A Comprehensive and Efficient Handoff Procedure for IPv6 Mobility Support

Christian Vogt
Institute of Telematics, Universität Karlsruhe (TH), Germany
chvogt@tm.uka.de

Abstract

Handoff performance with Mobile IPv6 Route Optimization strongly depends on the efficiency of IP-layer autoconfiguration mechanisms as well as the flexibility of mobile nodes to schedule and parallelize their signaling. This paper provides a comprehensive analysis of the handoff performance with the standard IPv6 protocol suite and Mobile IPv6, and it identifies several sources for delay. While some of the delays are already well known, an optimized and widely applicable handoff approach is yet to be found. The paper hence proceeds to discuss existing and new optimization proposals, some of which are currently under standardization within the IETF, and elaborates how a combination of those can significantly improve handoff experience.

1 Introduction

As Internet-based services pervade daily life more and more, users increasingly desire them to be accessible at any place and any time. At the same time grows the importance of real-time communications [19] such as audio and video streaming, IP telephony, or video conferencing. Real-time communications are highly delay-sensitive and exhibit a susceptibility to long propagation latencies and handoff delays. Efficient mobility support was hence amongst the primary objectives during the design of the next-generation Internet, and a mode for *Route Optimization* was incorporated into the Mobile IPv6 [9] mobility protocol. Route Optimization allows peers to communicate via a direct path. This complements the classic approach of routing a mobile node's traffic through a stationary proxy, its *home agent*.

While Route Optimization mitigates the problem with propagation latencies, handoff delays are still substantial enough to effectively *preclude* meaningful real-time support [2, 12, 13]. In fact, handoff delays in a standard IPv6 deployment are in the order of seconds. This is not only due to Mobile IPv6, but also affects standard IPv6 configuration and movement-detection mechanisms [15, 24]. Very fortunately, a multitude of optimization techniques [3, 5, 10, 14, 18] have recently been put forth to streamline individual handoff-related activities. Measurement data is typically available to corroborate the benefits of any specific technique. But a study of how well the optimizations integrate has so far been largely neglected [1].

This paper examines the challenges with mobility from a higher perspective: It explains the overall handoff procedure in a standard IPv6 deployment from an IP layer's perspective and analyzes inhowfar it falls short of expectations. Since the results strongly advise optimization, the paper proceeds to explore promising existing and new proposals that have recently gained momentum in both in the Internet Engineering Task Force (IETF) and the academic research community. The optimizations are also evaluated with respect to their interactions. The paper finally proposes an integrated solution for improved handoff performance.

2 Standard Handoff Procedure

A mobile node undergoes an IP-layer handoff, or simply a *handoff*, when it changes IP connectivity. This begins with a change in link-layer attachment, also referred to as a *link-layer handoff*, and is followed by the discovery of new routers, address configuration, movement detection, and finally Mobile IPv6 registrations. Figure 1 illustrates these handoff steps, which are separately discussed next.

2.1 Router Discovery

A mobile node learns about local routers and on-link prefixes during router discovery. This process is facilitated through Router Advertisement messages, which routers multicast to link-local nodes on a loosely periodic basis. The IPv6 Neighbor Discovery RFC [16] states that unsolicited Router Advertisement messages are to be sent in random intervals of between 3 and 4 seconds at least and between 1350 and 1800 seconds at most. Since these conservative limits are tailored towards stationary nodes and fail to meaningfully support mobility, the Mobile IPv6 RFC decreases the lower bound to one beacon every 30 to 70 milliseconds. This reduces the mean time between successive

