Mobile IPv6 Security

Tuomas Aura

Microsoft Research Ltd.
Roger Needham Building, 7 JJ Thomson Avenue, Cambridge, CB3 0FB, UK
tuomaura@microsoft.com

Abstract. This paper presents a case study of security protocol design: authentication of binding updates in Mobile IPv6. We go step by step through the threat analysis and show how each threat is addressed in the protocol design. The goal is to solve any new security issues caused by the introduction of mobility without requiring any new security infrastructure.

1 Introduction

This paper describes the Mobile IPv6 security protocol. The focus is on the authentication of binding updates, i.e. location information sent by the mobile to its correspondents.

We explain the security threats created by the introduction of mobility and the mechanisms that have been used to prevent the attacks. The protocol design is unusual and it would not be considered secure by the measures of traditional security protocol analysis. The security of the protocol depends on the partial reliability of the Internet routing infrastructure. The reason is that the protocol must work between any mobile node and any other Internet node that have no previous relationship, and we cannot assume the existence of a global PKI or other global security infrastructure. On the other hand, the only security requirement was to counter the new threats created by mobility. The protocol does exactly that. This is a pragmatic way of thinking when introducing new technology such as mobility.

The return-routability protocol described in this paper was originally a part of a protocol family designed by Michael Roe, the current author, Greg O'Shea and Jari Arkko [RAOA01]. Many of the protocol details have since been refined by the IETF Mobile IP working group [JPA03]. The solution enabled the Mobile IPv6 standardization process to continue after it had halted because of security concerns. When designing the protocol, we discovered a new class of attacks and introduced new defense mechanisms that have since been copied to other mobility protocols. The protocol described in this paper is a slightly simplified version of the actual Mobile IP6 protocol. We concetrate on the binding-update messages sent by the mobile to its correspondents and ignore some details that are intended to protect messages sent in the other direction.

The paper is organized roughly to follow the design process. We first introduce the Mobile IPv6 protocol and route optimization in Section 2. Section 3 describes the basic binding-update authentication protocol. Section 4 explains how even authenticated binding updates can be used for denial-of-service attacks and how these attacks are prevented. Section 5 explains some less serious threats and how the protocol was enhanced to mitigate them. Section 6 concludes the paper.

2 Mobile IPv6

Mobile IPv6 is an IP-layer mobility protocol for the IPv6 Internet. It is being standardized by the IETF. The idea is that when mobility, like any other functionality, is implemented in the network layer, it needs to be implemented only once and will then be transparently available for all higher-layer protocols. It remains to be seen how well this promise is fulfilled in practice. There are, however, some applications like mobile VPN access, for which Mobile IP is clearly a good solution. This section describes the Mobile IPv6 architecture and protocol.

2.1 The Mobile Network Architecture

The first half of an IPv6 address indicates the subnet to which the address belongs and it is used for routing IP packets across the Internet. Thus, when a mobile Internet host (called *mobile node (MN)* in the Mobile IPv6 terminology) moves to a different place in the network topology, its subnet and, thus, IP address necessarily change. This creates two kinds of problems: existing connections (e.g. TCP connections and IPSec security associations) between the mobile and other hosts (called *correspondent nodes (CN)*) become invalid, and the mobile is no more reachable in its old address for new connections. The former problem is important in stateful protocols and has little effect of stateless protocols such as HTTP. The latter problem typically concerns servers and not client computers.

Mobile IPv6 has two basic goals: all transport-layer and higher-layer connections and security associations between the mobile and its correspondents should survive the address change, and themobile host should be reachable as long as it is connected to the Internet somewhere in the world.

Mobile IP makes some quite strong assumptions about the environment in which it is used. First, all mobile hosts have a home network and a *home address (HoA)* on that network. This is a reflection from a time when mobility was an exception: few Internet nodes would be mobile and even they would for most of the time remain stationary at home. In any case, Mobile IP solves the reachability problem by ensuring that the mobile is always able to receive packets sent to its home address.

