

Network Working Group
Request for Comments: 4651
Category: Informational

C. Vogt
Universitaet Karlsruhe (TH)
J. Arkko
Ericsson Research NomadicLab
February 2007

A Taxonomy and Analysis of Enhancements to
Mobile IPv6 Route Optimization

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

IESG Note:

This RFC is a product of the Internet Research Task Force and is not a candidate for any level of Internet Standard. The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment.

Abstract

This document describes and evaluates strategies to enhance Mobile IPv6 Route Optimization, on the basis of existing proposals, in order to motivate and guide further research in this context. This document is a product of the IP Mobility Optimizations (MobOpts) Research Group.

Table of Contents

1. Introduction	3
1.1. A Note on Public-Key Infrastructures	4
1.2. A Note on Source Address Filtering	5
2. Objectives for Route Optimization Enhancement	7
2.1. Latency Optimizations	8
2.2. Security Enhancements	8
2.3. Signaling Optimizations	9
2.4. Robustness Enhancements	9
3. Enhancements Toolbox	9
3.1. IP Address Tests	10
3.2. Protected Tunnels	10
3.3. Optimistic Behavior	11
3.4. Proactive IP Address Tests	11
3.5. Concurrent Care-of Address Tests	12
3.6. Diverted Routing	13
3.7. Credit-Based Authorization	14
3.8. Heuristic Monitoring	17
3.9. Crypto-Based Identifiers	18
3.10. Pre-Configuration	19
3.11. Semi-Permanent Security Associations	20
3.12. Delegation	21
3.13. Mobile Networks	21
3.14. Location Privacy	22
4. Discussion	22
4.1. Cross-Layer Interactions	23
4.2. Experimentation and Measurements	23
4.3. Future Research	24
5. Security Considerations	24
6. Conclusions	25
7. Acknowledgments	25
8. References	26
8.1. Normative References	26
8.2. Informative References	26

1. Introduction

Mobility support for IPv6, or Mobile IPv6, enables mobile nodes to migrate active transport connections and application sessions from one IPv6 address to another. The Mobile IPv6 specification, RFC 3775 [1], introduces a "home agent", which proxies a mobile node at a permanent "home address". A roaming mobile node connects to the home agent through a bidirectional tunnel and can so communicate, from its local "care-of address", as if it was present at the home address. The mobile node keeps the home agent updated on its current care-of address via IPsec-protected signaling messages [40].

In case the correspondent node lacks appropriate mobility support, it communicates with the mobile node's home address, and thus all data packets are routed via the home agent. This mode, Bidirectional Tunneling, increases packet-propagation delays. RFC 3775 hence defines an additional mode for Route Optimization, which allows peers to communicate on the direct path. It requires that the correspondent node can cache a binding between the mobile node's home address and current care-of address. The challenge with Route Optimization is that an administrative relationship between the mobile node and the correspondent node can generally not be presupposed. So how can the two authenticate and authorize the signaling messages that they exchange?

Mobile IPv6 solves this problem by verifying a routing property of the mobile node. Specifically, the mobile node is checked to be reachable at its home address and current care-of address in that it must prove the reception of a home and care-of keygen token, respectively. This is called the "return-routability procedure". It takes place right before a mobile node registers a new care-of address with a correspondent node and is periodically repeated in case the mobile node does not move for a while.

The advantage of the return-routability procedure is that it is lightweight and does not require pre-shared authentication material. It also requires no state at the correspondent node. On the other hand, the two reachability tests can lead to a handoff delay unacceptable for many real-time or interactive applications such as Voice over IP (VoIP) and video conferencing. Also, the security that the return-routability procedure guarantees might not be sufficient for security-sensitive applications. And finally, periodically refreshing a registration at a correspondent node implies a hidden signaling overhead that may prevent mobile nodes from hibernation during times of inactivity.

Manifold enhancements for Route Optimizations have hence been suggested. This document describes and evaluates various strategies

on the basis of existing proposals. It is meant to provide a conceptual framework for further work, which was found to be inevitable in the context of Route Optimization. Many scientists volunteered to review this document. Their names are duly recorded in Section 7. Section 2 analyzes the strengths and weaknesses of Route Optimization and identifies potential objectives for enhancement. Different enhancement strategies are discussed, based on existing proposals, in Section 3. Section 4 discusses the different approaches and identifies opportunities for further research. Section 5 and Section 6 conclude the document.

This document represents the consensus of the MobOpts Research Group. It has been reviewed by the Research Group members active in the specific area of work. At the request of their chairs, this document has been comprehensively reviewed by multiple active contributors to the IETF MIP6 Working Group. At the time of this writing, some of the ideas presented in this document have been adopted by the Mobility for IP: Performance, Signaling and Handoff Optimization (mipshop) Working Group in the IETF.

1.1. A Note on Public-Key Infrastructures

Mobile IPv6 Route Optimization verifies a mobile node's authenticity through a routing property. An alternative is cryptographic authentication, which requires a binding between a node's identity and some sort of secret information. Although some proposals suggest installing shared secrets into end nodes when possible (see Section 3.10), pre-configuration is not an option for general Internet use for scalability reasons. Authentication based on a Public-Key Infrastructure (PKI) does not require pair-wise pre-configuration. Here, the secret information is the private component of a public/private-key pair, and the binding between a node's identity and private key exists indirectly through the cryptographic properties of public/private-key pairs and a binding between the identity and the public key. An authority trusted by both end nodes issues a certificate that effects this latter binding.

Large-scale use of a PKI, however, was considered unsuitable for mobility management due to the following reasons.

- o There are differing opinions on whether a PKI could scale up to hundreds of millions of mobile nodes. Some people argue they do, as there are already examples of certification authorities responsible for millions of certificates. But more important than the expected increase in the number of certificates would be a shift in application patterns. Nowadays, public-key cryptography is used only for those applications that require strong, cryptographic authentication.

If it was used for mobility management as well, certificate checks would become mandatory for any type of application, leading to more checks per user. Busy servers with mobility support might be unwilling to spend the processing resources required for this depending on the service they provide.

- o Revoked certificates are identified on Certificate Revocation Lists (CRLs), which correspondent nodes with mobility support would have to acquire from certification authorities. CRLs must be kept up to date, requiring periodic downloads. This and the act of checking a certificate against a CRL create overhead that some correspondent nodes might be unwilling to spend.
- o Certificate verification may take some time and hence interrupt ongoing applications. This can be disturbing from the user's perspective, especially when Route Optimization starts in the middle of a session, or the session is very short-term anyway.
- o The bigger a PKI grows, the more attractive it becomes as an attack target, endangering the Internet as a whole.
- o There is little experience with using home addresses as identifiers in certificates. Although the home address could theoretically be placed into a certificate's Subject Alternate Name field, the entities responsible for IP-address assignment and certification are usually not the same, and it may not be easy to coordinate the two.