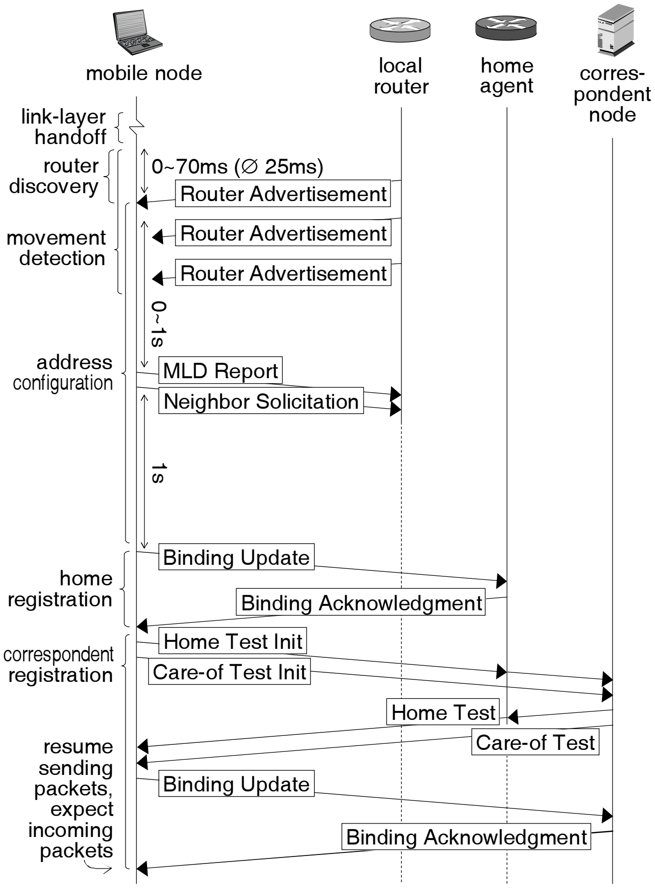

Figure 1. The standard handoff procedure

advertisements to 50 milliseconds so that a mobile node can expect to receive the first post-handoff advertisement after 25 milliseconds. On the other hand, high frequencies for multicast advertisements may be an issue in low-bandwidth, wide-area networks, where many users may not frequently leave the geographic area covered by the same IP subnet. Figure 1 shows the advertisements that are relevant for the handoff procedure; other advertisements are concealed.

2.2 Address Configuration

A mobile node configures a new global IP address upon receipt of a Router Advertisement message with an unknown prefix. This typically happens in compliance with Stateless Address Autoconfiguration [22]: The mobile node chooses an interface identifier, either randomly or based on the interface's MAC address, and prepends to this the obtained prefix. It then sends a Multicast Listener Report message [4] to subscribe to the solicited-node multicast group corresponding to the new address. If the Router Advertisement message was a multicast transmission, which usually is the case, the Multicast Listener Report message is delayed by up to 1 second to desynchronize with neighboring nodes that may be reacting to the same advertisement. The mobile node then runs the Duplicate Address Detection protocol to verify whether the new address is unique: It transmits a Neighbor Solicitation message for the address and, if no responses are received within a period of 1 second, assigns the address to the interface. The total configuration period hence ranges between 1 and 2 seconds if the address is unique. The probability for an IPv6 address to already be in use by another node is small enough to make it negligible.

Even though the link-local address keeps its prefix during handoff, the mobile node must still re-verify uniqueness of this address when IP connectivity changes, because a node on the new link may already be using the same link-local address. This is done through another run of Duplicate Address Detection.[1] Since only movement detection can establish whether IP connectivity has changed, re-verification of the link-local address typically begins after movement detection. This is not shown in figure 1, however, given that the availability of the link-local address does not influence the schedule for other handoff-related activities.

2.3 Movement Detection

Mobile nodes implement movement detection to recognize changes in IP connectivity. Such a change implies that a mobile node chooses a new default router, invalidates stale global addresses, re-verifies uniqueness of its link-local address, and initiates Mobile IPv6 registrations. Movement detection relies on analyzing the on-link prefixes advertised in Router Advertisement messages and possibly also probing reachability of routers considered off-link. When the prefixes in use by the mobile node are no longer seen to be advertised, but new prefixes show up instead, the mobile node typically decides that it has moved to a different network. On the other hand, received prefixes may also indicate that IP connectivity did not change in spite of a link-layer handoff, e.g., when the mobile node switches access points that connect to the same subnet.