The IP address of a stationary IP node normally serves two purposes: it is both an identifier for the node and an address that is used for routing messages to the node.

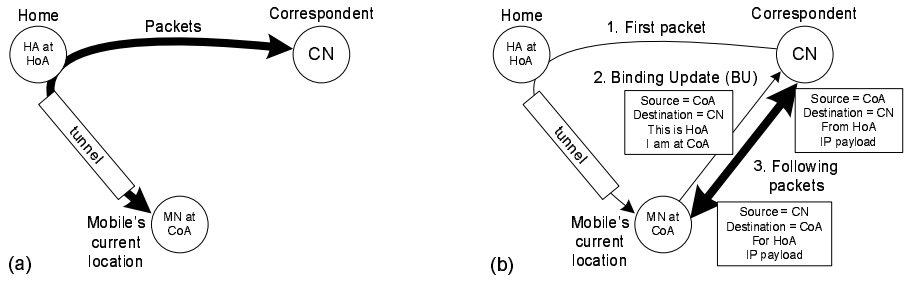

Fig. 1. Mobile IPv6 tunneling (a) and route optimization (b)

Mobile IP preserves this dual use of home addresses. The home address is an identifier for the mobile, as well as an address to which correspondents can send packets. The mobile's current location, called *care-of address (CoA)*, on the other hand, is a pure address and serves no identification purposes.

Any IPv6 address can be or become mobile and there is no way of distinguishing a mobile and stationary host by just looking at its address. This is because the Mobile IP protocol was originally designed to be transparent to the mobile's correspondents and the correspondent did not need to know that the mobile, in fact, was a mobile.

2.2 The Mobility Protocol

The transparent mode of operation is shown in Figure 1(a). At its home network, the mobile has an agent, called *home agent (HA)*. The home agent is a router that tunnels packets to and from the mobile. It intercepts packets sent by correspondents to the mobile's home address and forwards them to the mobile at its current location over an IPIP or IPSec tunnel. When the mobile wants to send packets to a correspondent, it sends them to the home agent over the reverse tunnel. The home agent un-encapsulates the packets and forwards them to the correspondent. When the mobile moves to a new location, it tells the home agent its new care-of address.

The tunnelling protocol is sufficient to enable mobility but it results in suboptimal routing. Packets between the mobile and its correspondents have to travel via the home network, which may be far away. To rectify this problem, Mobile IPv6 defines a mechanism called route optimization. The optimization requires changes to the correspondent but it is seen as so important that every IPv6 host has to support the protocol.

Route optimization typically works as shown in Figure 1(b). When the mobile receives a tunnelled packet, it initiates the route optimization protocol. The mobile sends to the correspondent a message called *binding update (BU)*. The binding update contains the mobile's home address and current care-of address. The correspondent stores this information in its binding cache, which is effectively a routing table: it tells that packets

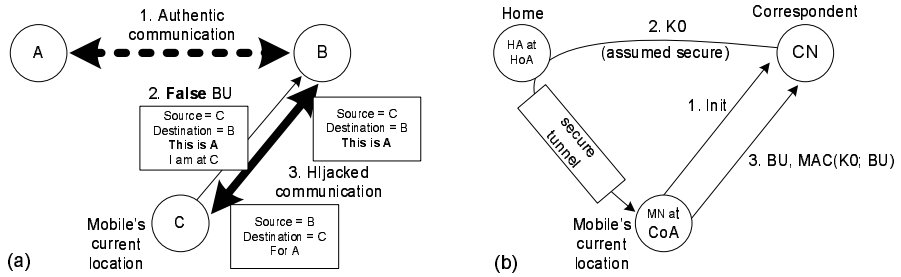

Fig. 2. False BU (a) and routing-based authentication protocol v.1 (b)

destined to HoA should instead be sent to CoA. Finally, the correspondent acknowledges the binding update. (For simplicity, the acknowledgements are not shown in the figures and we ignore the authentication of messages from the correspondent to the mobile.)