For these reasons, this document does not consider direct authentication of mobile nodes based on a PKI. Nevertheless, it does evaluate certificate-based techniques that make the problems identified above more tractable (see Section 3.12).

1.2. A Note on Source Address Filtering

RFC 3775 uses care-of-address tests to probe a mobile node's presence at its claimed location. Alternatively, verification of care-of addresses may be based on infrastructure in the mobile node's local access network. For instance, the infrastructure can verify that the IP source addresses of all packets leaving the network are correct. "Ingress filtering" [38][43] provides this feature to the extent that it inspects the prefix of IP source addresses and ensures topological correctness. Network-access providers that use ingress filtering normally deploy the technique in their first-hop and site-exit routers. Similarly, ISPs may filter packets originating from a downstream network.

Ingress filtering may eventually provide a way to replace care-of-address tests. But there are still a number of uncertainties today:

- o By definition, ingress filtering can prevent source-address spoofing only from those networks that do deploy the technique. As a consequence, ingress filtering needs to be widely, preferably universally, deployed in order to constitute Internet-wide protection. As long as an attacker can get network access without filters, all Internet nodes remain vulnerable.
- o There is little incentive for ISPs to deploy ingress filtering other than conscientiousness. Legal or regulatory prescription as well as financial motivation does not exist. A corrupt ISP might even have a financial incentive not to deploy the technique, if redirection-based denial-of-service (DoS) attacks using Route Optimization ever become possible and are exploited for financial gain. A similar issue was observed with, for example, email spam.
- o Ingress filtering is most effective, and easiest to configure, at the first-hop router. However, since only prefixes are checked, the filters inevitably get less precise the further upstream they are enforced. This issue is inherent in the technique, so the best solution is checking packets as close to the originating nodes as possible, preferably in the first-hop routers themselves.
- o A popular implementation of ingress filtering is "Reverse Path Forwarding" (RPF). This technique relies on routes to be symmetric, which is oftentimes the case between edge networks and ISPs, but far less often between peering ISPs. Alternatives to RPF are either manually configured access lists or dynamic approaches that are more relaxed, and thereby less secure, than RPF [43].
- o Another problem with ingress filtering is multi-homing. When a router attempts to forward to one ISP a packet with a source-address prefix from another ISP, filters at the second ISP would block the packet. The IETF seeks to find a way around this [39]. For instance, one could tunnel the packet to the topologically correct ISP, or one could allow source-address changes by means of a locator-identifier split [45].
- o Finally, RFC 3775 defines an Alternative Care-of Address option that mobile nodes can use to carry a care-of address within a Binding Update message outside of the IPv6 header. Such an address is not subject to inspection by ingress filtering and would have to be verified through other means [14].

Although these problems are expected to get solved eventually, there is currently little knowledge on how applicable and deployable, as a candidate for care-of-address verification, ingress filtering will be. High investments or administrative hurdles could prevent a large, preferably universal deployment of ingress filtering, which would hinder Internet-wide protection, as mentioned in the first bullet. For these reasons, this document does not consider ingress filtering as a viable alternative to care-of-address tests, although things may be different in the future.

2. Objectives for Route Optimization Enhancement

Wireless environments with frequently moving nodes feature a number of salient properties that distinguish them from environments with stationary nodes or nodes that move only occasionally. One important aspect is the efficiency of mobility management. Nodes may not bother about a few round-trip times of handoff latency if they do not change their point of IP attachment often. But the negative impact that a mobility protocol can have on application performance increases with the level of mobility. Therefore, in order to maximize user satisfaction, it is important to reduce the handoff latency that the mobility protocol adds to existing delays in other places of the network stack. A related issue is the robustness of the mobility protocol, given that temporary outage of mobility support can render mobile nodes incapable of continuing to communicate.

Furthermore, the wireless nature of data transmissions makes it potentially easier for an attacker to eavesdrop on other nodes' data or send data on behalf of other nodes. While applications can usually authenticate and encrypt their payload if need be, similar security measures may not be feasible for signaling packets of a mobility protocol, in particular if communicating end nodes have no pre-existing relationship.

Given the typically limited bandwidth in a wireless medium, resources ought to be spent in an economic matter. This is especially important for the amount of signaling that a mobility protocol requires.

Endeavors to enhance RFC 3775 Route Optimization generally strive for reduced handoff latency, higher security, lower signaling overhead, or increased protocol robustness. These objectives are herein discussed from a requirements perspective; the technical means to reach the objectives is not considered, nor is the feasibility of achieving them.

2.1. Latency Optimizations

One important objective for improving Route Optimization is to reduce handoff latencies. Assuming that the home-address test dominates the care-of-address test in terms of latency, a Mobile IPv6 handoff takes one round-trip time between the mobile node and the home agent for the home registration, a round-trip time between the mobile node and the home agent plus a round-trip time between the home agent and the correspondent node for the home-address test, and a one-way time from the mobile node to the correspondent node for the propagation of the Binding Update message. The first packet sent to the new care-of address requires an additional one-way time to propagate from the correspondent node to the mobile node. The mobile node can resume communications 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.

These delays are additive and are not subsumed by 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.

2.2. Security Enhancements

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 [46]. As such, it protects against impersonation, denial of service, and redirection-based flooding attacks that would not be possible without Route Optimization. This approach is based on an assumption that a mobile Internet cannot become any safer than the non-mobile Internet.

Applications that require a security level higher than what the return-routability procedure can provide are generally advised to use end-to-end protection such as IPsec or Transport Layer Security (TLS). But even then they are vulnerable to denial of service. This motivates research for stronger Route Optimization security. Security enhancements may also become necessary if future technological improvements mitigate some of the existing mobility-unrelated vulnerabilities.

One particular issue with Route Optimization is location privacy because route-optimized packets carry both home and care-of addresses in plaintext. A standard workaround is to fall back to Bidirectional Tunneling when location privacy is needed. Packets with the care-of address are then transferred only between the mobile node and the

home agent, where they can be encrypted through IPsec Encapsulating Security Payload (ESP) [42]. But even Bidirectional Tunneling requires the mobile node to periodically re-establish IPsec security associations with the home agent so as to become untraceable through Security Parameter Indexes (SPIs).

2.3. Signaling Optimizations

Route Optimization requires periodic signaling even when the mobile node does not move. The signaling overhead amounts to 7.16 bits per second if the mobile node communicates with a stationary node [6]. 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 refreshes consume a fraction of the wireless bandwidth that one could use more efficiently. Optimizations for reduced signaling overhead could mitigate these issues.

