authenticated by a proof of the mobile node's knowledge of a permanent home keygen token, the home nonce index MUST be set to zero.
- o If the Binding Update message is to be authenticated through a proof of the mobile node's reachability at the home address, the mobile node uses a temporary home keygen token to calculate the authenticator for the Binding Update message, and the associated home nonce index MUST be taken from the Home Test message with which the home keygen token was obtained.

The care-of nonce index is determined according to the following rules:

- o If the Binding Update message is complete, the care-of nonce index is taken from the Care-of Test option or Care-of Test message with which the care-of keygen token (used to calculate the authenticator for the Binding Update message) was obtained.
- o If the Binding Update message is early, the care-of nonce index MUST be set to zero.

- o If the purpose of the Binding Update message is to delete a binding at the correspondent node, the care-of nonce index MUST be set to zero.

The Nonce Indices option follows the CGA Parameters, Signature, Care-of Test Init, and CGA Parameters Request options if those are included in the Binding Update message as well.

The mobile node finally calculates an authenticator for the Binding Update message based on the selected home and care-of keygen tokens, following the rules described in Section 5.2 and Section 6.2.7 of [1]. For a Binding Update message that requests the deletion of an existing binding at the correspondent node, the authenticator is calculated based on only a home keygen token, and it does not incorporate a care-of keygen token. The authenticator is placed into the Authenticator field of a Binding Authorization Data option [1], which the mobile node adds to the Binding Update message as the last option.

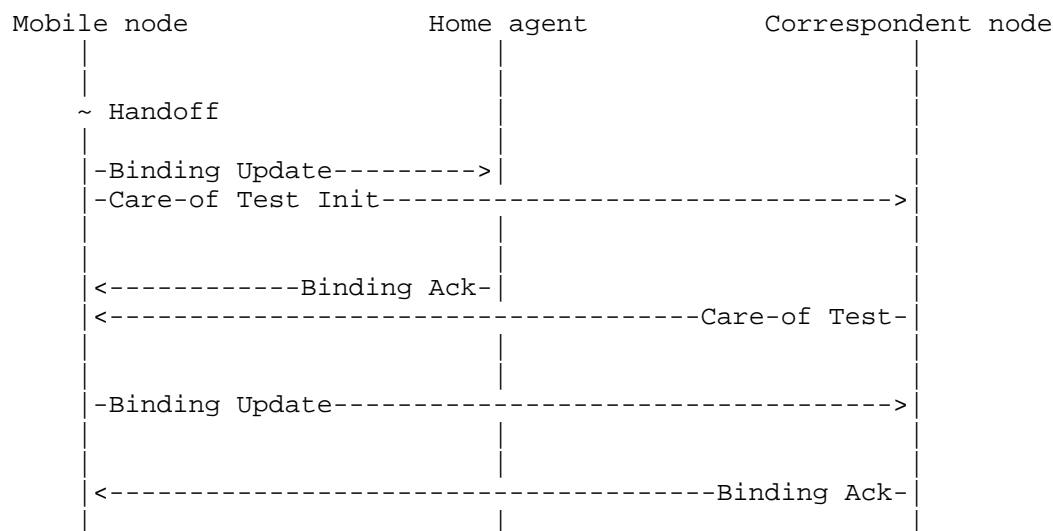


Figure 3: Correspondent registration with authentication by a proof of the mobile node's knowledge of a permanent home keygen token; explicit care-of address test

The time-sequence diagrams in Figure 1 through Figure 3 illustrate the operation of Enhanced Route Optimization based on a few selected message exchanges. Figure 1 shows the messages exchanged for a correspondent registration where an early Binding Update message is authenticated by a proof of the mobile node's knowledge of a permanent home keygen token. A Care-of Test Init option in the early

Binding Update message requests the correspondent node to add to the Binding Acknowledgment message a fresh care-of keygen token in a Care-of Test option. The mobile node finally concludes the correspondent registration with a complete Binding Update message. Figure 2 shows the procedure of a correspondent registration where the Binding Update message is authenticated with a proof of reachability at the home address. The home address test is proactively performed prior to handoff, permitting the mobile node to issue a Binding Update message directly after the handoff. The Binding Update message is again early, and a care-of keygen token is delivered to the mobile node along with the Binding Acknowledgment message. Figure 3 depicts a correspondent registration where the mobile node initially obtains a fresh care-of keygen token through the dedicated exchange of Care-of Test Init and Care-of Test messages. It subsequently issues a complete Binding Update message that is authenticated with the CGA property of the home address.

4.2. Receiving Binding Update Messages

When the correspondent node receives a Binding Update message, it must first verify whether the sending mobile node is the legitimate owner of the home address specified in the message. The correspondent node selects the authentication method based on the home nonce index given in the Nonce Indices option of the Binding Update message, and on the existence of CGA Parameters and Signature options in the Binding Update message:

- o If the home nonce index is set to a non-null value and the Binding Update message includes one or more CGA Parameters options followed by a Signature option, the correspondent node MUST authenticate the Binding Update message based on the CGA property of the mobile node's home address.
- o If the home nonce index is zero and the Binding Update message does not include one or more CGA Parameters options followed by a Signature option, the correspondent node MUST authenticate the Binding Update message by a proof of the mobile node's knowledge of a permanent home keygen token.
- o If the home nonce index is set to a non-null value and the Binding Update message does not include one or more CGA Parameters options followed by a Signature option, the correspondent node MUST authenticate the Binding Update message through a proof of the mobile node's reachability at the home address.

In addition to the validation procedure for Binding Update messages specified in [1], the correspondent node must take the following additional steps to reject Binding Update messages that are inappropriately authenticated:

- o If the Binding Update message includes one or more CGA Parameters options followed by a Signature option and the home nonce index is zero, the correspondent node MUST send a Binding Acknowledgment message with status code 150 ("Non-null home nonce index expected"). This ensures that a Binding Update message that is authenticated based on the CGA property of the mobile node's home address must also provide a proof of the mobile node's reachability at the home address.
- o If the Binding Update message is to be authenticated by a proof of the mobile node's knowledge of a permanent home keygen token, the correspondent node MUST verify that it has a Binding Cache entry for the mobile node that includes a permanent home keygen token. In case the correspondent node does not have a Binding Cache entry for the mobile node, or if the existing Binding Cache entry for the mobile node does not include a permanent home keygen token, the correspondent node MUST reject the Binding Update message by sending a Binding Acknowledgment message with status code 147 ("Permanent home keygen token unavailable").
- o If the Binding Update message is to be authenticated through a proof of the mobile node's reachability at the home address, the correspondent node MUST verify that it does not have a permanent home keygen token in its Binding Cache entry for the mobile node. If the correspondent node has a permanent home keygen token in its Binding Cache entry for the mobile node, it MUST reject the Binding Update message by sending a Binding Acknowledgment message with status code 149 ("Permanent home keygen token exists"). This ensures that an attacker cannot downgrade the authentication method to hijack the binding of a legitimate mobile node.

The authenticator for the Binding Update message is calculated based on a permanent or temporary home keygen token. Which type of home keygen token the correspondent node uses in validating the authenticator, and how it retrieves or recomputes the home keygen token, depends on the authentication method:

- o If the Binding Update message is to be authenticated based on the CGA property of the mobile node's home address, the correspondent node MUST recompute the temporary home keygen token defined by the (non-null) home nonce index in the Nonce Indices option of the Binding Update message, and it MUST use this recomputed token in validating the authenticator of the message.

- o If the Binding Update message is to be authenticated by a proof of the mobile node's knowledge of a permanent home keygen token, the correspondent node MUST use the permanent home keygen token that it has in its Binding Cache entry for the mobile node in validating the authenticator of the Binding Update message.
- o If the Binding Update message is to be authenticated through verification of the mobile node's reachability at the home address, the correspondent node MUST recompute the temporary home keygen token defined by the (non-null) home nonce index in the Nonce Indices option of the Binding Update message, and it MUST use this recomputed token in validating the authenticator of the message.

Unless the purpose of the Binding Update message is to delete an existing binding at the correspondent node, the authenticator is also calculated based on a care-of keygen token. Which care-of keygen token the correspondent node uses in validating the authenticator depends on whether the Binding Update message is complete or early:

- o If the care-of nonce index in the Nonce Indices option of the Binding Update message is set to a non-null value, the Binding Update message is complete. In this case, the correspondent node MUST recompute the care-of keygen token that is identified by the care-of nonce index, and it MUST use this recomputed token in validating the authenticator of the message.
- o If the care-of nonce index in the Nonce Indices option of the Binding Update message is zero, the Binding Update message is early. The care-of keygen token to be used by the correspondent node in validating the authenticator of the Binding Update message is zero in this case.

The correspondent node finally validates the authenticator in the Binding Update message based on the selected home and care-of keygen tokens, following the algorithm described in Section 9.5.1 of [1].

If the validation fails, the correspondent node MUST discard the Binding Update message. The correspondent node may have to send a Binding Acknowledgment message with a status code indicating the failure, as described in [1].