The route optimization is a voluntary in the sense that either the mobile or the correspondent can refuse to do it and they can continue communicating via the home agent. It is, however, and important optimization because sending packets via the home agent can be very inefficient.

After the binding update, the following packets between the mobile and the correspondent are sent directly. The packets from the mobile to the correspondent contain a header field called *home address option (HAO)*, which tells the correspondent that although the source address is CoA, the packet is actually from the node whose address is HoA. The packets from the correspondent to the mobile contain a *routing header*, which tells the mobile that although the packet is destined to CoA, it is really intended to HoA. The headers are, in effect, a kind of tunnel between the mobile and the correspondent. Every few minutes, the mobile needs to send another binding update to refresh the binding cache entry even if the care-of address has not changed.

The assumptions and design choices made in the Mobile IP protocol are not necessarily the same that would be made if the Internet mobility protocol were designed again from scratch. Nevertheless, these are the protocol and assumptions that we had to live with when considering Mobile IPv6 security.

3 BU Authentication

It is quite obvious that the binding update protocol, if implemented as described above, would create serious new security vulnerabilities. The first thing that one notices is that the binding updates are not authenticated. This section describes the basic attacks using unauthentic BUs and a BU authentication mechanism.

3.1 The Need for Infrastructureless Authentication

A possible attack is shown in Figure 2(a). An attacker at address C sends a false binding update to an Internet node B, claiming to be a mobile with home address A. If B, acting in the role of a correspondent, believes this binding update, it will redirect to C all packets that are intended for A. Thus, the attacker can intercept packets sent by B to A. The attacker can also spoof data packets from A by inserting a false home-address option in them. This enables the attacker to hijack existing connections between A and B, and to open new ones pretending to be A. The attacker can also redirect the packets to somewhere else than C, which prevents A and B from communicating with each other. End-to-end data protection, e.g. IPSec or SSL, prevents most of the attacks but not denial of service (DoS).

These attacks are serious because A, B and C can be any IPv6 addresses anywhere on the Internet. All the attacker needs to know is the IPv6 addresses of A and B. Since there is no visible difference between a mobile home address and a stationary IPv6 address, the attacks work as well against stationary Internet nodes as against mobile ones. The possibility of these attacks caused IETF to halt the Mobile IPv6 standardization process until a solution for authenticating binding updates was found. It is believed that deployment of the protocol without security could result in a break-down of the entire Internet.

Obviously, the solution is to authenticate the binding updates. A typical authentication mechanism would involve a trusted online server or a public-key infrastructure (PKI). The problem is that the authentication needs to work between any mobile Internet node and any correspondent. There does not currently exist any authentication infrastructure that could be used for such global authentication between any two IPv6 nodes. Neither is it realistic to suggest creating such infrastucture for the needs of Mobile IPv6. Hence, using the conventional authentication mechanism would confine route optimization to intra-organizational use where the required security services are in place.

For the above reason, we were forced to consider unconventional authentication methods. The advantage we had on our side is that the security requirements for BU authentication are unusually weak. The stated goal in the IETF working group was that the Mobile IPv6 protocol should be at least as secure as the current non-mobile IPv4 Internet. This means that we were not confined to designing a traditional strong security protocol. Our ambition was limited to making sure that Mobile IPv6 does not introduce any new major vulnerabilities to the Internet. The goal was not to create a strong general-purpose authentication protocol.

The IP layer provides two kinds of infrastructure. First, the addressing architecture [HD98] provides Internet nodes unique, globally routable IPv6 addresses. Second, the routing infrastructure [Hui95] delivers packets across the Internet to their destination address. It turns out that both the addressing and the routing can be used to bootstrap some form of authentication, not necessarily as strong as a PKI would enable in closed networks but nevertheless better than no authentication. Since these techniques do not

require any special security infrastructure, they are, somewhat misleadingly, called *infrastructureless authentication*.