2.4. Robustness Enhancements

Route Optimization could conceptually enable continued communications during periods of temporary home-agent unavailability. The protocol defined in RFC 3775 does not achieve this independence, however, as the home agent plays an active role in the return-routability procedure. Appropriate enhancements could increase the independence from the home agent and thus enable robust Route Optimization even in the absence of the home agent.

3. Enhancements Toolbox

A large body of effort has recently gone into improving Mobile IPv6 Route Optimization. Some of the proposed techniques are modifications to the return-routability procedure, while others replace the procedure by alternative mechanisms. Some of them operate end-to-end; others introduce network-side mobility support. In most cases, it is the combination of a set of techniques that is required to gain a complete -- that is, efficient and secure -- route-optimization mechanism.

3.1. IP Address Tests

RFC 3775 uses IP-address tests to ensure that a mobile node is live and on the path to a specific destination address: The home-address test provides evidence that the mobile node is the legitimate owner of its home address; the care-of-address test detects spoofed care-of addresses and prevents redirection-based flooding attacks. Both tests can be performed in parallel.

A home-address test should be initiated by the mobile node so that the correspondent node can delay state creation until the mobile node has authenticated. The care-of-address test can conceptually be initiated by either side. It originates with the mobile node in RFC 3775, but with the correspondent node in [16] and [22]. The correspondent-node-driven approach suggests itself when authentication is done through other means than a home-address test.

Important advantages of IP-address tests are zero-configurability and the independence of ancillary infrastructure. As a disadvantage, IP-address tests can only guarantee that a node is on the path to the probed address, not that the node truly owns this address. This does not lead to new security threats, however, because the types of attacks that an on-path attacker can do with Route Optimization are already possible in the non-mobile Internet [46].

3.2. Protected Tunnels

RFC 3775 protects certain signaling messages, exchanged between a mobile node and its home agent, through an authenticated and encrypted tunnel. This prevents unauthorized nodes on that path, including eavesdroppers in the mobile node's wireless access network, from listening in on these messages.

Given that a pre-existing end-to-end security relationship between the mobile node and the correspondent node cannot generally be assumed, this protection exists only for the mobile node's side. If the correspondent node is immobile, the path between the home agent and the correspondent node remains unprotected. This is a path between two stationary nodes, so all types of attacks that a villain could wage on this path are already possible in the non-mobile Internet. In case the correspondent node is mobile, it has its own home agent, and only the path between the two (stationary) home agents remains unprotected.

3.3. Optimistic Behavior

Many Mobile IPv6 implementations [29][31] defer a correspondent registration until the associated home registration has been completed successfully. In contrast to such "conservative" behavior, a more "optimistic" approach is to begin the return-routability procedure in parallel with the home registration [52]. Conservative behavior avoids a useless return-routability procedure in case the home registration fails. This comes at the cost of additional handoff delay when the home registration is successful. Optimistic behavior saves this delay, but the return-routability procedure will be in vain should the corresponding home registration be unsuccessful.

While a parallelization of the home registration and the return-routability procedure is feasible within the bounds of RFC 3775, the specification does not permit mobile nodes to continue with the correspondent registration, by sending a Binding Update message to the correspondent node, until a Binding Acknowledgment message indicating successful home registration has been received. This is usually not a problem because the return-routability procedure is likely to take longer than the home registration anyway. However, some optimizations (see Section 3.4) reduce the delay caused by the return-routability procedure. A useful improvement is then to allow Binding Update messages to be sent to correspondent nodes even before the home registration has been acknowledged.

The drawback of optimistic behavior is that a lost, reordered, or rejected Binding Update message can cause data packets to be discarded. Nevertheless, packet loss would have similar negative impacts on conservative approaches, so the mobile node needs to be prepared for the possible loss of these packets in any case.

3.4. Proactive IP Address Tests

The critical handoff phase, during which the mobile node and the correspondent node cannot fully communicate, spans the home registration and the correspondent registration, including the return-routability procedure. One technique to shorten this phase is to accomplish part of the signaling proactively before the handoff. In particular, the home-address test can be done in advance without violating the specifications of RFC 3775 [52][51].

In order to have a fresh home keygen token ready for a future handoff, the mobile node should initiate a proactive home-address test at least once per token lifetime, that is, every 3.5 minutes. This does at most double the signaling overhead spent on home-address tests given that correspondent registrations must be refreshed every

7 minutes even when the mobile node does not move for a while. An optimization is possible where the mobile node's local link layer can anticipate handoffs and trigger the home-address test in such a case. [6] or [54] reduce the frequency of home-address tests even further. Proactive care-of-address tests are possible only if the mobile node is capable of attaching to two networks simultaneously. Dual attachment is possible if the link-layer technology enables it with a single interface [10], or if the mobile node is endowed with multiple interfaces [7].

3.5. Concurrent Care-of Address Tests

Without the assumption that a mobile node can simultaneously attach to multiple networks, proactive care-of-address tests, executed prior to handoff, are not an option. A correspondent node may instead authorize a mobile node to defer the care-of-address test until an early, tentative binding has been registered [52][51]. This in combination with a technique to eliminate the handoff delay of home-address tests (see Section 3.4 and Section 3.9) facilitates early resumption of bidirectional communications subsequent to handoff. The care-of address is called "unverified" during the concurrent care-of-address test, and it is said to be "verified" once the correspondent node has obtained evidence that the mobile node is present at the address. A tentative binding's lifetime can be limited to a few seconds.

Home-address tests must not be accomplished concurrently, however, given that they serve the purpose of authentication. They guarantee that only the legitimate mobile node can create or update a binding pertaining to a particular home address.