Provided that the validation of the authenticator in the Binding Update message succeeds, the correspondent node registers the mobile node's new care-of address, either updating an existing Binding Cache entry, if one exists, or creating a new Binding Cache entry. The lifetime granted for the binding depends on the lifetime requested by the mobile node in the Lifetime field of the Binding Update message

and the method by which the Binding Update message is authenticated. If the Binding Update message is authenticated based on the CGA property of the mobile node's home address or by a proof of the mobile node's knowledge of a permanent home keygen token, the lifetime for the binding SHOULD be set to the maximum of MAX_CGA_BINDING_LIFETIME and the value specified in the Lifetime field of the Binding Update message. If the Binding Update message is authenticated through a proof of the mobile node's reachability at the home address, then the lifetime for the binding SHOULD be set to the maximum of MAX_RR_BINDING_LIFETIME [1] and the value specified in the Lifetime field of the Binding Update message. The correspondent node may in either case grant a further reduced lifetime, but it MUST NOT accept a higher lifetime.

The state of the new care-of address depends on whether the Binding Update message is complete or early:

- o If the Binding Update message is complete, the new care-of address is set to VERIFIED state. The correspondent node may then immediately send packets to the new care-of address without restrictions.
- o If the Binding Update message is early, the new care-of address is set to UNVERIFIED state. The correspondent node MUST then follow the rules defined in Section 4.10 for sending packets to this care-of address until the care-of address is set in VERIFIED state.

If the Binding Update message contains one or multiple CGA Parameters options, the mobile node is requesting the correspondent node to accept the included CGA parameters either for establishing a new, or for renewing an existing permanent home keygen token shared between the mobile node and the correspondent node. The correspondent node MUST in this case check if the CGA Parameters options are directly followed by a Signature option and, if so, validate the CGA parameters and signature as described in Section 4.6.

If the CGA Parameters option is not directly followed by a Signature option, or the validation of the included CGA parameters and signature fails, the correspondent node MUST discard the Binding Update message and send a Binding Acknowledgment message with status code 148 ("CGA and signature verification failed") to the mobile node.

Provided that the signature included in the Signature option is correct, the correspondent node generates a permanent home keygen token to be shared with the mobile node and stores it in its Binding Cache entry for the mobile node. The permanent home keygen token is

sent to the mobile node within a Binding Acknowledgment message as described in Section 4.3.

4.3. Sending Binding Acknowledgment Messages

Upon receipt of a valid Binding Update message, the correspondent node returns to the mobile node a Binding Acknowledgment message in any of the following cases:

- o The Acknowledge flag in the Binding Update message is set.
- o The Binding Update message contains one or multiple CGA Parameters options directly followed by a Signature option, and the signature included in the latter was determined to be correct.
- o The Binding Update message is early and includes a Care-of Test Init option.

If the Binding Update message further contains a CGA Parameters Request option and the correspondent node's IP address is a CGA, the correspondent node MUST include its CGA parameters and signature in the Binding Acknowledgment message by adding one or more CGA Parameters options directly followed by a Signature option. The correspondent node's CGA parameters and signature enable the mobile node to verify that the permanent home keygen token received in the Binding Acknowledgment message was generated by the right correspondent node. If the Binding Update message contains a CGA Parameters Request option, but the correspondent node's IP address is not a CGA, the correspondent node ignores the CGA Parameters Request option and processes the Binding Update message further as described below.

If the Binding Update message contains one or multiple CGA Parameters options directly followed by a Signature option, and the signature included in the latter was determined to be correct, the correspondent node MUST add a Permanent Home Keygen Token option (see Section 5.3) with a new permanent home keygen token to the Binding Acknowledgment message. The correspondent node also stores this permanent home keygen token in its Binding Cache entry for the mobile node.

If the Binding Update message includes a Care-of Test Init option, the correspondent node MUST append to the Binding Acknowledgment message a Care-of Test option with a pseudo-random value in the Care-of Keygen Token field. The Care-of Test option MUST appear after the Permanent Home Keygen Token option in case both options are present in the Binding Acknowledgment message.

A Binding Authorization Data option must be added to the Binding Acknowledgment message as a last option, as described in Section 5.2 and Section 6.2.7 of [1].

4.4. Receiving Binding Acknowledgment Messages

A mobile node first verifies a received Binding Acknowledgment message according to the rules specified in [1]. Provided that the Binding Acknowledgment message is not rejected based on these rules, the mobile node takes the following additional steps.

If the mobile node included a CGA Parameters Request option in the Binding Update message and the Binding Acknowledgment message contains a Permanent Home Keygen Token option, the mobile node first processes any CGA Parameters and Signature options in the Binding Acknowledgment message in the following manner. If the Binding Acknowledgment message contains one or more CGA Parameters options that are directly followed by a Signature option, the mobile node MUST check the ownership of the correspondent node's IP address by verifying the included CGA parameters and signature as described in Section 4.6. If the validation of the CGA parameters and signature fails, the mobile node MUST silently discard the Binding Acknowledgment message. The mobile node MUST also silently discard the Binding Acknowledgment message if the message includes one or more CGA Parameters options that are not directly followed by a Signature option, or if the Binding Acknowledgment message lacks any CGA Parameters options in the presence of a Signature option.

If the mobile node did not include a CGA Parameters Request option in the Binding Update message or the Binding Acknowledgment message does not contain a Permanent Home Keygen Token option, the mobile node ignores any CGA Parameters and Signature options that the Binding Acknowledgment message may contain. Careful use of the CGA Parameters Request option in Binding Update messages enables the mobile node to control the processing resources it spends on the verification of a correspondent node's CGA as well as to disable such verification in the case of persistent verification failures, which may be due to misconfigured or outdated CGA software [12] on the correspondent node side or at the mobile node itself. Specifically, if the mobile node repeatedly fails to receive a Binding Acknowledgment message including valid CGA Parameters and Signature options in response to sending a Binding Update message with a CGA Parameters Request option, the mobile node SHOULD refrain from including a CGA Parameters Request option in future Binding Update messages for the same correspondent node.

If the mobile node included a CGA Parameters Request option in the Binding Update message, but the Binding Acknowledgment message does not contain any CGA Parameters or Signature options, the mobile node cannot be sure if the correspondent node's IP address is simply not a CGA, or if the Binding Acknowledgment message originates from an attacker on the path from the mobile node to the correspondent node. To avoid accepting a permanent home keygen token from an on-path attacker, the mobile node MUST give precedence to Binding Acknowledgment messages that include valid CGA Parameters and Signature options over Binding Acknowledgment messages without such options. One possible algorithm for the mobile node to follow in this regard is to always accept the Binding Acknowledgment message received first, and if this message does not contain valid CGA Parameters or Signature options and another Binding Acknowledgment message including such options is received later on, to revert any state changes involved in accepting the first Binding Acknowledgment in favor of this subsequent Binding Acknowledgment message. Giving precedence to Binding Acknowledgment messages with valid CGA Parameters and Signature options over Binding Acknowledgment messages without such options enables the mobile node to communicate with correspondent nodes that do not use a CGA, and at the same time protects against most on-path attackers. The strategy does not protect against an attacker that can intercept Binding Acknowledgment messages from the correspondent node, but such an attacker could preclude mobility management between the mobile node and the correspondent node anyway. When the mobile node has permanently accepted a Binding Acknowledgment message without valid CGA Parameters and Signature options, the mobile node SHOULD refrain from including a CGA Parameters Request option in future Binding Update messages for the same correspondent node.

If the Binding Acknowledgment message contains a Permanent Home Keygen Token option, the mobile node extracts the permanent home keygen token included in this option and stores it in its Binding Update List entry for the correspondent node. Future Binding Update messages will then be authenticated by a proof of the mobile node's knowledge of this permanent home keygen token.

If the Binding Acknowledgment message contains a Care-of Test option, the mobile node extracts the care-of keygen token included in this option, stores the token in its Binding Update List entry for the correspondent node, and sends the correspondent node a complete Binding Update message as defined in Section 4.1. Note that the complete Binding Update message will be authenticated based on the CGA property of the mobile node's home address if the Binding Acknowledgment message also includes a Permanent Home Keygen Token option. This is independent of the authentication method that was used for the corresponding early Binding Update message.

A mobile node MUST ensure that, while it has a binding for a certain home address at a correspondent node, it also has a valid binding at its home agent for the same home address. This may at times require the mobile node to extend the binding lifetime at the home agent, request a correspondent node to use a binding lifetime less than the permitted maximum, or explicitly deregister an existing binding at a correspondent node.

If the mobile node authenticates Binding Update messages for a particular correspondent node by proving its knowledge of a permanent home keygen token, but registrations at this correspondent node persistently fail, the mobile node SHOULD renew the permanent home keygen token by sending a Binding Update message that is authenticated based on the CGA property of its home address. This Binding Update message includes the mobile node's CGA parameters and signature, and it requests the correspondent node to generate a new permanent home keygen token and send this to the mobile node within a Binding Acknowledgment message.

If the mobile node persistently receives Binding Acknowledgment messages with status code 148 ("CGA and signature verification failed") from a correspondent node, the mobile node SHOULD authenticate future Binding Update messages for the same correspondent nodes through a proof of its reachability at the home address. This enables the mobile node to recover from misconfigured or outdated CGA software [12] on the correspondent node side or at the mobile node itself.