The address-based technique was first proposed in the CAM protocol for binding update authentication by O'Shea and Roe [OR01]. The basic idea was to create the second half of the home address as a one-way hash of the mobile node's public key. Such address are called cryptographically generated addresses (CGA). The mobile signs the binding update and attaches its public key to the message. The correspondent can verify without any additional infrastructure that the binding update was signed by the owner of the home address. While the authentication of the sender's IPv6 address would be of little value in most applications, it is exactly what is needed to authorize the binding update. There were several further proposal for using this technique for protecting mobility protocols [Nik01,MC01].

The routing-based authentication will be covered in detail in the rest of the paper. Our original protocol design allowed both types of infrastructureless authentication but the IETF working group chose to standardize only the simpler routing-based technique.

3.2 Routing-Based Authentication

The second infrastructureless authentication method is based on the fact that routing in the Internet is semi-reliable. It is difficult for a remote attacker to change the route of packets that do not travel via the attacker's network. Thus, in order to sniff or intercept a packet, the attacker needs to be on its route.

The first version of the BU authentication protocol is shown in Figure 2(b). The idea is that after the mobile initiates the BU protocol (message 1), the correspondent sends a secret key as plaintext to the mobile's home address (message 2). The home agent intercepts the message and forwards it to the mobile via a secure tunnel. The mobile then uses the key to compute a message authentication code for the binding update (message 3). This mechanism is called *return-routability test for the home address* because the mobile must return to the correspondent (a function of) a value sent by the correspondent to the home address. In effect, the correspondent verifies that the mobile is able to receive messages at the home address.

In order to break the protocol, i.e. to spoof a binding update, the attacker needs to be on the route between the correspondent and the mobile. Thus, the protocol is not secure against the standard attacker model where the attacker can sniff and intercept all messages on the network. It is natural that most readers previously unfamiliar with the protocol will at this point object to the idea of sending a key in plaintext. There are, nevertheless, strong arguments in favor of the design.

First, the number of potential attackers and targets is dramatically reduced. Without authentication, any Internet node C could spoof binding updates from any Internet node A to any Internet node B. In our protocol, the attacker C must be on the route from B to A, which means that there are typically only tens or hundreds of nodes that can execute

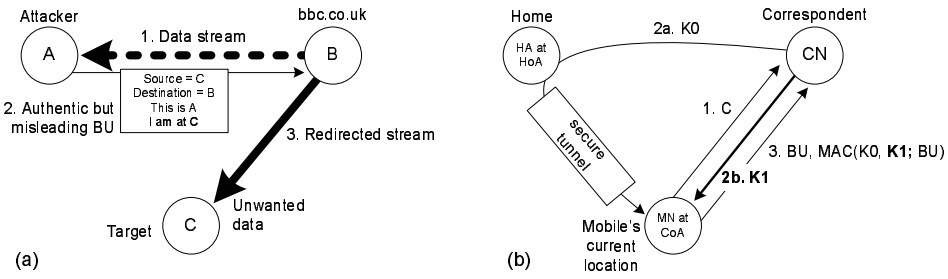

Fig. 3. Bombing attack (a) and verifying the care-of address (b)

that attack, i.e., the routers between A and B and the hosts on the local networks of A and B. On the other hand, a malicious node is able to target only the connections that pass though it its local network. For a typical attack, such as a compromised host, the number of such connections is small. This reduction in the scale of the potential damage alone means that deployment of the Mobile IPv6 would no longer be an danger to the Internet's stability.

Second, the protocol fulfills the explicit design goal of being as secure as the current Internet without mobility. Assume that the mobile node never leaves its home network and always communicates directly from the home address. In that case, an attacker on the route between A and B can spoof, intercept and sniff packets between them, and it can execute all the same attacks that were possible by exploiting the weaknesses of our BU authentication protocol. Therefore, we argue that the simple protocol of Figure 2(b) is sufficient for authenticating the sender of a binding update.