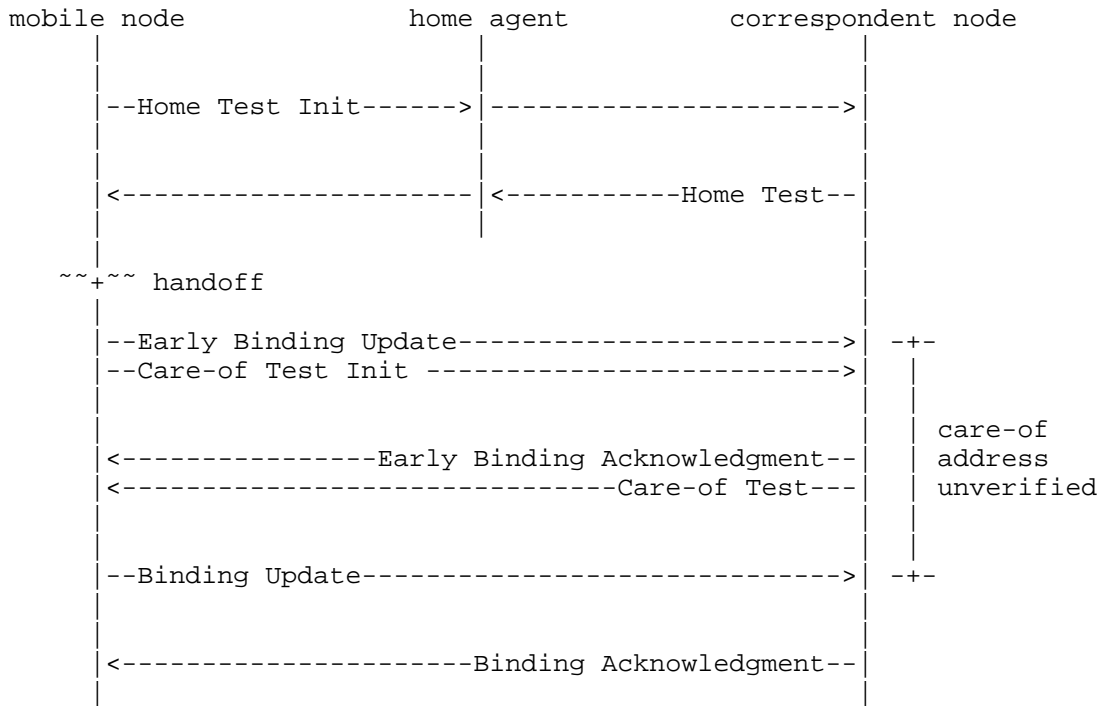


Figure 1: Concurrent Care-of Address Tests

Figure 1 illustrates how concurrent care-of-address tests are used in [52][51]: As soon as the mobile node has configured a new care-of address after a handoff, it sends to the correspondent node an Early Binding Update message. Only a home keygen token, obtained from a proactive home-address test, is required to sign this message. The correspondent node creates a tentative binding for the new, unverified care-of address when it receives the Early Binding Update message. This address can be used immediately. The mobile node finally sends a (standard) Binding Update message to the correspondent node when the concurrent care-of-address test is complete. Credit-Based Authorization (see Section 3.7) prevents misuse of care-of addresses while they are unverified.

3.6. Diverted Routing

Given that a home registration is faster than a correspondent registration in the absence of additional optimizations, the mobile node may request its traffic to be routed through the home address until a new binding has been set up at the correspondent node [52][51]. The performance of such diverted routing depends on the propagation properties of the involved routes, however.

For packets to be diverted via the home address, signaling is necessary with both the home agent and the correspondent node. The home agent must be informed about the new care-of address so that it can correctly forward packets intercepted at the home address. The correspondent node continues to send packets to the old care-of address until it receives a Binding Update message indicating that the current binding is no longer valid and ought to be removed. This request requires authentication through a home-address test in order to prevent denial of service by unauthorized nodes. The test can be accomplished in a proactive way (see Section 3.4).

The mobile node may send packets via the home address as soon as it has dispatched the Binding Update message to the home agent. It may send outgoing packets along the direct path once a Binding Update message for the new care-of address has been sent to the correspondent node.

It depends on the propagation latency on the end-to-end path via the home agent relative to the latency on the direct path for how long the correspondent node should continue to send packets to the home address. If the former path is slow, it may be better to queue some of the packets until the correspondent registration is complete and packets can be sent along the direct route.

3.7. Credit-Based Authorization

Concurrent care-of-address tests (see Section 3.5) require protection against spoofed unverified care-of addresses and redirection-based flooding attacks. Credit-Based Authorization [50] is a technique that provides such protection based on the following three hypotheses:

1. A flooding attacker typically seeks to somehow multiply the packets it assembles for the purpose of the attack because bandwidth is an ample resource for many attractive victims.
2. An attacker can always cause unamplified flooding by generating bogus packets itself and sending them to its victim directly.
3. Consequently, the additional effort required to set up a redirection-based flooding attack pays off for the attacker only if amplification can be obtained this way.

On this basis, rather than eliminating malicious packet redirection in the first place, Credit-Based Authorization prevents any amplification that can be reached through it. This is accomplished by limiting the data a correspondent node can send to an unverified care-of address of a mobile node by the data that the correspondent

node has recently received from that mobile node. (See Section 3.5 for a definition on when a care-of address is verified and when it is unverified.) A redirection-based flooding attack is thus 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 must keep Mobile IPv6 state to coordinate the redirection.

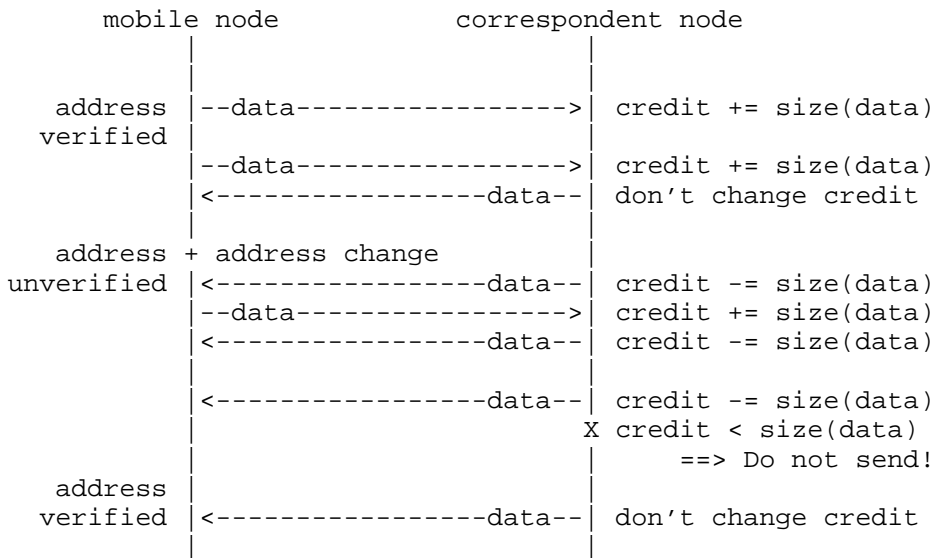


Figure 2: Credit-Based Authorization

Figure 2 illustrates Credit-Based Authorization for an exemplifying exchange of data packets: The correspondent node measures the bytes received from the mobile node. When the mobile node registers a new care-of address, the correspondent node labels this address "unverified" and sends packets there as long as the sum of the packet sizes does not exceed the measured, received data volume. A concurrent care-of-address test is meanwhile performed. Once the care-of address has been verified, the correspondent node relabels the address from "unverified" to "verified". Packets can then be sent to the new care-of address without restrictions. When insufficient credit is left while the care-of address is still "unverified", the correspondent node stops sending further packets to the address until the verification completes. The correspondent node may drop these packets, direct them to the mobile node's home address, or buffer them for later transmission when the care-of address is verified. Figure 2 does not show Mobile IPv6 signaling packets.