4.5. Sending CGA Parameters

A mobile node includes its CGA parameters and signature in a Binding Update message for a correspondent node in any of the following situations:

- o To acquire a permanent home keygen token if the mobile node's home address is a CGA, and the mobile node does not yet have a permanent home keygen token from the correspondent node.
- o To extend the lifetime of an existing binding if the mobile node already has a permanent home keygen token from the correspondent node, and the lifetime of the binding at the correspondent node is about to expire.
- o To renew an existing permanent home keygen token to prevent replay attacks in the imminent event of a sequence number rollover, or for improved protection against cryptanalysis.

A correspondent node whose IP address is a CGA includes its CGA parameters and signature in a Binding Acknowledgment message for the mobile node when it receives a Binding Update message with a CGA Parameters Request option.

CGA parameters are transmitted in the format of the CGA Parameters data structure defined in [2]. The CGA Parameters data structure is split over one or more CGA Parameters options as described in Section 5.1. The last CGA Parameters option MUST be directly followed by a Signature option.

The value for the Signature field in the Signature option is calculated according to the signature generation algorithm defined in Section 6 of [2]. The value is calculated with the mobile or correspondent node's private key over the following sequence of octets:

```
mobility data =  
  care-of address | correspondent node IP address | MH data
```

where "|" denotes concatenation. "Care-of address" is the mobile node's care-of address, and "correspondent node IP address" is the IP address of the correspondent node that is visible to protocol layers above IP. In case the correspondent node is mobile, "correspondent node IP address" refers to the correspondent node's home address. "MH data" is the content of the Binding Update or Binding Acknowledgment message including the mobility header and all options up to the last CGA Parameters option. That is, "MH data" excludes the IPv6 header and any IPv6 extension headers other than the mobility header itself. The "mobility data" constitutes what is referred to as the "message" in Section 6 of [2].

The value for the Signature field is calculated as if the Checksum field in the mobility header was zero. The Checksum field in the transmitted packet is still calculated in the usual manner, with the calculated value in the Signature field being a part of the packet protected by the checksum.

4.6. Receiving CGA Parameters

Mobile and correspondent nodes that receive a Binding Update or Binding Acknowledgment message including one or more CGA Parameters options directly followed by a Signature option first process the message as described in [1]. This includes a verification of the authenticator in the Authenticator field of the Binding Authorization Data option. If the Binding Update or Binding Acknowledgment message is rejected due to an incorrect authenticator or for any other reason, the message is not processed further.

Otherwise, if the validation of the Binding Update or Binding Acknowledgment message succeeds, the mobile or correspondent node reassembles the CGA Parameters data structure from the CGA Parameters options included in the message as described in Section 5.1, and executes the CGA verification algorithm defined in Section 5 of [2]. The CGA verification algorithm takes the to-be-verified CGA and the reassembled CGA Parameters data structure as input. The to-be-verified CGA is the mobile node's home address when the CGA verification algorithm is executed by the correspondent node. When the mobile node executes the CGA verification algorithm, the to-be-verified CGA is the correspondent node's IP address that is visible to protocol layers above IP. This is the correspondent node's home address in case the correspondent node is mobile. The following steps are skipped if the CGA verification fails.

If the CGA verification succeeds, the mobile or correspondent node performs a more time-consuming check of the signature. It extracts the signature from the Signature field in the Signature option and executes the signature verification algorithm defined in Section 6 of [2]. The signature verification algorithm takes as input the to-be-verified CGA as defined above, the reassembled CGA Parameters data structure, the MH data as defined in Section 4.5, the CGA Message Type tag of Enhanced Route Optimization as defined in Section 7, and the signature itself.

4.7. Sending Permanent Home Keygen Tokens

A correspondent node assigns a mobile node a new permanent home keygen token after it has received from the mobile node a Binding Update message with included CGA Parameters and Signature options, and these options have been successfully validated as described in Section 4.6. The permanent home keygen token is a 64-bit value randomly generated by the correspondent node. The correspondent node stores the permanent home keygen token in the binding cache entry that it maintains for the mobile node.

The correspondent node sends the permanent home keygen token to the mobile node in encrypted form within a Permanent Home Keygen Token option in a Binding Acknowledgment message. It sends this message even if the Acknowledge flag in the corresponding Binding Update message was clear. The correspondent node encrypts the permanent home keygen token with the mobile node's public key using the RSAES-PKCS1-v1_5 format [4], and places the ciphertext into the Permanent Home Keygen Token field of the Permanent Home Keygen Token option.

The Binding Authorization Data option MUST be the last option in the Binding Acknowledgment message. That is, the authenticator in the

Binding Authorization Data option covers the Permanent Home Keygen Token option.

4.8. Receiving Permanent Home Keygen Tokens

A mobile node that receives a Binding Acknowledgment message first processes the message as described in [1], independent of whether the message includes a Permanent Home Keygen Token option. This includes a verification of the authenticator in the Authenticator field of the Binding Authorization Data option. If the Binding Acknowledgment message is rejected due to an incorrect authenticator or for any other reason, the mobile node does not process the message further.

Otherwise, if the mobile node accepts the Binding Acknowledgment message and the message includes a Permanent Home Keygen Token option, the mobile node extracts the ciphertext from the Permanent Home Keygen Token field in this option and decrypts it with its private key using the RSAES-PKCS1-v1_5 format [4]. The result of the encryption is the permanent home keygen token to be used in further registrations with the correspondent node. The mobile node stores the permanent home keygen token in the Binding Update List entry that it maintains for the correspondent node.

4.9. Renewing Permanent Home Keygen Tokens

A mobile node that shares a permanent home keygen token with a correspondent node MUST NOT use the same sequence number twice with this permanent home keygen token in order to protect against replay attacks. The mobile node MUST renew the permanent home keygen token by including its CGA parameters and signature in a Binding Update message for the correspondent node when a sequence number rollover is imminent. In addition, the mobile node MAY renew its permanent home keygen token at any time. Periodic renewal of the permanent home keygen token provides increased protection against cryptanalysis. Finally, the mobile node may in most cases want to renew the permanent home keygen token when the lifetime of its binding at the correspondent node expires.

4.10. Handling Payload Packets

The immediate exchange of an early Binding Update message after a handoff on the mobile node side enables mobile and correspondent nodes to quickly reestablish route-optimized communications via the mobile node's new care-of address. The mobile node may send payload packets to the correspondent node from the new care-of address as soon as it has dispatched the early Binding Update message. The correspondent node redirects outgoing payload packets for the mobile node to the new care-of address once it has received the early

Binding Update message and registered the new care-of address. Here, a "payload packet" is defined as a packet that originates at a protocol layer above IP.

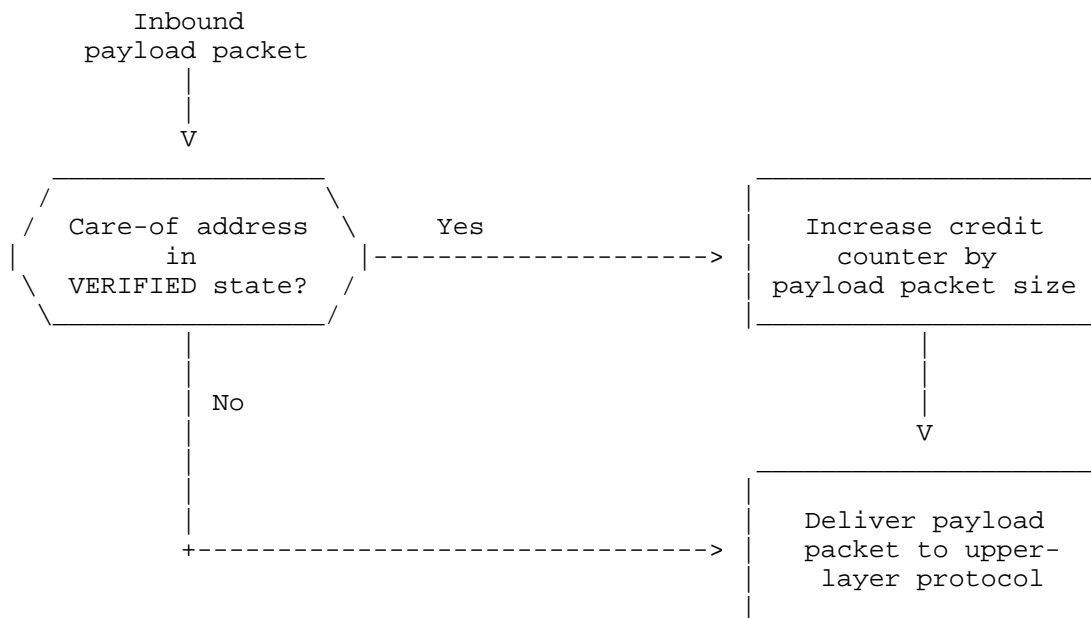


Figure 4: Handling outbound payload packets

A new care-of address that was registered with an early Binding Update message is maintained in UNVERIFIED state by the correspondent node until the correspondent node receives a complete Binding Update message from the mobile node. The correspondent node then sets the care-of address to VERIFIED state. The state of the care-of address determines the maximum amount of data that the correspondent node is allowed to send to the care-of address, as is necessary to prevent amplified, redirection-based flooding attacks. For this purpose, the correspondent node maintains a "credit counter" for each mobile node with an entry in its Binding Cache. Whenever a payload packet arrives from a mobile node with a care-of address in VERIFIED state, the correspondent node SHOULD increase the mobile node's credit counter by the size of the received payload packet. The correspondent node MAY be restricted by policy to increase the credit counter by a lower value or not to increase the credit at all. The credit counter does not change when an inbound payload packet is received from a care-of address in UNVERIFIED state. Figure 4 shows a flow chart of this procedure.