The way the mobile and the home agent establish the secure tunnel between themselves is beyond the scope of this paper and currently depends on the implementation. We assume that the mobile and the correspondent are in a close alliance and, thus, have a pre-shared key or another method for authenticating each other. Similarly, the mobile can send binding updates to the home agent via a secure tunnel, so that the home agent knows where to forward the packets. (This is theoretically straightforward but relatively cumbersome in practice because some authentication protocols use the source address of packets as the endpoint identifier.) On the other hand, if the mobile and the home agent do not have a secure mechanism for authenticating each other or if they do not trust each other, then the assumptions of our protocol do not apply and some other kind of protocol is needed.

4 Verifying the CoA

The protocol described above is sufficient to authenticate the sender of the binding update and, thus, solves the problem that we started with. In the process of designing

the protocol, we had to develop an in-depth understanding of the threats against Mobile IPv6. As a result, we discovered another type of treat that is at least as serious as spoofed binding updates. This section explains how even authenticated binding updates can be used to amplify a packet flooding attack, and how the protocol was modified to prevent the attack.

4.1 Bombing Attacks

The key observation is that a binding update contains two pieces of information, HoA and CoA, and the protocol described above only verifies the correctness of the HoA. Even if the binding update is authentic in the sense that it was sent by a mobile node whose home address is the one stated on the packet, the mobile might provide a false value for the care-of address. In other words, the mobile may be lying about its own location.

Once we made the above observation, it is easy to come up with an attack. Figure 3(a) shows a scenario where the attacker A tricks a public web site B into sending a flood of unwanted packets to a third party C. The attacker A first starts to download a stream of data, such as a long TCP stream, from a public server B. It then sends an authenticated binding update to the server claiming to be at the care-of address C. The server accept the binding update because A used an authentic home address. (A does not need to be mobile. It can use its own stationary address A as the home address and act both as the home agent and as the mobile node in the binding update protocol.) As a result, the server redirects the data stream to the false care-of address C.

Readers familiar with communications protocols will have noticed by now that B will soon stop transmitting the data stream because it does not receive acknowledgments from C. Unfortunately, this does not help much because the attacker can spoof the acknowledgments. Since the attacker received the first packets of the data stream, it knows the initial TCP sequence numbers and can spoof TCP acknowledgments. Moreover, the attacker only needs to send one acknowledgment per TCP window, which means that by spoofing only a few packets it can get B to send a large data stream to C.

Alert readers might also note that the recipient of unwanted TCP packets usually sends a TCP Reset signal to the source of the packets, which puts in immediate stop to the data stream. Thus, one might assume (as we did for quite a while) that the target node C send a TCP Reset signal to B. Unfortunately, this does not quite work in our case. The packets sent by B to C have a routing header that says the packets are intended for A. When the IP layer in C's stack processes the routing header, it encounters a strange address A and drops the packet without ever processing the following TCP header. Thus, no TCP Reset will ever be sent.

The attack is serious because it can be used to bomb any Internet node with data and the target node cannot do anything to protect itself. If used in combination with distributed denial-of-service (DDoS) attacks, the bombing attack can cause serious problems to the stability of the Internet.

4.2 A Bombing-Resistant Protocol

Out first reaction to the discovery of the bombing attack was that, before sending data to the new care-of address, the correspondent should somehow verify that the mobile is in that address. That is, the correctness of both the HoA and the CoA in the binding update needs to be checked. But secure verification of the physical location of a mobile node turns out to be extremely difficult, especially when the only thing the correspondent knows about the mobile is the home address, and when the access network does not participate in the protocol at all. The solution is to relax the requirements a bit, as we'll explain below.

Figure 3(b) shows how we the improved protocol. The correspondent sends a second secret key to the mobile. This key is sent directly to the mobile's care-of address. The correspondent uses both keys to compute the MAC on the binding update. This proves to the correspondent that the mobile is able to receive messages sent to the new care-of address. This mechanism is called *return-routability (RR) test for the care-of address*.