The correspondent node ensures that the mobile node's acquired credit gradually decreases over time. This "aging" prevents the mobile node from building up credit over a long time. A malicious node with a slow Internet connection could otherwise provision for a burst of redirected packets that does not relate to its own upstream capacity.

Allocating the mobile node's credit based on the packets that the mobile node sends and reducing the credit based on packets that the mobile node receives is defined as "Inbound Mode". (The correspondent node is in control of credit allocation, and it computes the credit based on inbound packets received from the mobile node.) A nice property of Inbound Mode is that it does not require support from the mobile node. The mobile node neither needs to understand that Credit-Based Authorization is effective at the correspondent node, nor does it have to have an idea of how much credit it has at a particular point in time.

Inbound Mode works fine with applications that send comparable data volumes into both directions. On the other hand, the mode may prevent the mobile node from collecting the amount of credit it needs for a handoff when applications with asymmetric traffic patterns are in use. For instance, file transfers and media streaming are characterized by high throughput towards the client, typically the mobile node, and comparably little throughput towards the serving correspondent node.

An additional "Outbound Mode" was designed to better accommodate applications with asymmetric traffic patterns. In Outbound Mode, packets that the correspondent node sends to the mobile node determine both, how much the credit increases while the current care-of address is verified, and how much the credit shrinks while the care-of address is unverified. This resolves the issue with asymmetric traffic patterns.

The security of Outbound Mode is based on the further hypothesis that the mobile node invests comparable effort for packet reception and transmission in terms of bandwidth, memory, and processing capacity. This justifies why credit, allocated for packets received by the mobile node, can be turned into packets that the correspondent node sends. The question is, though, how the correspondent node can determine how many of the packets sent to a mobile node are actually received and processed by that mobile node. Relying on transport-layer acknowledgments is not an option as such messages can easily be faked. Outbound Mode hence defines its own feedback mechanism, Care-of Address Spot Checks, which is robust to spoofing. The correspondent node periodically tags packets that it sends to the mobile node with a random, unguessable number, a so-called Spot Check Token. When the mobile node receives a packet with an attached Spot

Check Token, it buffers the token until it sends the next packet to the correspondent node. The Spot Check Token is then included in this packet. Upon reception, the correspondent node verifies whether the returned Spot Check Token matches a token recently sent to the mobile node. New credit is allocated in proportion to the ratio between the number of successfully returned Spot Check Tokens and the total number of tokens sent. This implies that new credit is approximately proportional to the fraction of packets that have made their way at least up to the mobile node's IP stack. The preciseness of Care-of Address Spot Checks can be traded with overhead through the frequency with which packets are tagged with Spot Check Tokens.

An interesting question is whether Outbound Mode could be misused by an attacker with asymmetric Internet connection. Widespread digital subscriber lines (DSL), for example, typically have a much higher download rate than upload rate. The limited upload rate would render most denial-of-service attempts through direct flooding meaningless. But the attacker could leverage the strong download rate to build up credit at one or multiple correspondent nodes. It could then illegitimately spend the credit on a stronger, redirection-based flooding attack. The reason why this has so far not been considered an issue is that, in order to accumulate enough credit at the remote end, the attacker would first have to expose itself to the same packet flood that it could then redirect towards the victim.

3.8. Heuristic Monitoring

Heuristic approaches to prevent misuse of unverified care-of addresses (see Section 3.5) are conceivable as well. A heuristic, implemented at the correspondent node and possibly supplemented by a restrictive lifetime limit for tentative bindings, can prevent, or at least effectually discourage such misuse. The challenge here seems to be a feasible heuristic: On one hand, the heuristic must be sufficiently rigid to quickly respond to malicious intents at the other side. On the other hand, it should not have a negative impact on a fair-minded mobile node's communications.

Another problem with heuristics is that they are usually reactive. The correspondent node can only respond to misbehavior after it appeared. If sanctions are imposed quickly, attacks may simply not be worthwhile. Yet premature measures should be avoided. One must also bear in mind that an attacker may be able to use different home addresses, and it is in general impossible for the correspondent node to see that the set of home addresses belongs to the same node. The attacker may furthermore exploit multiple correspondent nodes for its attack in an attempt to amplify the result.

3.9. Crypto-Based Identifiers

A Crypto-Based Identifier (CBID) is an identifier with a strong cryptographic binding to the public component of its owner's public/private-key pair [33]. This allows the owner to prove its claim on the CBID: It signs a piece of data with its private key and sends this to the verifier along with its public key and the parameters necessary to recompute the CBID. The verifier recomputes the CBID and checks the owner's knowledge of the corresponding private key.

CBIDs offer three main advantages: First, spoofing attacks against a CBID are much harder than attacks against a non-cryptographic identifier like a domain name or a Mobile IPv6 home address. Though an attacker can always create its own CBID, it is unlikely to find a public/private-key pair that produces someone else's. Second, a CBID does not depend on a PKI given its inherent binding to the owner's public key. Third, a CBID can be used to bind a public key to an IP address, in which case it is called a Cryptographically Generated Address (CGA) [44][34][47]. A CGA is syntactically just an ordinary IPv6 address. It has a standard routing prefix and an interface identifier generated from a hash on the CGA owner's public key and additional parameters.

Many applications are conceivable where CGAs are advantageous. In Mobile IPv6, CGAs can bind a mobile node's home address to its public key [35][5] and so avoid the home-address test in most correspondent registrations. This accelerates the registration process and allows the peers to communicate independently of home-agent availability.

Since only the interface identifier of a CGA is cryptographically protected, its network prefix can be spoofed, and flooding attacks against networks are still an issue. An initial home-address test is hence required to validate the network prefix even when the home address is a CGA. For the same reason, CGAs are rarely used as care-of addresses.

One limitation of CGAs compared to other types of CBIDs is that the cryptographically protected portion is only at most 62 bits long. The rest of the address is occupied by a 64-bit network prefix as well as the universal/local and individual/group bits. (The specification in [44] further hard-codes a 3-bit security parameter into the address, reducing the cryptographically protected portion to 59 bits.) A brute-force attack might thus reveal a public/private key public/private-key pair that produces a certain CGA. This vulnerability can be contained by including the network prefix in the hash computation for the interface identifier so that an attacker, in

case it did find the right public/private key public/private-key pair, could not form CGAs for multiple networks from it.

To resolve collisions in generating CGAs, a collision count is part of the input to the hash function. Changing this produces a different CGA. Unfortunately, the collision count also reduces the complexity of a brute-force attack against a CGA because it allows the same private/public-key pair to be used to generate multiple CGAs. The collision count is therefore limited to a few values only.