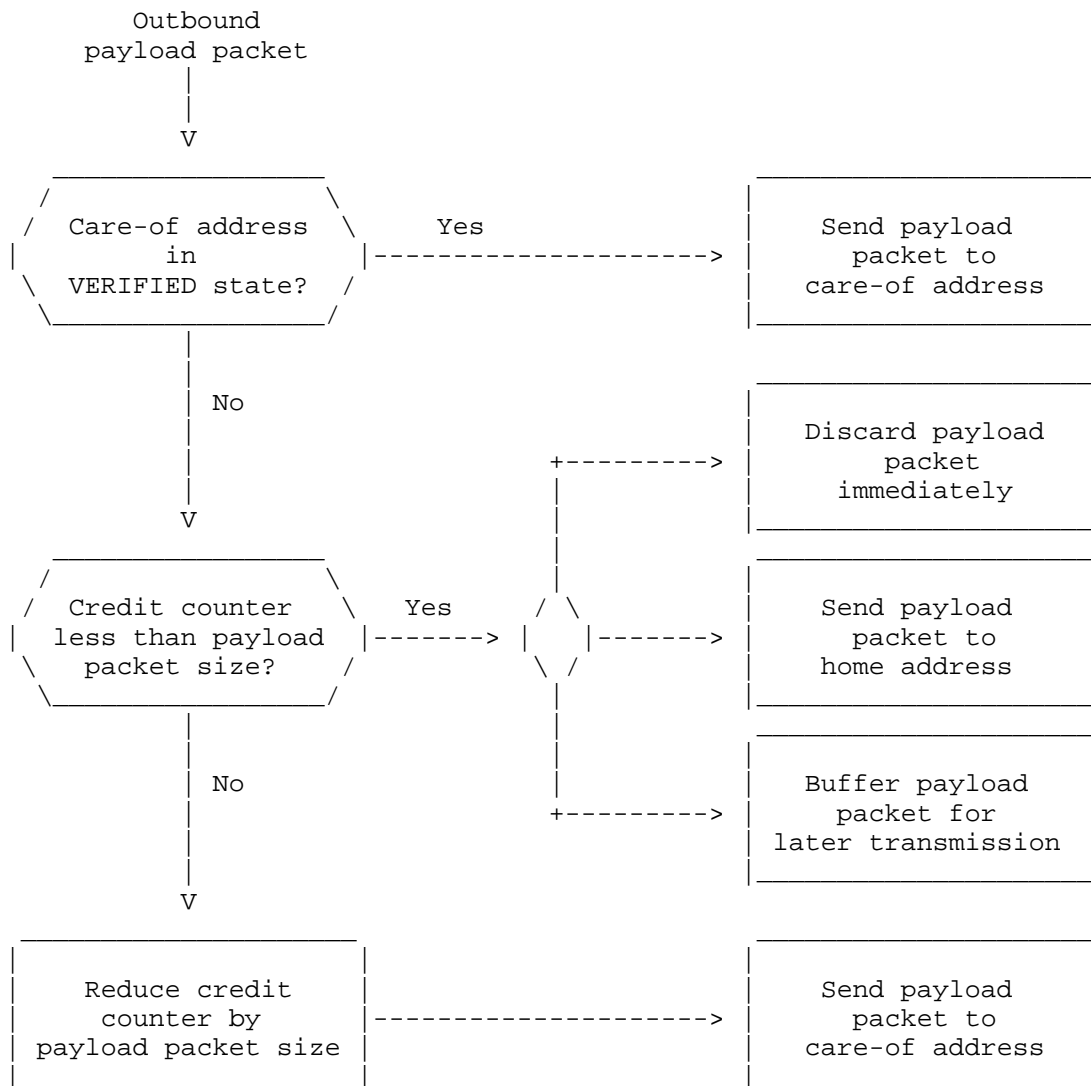


Figure 5: Handling outbound payload packets

When the correspondent node has a payload packet to send to the mobile node, further treatment of the payload packet depends on the state of the mobile node's care-of address and the current value of the mobile node's credit counter, as illustrated in Figure 5: The correspondent node MUST send the payload packet to the mobile node's care-of address if the care-of address is in VERIFIED state. If the care-of address is in UNVERIFIED state and the value of the credit counter is higher than or equal to the size of the payload packet,

the correspondent node MUST reduce the mobile node's credit counter by the size of the payload packet and send the payload packet to the care-of address as well. However, if the care-of address is in UNVERIFIED state and the credit counter is less than the size of the payload packet, the payload packet MUST NOT be sent to the mobile node's care-of address. The correspondent node SHOULD then discard the payload packet, although it MAY alternatively buffer the payload packet until the care-of address moves to VERIFIED state, or send the payload packet to the mobile node's home address. The credit counter of the mobile node does not change when the correspondent node sends a payload packet to the mobile node's care-of address while the care-of address is in VERIFIED state.

The amount of data that the mobile node may send to the correspondent node is never restricted due to the state of the mobile node's care-of address. The care-of address state also does not change the addressing and routing of payload packets in either traffic direction: All payload packets that originate from the mobile node have the care-of address in the Source Address field of the IPv6 header and the home address in the Home Address option of the IPv6 Destination Options extension header. Vice versa, all payload packets from the correspondent node have the care-of address in the Destination Address field of the IPv6 header and the home address in the IPv6 Routing extension header.

4.11. Credit Aging

A correspondent node ensures that all credit counters that it maintains gradually decrease over time. Each credit counter is multiplied with a factor, `CreditAgingFactor`, of less than one in fixed time intervals of `CreditAgingInterval` length. Such "credit aging" limits the total credit that a mobile node can earn, provided that the replenishing rate for the credit is constant or nearly constant. It thereby enforces an upper bound on the rate at which the correspondent node can durably send to the mobile node's care-of address while the care-of address is in UNVERIFIED state. In the absence of credit aging, a malicious node with poor up-link capacity could adopt the role of a mobile node, build up credit at a very slow speed and over a long period, and spend this credit during a much shorter period on redirecting a burst of payload packets to the IP address of a victim.

Choosing appropriate values for `CreditAgingFactor` and `CreditAgingInterval` is important to facilitate applications where the correspondent node sends at a higher rate than the mobile node. If `CreditAgingFactor` or `CreditAgingInterval` is too small, the credit counter might persistently prevent the transmission of payload packets to a care-of address in UNVERIFIED state. The values given

in Section 7 are RECOMMENDED as they work well when the correspondent node transfers a file to the mobile node via a TCP connection and the end-to-end round-trip time does not exceed 500 milliseconds.

4.12. Simultaneous Movements

As specified in [1], Binding Update messages are sent to a mobile correspondent node's home address. This makes it possible for two mobile nodes to continue communications even if both of them change IP connectivity at the same time.

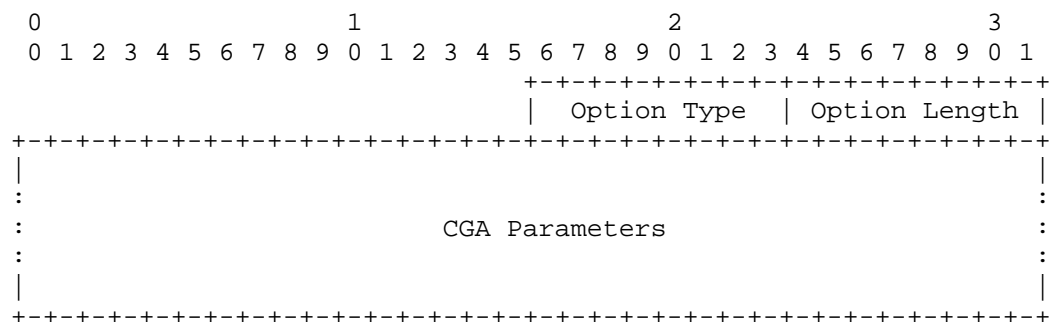
5. Option Formats and Status Codes

Enhanced Route Optimization uses a set of new mobility options and status codes in addition to the mobility options and status codes defined in [1]. These are described below.

5.1. CGA Parameters Option

The CGA Parameters option is used in Binding Update and Binding Acknowledgment messages. It contains part of the mobile or correspondent node's CGA parameters. [1] limits mobility header options to a maximum length of 255 bytes, excluding the Option Type and Option Length fields. Since the CGA parameters are likely to exceed this limit, multiple CGA Parameters options may have to be concatenated to carry all CGA parameters.