The above protocol does not strictly prove that the mobile is located at the new care-of address. The attacker could capture the second key at the CoA or on the route from the correspondent to the CoA. But if that is the case, the stream of unwanted packets will flow to or through the attacker's network. Thus, the attacker will suffer from the bombing as much as the target node, and the attacker could just as easily send the flood of packets itself, rather than using the correspondent as the sender.

Effectively, the correspondent in the improved protocol verifies that someone on the new route wants to receive the data. This is sufficient to make packet-flooding attacks unattractive. (The reader may have noted that the return-routability test for the CoA is similar to transport-layer acknowledgment and the same effect could have been achieved by enhancing the transport-layer acknowledgement mechanisms.)

Although we have used the terms authentication and verification, both of the return-routability tests can be seem as forms of authorization. The RR test for the HoA authorizes the sender of the binding update to change the binding for the home address. The RR test for the CoA authorizes the sender to redirect data to the care-of-address. These are quite different security goals and we have achieved them using curiously symmetric mechanisms. This is perhaps best explained by viewing the two tests as a way to verify that the sender is authorized to control the use of the two addresses.

5 Other attacks

Verification of the home and care-of addresses is sufficient to prevent most attacks that exploit weaknesses of the Mobile IPv6 route optimization. The return-routability protocol does this and, thus, protects the Internet from the new vulnerabilities introduced by the mobility mechanism.

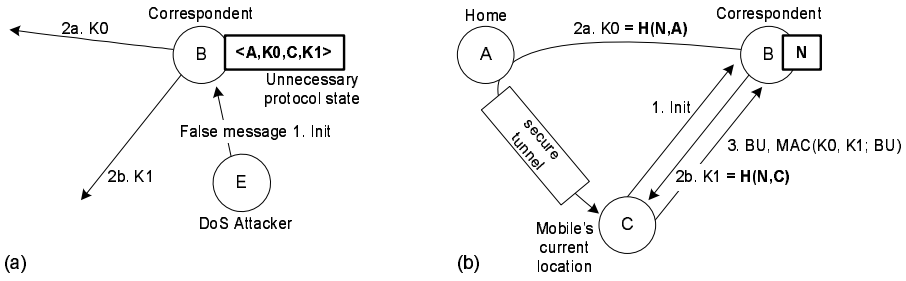

Fig. 4. State-storage exhaustion attack (a) and stateless correspondent (b)

But like in all security protocols, there are a number of potential attacks against the security protocol itself that need to be considered. In this section, we make small changes to the protocol of Figure 3(b) to prevent state-storage exhaustion and reflection attacks. We also describe a class of attack that cannot be prevented by any security protocol.

5.1 Stateless Correspondent

It is well-known that stateful protocols expose the protocol participants to denial of service attacks. In particular, if a host stores a state in a protocol run that someone else has initiated and before authenticating the initiator, an attacker can initiate the protocol many times and cause the host to store a large number of unnecessary protocol states.

Figure 4(a) shows how this attack works against our protocol. The attacker sends a spoofed initial message with a false home address A and false care-of address C. The correspondent responds with two randomly generated secret keys, which it has to remember until it receives the authenticated BU. If the attacker repeats this many times, the correspondent will not be able to store all the state data and may drop some initial messages. This may prevent legitimate mobiles from using route optimization with the correspondent.

While this attack is not nearly as serious as the one described earlier and it could be prevented by adding memory and managing the state storage carefully, it much easier to design the protocol to be stateless. (For stateless design techniques, see e.g. [AN97]). In Figure 4(b), the correspondent does not store a separate state for each mobile. Instead, it stores a single periodically-changing randomly-generated master secret N and computes the two secret keys K0 and K1 with a one-way hash function from the master secret and from the addresses (HoA and CoA). The correspondent does not remember the keys but instead recomputes them when it receives the authenticated binding update. This means that the correspondent can remain stateless until it has authenticated the mobile.