Movement detection is complicated by the fact that Router Advertisement messages may include incomplete sets of prefixes. Reception of a single advertisement is therefore usually insufficient to decide whether IP connectivity has changed. It is also generally impossible to determine when an advertisement should have been received, but did not appear, due to the lack of a guaranteed advertisement interval. The Mobile IPv6 RFC helps in this respect in that it introduces an Advertisement Interval option for Router Advertisement messages. Routers use this option to indicate an upper bound on their beaconing periods. Where

[1] Transmission of another Multicast Listener Report message can be spared if the multicast group corresponding to the link-local address is the same as that of a global address configured on the new link.

this is as low as 70 milliseconds (cf. section 2.1), an extra 20 milliseconds are added in order to account for scheduling granularities in mobile nodes and routers. Mobile nodes then expect Router Advertisement messages to arrive in intervals of at most 90 milliseconds.

Nevertheless, the absence of a single expected advertisement still does not imply a change in IP connectivity given the potential for packet loss. Three missing advertisements indicate movement more reliably. A decision can then be made at most 270 milliseconds after the last advertisement was received from the old default router. The actual link-layer handoff may occur up to 70 milliseconds later, so movement detection can take any time between 200 and 270 milliseconds. On average, the period between reception of the last advertisement from the old default router and the link-layer handoff is 25 milliseconds, yielding a mean movement-detection delay of 245 milliseconds.

2.4 Mobile IPv6 Registration

After address configuration and movement detection, the mobile node selects one of its new global addresses to be registered as a *care-of address* with its home agent and correspondent nodes. This establishes a binding between the care-of address and the mobile node's *home address*, which has a prefix from the home agent's network and remains stable across movements. The home address is used at stack layers above IP as part of end-point identification. Data packets that a mobile node exchanges with a peer have the care-of address in the IP header and the home address in an IPv6 extension header while on the wire. Both end nodes swap the addresses when a packet traverses the IP layer so that transport protocols and applications can access the home address as usual.

Figure 1 illustrates the Mobile IPv6 registration procedure for the home agent and a single correspondent node. The *home registration* consists of a Binding Update message which notifies the home agent of the new care-of address, and a Binding Acknowledgment message indicating success or failure. Care must be taken to preclude illegitimate bindings [17], which malicious nodes could attempt to establish for the purpose of impersonation or redirection-based flooding. The mobile node and the home agent are typically under the same administration and pre-share credentials to bootstrap an IPsec security association. The home registration can so be authenticated and encrypted.

The *correspondent registration* permits Route Optimization. It includes a Binding Update message that conveys the new care-of address to the correspondent node, and a responding Binding Acknowledgment message[2]. These cannot generally be protected through IPsec, however, because mobile nodes are neither likely to share authentication credentials with all correspondent nodes they may at some point communicate with, nor is a ”global” public-key infrastructure, available for arbitrary pairs of nodes, expected to come into existence any time soon [17]. Correspondent registrations are instead authenticated and authorized through a *return-routability procedure*, based on non-cryptographic verification of a mobile node's reachability at the home and care-of addresses. Reachability at both addresses entitles the mobile node to initiate a binding between the addresses.

For the *home-address test*, the mobile node tunnels a Home Test Init message to the home agent, which forwards the message to the correspondent node. The correspondent node returns an unpredictable *home keygen token* to the home address within a Home Test message. The home agent intercepts this message and tunnels it to the mobile node. The *care-of-address test* is a direct exchange between the mobile node and the correspondent node. It consists of a Care-of Test Init message and a Care-of Test message with an unpredictable *care-of keygen token*. Knowledge of the home and care-of keygen tokens proves the mobile node's ability to receive packets at the home address and care-of address, respectively. The mobile node demonstrates this knowledge by authenticating the Binding Update message for the correspondent node with a key derived from both tokens. The correspondent node uses the same key to authenticate the final Binding Acknowledgment message.