Higher security can be achieved through longer CBIDs. For example, a node's primary identifier in the Host Identity Protocol [21] is a 128-bit hash on the node's public key. It is used as an IP-address replacement at stack layers above IP. This CBID is not routable, so there needs to be some external localization mechanism if a node wants to contact a peer of which it only knows the identifier.

3.10. Pre-Configuration

Where mobile and correspondent nodes can be pre-configured with a shared key, bound to the mobile node's home address, authentication through a home-address test can be replaced by a cryptographic mechanism. This has three advantages. First, cryptography allows for stronger authentication than address tests. Second, strong authentication facilitates binding lifetimes longer than the 7-minute limit that RFC 3775 defines for correspondent registrations. Third, handoff delays are usually shorter with cryptographic approaches because the round-trips of the home-address test can be spared. The disadvantage of pre-configuration is its limited applicability.

Two proposals for pre-configuration are currently under discussion within the IETF. [25] endows mobile nodes with the information they need to compute home and care-of keygen tokens themselves rather than having to obtain them through the return-routability procedure. [15] uses the Internet Key Exchange protocol to establish an IPsec security association between the peers.

From a technical standpoint, pre-configuration can only replace a home-address test. A test of the care-of address is still necessary to verify the mobile node's presence at that address. The problem is circumvented in [25] by postulating that the correspondent node has sufficient trust in the mobile node to believe that the care-of address is correct. This assumption discourages the use of pre-configuration in scenarios where such trust is unavailable, however. For example, a mobile-phone operator may be able to configure subscribers with secret keys for authorization to a particular service, but it may not be able to vouch that all subscribers use

this service in a responsible manner. And even if users are trustworthy, their mobile nodes may become infected with malware and start behaving unreliably.

Another way to avoid care-of-address verification is to rely on access networks to filter out packets with incorrect IP source addresses [38][43]. This approach is taken in [15]. The problem with local filtering is that it can only protect a network from becoming the source of an attack, not from falling victim to an attack. The technique is hence potentially unreliable unless deployed in access networks worldwide (see Section 1.2).

Care-of-address tests facilitate the use of pre-configuration in spite of lacking trust relationships or the existence of access networks without local filtering techniques. For increased performance, concurrent care-of-address tests can be used in combination with Credit-Based Authorization or heuristic monitoring.

3.11. Semi-Permanent Security Associations

A compromise between the return-routability procedure and pre-configuration are semi-permanent security associations. A semi-permanent security association is established between a mobile node and a correspondent node upon first contact, and it is used to authenticate the mobile node during subsequent correspondent registrations. Semi-permanent security associations eliminate the need for periodic home-address tests and permit correspondent registrations with lifetimes longer than the 7-minute limit specified in RFC 3775.

It is important to verify a mobile node's home address before a security association is bound to it. An impersonator could otherwise create a security association for a victim's IP address and then redirect the victim's traffic at will until the security association expires. An initial home-address test mitigates this vulnerability because it requires the attacker to be on the path between the victim and the victim's peer at least while the security association is being established. Stronger security can be obtained through cryptographically generated home addresses (see Section 3.9).

Semi-permanent security associations alone provide no verification of care-of addresses and must therefore be supplemented by care-of-address tests. These may be performed concurrently for reduced handoff delays. Semi-permanent security associations were first developed in [8] where they were called "purpose-built keys".

3.12. Delegation

Section 1.1 lists numerous problems of PKIs with respect to authentication of mobile nodes. These problems become more tractable, however, if correspondent nodes authenticate home agents rather than mobile nodes, and the home agents vouch for the authenticity and trustworthiness of the mobile nodes [37]. Such delegation of responsibilities solves the scalability issue with PKIs given that home agents can be expected to be much less numerous than mobile nodes. Certificate revocation becomes less delicate as well because home agents are commonly administrated by a mobility provider and should as such be more accountable than mobile nodes.

Another advantage of delegation is that it avoids public-key computations at mobile nodes. On the other hand, the processing overhead at correspondent nodes increases. This may or may not be an issue depending on resources available at the correspondent node relative to the services that the correspondent node provides. The correspondent node may also be mobile itself, in which case cryptographic operations would be problematic. Furthermore, the increased overhead implies a higher risk to resource-exhaustion attacks.

3.13. Mobile Networks

Mobile nodes may move as a group and attach to the Internet via a "mobile router" that stays with the group. This happens, for example, in trains or aircraft where passengers communicate via a local wireless network that is globally interconnected through a satellite link.

It is straightforward to support such network mobility [41] with a single home agent and a tunnel between the mobile router and this home agent. The mobile nodes themselves then do not have to be mobility-aware. However, Route Optimization for moving networks [36][26][27][55] is more complicated. One possibility is to have the mobile router handle Route Optimization on behalf of the mobile nodes. This requires the mobile router to modify incoming and outgoing packets such that they can be routed on the direct path between the end nodes. The mobile router would also have to perform Mobile IPv6 signaling on behalf of the mobile nodes. Similarly, a network of correspondent nodes can communicate with mobile nodes, through a "correspondent router", in a route-optimized way without providing mobility support themselves.

3.14. Location Privacy

RFC 3775 fails to conceal a mobile node's current position as route-optimized packets always carry both home and care-of addresses. Both the correspondent node and a third party can therefore track the mobile node's whereabouts. A workaround is to fall back to bidirectional tunneling where location privacy is needed. Packets carrying the mobile node's care-of address are thus only transferred between the mobile node and the home agent, where they can be encrypted through IPsec ESP [42]. But even then should the mobile node periodically re-establish its IPsec security associations so as to become untraceable through its SPIs. Early efforts on location privacy in Route Optimization include [17][13][24][30].

4. Discussion

Common to the proposals discussed in Section 3 is that all of them affect a trade-off between effectiveness, on one hand, and economical deployability, administrative overhead, and wide applicability, on the other. Effectiveness may be equated with low latency, strong security, reduced signaling, or increased robustness. Economy implies no, or only moderate requirements in terms of hardware upgrades and software modifications. Administrative overhead relates to the amount of manual configuration and intervention that a technique needs.

The standard return-routability procedure avoids costly pre-configuration or new network entities. This minimizes both deployment investments as well as administrative expenses. Variants with optimistic behavior and proactive or concurrent IP-address tests have these advantages as well. CBIDs allow for public-key authentication without a PKI. They constitute a more secure alternative to home-address tests and are as such most effective when combined with concurrent reachability verification. CBID-based authentication may require nodes to be programmed with a mapping between human-readable identifiers and the corresponding CBIDs. Pre-configuration is another approach to avoid home-address tests. It does without computationally expensive public-key algorithms, but requires pair-wise credentials and, therefore, administrative maintenance. Where suitable infrastructure is available, end nodes may delegate authentication and encryption tasks to trusted network entities which, in turn, vouch for the end nodes. Delegation could resurrect the use of certificates for the purpose of mobility support. But it introduces a dependency on the delegates, adds the provisioning costs for new network entities, and is likely to be limited to communities of authorized nodes.