The attacks is similar to the SYN flooding attack against the TCP protocol [SKK+97]. The two messages sent by the correspondent to the mobile's home address and care-of

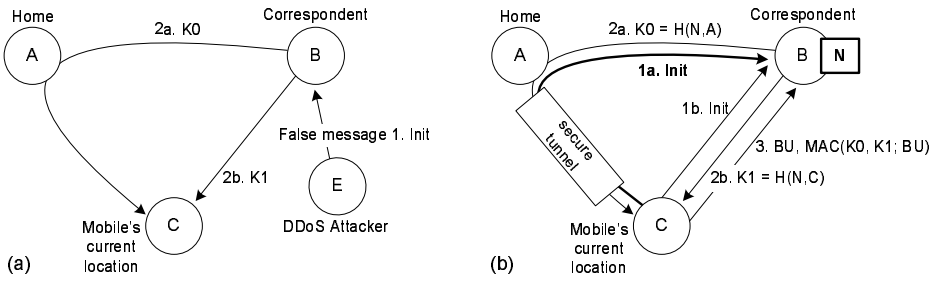

Fig. 5. Reflection attack (a) and balanced messages (b)

address are similar to the SYN cookies that have been used to prevent the SYN flooding attack.

Careful readers might suggest two further improvements: binding the two secret keys (K0 and K1) together so that keys from different protocol runs cannot be mixed and matched, and making the two key formats dfferent so that K0 cannot be replayed in the role of K1 and visa versa. While both improvements are unnecessary, they might increase the robustness of the protocol if the assumptions or goals change in the future. The idea of binding the keys together was rejected because it is useful to be able to reuse K0 within a short time even if the care-of address changes. The idea of making the key formats slightly different, on the other hand, was adopted to the standard protocol. This is done by inserting octets 0 and 1 into the respective hash inputs.

6 Preventing Reflection and Amplification

Another attack that takes advantage of the security protocol is shown in Figure 5(a). The attacker E spoofs the initial message, which induces the correspondent to send two messages to the mobile. This causes two problems. First, the attacker sends only one packet but two arrive at the mobile. Thus, the attacker can use the binding-update authentication protocol to amplify a packet flooding attack against a mobile node by a factor of two. Second, the sender address of the two messages arriving at the target of the flooding attacks have the correspondent's address as their source address. Any efficient mechanism for tracing the source of the attacks probably won't be able to trace the attack back to its real origin. (For a detailed discussion of the problems caused by reflection, see [Pax01].)

While these attacks may not seem very serious, it is hard to justify a security protocol that creates new vulnerabilities. The problem was solved by duplicating the initial message. In the protocol of Figure 5(b), the mobile sends one initial message via its home agent and another one directly to the correspondent. The correspondent responds to both initial messages independently by sending a secret key to the same address that

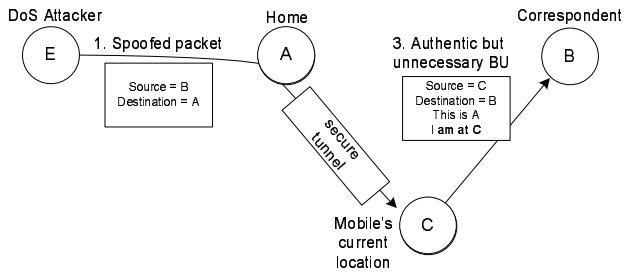

Fig. 6. Inducing unnecessary BUs

the initial message came from. The mobile needs to send both initial messages to receive both keys. The result is that the correspondent sends only as many messages as it receives, thus eliminating the amplification problem, and the correspondent always responds to the same address from which the initial message appears to come, which may make it easier to trace the origin attack using standard methods.

7 Avoiding Unnecessary Authentication

There is one last attack that needs to be considered. This attack is possible regardless of what kind of authentication is used for the binding updates. In fact, the stronger and the more expensive the authentication protocol, the more serous this attack becomes.

