

A new Route Optimization Protocol for Mobile IPv6 (ROM)

Christian Veigner¹ and Chunming Rong
University of Stavanger
christian.veigner@his.no, chunming.rong@his.no

Abstract-Mobile IPv6 is the mobility protocol of the IPv6 protocol. The mobility feature in IPv6 is supposed to be default, but there are still obstacles carrying out deployment. Great needs for this new Internet Protocol are security, scalability, efficiency and the lack of IPv4 addresses.

Deployment of this long needed new protocol has been delayed partly due to the route optimization problem of mobile nodes. Route optimization is the property of sending data via the shortest route, even when the mobile node is roaming, thereby making packet sending between a mobile node and a corresponding node through the mobile node's home network for further redirection obsolete. Any other nodes communicating with a mobile node are referred to as corresponding nodes. Routing packets over the shortest route will both decrease the latency of data transmission and bandwidth consumption.

In this paper, we suggest a new route optimization protocol for Mobile IPv6 (ROM), which we will show as a more efficient solution to this problem than other existing schemes.

Keywords: Mobile IPv6, route optimization, authentication.

I. Introduction

Mobile IPv6 (MIPv6) [1] is supposed to be a part of the Internet Protocol version 6 (IPv6) [2], and has a mandatory feature permitting MIPv6 nodes roaming from one subnet to another without disrupting their sessions during handover, where a handover is the change of connection from one network to another. It is crucial that handovers are authentic. IPv6 was supposed to be deployed in 2001, but partly due to authentication issues in MIPv6 route optimization, as shown in figure 1 and 2, the new IP protocol has not yet been deployed.

Consider a mobile node (MN) as shown in figure 1, having an ongoing session with a corresponding node (CN).




Figure 1. Movement of mobile node




Figure 2. Route optimization

If MN changes its point of location in the network, it is desirable to make the communication proceed between the MN and CN as shown in figure 2, without sending the packets through MN's home agent (HA) for further redirection to the MN. The HA is a node at MN's home link cooperating with MN in packet routing when MN is visiting a foreign network. In MIPv6, route optimization as shown in figure 2 is the default mode; every IPv6 node has to support it.

There are actually no distinctions between a MN and a fixed node in IPv6, which enhance the necessity of completing the work of MIPv6 even more.

Making the flow of data from a MN to a CN and vice versa without sending the packets through the HA has been a topic of recent research. There are several proposals for this type of management, e.g. RR [1] CGA [3], CAM [4], HIP [5], ABK [6], BAKE/2 [7] and IPsec [8], where each of these proposals suffers from different weaknesses. CGA is patented, CAM is similar in many ways to CGA, but has security $O(2^{62})$ assuming a brute force attack. This level of security may be enough considering most attackers currently available and affordable technology, but the design of IPv6 should be secure for more than the next decade. HIP, ABK, BAKE/2

¹This work was supported by Rogaland University Fund, UiS 95310

and IPsec assumes the existence of additional infrastructure, and will not easily integrate into existing networks for managing Mobile IPv6 handovers in a secure and scalable way. The protocol chosen by IETF is the Return Routability (RR). It has some known weaknesses, though it does not rely on any additional infrastructure.

Considering the RRv3 protocol, we will show that it might be too time consuming when a MN suddenly has to initiate a handover to a new subnet, which may result in the user suffering from lack of seamless handover.

ROM is a protocol initiated by a mobile node (MN). Its goal is to prepare the corresponding nodes (CNs), which the MN is currently communicating with, of MN's change in location prior to the actual movement. When MN changes its point of location, only one single message to each of the CNs is required to make verifiable new bindings. This will increase the probability of achieving a seamless handover compared to RRv3, where several messages must be sent and received before the handover is complete. A binding update message in MIPv6 is a message notifying the receiving node of the sender's new location.

In Mobile IPv6 the MNs must also be reachable by their IPv6 home addresses. CNs are initially unaware of the current location of a MN, and will use the MN's home address (HoA) when transmitting an initial packet. In this way, packets will be routed through the HA. The HA reroutes the packets intended for a MN to its new location. To manage this, the MN has to register its current foreign address at the HA before the HA is able to perform rerouting of incoming packets.

In time users will demand uninterrupted sessions when moving securely from one subnet to another. This is a motivation for implementing ROM.

Constructing a location management protocol we believe in the importance of scalability, entailing the importance of not making the protocol rely on additional infrastructure as certificate nodes, trusted third parties or any other infrastructure that will be both cost consuming and not easily scalable. Security issues of the protocol regarding authenticity are also very important.