The Mobile IPv6 RFC leaves mobile nodes liberties with respect to scheduling signaling and data packets. Figure 1 shows a *conservative* mobile node, which waits for the Binding Acknowledgment message from the home agent before it initiates the return-routability procedure. In contrast, an *optimistic* mobile node executes the home registration and the return-routability procedure in parallel. An optimistic mobile node furthermore starts sending packets to the correspondent node as soon as the Binding Update message for the correspondent node has been brought on way, whereas a conservative mobile node uses the new care-of address only after reception of an acknowledgment. In either case, the correspondent node is unaware of the new care-of address until it receives the Binding Update message. Its first packet sent to the new care-of address will hence be delivered to the mobile node roughly along with the Binding Acknowledgment message, assuming that one was requested by the mobile node.

Conservative mobile nodes avoid a useless return-routability procedure in case the home registration fails. They also do not risk loss of packets sent shortly after a lost or rejected Binding Update message. The correspondent node would discard these packets in the face of a mismatching binding due to security measures. This comes at the cost of additional handoff latency when both registra-

[2] Whether or not the correspondent node sends an acknowledgment is left to the discretion of the mobile node. The mobile node can request one by setting a flag in the Binding Update message.

tions are successful. For outgoing route-optimized packets, this is a round-trip time between the mobile node and the home agent plus a round-trip time between the mobile node and the correspondent node. For incoming packets, the additional handoff latency is a round-trip time between the mobile node and the home agent. Optimistic mobile nodes perform better in the general case. But they may attempt a return-routability procedure in vain or suffer packet loss should the home or correspondent registration fail.

3 Existing and Proposed Approaches to Improve Handoff Performance

The delays of the standard handoff procedure can significantly impair the quality of real-time applications, even though Route Optimization was designed with an intention to improve support for these applications. The research community has been working to decrease handoff delays for some time now, and a number of proposals have been made. Particularly promising are the following approaches.

3.1 Router Discovery

More sophisticated scheduling intervals in routers can improve router discovery with respect to both bandwidth consumption and efficiency. FastRA [5] permits a mobile node to solicit an immediate advertisement. This is useful when the mobile node's link layer can indicate changes in network attachment. Based on on-link routers' link-local addresses and the source address of the solicitation, each router autonomously computes a dynamic ranking indicating which router should respond immediately, and possibly which other routers should send additional advertisements shortly thereafter. The Fast Router Discovery [3] proposal suggests that access points replay a cached Router Advertisement message once a node has been associated. Such link-layer support on the network side eliminates the requirement for link-layer triggers in mobile nodes.

3.2 Address Configuration

Different proposals have been made to avoid the handoff delays caused by standard Duplicate Address Detection. The IPv6 working group within the IETF is developing Optimistic Duplicate Address Detection [14], which allows for limited use of a potentially duplicate IP address. Mobile nodes temporarily change the rules by which they do IPv6 Neighbor Discovery signaling so as to avoid pollution of other nodes' neighbor caches with possibly illegitimate address-resolution information.

With Advanced Duplicate Address Detection [8], routers generate a pool of unique addresses which they then assign to mobile nodes. Duplicate Address Detection is performed on the addresses in advance so that mobile nodes can configure them instantly without first having to verify uniqueness themselves.

3.3 Movement Detection

The DNA working group within the IETF tackles the problem of slow movement detection with two complementary approaches. The Complete Prefix List protocol [18] works with unmodified routers. A mobile node maintains a list of learned on-link prefixes, possibly obtained by reception of multiple Router Advertisement messages. After the list has matured for a while, the mobile node can assume a change in IP connectivity with high probability when a newly received advertisement exclusively contains prefixes not in the list. Such predictions are based on potentially incomplete information, so the mobile node might assert movement even when none actually occurred.

The DNA protocol [10] uses FastRA for timely transmissions of solicited Router Advertisement messages. Routers choose a certain prefix to serve as a *link identifier* and be as such indicated in all transmitted advertisements. This allows a mobile node to reliably detect changes in IP connectivity based on a single advertisement. Alternatively, the mobile node can explictly check with routers as part of the solicitation-advertisement exchange whether a network prefix used before a link-layer handoff, as such called a *landmark*, is still valid on the possibly new link. The DNA protocol integrates Complete Prefix List as a fall-back mechanism for links with legacy routers.