4.1. Cross-Layer Interactions

The performance of Route Optimization, as evaluated in this document, should be put into perspective of handoff-related activities in other parts of the network stack. These include link-layer attachment procedures; link-layer security mechanisms such as negotiation, authentication, and key agreement; as well as IPv6 router discovery, address configuration, and movement detection. A complete network attachment in a typical IEEE 802.11 commercial deployment requires over twenty link- and IP-layer messages. Current protocol stacks also have a number of limitations in addition to long attachment delays, such as denial-of-service vulnerabilities, difficulties in trusting a set of access nodes distributed to physically insecure locations, or the inability to retrieve sufficient information for making a handoff decision [2].

A number of proposals have been put forth to improve handoff performance on different parts of the network stack, mostly focusing on handoff performance. These include link-layer parameter tuning [49] and network-access authentication [18][2][32], as well as IPv6 router discovery [11][12], address configuration [23], and movement detection [19][20]. It is uncertain how far this optimization can be taken by only looking at the different parts individually. An integrated approach may eventually become necessary [4][53].

4.2. Experimentation and Measurements

The number and diversity of mobility-related activities within a typical network stack oftentimes render theoretical analyses insufficient and call for additional, extensive experimentation or simulation. The following is a non-exhaustive list of areas where practical experience is likely to yield valuable insight.

- o Conception of a set of standard scenarios that can be used as a reference for comparable measurements and experimentation. Ideally, such standard scenarios ought to be derived from real-world environments, and they should include all features that would likely be needed in a commercial deployment. These features include link-layer access control, for instance.
- o Measurements of the performance impacts that existing enhancement proposals have on the different parts of the stack.
- o Comparisons of different implementations that are based on the same specification. For instance, it would be valuable to know how much implementations differ with regards to the use of parallelism that RFC 3775 allows in home and correspondent registrations.

- o Measurements of the impact that network conditions such as packet loss can have on existing and new optimizations.
- o Statistical data collection on the behavior of mobile nodes in different networks. Several Route Optimization techniques behave differently depending on the degree of mobility.
- o Measurements of the performance that Route Optimization schemes show under different application scenarios, such as the use of applications with symmetric vs. asymmetric traffic patterns.

4.3. Future Research

Future research that goes beyond the techniques discussed in this document may consider the following items.

- o Local mobility support or local route-repair mechanisms that do not require expensive configuration. This includes infrastructure-based Route Optimization like [48].
- o Care-of-address verification mechanisms that are based on Secure Neighbor Discovery.
- o The introduction of optimizations developed in the context of Mobile IPv6 to other mobility protocols, such as the Host Identity Protocol, the Stream Control Transmission Protocol, the Datagram Congestion Control Protocol, or link-layer mobility solutions.
- o The extension of the developed mobility techniques to full multi-addressing, including multi-homing.
- o Further strategies that are based on "asymmetric cost wars" [3], such as Credit-Based Authorization.
- o Integrated techniques taking into account both link- and IP-layer mobility tasks.

5. Security Considerations

Standard Route Optimization enables mobile nodes to redirect IP packets at a remote peer from one IP address to another IP address. This ability introduces new security issues, which are explained and discussed in depth in [46]. The alternative Route Optimization techniques described in this document may introduce new security threats that go beyond those identified in [46]. Where such new threats exist, they are discussed and analyzed along with the description of the respective technique in Section 3.

6. Conclusions

Mobile IPv6 Route Optimization reduces packet-propagation latencies so as to facilitate interactive and real-time applications in mobile environments. Unfortunately, the end-to-end protocol's high handoff latencies hinder exactly these applications. A large body of effort has therefore recently been dedicated to Route Optimization improvements. Some of the proposed techniques operate on an end-to-end basis, others require new or extended infrastructure in the network; some need pre-configuration, others are zero-configurable. This document has compared and evaluated the different strategies based on a selected set of enhancement proposals. It stands out that all proposals make a trade-off between effectiveness, on one hand -- be it in terms of reduced handoff latency, increased security, or lower signaling overhead -- and pre-configuration costs or requisite network upgrades, on the other. An optimization's investment requirements, in turn, are in relation to its suitability for widespread deployment.

However, the real-life performance of end-to-end mobility does not only depend on enhancements of Route Optimization, but ultimately on all parts of the protocol stack [2]. Related optimization endeavors are in fact gaining momentum, and a comprehensive approach towards Route Optimization must incorporate the most suitable solutions amongst them [4]. Whichever proposals will eventually reach a maturity level sufficient for standardization, any effort should be expended to arrive at that point within the foreseeable future. Route Optimization requires support from both peers and depends on a solid basis of installed implementations in correspondent nodes. This should hence be included in emerging IPv6 stacks early on. Although IPv6 deployment is yet far away from becoming widespread, the sooner efficient Route Optimization will be available, the more likely that it will in the end be ubiquitously supported.

7. Acknowledgments

This document was thoroughly reviewed, in alphabetical order, by Samita Chakrabarti, Francis Dupont, Thierry Ernst, Gerardo Giaretta, James Kempf, Rajeev Koodli, Gabriel Montenegro, Vidya Narayanan, and Fan Zhao. The authors wish to thank these folks for their valuable comments and suggestions.

8. References

8.1. Normative References

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

8.2. Informative References

- [2] Alimian, A. and B. Aboba, "Analysis of Roaming Techniques", IEEE Contribution 802.11-04/0377r1, March 2004.
- [3] Arkko, J. and P. Nikander, "Weak Authentication: How to Authenticate Unknown Principals without Trusted Parties", Proceedings of Security Protocols Workshop 2002, Cambridge, UK, April 2002.
- [4] Arkko, J., Eronen, P., Nikander, P., and V. Torvinen, "Secure and Efficient Network Access", Proceedings of the DIMACS Workshop on Mobile and Wireless Security, November 2004.
- [5] Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6", Work in Progress, August 2006.
- [6] Arkko, J. and C. Vogt, "Credit-Based Authorization for Binding Lifetime Extension", Work in Progress, May 2004.
- [7] Bahl, P., Adya, A., Padhye, J., and A. Walman, "Reconsidering Wireless Systems With Multiple Radios", ACM SIGCOMM Computer Communication Review, ACM Press, Vol. 34, No. 5, October 2004.
- [8] Bradner, S., Mankin, A., and J. Schiller, "A Framework for Purpose-Built Keys (PBK)", Work in Progress, June 2003.
- [9] Castellucia, C., Montenegro, G., Laganier, J., and C. Neumann, "Hindering Eavesdropping via IPv6 Opportunistic Encryption", Proceedings of the European Symposium on Research in Computer Security, Lecture Notes in Computer Science, Springer-Verlag, September 2004.
- [10] Chandra, R., Bahl, P., and P. Bahl, "MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card", Proceedings of the IEEE INFOCOM, Vol. 2, August 2004.
- [11] Daley, G., Pentland, B., and R. Nelson, "Effects of Fast Routers Advertisement on Mobile IPv6 Handovers", Proceedings of the IEEE International Symposium on Computers and Communication, Vol. 1, June 2003.