The format of the CGA Parameters option is as follows:



Option Type

8-bit identifier of the type of this mobility option. Its value is 12.

Option Length

8-bit unsigned integer representing the length of the CGA Parameters field in octets.

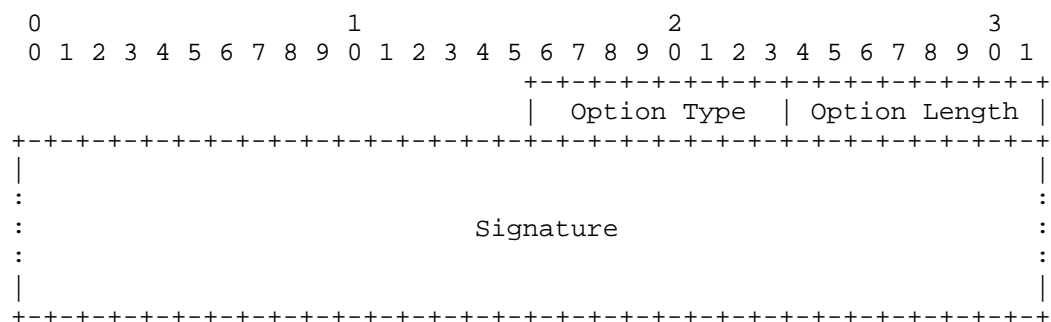
CGA Parameters

This field contains up to 255 bytes of the CGA Parameters data structure defined in [2]. The concatenation of all CGA Parameters options in the order they appear in the Binding Update message MUST result in the original CGA Parameters data structure. All CGA Parameters options in the Binding Update message except the last one MUST contain exactly 255 bytes in the CGA Parameters field, and the Option Length field MUST be set to 255 accordingly. All CGA Parameters options MUST appear directly one after another, that is, a mobility option of a different type MUST NOT be placed in between two CGA Parameters options.

5.2. Signature Option

The Signature option is used in Binding and Binding Acknowledgment Update messages. It contains a signature that the mobile or correspondent node generates with its private key over one or more preceding CGA Parameters options.

The format of the Signature option is as follows:



Option Type

8-bit identifier of the type of this mobility option. Its value is 13.

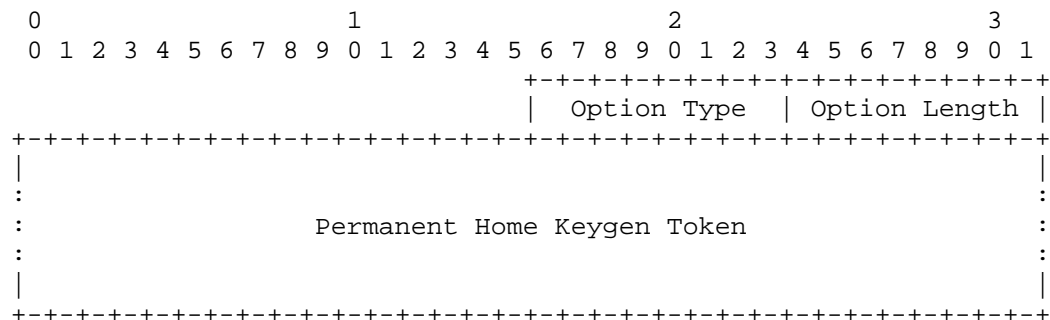
Option Length

8-bit unsigned integer representing the length of the Signature field in octets.

This field contains the mobile or correspondent node's signature, generated with the mobile or correspondent node's private key as specified in Section 4.5.

5.3. Permanent Home Keygen Token Option

The format of the Permanent Home Keygen Token option is as follows:



8-bit identifier of the type of this mobility option. Its value is 14.

8-bit unsigned integer representing the length of the Permanent Home Keygen Token field in octets.

This field contains the permanent home keygen token generated by the correspondent node. The content of this field MUST be encrypted with the mobile node's public key as defined in Section 4.7. The length of the permanent home keygen token is 8 octets before encryption, though the ciphertext [4] and, hence, the Permanent Home Keygen Token field may be longer.

5.4. Care-of Test Init Option

The Care-of Test Init option is included in Binding Update messages. It requests a correspondent node to return a Care-of Test option with a fresh care-of keygen token in the Binding Acknowledgment message.

The format of the Care-of Test Init option is 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
                                     +-----+
                                     | Option Type | Option Length |
                                     +-----+
```

Option Type

8-bit identifier of the type of this mobility option. Its value is 15.

Option Length

This field MUST be set to zero.

5.5. Care-of Test Option

The Care-of Test option is used in Binding Acknowledgment messages. It contains a fresh care-of keygen token, which the correspondent node sends to the mobile node after it has received a Care-of Test Init option in a Binding Update message.

The format of the Care-of Test option is 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
                                     +-----+
                                     | Option Type | Option Length |
+-----+
|                                     |
+                                     +
|                                     |
+-----+
                                     +-----+
                                     | Option Type | Option Length |
+-----+
```

Option Type

8-bit identifier of the type of this mobility option. Its value is 16.

Option Length

This field MUST be set to 8. It represents the length of the Care-of Keygen Token field in octets.

Care-of Keygen Token

This field contains the care-of keygen token generated by the correspondent node, as specified in Section 4.3.

5.6. CGA Parameters Request Option

The CGA Parameters Request option is included in Binding Update messages that are authenticated based on the CGA property of the mobile node's home address. It requests a correspondent node to return its CGA parameters and signature in the Binding Acknowledgment message, enabling the mobile node to verify that the permanent home keygen token returned in the Binding Acknowledgment message was generated by the right correspondent node.