3.4 Mobile IPv6 Optimizations

Many Mobile IPv6 optimizations reduce the handoff delays of Route Optimization through modifications of the return-routability procedure. A combination of Early Binding Updates [25] and Credit-Based Authorization [23] achieves this, on a purely end-to-end basis, with the following four constituent optimizations:

1. **Proactive home-address tests:** A mobile node acquires a home keygen token for a future handoff during a proactive home-address test. This saves a possibly long round trip through the home agent during the critical handoff period. The mobile node can invoke proactive home-address tests on a just-in-time basis, if its link layer provides a trigger indicating imminent handoff, or periodically whenever the most recently obtained home keygen token is about to expire.

2. **Concurrent care-of-address tests:** Data packets can already be exchanged, to a limited extent, via a new care-of address, while the mobile node's reachability at that care-of address is being verified.

3. **Tentative bindings:** The mobile node registers a tentative binding between its home address and an *unverified* care-of address by exchanging Early Binding Update and Early Binding Acknowledgment messages with a correspondent node. The messages are authenticated only with the home keygen token obtained from a recent proactive home-address test, thus facilitating a subsequent, concurrent care-of-address test. Once the mobile has executed the concurrent care-of-address test, it authenticates a standard Binding Update message and registers a *verified* care-of address with the correspondent node.

4. **Parallel home and correspondent registrations:** The Mobile IPv6 specification does not permit the mobile node to send a Binding Update message to a correspondent node before it receives an acknowledgment from the home agent. This becomes a performance issue if the combination of proactive home-address tests and concurrent care-of-address tests hides the latency of the return-routability procedure. The rules of Mobile IPv6 are hence relaxed so as to allow a mobile node to send an Early Binding Update message when the home registration is still pending.

Well-known security guidelines [17] prohibit a correspondent node to send packets to a care-of address for which reachability has not yet been verified. This is a precaution against malicious nodes which could otherwise trick correspondent nodes into flooding a third party with unrequested packets. The appeal of such *redirection-based flooding attacks* is the potential for significant amplification. E.g., an attacker could accomplish the initial TCP handshake for a voluminous file download through its own address (or home address, for that matter), and then redirect the flow to the address of its victim. The attacker could, and would have to, spoof acknowledgments on behalf of the victim based on the sequence numbers it learned during the initial handshake. But the acknowledgments would be small compared to the data segments that the correspondent node generates. Credit-Based Authorization prevents redirection-based amplified flooding, yet enables bidirectional communications via unverified care-of addresses. The correspondent node maintains a byte counter for the mobile node, also called the mobile node's *credit*, which increases with the data volume received from the mobile node and decreases with the data volume sent to the mobile node while the care-of address is unverified. Exponential aging assures that existing credit represents only data recently received from the mobile node. When the correspondent node has a packet for the mobile node, it sends it to the care-of address if the address is verified, or if the address is unverified, but the packet size does not exceed the currently available credit. Otherwise, the correspondent node may drop the packet, buffer it until the care-of address becomes verified[3], or send it to the home address.

Other Route Optimization enhancements require some form of pre-configuration: Where end nodes share a key or the credentials to bootstrap a security association, more efficient, cryptographic authentication can replace the home address test. Two such proposals are currently under discussion in the IETF. In [20], mobile nodes and correspondent nodes are pre-configured with a shared, secret authentication key. [6] uses IPsec and the Internet Key Exchange protocol. These techniques suffer from a scalability problem, however, given that end nodes must be set up with pairwise credentials. Also, neither of the techniques provides verification of a mobile node's reachability, so both cannot technically do without the care-of-address test. End nodes which trust in the peer's reachability may further omit the care-of-address test, but such trust is unavailable in many important business models. E.g., a mobile-phone operator may be able to configure subscribers with secret authentication keys, but it may not be able to vow that all subscribers use these keys in a trustworthy manner.

Another family of Mobile IPv6 optimizations is based on router support in the mobile nodes' visited networks. Where Fast Handovers for Mobile IPv6 [11] are deployed, a mobile node can request its current default router to establish a bidirectional tunnel to a new care-of address. This allows the mobile node to temporarily communicate through its old care-of address after a handoff, and to register the new care-of address with its home agent and correspondent nodes in the meantime. By help of proxy router discovery and assisted address configuration, the mobile node may request the tunnel prior to handoff, provided that it can anticipate movements. Additional capabilities optimize reactive handoff management for cases of unexpected link breaks.