- [12] Daley, G., Pentland, B., and R. Nelson, "Movement Detection Optimizations in Mobile IPv6", Proceedings of the IEEE International Conference on Networks, September 2003.
- [13] Daley, G., "Location Privacy and Mobile IPv6", Work in Progress, January 2004.
- [14] Dupont, F., "A Note about 3rd Party Bombing in Mobile IPv6", Work in Progress, July 2006.
- [15] Dupont, F. and J. Combes, "Using IPsec between Mobile and Correspondent IPv6 Nodes", Work in Progress, August 2004.
- [16] Dupont, F. and J. Combes, "Care-of Address Test for MIPv6 using a State Cookie", Work in Progress, July 2006.
- [17] Haddad, W., Nordmark, E., Dupont, F., Bagnulo, M., and B. Patil, "Privacy for Mobile and Multi-homed Nodes: MoMiPriv Problem Statement", Work in Progress, June 2006.
- [18] "IEEE Standard for Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, December 2004.
- [19] Choi, J. and E. Nordmark, "DNA with Unmodified Routers: Prefix List Based Approach", Work in Progress, January 2006.
- [20] Narayanan, S., Ed., "Detecting Network Attachment in IPv6 Networks (DNAv6)", Work in Progress, October 2006.
- [21] Moskowitz, R., Nikander, P., Jokela, Ed., P., and T. Henderson, "Host Identity Protocol", Work in Progress, June 2006.
- [22] Henderson, T., Ed., "End-Host Mobility and Multihoming with the Host Identity Protocol", Work in Progress, June 2006.
- [23] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 2006.
- [24] Koodli, R., "IP Address Location Privacy and Mobile IPv6: Problem Statement", Work in Progress, October 2006.
- [25] Perkins, C., "Securing Mobile IPv6 Route Optimization Using a Static Shared Key", RFC 4449, June 2006.
- [26] Ng, C., Thubert, P., Watari, M., and F. Zhao, "Network Mobility Route Optimization Problem Statement", Work in Progress, September 2006.

- [27] Ng, C., Zhao, F., Watari, M., and P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", Work in Progress, September 2006.
- [28] Arbaugh, W. and B. Aboba, "Handoff Extension to RADIUS", Work in Progress, October 2003.
- [29] "Kame-Shisa", Mobile IPv6 for FreeBSD.
- [30] Koodli, R., Devarapalli, V., Flinck, H., and C. Perkins, "Solutions for IP Address Location Privacy in the Presence of IP Mobility", Work in Progress, February 2005.
- [31] Nuorvala, V., Petander, H., and A. Tuominen, "Mobile IPv6 for Linux (MIPL)".
- [32] Mishra, A., Shin, M., Petroni Jr., N., Clancy, C., and W. Arbaugh, "Proactive Key Distribution Using Neighbor Graphs", IEEE Wireless Communications, Vol. 11, No. 1, February 2004.
- [33] Montenegro, G. and Claude. Castelluccia, "Crypto-Based Identifiers (CBIDs): Concepts and Applications", ACM Transactions on Information and System Security Vol.7, No. 1, February 2004.
- [34] 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.
- [35] O'Shea, G. and M. Roe, "Child-proof Authentication for MIPv6", ACM SIGCOMM Computer Communication Review, April 2001.
- [36] Perera, E., Sivaraman, V., and A. Seneviratne, "Survey on Network Mobility Support", ACM SIGCOMM Computer Communication Review, Vol. 8, No. 2, ACM Press, April 2004.
- [37] Bao, F., Deng, R., Qiu, Y., and J. Zhou, "Certificate-based Binding Update Protocol (CBU)", Work in Progress, March 2005.
- [38] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [39] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, August 2003.

- [40] Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", RFC 3776, June 2004.
- [41] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.
- [42] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [43] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [44] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [45] Huston, G., "Architectural Approaches to Multi-homing for IPv6", RFC 4177, September 2005.
- [46] Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background", RFC 4225, December 2005.
- [47] Roe, M., Aura, T., O'Shea, G., and J. Arkko, "Authentication of Mobile IPv6 Binding Updates and Acknowledgments", Work in Progress, February 2002.
- [48] Vadali, R., Li, J., Wu, Y., and G. Cao, "Agent-Based Route Optimization for Mobile IP", Proceedings of the IEEE Vehicular Technology Conference, October 2001.
- [49] Velayos, H. and G. Karlsson, "Techniques to Reduce IEEE 802.11b MAC Layer Handoff Time", Laboratory for Communication Networks, KTH, Royal Institute of Technology, Stockholm, Sweden, TRITA-IMIT-LCN R 03:02, April 2003.
- [50] Vogt, C., "Credit-Based Authorization for Concurrent IP-Address Tests", Proceedings of the IST Mobile and Wireless Communications Summit, June 2005.
- [51] Vogt, C., Bless, R., Doll, M., and T. Kuefner, "Early Binding Updates for Mobile IPv6", Proceedings of the IEEE Wireless Communications and Networking Conference, IEEE, Vol. 3, March 2005.

- [52] Vogt, C. and M. Doll, "Efficient End-to-End Mobility Support in IPv6", Proceedings of the IEEE Wireless Communications and Networking Conference, April 2006.
- [53] Vogt, C., "A Comprehensive Delay Analysis for Reactive and Proactive Handoffs with Mobile IPv6 Route Optimization", Institute of Telematics, Universitaet Karlsruhe (TH), Karlsruhe, Germany, TM-2006-1, January 2006.
- [54] Zhao, F., Wu, F., and S. Jung, "Extensions to Return Routability Test in MIPv6", Work in Progress, February 2005.
- [55] Calderon, M., Bernardos, C., Bagnulo, M., Soto, I., and A. de la Oliva, "Design and Experimental Evaluation of a Route Optimisation Solution for NEMO", IEEE Journal on Selected Areas in Communications, Vol. 24, No. 9, September 2006.

Authors' Addresses

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

E-Mail: chvogt@tm.uka.de

Jari Arkko
Ericsson Research NomadicLab
FI-02420 Jorvas
Finland

E-Mail: jari.arkko@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.