The format of the CGA Parameters Request option is 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
                                     +-----+-----+-----+-----+
                                     | Option Type | Option Length |
                                     +-----+-----+-----+-----+
```

Option Type

8-bit identifier of the type of this mobility option. Its value is 11.

Option Length

This field MUST be set to zero.

5.7. Status Codes

Enhanced Route Optimization uses the following four new status codes for Binding Acknowledgment messages in addition to the status codes defined in [1]:

Permanent home keygen token unavailable (147)

A correspondent node returns a Binding Acknowledgment message with status code 147 to a mobile node if it has received from the mobile node a Binding Update message that was authenticated

through the CGA property of the mobile node's home address, but the correspondent node either does not have a Binding Cache entry for the mobile node, or the existing Binding Cache entry for the mobile node does not contain a permanent home keygen token. A Binding Acknowledgment message with status code 147 indicates to the mobile node that it should request a new permanent home keygen token from the correspondent node by sending the correspondent node a Binding Update message including its CGA parameters and signature. This in particular enables the mobile node to quickly recover from state loss at the correspondent node.

[1] does not allow a correspondent node to send a Binding Acknowledgment message with a status code indicating failure when the authenticator of a received Binding Update message turns out to be incorrect. This causes additional handoff latency with high probability because the mobile node can detect the problem only after the expiration of a retransmission timer. The mobile node is furthermore likely to assume packet loss and resend the incorrectly authenticated Binding Update message additional times. A Binding Acknowledgment message with status code 147 helps the mobile node to identify the underlying problem more efficiently when the correspondent node could not verify the CGA property of the mobile node's home address.

CGA and signature verification failed (148)

A correspondent node returns a Binding Acknowledgment message with status code 148 to a mobile node if it has received from the mobile node a Binding Update message that includes one or more CGA Parameters options directly followed by a Signature option, but either the CGA property of the home address cannot be verified based on the contents of the CGA Parameters options, or the verification of the signature in the Signature option has failed.

Permanent home keygen token exists (149)

A correspondent node returns a Binding Acknowledgment message with status code 149 to a mobile node if it has received from the mobile node a Binding Update message that was authenticated through verification of the mobile node's reachability at the home address and does not include one or more CGA Parameters options directly followed by a Signature option, but the correspondent node has a permanent home keygen token in its Binding Cache entry for the mobile node. The Binding Update message is processed further if it includes one or more CGA Parameters options directly followed by a Signature option. This enables a mobile node to obtain a new permanent home keygen token from the correspondent node in case it has lost the existing one, for instance, due to a

reboot. Whether the correspondent node accepts the Binding Update message in this case depends on the verification of the CGA parameters and the signature provided in the Binding Update message.

Non-null home nonce index expected (150)

A correspondent node returns a Binding Acknowledgment message with status code 150 to a mobile node if it has received from the mobile node a Binding Update message that includes one or more CGA Parameters options directly followed by a Signature option, but the home nonce index specified in the Nonce Indices option is zero. This behavior ensures that a Binding Update message that is authenticated based on the CGA property of the mobile node's home address must also provide a proof of the mobile node's reachability at the home address.

6. Security Considerations

Enhanced Route Optimization differs from base Mobile IPv6 in that it applies a set of optimizations for increased handoff performance, stronger security, and reduced signaling overhead. These optimizations entail the following conceptual changes to the security model [5] of base Mobile IPv6:

- o Base Mobile IPv6 conducts periodic tests of a mobile node's reachability at the home address as a proof of home address ownership. Enhanced Route Optimization applies an initial cryptographic home address ownership proof in combination with a verification of the mobile node's reachability at the home address in order to securely exchange a secret permanent home keygen token. The permanent home keygen token is used for cryptographic authentication of the mobile node during subsequent correspondent registrations, so that these later correspondent registrations can be securely bound to the initial home address ownership proof. No further periodic reachability verification at the home address tests is performed.
- o Base Mobile IPv6 requires a mobile node to prove its reachability at a new care-of address during a correspondent registration. This implies that the mobile node and the correspondent node must exchange Care-of Test Init and Care-of Test messages before the mobile node can initiate the binding update proper. Enhanced Route Optimization allows the mobile node to initiate the binding update first and follow up with a proof of reachability at the care-of address. Mobile and correspondent nodes can so resume communications early on after a handoff, while reachability verification proceeds concurrently. The amount of data that the

correspondent node is permitted to send to the care-of address until reachability verification completes is governed by Credit-Based Authorization.

- o The maximum binding lifetime for correspondent registrations is 7 minutes in base Mobile IPv6. A mobile node must hence periodically refresh a correspondent registration in cases where it does not change IP connectivity for a while. This protocol increases the maximum binding lifetime to 24 hours, reducing the need for periodic refreshes to a negligible degree.

The ensuing discussion addresses the implications that these conceptual changes of the Mobile IPv6 security model have. The discussion ought to be seen in context with the security considerations of [1], [2], and [5].

6.1. Home Address Ownership

Enhanced Route Optimization requires a mobile node to deliver a strong cryptographic proof [2] that it is the legitimate owner of the home address it wishes to use. The proof is based on the true home address owner's knowledge of the private component in a public/private-key pair with the following two properties:

- o As an input to an irreversible CGA generation function along with a set of auxiliary CGA parameters, the public key results in the mobile node's home address.
- o Among the CGA parameters that are fed into the CGA generation function is a modifier that, as an input to an irreversible hash extension function along with the public key, results in a string with a certain minimum number of leading zeroes. Three reserved bits in the home address encode this minimum number.

The first property cryptographically binds the home address to the mobile node's public key and, by virtue of public-key cryptography, to the private key. It allows the mobile node to claim ownership of the home address by proving its knowledge of the private key. The second property increases the cost of searching in brute-force manner for a public/private-key pair that suffices the first property. This increases the security of a cryptographically generated home address despite its limitation to 59 bits with cryptographic significance. Solely enforcing the first property would otherwise allow an attacker to find a suitable public/private-key pair in $O(2^{59})$ steps. By addition of the second property, the complexity of a brute-force search can be increased to $O(2^{(59+N)})$ steps, where N is the minimum number of leading zeroes that the result of the hash extension function is required to have.

In practice, for a legitimate mobile node to cryptographically generate a home address, the mobile node must first accomplish a brute-force search for a suitable modifier, and then use this modifier to execute the CGA generation function. An attacker who is willing to spoof the mobile node's home address, so-called "IP address stealing" [5], then has two options: It could either generate its own public/private-key pair and perform a brute-force search for a modifier which, in combination with the generated public key, suffices the initially described two properties; or it could integer-factor the mobile node's public key, deduce the corresponding private key, and copy the mobile node's modifier without a brute-force search. The cost of the attack can be determined by the mobile node in either case: Integer-factoring a public key becomes increasingly complex as the length of the public key grows, and the key length is at the discretion of the mobile node. The cost of a brute-force search for a suitable modifier increases with the number of leading zeroes that the result of the hash extension function is required to have. This number, too, is a parameter that the mobile node can choose. Downgrading attacks, where the attacker reduces the cost of spoofing a cryptographically generated home address by choosing a set of CGA parameters that are less secure than the CGA parameters the mobile node has used to generate the home address, are hence impossible.

The CGA specification [2] requires the use of RSA public and private keys, and it stipulates a minimum key length of 384 bits. This requirement that was tailored to Secure Neighbor Discovery for IPv6 [13], the original CGA application. Enhanced Route Optimization does not increase the minimum key length because, in the absence of downgrading attacks as explained before, the ability to use short keys does not compromise the security of home addresses that were cryptographically generated using longer keys. Moreover, extensions to [2] may eventually permit the use of public/private-key classes other than RSA. Such extensions are compatible with the CGA application of Enhanced Route Optimization. Care must be taken in selecting an appropriate key class and length, however. Home addresses are typically rather stable in nature, so the chosen parameters must be secure for a potentially long home address lifetime. Where RSA keys are used, a minimum key length of 1024 bits is therefore RECOMMENDED.

While the CGA generation function cryptographically ties the interface identifier of a home address to the subnet prefix of the home address, the function accepts any subnet prefix and hence does not prevent a node from cryptographically generating a home address with a spoofed subnet prefix. As a consequence, the CGA property of a home address does not guarantee the owner's reachability at the home address. This could be misused for a "return-to-home flooding

attack" [5], where the attacker uses its own public key to cryptographically generate a home address with a subnet prefix from a victim network, requests a correspondent node to bind this to the attacker's current care-of address, initiates the download of a large file via the care-of address, and finally deregisters the binding or lets it expire. The correspondent node would then redirect the packets being downloaded to the victim network identified by the subnet prefix of the attacker's spoofed home address. The protocol defined in this document performs a reachability test for the home address at the time the home address is first registered with the correspondent node. This precludes return-to-home flooding.

The verification of the CGA property of a mobile node's home address involves asymmetric public-key cryptography, which is relatively complex compared to symmetric cryptography. Enhanced Route Optimization mitigates this disadvantage through the use of symmetric cryptography after an initial public-key-based verification of the mobile node's home address has been performed. Specifically, the correspondent node assigns the mobile node a permanent home keygen token during the initial correspondent registration based on which the mobile node can authenticate to the correspondent node during subsequent correspondent registrations. Such authentication enables the correspondent node to bind a subsequent correspondent registration back to the initial public-key-based verification of the mobile node's home address. The permanent home keygen token is never sent in plain text; it is encrypted with the mobile node's public key when initially assigned, and irreversibly hashed during subsequent correspondent registrations.

6.2. Care-of Address Ownership

A secure proof of home address ownership can mitigate the threat of IP address stealing, but an attacker may still bind a correct home address to a false care-of address and thereby trick a correspondent node into redirecting packets, which would otherwise be delivered to the attacker itself, to a third party. Neglecting to verify a mobile node's reachability at its claimed care-of address could therefore cause one or multiple correspondent nodes to unknowingly contribute to a redirection-based flooding attack against a victim chosen by the attacker.

Redirection-based flooding attacks may target a single node, a link, or a router or other critical network device upstream of an entire network. Accordingly, the attacker's spoofed care-of address may be the IP address of a node, a random IP address from a subnet prefix of a particular link, or the IP address of a router or other network device. An attack against a network potentially impacts a larger number of nodes than an attack against a specific node, although

neighbors of a victim node on a broadcast link typically suffer the same damage as the victim itself.

Requiring mobile nodes to cryptographically generate care-of addresses in the same way as they generate home addresses would mitigate the threat of redirection-based flooding only marginally. While it would prevent an attacker from registering as its care-of address the IP address of a specific victim node, the attacker could still generate a different CGA-based care-of address with the same subnet prefix as that of the victim's IP address. Flooding packets redirected towards this care-of address would then not have to be received and processed by any specific node, but they would impact an entire link or network and thus cause comparable damage. CGA-based care-of addresses therefore have little effectiveness with respect to flooding protection. On the other hand, they would require a computationally expensive, public-key-based ownership proof whenever the care-of address changes. For these reasons, Enhanced Route Optimization uses regular IPv6 care-of addresses.

A common misconception is that a strong proof of home address ownership would mitigate the threat of redirection-based flooding and consequently eliminate the need to verify a mobile node's reachability at a new care-of address. This notion may originate from the specification of a base Mobile IPv6 home registration in [1], which calls for the authentication of a mobile node based on an IPsec security association, but does not require this to be supplemented by a verification of the mobile node's reachability at the care-of address. However, the reason not to mandate reachability verification for a home registration is in this case the existence of an administrative relationship between the home agent and the mobile node, rather than the fact that the home agent can securely verify the mobile node's home address ownership, or that the home registration is IPsec-protected. The administrative relationship with the mobile node allows the home agent, first, to trust in the correctness of a mobile node's care-of address and, second, to quickly identify the mobile node should it still start behaving maliciously, for example, due to infection by malware. Section 15.3 in [1] and Section 1.3.2 in [5] explain these prerequisites.

Assuming trust, an administrative relationship between the mobile node and its home agent is viable, given that the home agent is an integral part of the mobility services that a mobile user typically subscribes to, sets up her- or himself, or receives based on a business relationship. A Mobile IPv6 extension [14] that leverages a shared authentication key, preconfigured on the mobile node and the correspondent node, preassumes the same relationship between the mobile node and a correspondent node. While this assumption limits the applicability of the protocol (Section 2 of [14] acknowledges