The rest of this paper is organized as follows. The next section gives a short introduction to the current solution adopted by IETF, RRv3. Section III explains the properties of our solution of the same problem and gives a short introduction of the main advantages of the new protocol. In section IV we analyse different attacks possible when offering route optimization, and show to which extent our protocol withstand these attacks. Finally, we conclude the paper in section V.

II. Return Routability (RR)

Return Routability (RR) is the protocol adopted by IETF for managing mobile nodes (MNs) binding

updates (BUs) sent to corresponding nodes (CNs) and home agents (HAs). RR assumes the existence of IPsec Security Associations (SAs) [8] between a MN and its HA. The communication of mobility management messages between MN and its HA must be secured by IPsec ESP [8] using these associations.

A. Return Routability version 3

The most recent version of RR is version 3 (RRv3) [1]. In this scheme, when updating a binding at a CN, the MN has to send and receive messages from CN in three different stages, as shown in figure 3. This is all done in turn after MN has sent a binding update (BU) message to its HA and received binding update acknowledgment (BUAck) in return. In RRv3, all these messages must be sent when the MN has arrived at its new location. Sending so many messages back and forth, where some of them depends on the cooperation of the HA must be considered a great threat to the seamless handover procedure. The reason for sending messages via different routes is the belief in the hardness of an attacker monitoring two different dynamically changing routes at the same time. If the MN when sending a BU message knows the two values sent to it via different routes from CN (tokens in message 4a and 4b), the CN consider the BU authentic.




Figure 3. Return Routability version 3

If a MN moves quickly from one subnet to another, and suddenly its connection with the old subnet is disrupted, it will be out of connection until the RRv3 protocol finishes executing at the new network. Constructing a protocol reducing the binding latency reduces the possibilities of noticeable handover procedures. We believe, in time, users are going to demand non-interrupted communication with their CNs even when they are in transit, which motivates the construction of a faster handover protocol as ROM. Another major problem in RRv3 occurs if the HA is off-line at the moment MN tries to initiate a handover. If the HA is down, there is no way MN can continue its session with a CN from the new location without subscribing to a new HA before

initiating the RRV3 protocol. This problem is eluded in the ROM protocol design.

III. ROM

A. Overview

The overall goal of the ROM protocol is to minimize the elapsed time from a MN in subnet transit realizes it has to connect to its current CNs via a new subnet, to the new connection is ready for use. If the MN is moving quickly from one subnet to another and/or the radio signal strength from the old subnet is rapidly decreasing, there may be no time sending messages back and forth establishing a new verifiable connection with the CNs from the new subnet, and still provide seamless handover to the user. Using the ROM protocol, a MN only has to send one message to each of the CNs which it has ongoing sessions to re-establish a verifiable binding when changing its location. This is a huge improvement over RRV3 where MN first has to bind to its HA, which may be off-line, then exchange four messages with each CN of current interest, and then finally send a BU message to each of these CNs.

Briefly, the ROM protocol is to be used by a MN for sending a hash value via the HA to each of its currently used CNs. This is done before movement. Each of these hash values should be unique. When changing location, MN must send a binding update (BU) message directly to each of these CNs if it wants to continue its sessions. A CN is now capable of verifying the new binding due to MN's knowledge of the nonce value used to generate the hash. This nonce value must be included in the BU message. We will now explain the ROM protocol in more detail.

In ROM, as in RRV3, we assume the existence of IPsec SAs between MN and HA.

A MN may want to establish a binding with one or several previously unused CNs. If the MN is away from home, our solution is to make MN send a list to its HA containing IPv6 addresses of these CNs and corresponding hashed nonce values, see message 1 in figure 4. These hash values are generated by MN hashing different nonce values. HA sends the correct hash value to the correct CN, message 2a in figure 4, one hash value to each CN, explained in further details later. MN will then, when binding, send BU messages directly to the CNs it wants to bind with. A BU message will be verified at a CN by the nonce value included in the message. A one-way hash function is needed.

The 2a-2c sequence of messages is only intended to prepare the CNs of MN's next BU message. A BU message may be sent by MN from the current location establishing or renewing its binding, or it may be sent from MN's new location.