Conversely, Media Independent Pre-Authentication [7] uses a bidirectional tunnel between the old care-of address and a new default router. A mobile node is assigned a new care-of address from remote and effects home and correspondent registrations before it changes links. The benefits of this approach are similar to those of Fast Handovers if the overlap between neighboring cells is sufficiently large to permit timely completion of handoff preparations. However, where cell overlaps are small relative to node velocities, postponing global signaling to a stage after handoff is advantageous, since the wireless signal quality is generally higher and more durable then. The strength of Media Independent Pre-Authentication is its ability to pre-authenticate a mobile node to the new network prior to handoff.

[3] Packet buffering makes the correspondent node susceptible to memory-overflow attacks and may hence represent a denial-of-service vulnerability on its own. However, where the correspondent node can identify trustworthy mobile nodes based on their (authenticated) home addresses, packet buffering could be an option.

Hierarchical Mobile IPv6 [21] enables a mobile node to bind its current on-link care-of address to a more stable *regional care-of address* from a Mobility Anchor Point's network located elsewhere in the visited domain. The mobile node sends and receives packets via the regional care-of address through a bidirectional tunnel between itself and the Mobility Anchor Point. It registers the regional care-of address with the home agent and correspondent nodes and updates the Mobility Anchor Point whenever its on-link care-of address changes in the wake of a movement. Movements can so be concealed from the home agent and correspondent nodes as long as the mobile node roves within the same Mobility Anchor Point's realm.

4 Proposed Handoff Procedure

Enhancements to Mobile IPv6 in conjunction with optimizations for router discovery, address configuration, and movement detection facilitate separation of a timely correspondent registration from less critical handoff tasks. The optimized handoff procedure proposed below uses a subset of the techniques described in section 3. Figure 2 illustrates it for communications with a single correspondent node.

The mobile node executes a proactive home-address test with the correspondent node prior to handoff. It may do this periodically whenever the previously acquired home keygen token is about to expire, or on-demand if forthcoming handoffs can be anticipated. A fresh home keygen token should thus be available when the mobile node is eventually notified about a change in link-layer attachment. The mobile node then sends an immediate Router Solicitation message and promptly receives a Router Advertisement message which allows it to review its current IP configuration. Where routers implement FastRA, the mobile node must use Complete Prefix List logic to quickly derive movement decisions. Alternatively, routers and the mobile node may implement the DNA protocol. Router discovery and movement detection can be achieved almost instantly in either case. Besides, solicited advertisements are generally more rate-economic than frequently transmitted, unsolicited multicast advertisements. The mobile node uses Optimistic Duplicate Address Detection to configure new and re-verify existing addresses. This proceeds in background mode and does not cause additional handoff latencies.