this), it permits omission of care-of address reachability verification as in the case of the home registration. Enhanced Router Optimization does not make assumptions on the relationship between mobile and correspondent nodes. This renders the protocol applicable to arbitrary scenarios, but necessitates that correspondent nodes must verify a mobile node's reachability at every new care-of address.

6.3. Credit-Based Authorization

Enhanced Route Optimization enables mobile and correspondent nodes to resume bidirectional communications after a handoff on the mobile-node side before the mobile node's reachability at the new care-of address has been verified by the correspondent node. Such concurrency would in the absence of appropriate protection reintroduce the threat of redirection-based flooding, which reachability verification was originally designed to eliminate: Given that the correspondent node is in general unaware of the round-trip time to the mobile node, and since reachability verification may fail due to packet loss, the correspondent node must accept a sufficiently long concurrency period for reachability verification to complete. An attacker could misuse this to temporarily trick the correspondent node into redirecting packets to the IP address of a victim. The attacker may also successively postpone reachability verification in that it registers with the correspondent node anew, possibly with a different spoofed care-of address, shortly before the correspondent node's maximum permitted concurrency period elapses and the correspondent node switches to waiting for the completion of reachability verification without sending further packets. This behavior cannot necessarily be considered malicious on the correspondent node side since even a legitimate mobile node's reachability may fail to become verified before the mobile node's care-of address changes again. This may be due to high mobility on the mobile node side, or to persistent packet loss on the path between the mobile node and the correspondent node. It is generally non-trivial to decide on the correspondent node side whether the party at the other end behaves legitimately under adverse conditions or maliciously.

Enhanced Route Optimization eliminates the threat of redirection-based flooding despite concurrent reachability verification through the use of Credit-Based Authorization. Credit-Based Authorization manages the effort that a correspondent node expends in sending payload packets to a care-of address in UNVERIFIED state. This is accomplished based on the following three hypotheses:

1. A flooding attacker typically seeks to shift the burden of assembling and sending flooding packets to a third party. Bandwidth is an ample resource for many attractive victims, so the effort for sending the high rate of flooding packets required to impair the victim's ability to communicate may exceed the attacker's own capacities.
2. The attacker can always flood a victim directly by generating bogus packets itself and sending those to the victim. Such an attack is not amplified, so the attacker must be provisioned enough to generate a packet flood sufficient to bring the victim down.
3. Consequently, the additional effort required to set up and coordinate a redirection-based flooding attack pays off for the attacker only if the correspondent node can be tricked into contributing to and amplifying the attack.

Non-amplified redirection-based flooding is hence, from an attacker's perspective, no more attractive than pure direct flooding, where the attacker itself sends bogus packets to the victim. It is actually less attractive given that the attacker needs to maintain a context for mobility management in order to coordinate the redirection. On this basis, Credit-Based Authorization extinguishes the motivation for redirection-based flooding by preventing the amplification that could be reached through it, rather than eliminating malicious packet redirection in the first place. The ability to send unrequested packets is an inherent property of packet-oriented networks, and direct flooding is a threat that results from this. Since direct flooding exists with and without mobility support, it constitutes a reasonable measure in comparing the security provided by Enhanced Route Optimization to the security of the non-mobile Internet. Through the use of Credit-Based Authorization, Enhanced Route Optimization satisfies the objective to provide a security level comparable to that of the non-mobile Internet.

Since the perpetrator of a redirection-based flooding attack would take on the role of a mobile node, Credit-Based Authorization must be enforced on the correspondent node side. The correspondent node continuously monitors the effort that the mobile node spends in communicating with the correspondent node. The mobile node's effort is then taken as a limit on the effort that the correspondent node may spend in sending payload packets when the mobile node's care-of address is in UNVERIFIED state. The permission for the correspondent node to send a limited amount of payload packets to a care-of address in UNVERIFIED state enables immediate resumption of bidirectional communications once the mobile node has registered a new IP address with the correspondent node after a handoff.

If what appears to be a mobile node is in fact an attacker who tricks the correspondent node into redirecting payload packets to the IP address of a victim, Credit-Based Authorization ensures that the stream of flooding packets ceases before the effort that the correspondent node spends on generating the stream exceeds the effort that the attacker has recently spent itself. The flooding attack is therefore at most as effective as a direct flooding attack, and consequently fails to produce any amplification.

Another property of Credit-Based Authorization is that it does not assign a mobile node credit while its care-of address is in UNVERIFIED state. This deserves justification since it would technically be feasible to assign credit independent of the state of the mobile node's care-of address. However, the assignment of credit for packets received from a care-of address in UNVERIFIED state would introduce a vulnerability to sustained reflection attacks. Specifically, an attacker could cause a correspondent node to redirect packets for the attacker to the IP address of a victim, and sustain the packet flow towards the victim in that it continuously replenishes its credit by sending packets to the correspondent node. Although such a redirection-based reflection attack would fail to produce any amplification, it may still be appealing to an attacker who wishes to pursue an initial transport protocol handshake with the correspondent node -- which typically requires the attacker to receive some unguessable data -- and redirect the download to the victim's IP address afterwards. Credit-Based Authorization ensures that the attacker in this case cannot acquire additional credit once the download has been redirected, and thereby forces the attack to end quickly.

6.4. Time Shifting Attacks

Base Mobile IPv6 limits the lifetime of a correspondent registration to 7 minutes and so arranges that a mobile node's reachability at its home and care-of addresses is reverified periodically. This ensures that the return routability procedure's vulnerability to eavesdropping cannot be exploited by an attacker that is only temporarily on the path between the correspondent node and the spoofed home or care-of address. Such "time shifting attacks" [5] could otherwise be misused for off-path IP address stealing, return-to-home flooding, or flooding against care-of addresses.

Enhanced Route Optimization repeats neither the initial home address test nor any care-of address test in order to decrease handoff delays and signaling overhead. This does not limit the protocol's robustness to IP address stealing attacks because the required CGA-based ownership proof for home addresses already eliminates such attacks. Reachability verification does not add further protection in this regard. On the other hand, the restriction to an initial reachability verification facilitates time-shifted, off-path flooding attacks -- either against home addresses with incorrect prefixes or against spoofed care-of addresses -- if the perpetrator can interpose in the exchange before it moves to a different location.

The design choice against repeated home and care-of address tests was made based on the observation that time shifting attacks are already an existing threat in the non-mobile Internet of today. Specifically, an attacker can temporarily move onto the path between a victim and a correspondent node, request a stream of packets from the correspondent node on behalf of the victim, and then move to a different location. Most transport protocols do not verify an initiator's reachability at the claimed IP address after an initial verification during connection establishment. It enables an attacker to participate only in connection establishment and then move to an off-path position, from where it can spoof acknowledgments to feign continued presence at the victim's IP address. The threat of time shifting hence already applies to the non-mobile Internet.

It should still be acknowledged that the time at which Enhanced Route Optimization verifies a mobile node's reachability at a home or care-of address may well antecede the establishment of any transport layer connection. This gives an attacker more time to move away from the path between the correspondent node and the victim and so makes a time shifting attack more practicable. If the lack of periodic reachability verification is considered too risky, a correspondent node may enforce reruns of home or care-of address tests by limiting the registration lifetime, or by sending Binding Refresh Request messages to a mobile node.

6.5. Replay Attacks

The protocol specified in this document relies on 16-bit base Mobile IPv6 sequence numbers and periodic rekeying to avoid replay attacks. Rekeying allows mobile and correspondent nodes to reuse sequence numbers without exposing themselves to replay attacks. It must be pursued at least once every 24 hours due to the maximum permitted binding lifetime for correspondent registrations. Mobile and correspondent nodes also rekey whenever a rollover in sequence number space becomes imminent. This is unlikely to happen frequently, however, given that available sequence numbers are sufficient for up to 32768 correspondent registrations, each consisting of an early and a complete Binding Update message. The sequence number space thus permits an average rate of 22 correspondent registrations per minute without exposing a need to rekey throughout the 24-hour binding lifetime.

6.6. Resource Exhaustion

While a CGA-based home address ownership proof provides protection against unauthenticated Binding Update messages, it can expose a correspondent node to denial-of-service attacks since it requires computationally expensive public-key cryptography. Enhanced Route Optimization limits the use of public-key cryptography to only the first correspondent registration and if/when rekeying is needed. It is RECOMMENDED that correspondent nodes in addition track the amount of processing resources they spend on CGA-based home address ownership verification, and that they reject new correspondent registrations that involve public-key cryptography when these resources exceed a predefined limit. [2] discusses the feasibility of CGA-based resource exhaustion attacks in depth.

6.7. IP Address Ownership of Correspondent Node

Enhanced Route Optimization enables mobile nodes to authenticate a received Binding Acknowledgment message based on a CGA property of the correspondent node's IP address, provided that the correspondent node has a CGA. The mobile node requests this authentication by including a CGA Parameters Request option in the Binding Update message that it sends to the correspondent node, and the correspondent node responds by adding its CGA parameters and signature to the Binding Acknowledgment message within CGA Parameters and Signature options. Proving ownership of the correspondent node's IP address protects the mobile node from accepting a spoofed Binding Acknowledgment message and from storing the included permanent home keygen token for use during future correspondent registrations. Such an attack would result in denial of service against the mobile node because it would prevent the mobile node from transacting any binding

updates with the obtained permanent home keygen token. Enhanced Route Optimization recommends renewal of a permanent home keygen token in case of persistent correspondent registration failures, allowing mobile nodes to recover from denial-of-service attacks that involve spoofed permanent home keygen tokens.

The threat of the described denial-of-service attack is to some extent mitigated by requirements on the attacker's location: A Binding Update message that requests a correspondent node to provide a permanent home keygen token is authenticated based on the CGA property of the mobile node's home address. This authentication method involves a home address test, providing the mobile node with a home keygen token based on which it can calculate the authenticator of the Binding Update message. Since the mobile node expects the authenticator of the returning Binding Acknowledgment message to be calculated with the same home keygen token, an attacker that is willing to spoof a Binding Acknowledgment message that includes a permanent home keygen token must eavesdrop on the home address test. The attacker must hence be present on the path from the correspondent node to the mobile node's home agent while the home address test proceeds. Moreover, if the Binding Update message requesting the permanent home keygen token is complete, its authenticator is further calculated based on a care-of keygen token. The attacker must then also know this care-of keygen token to generate the authenticator of the Binding Acknowledgment message. This requires the attacker to be on the path from the correspondent node to the mobile node's current IP attachment at the time the correspondent node sends the care-of keygen token to the mobile node within a Care-of Test message or the Care-of Test option of a Binding Acknowledgment message.

Since a mobile node in general does not know whether a particular correspondent node's IP address is a CGA, the mobile node must be prepared to receive a Binding Acknowledgment message without CGA Parameters and Signature options in response to sending a Binding Update message with an included CGA Parameters Request option. Per se, this mandatory behavior may enable downgrading attacks where the attacker would send, on the correspondent node's behalf, a Binding Acknowledgment message without CGA Parameters and Signature options, claiming that the correspondent node's IP address is not a CGA. Enhanced Route Optimization mitigates this threat in that it calls for mobile nodes to prioritize Binding Acknowledgment messages with valid CGA Parameters and Signature options over Binding Acknowledgment messages without such options. This protects against downgrading attacks unless the attacker can intercept Binding Acknowledgment messages from the correspondent node. Given that the attacker must be on the path from the correspondent node to the mobile node's home agent at roughly the same time as explained above, the attacker may not be able to intercept the correspondent node's

Binding Acknowledgment messages. On the other hand, an attacker that can intercept Binding Acknowledgment messages from the correspondent node is anyway in a position where it can pursue denial of service against the mobile node and the correspondent node. This is a threat that already exists in the non-mobile Internet, and it is not specific to Enhanced Route Optimization.

External mechanisms may enable the mobile node to obtain certainty about whether a particular correspondent node's IP address is a CGA. The mobile node may then insist on an IP address ownership proof from the correspondent node, in which case it would discard any received Binding Acknowledgment messages that do not contain valid CGA Parameters and Signature options. One conceivable means for mobile nodes to distinguish between standard IPv6 addresses and CGAs might be an extension to the Domain Name System.

7. Protocol Constants and Configuration Variables

[2] defines a CGA Message Type namespace from which CGA applications draw CGA Message Type tags to be used in signature calculations. Enhanced Route Optimization uses the following constant, randomly generated CGA Message Type tag:

```
0x5F27 0586 8D6C 4C56 A246 9EBB 9B2A 2E13
```

[1] bounds the lifetime for bindings that were established with correspondent nodes by way of the return routability procedure to `MAX_RR_BINDING_LIFETIME`. Enhanced Route Optimization adopts this limit for bindings that are authenticated through a proof of the mobile node's reachability at the home address. However, the binding lifetime is limited to the more generous constant value of `MAX_CGA_BINDING_LIFETIME` when the binding is authenticated through the CGA property of the mobile node's home address:

```
MAX_CGA_BINDING_LIFETIME    86400 seconds
```

Credit aging incorporates two configuration variables to gradually decrease a mobile node's credit counter over time. It is RECOMMENDED that a correspondent node uses the following values:

```
CreditAgingFactor            7/8
CreditAgingInterval          5 seconds
```

8. IANA Considerations

This document defines the following six new mobility options, which must be assigned type values within the mobility option numbering space of [1]:

- o CGA Parameters Request mobility option (11)
- o CGA Parameters mobility option (12)
- o Signature mobility option (13)
- o Permanent Home Keygen Token mobility option (14)
- o Care-of Test Init mobility option (15)
- o Care-of Test mobility option (16)

This document allocates the following four new status codes for Binding Acknowledgment messages:

- o "Permanent home keygen token unavailable" (147)
- o "CGA and signature verification failed" (148)
- o "Permanent home keygen token exists" (149)
- o "Non-null home nonce index expected" (150)

The values to be assigned for these status codes must all be greater than or equal to 128, indicating that the respective Binding Update message was rejected by the receiving correspondent node.

This document also defines a new 128-bit value under the CGA Message Type namespace [2].

9. Acknowledgments

The authors would like to thank Tuomas Aura, Gabriel Montenegro, Pekka Nikander, Mike Roe, Greg O'Shea, Vesa Torvinen (in alphabetical order) for valuable and interesting discussions around cryptographically generated addresses.

The authors would also like to thank Marcelo Bagnulo, Roland Bless, Zhen Cao, Samita Chakrabarti, Greg Daley, Vijay Devarapalli, Mark Doll, Lakshminath Dondeti, Francis Dupont, Lars Eggert, Eric Gray, Manhee Jo, James Kempf, Suresh Krishnan, Tobias Kuefner, Lila Madour, Vidya Narayanan, Mohan Parthasarathy, Alice Qinxia, and Behcet

Sarikaya (in alphabetical order) for their reviews of and important comments on this document and the predecessors of this document.

Finally, the authors would also like to emphasize that [15] pioneered the use of cryptographically generated addresses in the context of Mobile IPv6 route optimization, and that this document consists largely of material from [16], [17], and [18] and the contributions of their authors.

10. References

10.1. Normative References

- [1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [2] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [3] Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels", IETF BCP 14, RFC 2119, March 1997.
- [4] Jonsson, J. and B. Kaliski, "Public-Key Cryptography Standards (PKCS) #1: RSA Cryptography Specifications Version 2.1", RFC 3447, February 2003.

10.2. Informative References

- [5] Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background", RFC 4225, December 2005.
- [6] Vogt, C. and J. Arkko, "A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization", RFC 4651, February 2007.
- [7] Vogt, C. and M. Doll, "Efficient End-to-End Mobility Support in IPv6", Proceedings of the IEEE Wireless Communications and Networking Conference, IEEE, April 2006.
- [8] Mirkovic, J. and P. Reiher, "A Taxonomy of DDoS Attack and DDoS Defense Mechanisms", ACM SIGCOMM Computer Communication Review, Vol. 34, No. 2, ACM Press, April 2004.
- [9] Arkko, J. and C. Vogt, "Credit-Based Authorization for Binding Lifetime Extension", Work in Progress, May 2004.

- [10] O'Shea, G. and M. Roe, "Child-Proof Authentication for MIPv6 (CAM)", ACM SIGCOMM Computer Communication Review, ACM Press, Vol. 31, No. 2, April 2001.
- [11] Nikander, P., "Denial-of-Service, Address Ownership, and Early Authentication in the IPv6 World", Revised papers from the International Workshop on Security Protocols, Springer-Verlag, April 2002.
- [12] Bagnulo, M. and J. Arkko, "Support for Multiple Hash Algorithms in Cryptographically Generated Addresses (CGAs)", Work in Progress, April 2007.
- [13] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [14] Perkins, C., "Securing Mobile IPv6 Route Optimization Using a Static Shared Key", RFC 4449, June 2006.
- [15] Roe, M., Aura, T., O'Shea, G., and J. Arkko, "Authentication of Mobile IPv6 Binding Updates and Acknowledgments", Work in Progress, March 2002.
- [16] Haddad, W., Madour, L., Arkko, J., and F. Dupont, "Applying Cryptographically Generated Addresses to Optimize MIPv6 (CGA-OMIPv6)", Work Progress, May 2005.
- [17] Vogt, C., Bless, R., Doll, M., and T. Kuefner, "Early Binding Updates for Mobile IPv6", Work in Progress, February 2004.
- [18] Vogt, C., Arkko, J., Bless, R., Doll, M., and T. Kuefner, "Credit-Based Authorization for Mobile IPv6 Early Binding Updates", Work in Progress, May 2004.

Authors' Addresses

Jari Arkko
Ericsson Research NomadicLab
FI-02420 Jorvas
Finland

EMail: jari.arkko@ericsson.com

Christian Vogt
Institute of Telematics
Universitaet Karlsruhe (TH)
P.O. Box 6980
76128 Karlsruhe
Germany

EMail: chvogt@tm.uka.de

Wassim Haddad
Ericsson Research
8400, Decarie Blvd
Town of Mount Royal
Quebec H4P 2N2, Canada

EMail: wassim.haddad@ericsson.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