Figure 6 shows how the attacker can induce authentic but unnecessary binding updates. When a spoofed packet sent by the attacker is tunneled to the mobile, the mobile typically responds by executing the binding update protocol with the claimed correspondent. The correspondent will eventually accept the binding update as both HoA and CoA are true. But the protocol execution is completely unnecessary. The attacker can repeat this with many different spoofed correspondent addresses to exhaust the resources of a single mobile, or with one spoofed correspondent address and many mobiles to attack a single correspondent.

Since the IP layer is stateless and BUs may be sent at any time, there is no practical way for the mobile or the correspondent to filter out the unnecessary binding updates without dropping also necessary ones. Therefore, the best defense against this attack is to limit the resources that the nodes allocates to processing binding updates to and from previously unknown mobiles. At worst, the attacker can prevent the use of route optimization.

8 Conclusion

We have described threats against the Mobile IPv6 route optimization protocol and protection mechanism used in the standard protocol. Some of the techniques are unconventional because the protocol needs to work globally without any global security infrastructure. Without such a solution, the Mobile IPv6 protocol would have been confined to intra-oraganizational use.

The experiences from the Mobiel IPv6 design process highlight the need to consider early the potential security threats created by new technology. The some of the solutions described in this paper have been found to be applicable to other mobility protocols such as HIP [NYW03]. It is promising that not only have the designers of newer protocols learned the specific protocol mechanisms but they have also started serious threat analysis and security design at an early stage in the design process.

References

[AN97] Tuomas Aura and Pekka Nikander. Stateless connections. In *Proc. International Conference on Information and Communications Security (ICICS'97)*, volume 1334 of *LNCS*, pages 87–97, Beijing, China, November 1997. Springer.

[HD98] Robert M. Hinden and Stephen E. Deering. IP version 6 addressing architecture. RFC 2373, IETF Network Working Group, July 1998.

[Hui95] Christian Huitema. *Routing in the Internet*. Prentice Hall, 1995.

[JPA03] David B. Johnson, Charles Perkins, and Jari Arkko. Mobility support in IPv6. Internet-Draft draft-ietf-mobileip-ipv6-24.txt, IETF Mobile IP Working Group, June 2003. Work in progress.

[MC01] Gabriel Montenegro and Claude Castelluccia. SUCV identifiers and addresses. Internet Draft draft-montenegro-sucv-02.txt, IETF, November 2001. Work in progress.

[Nik01] Pekka Nikander. A scaleable architecture for ipv6 address ownership. Internet-Draft draft-nikander-ipng-pbk-addresses-00.txt, March 2001. Work in progress.

[NYW03] Pekka Nikander, Jukka Ylitalo, and Jorma Wall. Integrating security, mobility, and multi-homing in a HIP way. In *Proc. Network and Distributed Systems Security Symposium (NDSS'03)*, pages 87–99, San Diego, CA USA, February 2003.

[OR01] Greg O'Shea and Michael Roe. Child-proof authentication for MIPv6 (CAM). *ACM Computer Communications Review*, 31(2), April 2001.

[Pax01] Vern Paxson. An analysis of using reflectors for distributed denial-of-service attacks. *ACM Computer Communications Review (CCR)*, 31(3), July 2001.

[RAOA01] Michael Roe, Tuomas Aura, Greg O'Shea, and Jari Arkko. Authentication of Mobile IPv6 binding updates and acknowledgments. Internet Draft draft-roe-mobileip-updateauth-01.txt, IETF Mobile IP Working Group, November 2001. Work in progress.

[SKK+97] Christoph L. Schuba, Ivan V. Krsul, Markus G. Kuhn, Eugene H. Spaffold, Aurobindo Sundaram, and Diego Zamboni. Analysis of a denial of service attack on TCP. In *Proc. 1997 IEEE Symposium on Security and Privacy*, pages 208–223, Oakland, CA USA, May 1997. IEEE Computer Society Press.