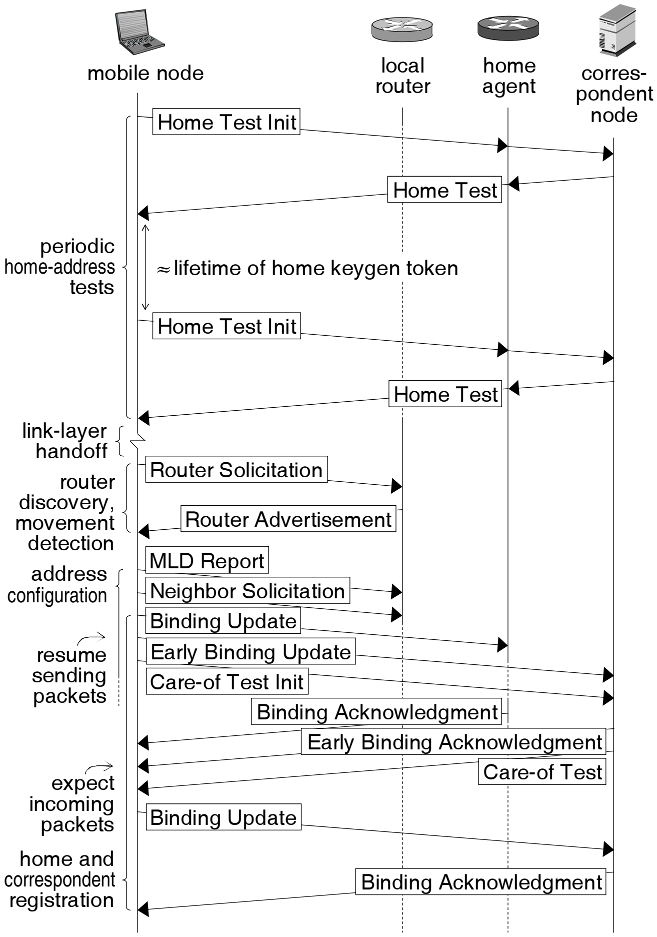

Figure 2. The proposed handoff procedure

When the mobile node detects a change in IP connectivity, it registers with its home agent and correspondent node a new care-of address for which uniqueness verification is still in progress. The home agent sends a Binding Acknowledgment message and subsequent data packets to the new care-of address once it has received the Binding Update message. The correspondent node registers a tentative binding for an unverified care-of address when it receives the Early Binding Update message. It sends an Early Binding Acknowledgment message if one was requested by the mobile node. The tentative binding conceals the latency of the return-routability procedure. The mobile node resumes communications with the correspondent node via the new care-of address immediately after it has sent the Early Binding Update message, and the correspondent node may send packets to the unverified care-of address, to the extent Credit-Based Authorization permits, as soon as it receives this message.

The mobile node initiates a concurrent care-of-address test with the correspondent node after it has sent the Early Binding Update message. When the test is complete, the mobile node sends an authenticated standard Binding Update message, causing the correspondent node to change the status of the care-of address from unverified to verified. Use of the care-of address is then no longer governed by Credit-Based Authorization. The correspondent node also sends a Binding Acknowledgment message if one was requested.

5 Conclusion

Efficient end-to-end IPv6 mobility support requires optimizations not only for the mobility protocol, but also for router discovery, address configuration, and movement detection. This paper has taken an in-depth look at the shortcomings of today's IPv6 protocol standards, explored own and various existing optimizations, examined how those interact, and how they can be combined into a complete and efficient mobility solution.

It is important to recognize that the basis for tomorrow's mobility support is to be laid today. This holds in particular for Route Optimization, which requires support from both peers and hence depends on a solid basis of implementations in correspondent nodes. Route Optimization functionality should therefore be included in emerging IPv6 stacks early on. Enhancements must also make their way into access routers, whose responsiveness is critical for efficient IPv6 configuration and movement detection. The sooner a set of necessary optimizations is widely accepted, the likelier that this set will in the end be ubiquitously supported.

6 Acknowledgment

The author thanks Ralf Beck, Roland Bless, Gregory Daley, Mark Doll, Daniel Jungbluth, and Max Laier (in alphabetical order) for valuable discussions on the topic of this paper.

References

[1] J. Arkko and C. Vogt. A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization. IETF Internet Draft draft-irtf-mobopts-ro-enhancements-07.txt (work in progress), Apr. 2006.

[2] M. Bernaschi, F. Cacace, G. Iannello, S. Za, and A. Pescape. Seamless Internetworking of WLANs and Cellular Networks: Architecture and Performance Issues in a Mobile IPv6 Scenario. *IEEE Wireless Communications*, 12(3), June 2005.

[3] J. Choi, D. Shin, and W. Haddad. Fast Router Discovery with L2 Support. IETF Internet Draft draft-ietf-dna-frd-00.txt (work in progress), Oct. 2005.

[4] M. J. Christensen, K. Kimball, and F. Solensky. Considerations for IGMP and MLD Snooping Switches. IETF Internet Draft draft-ietf-magma-snoop-12.txt (work in progress), Feb. 2005.

[5] G. Daley, B. Pentland, and R. Nelson. Movement Detection Optimizations in Mobile IPv6. In *Proceedings of the IEEE International Conference on Networks*, Sept. 2003.

[6] F. Dupont and J.-M. Combes. Using IPsec between Mobile and Correspondent IPv6 Nodes. IETF Internet Draft draft-ietf-mip6-cn-ipsec-02.txt (work in progress), Dec. 2005.

[7] A. Dutta, T. Zhang, K. Taniuchi, Y. Ohba, and H. Schulzrinne. MPA-Assisted Optimized Proactive Handoff Scheme. In *Proceedings of the International Conference on Mobile and Ubiquitous Systems*, July 2005.

[8] Y.-H. Han, J. Choi, H.-J. Jang, and S. D. Park. Advance Duplicate Address Detection. IETF Internet Draft draft-han-mobileip-adad-01.txt (work in progress), July 2003.

[9] D. Johnson, C. E. Perkins, and J. Arkko. Mobility Support in IPv6. IETF Request for Comments 3775, June 2004.

[10] J. Kempf. Detecting Network Attachment in IPv6 Networks (DNAv6). IETF Internet Draft draft-ietf-dna-protocol-00.txt (work in progress), Feb. 2006.

[11] R. Koodli. Fast Handovers for Mobile IPv6. IETF Request for Comments 4068, July 2005.

[12] N. Montavont and T. Noël. Handover Management for Mobile Nodes in IPv6 Networks. *IEEE Communications Magazine*, 40(8), Aug. 2002.

[13] N. Montavont and T. Noël. Analysis and Evaluation of Mobile IPv6 Handovers over Wireless LAN. *Kluwer Academic Publishers Mobile Networks and Applications*, 8(6), Dec. 2003.

[14] N. Moore. Optimistic Duplicate Address Detection for IPv6. IETF Internet Draft draft-ietf-ipv6-optimistic-dad-07.txt (work in progress), Dec. 2005.

[15] N. Moore and G. Daley. Fast Address Configuration Strategies for the Next-Generation Internet. In *Proceedings of the Australian Telecommunications, Networks, and Applications Conference*, Dec. 2003.

[16] T. Narten, E. Nordmark, W. A. Simpson, and H. Soliman. Neighbor Discovery for IP Version 6 (IPv6). IETF Internet Draft draft-ietf-ipv6-2461bis-06.txt (work in progress), Mar. 2006.

[17] P. Nikander, J. Arkko, T. Aura, G. Montenegro, and E. Nordmark. Mobile IP Version 6 Route Optimization Security Design Background. IETF Request for Comments 4225, Dec. 2005.

[18] E. Nordmark and J. Choi. DNA with Unmodified Routers: Prefix List Based Approach. IETF Internet Draft draft-ietf-dna-cpl-02.txt (work in progress), Jan. 2006.

[19] S. Ortiz Jr. Internet Telephony Jumps off the Wires. *IEEE Computer*, 37(12), Dec. 2004.

[20] C. Perkins. Securing Mobile IPv6 Route Optimization Using a Static Shared Key. IETF Internet Draft draft-ietf-mip6-precfgkbm-04.txt (work in progress), Dec. 2005.

[21] H. Soliman, C. Castelluccia, K. E. Malki, and L. Bellier. Hierarchical Mobile IPv6 Mobility Management (HMIPv6). IETF Request for Comments 4140, Aug. 2005.

[22] S. Thomson, T. Narten, and T. Jinmei. IPv6 Stateless Address Autoconfiguration. IETF Internet Draft draft-ietf-ipv6-rfc2462bis-08.txt (work in progress), May 2005.

[23] C. Vogt. Credit-Based Authorization for Concurrent IP-Address Tests. In *Proceedings of the IST Mobile and Wireless Communications Summit*, June 2005.

[24] C. Vogt, R. Bless, M. Doll, and G. Daley. Analysis of IPv6 Relocation Delays. Technical Report TM-2005-4, Institute of Telematics, Universität Karlsruhe (TH), Germany, Apr. 2005.

[25] C. Vogt and M. Doll. Efficient End-to-End Mobility Support in IPv6. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, Apr. 2006.