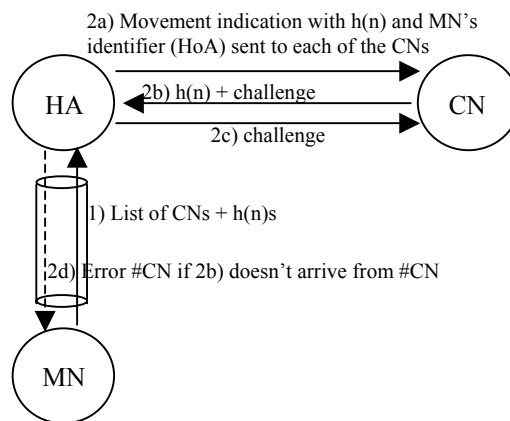


Figure 4. Preparation of CNs

When a CN gets a BU message from the MN (depicted later) containing a nonce value, the corresponding hash value used for verification will be obsolete and deleted. The CN must now be sent a new hash value as quickly as possible via MN's HA to be able of verifying a future BU message from MN.

New hash values must also be sent at given intervals to currently used CNs, even though they are in possession of hash values not yet used for verification and currently usable for verifying MN's next BU message. Sending of new hash values to such CNs could occur at given intervals, depending on the security of the hash function. If it is truly one-way the valid time may be extended. The reason for generating new nonce values and sending their corresponding hash values, even to a CN still in possession of a valid hash value usable for verification of the next BU message, is a security matter. The hash value can be detected by a passive attacker at the HA-CN link. This attacker can, if the nonce or a valid value is found, turn active and create a false BU, redirecting MN's incoming traffic.

In the general case, when the ROM protocol is used for speeding up a handover, message 2a in figure 4 is just an indication to a CN of MN's next movement. This message must also include MN's identifier (HoA), i.e. the home address of the MN, in addition to the hash value. The CN must store MN's home address, the hash value, and MN's verified care-of address (CoA) in a location table. The CoA is the IPv6 address a MN is associated with at its foreign subnet. The HoA and hash are stored at CN when receiving message 2a and 2c, figure 4, but the CoA is stored at CN when CN later receives a verified BU message from MN. CN does not change its location table for MN's CoA until a verified BU from the MN is received, thereby making it impossible for an attacker hijacking another MN's session without knowledge of the secret value used for generating the hash. Also; the hash value $h(n)$ at CN will not be changed until the 2a-2c procedure has been completed. This assures CN that the new $h(n)$ is sent from MN via its HA. It is the MN's home

network that is authenticated by this message exchange.

The 2b message in figure 4 must contain the $h(n)$ value, making the HA able to determine whether it has sent the 2a message to this node beforehand, or not. 2b should also include a challenge that must be returned by the HA. This authenticates the MN's home network. However there might be a man-in-the-middle attack, but this is as in the RRv3 protocol not considered a problem of the mobility protocol. The aim of this protocol is not solving the problems of the entire IPv6 protocol, but only to avoid new ones due to the mobility feature.

If HA doesn't receive 2b in return of 2a within a predefined short time, HA should send error message 2d as shown in figure 4 to the MN, notifying MN which CNs are unavailable/off-line, and not necessary sending BU messages to until online again.

The BU message (figure 5) from MN to CN will verify the previously received $h(n)$ at CN, by containing the correct nonce value. Inclusion of the HoA in the BU message tells the CN for which MN it should attempt verify the new binding.

The hashed nonce values sent to a CN via MN's HA should be unique. Each hash should also be different from the other hash values sent to the other CNs that MN has in its list of recently used CNs. This is a security property. An attacker may log all the different hash values sent via MN's HA, and over a long period of time find valid nonce values, or perhaps getting the nonce value sent from MN to a CN when MN is verifying its hash value. It is assumed hard for Eve monitoring both these links, but not necessarily if Eve is close to the CN we think, then the two links HA-CN and MN-CN may coincide. RRv3 bases its security on the hardness of this property. If MN uses the same hash value again, Eve can check its list, possibly finding a valid nonce for the hash value, and then hijack or redirect MN's session before MN has sent its binding update message to the CN.




Figure 5. BUs from MN

B. ROM improvements compared to RRv3

The messages in figure 4 may be sent before MN moves to its new location. If the MN moves to a new location and sends BU messages to its CNs, then the MN must once again send a list to HA providing the CNs with new hash values usable for verifying MN if it changes its point of location once again, or if the binding valid time is to expire and the need for sending a new verifiable BU message from MN's current location arises.

When MN has lost a connection or changes its point of location, it has to configure a new IPv6 address at the new subnet, as in RRv3, but is spared from sending the two messages via different routes (figure 3) and receiving the two messages also via different routes before it can send BU to its CN. In RRv3, all this has to be done after the MN successfully has updated its HA. The reason in RRv3 for sending the four messages (HoTI, CoTI, HoT and CoT) is for the CN to get an insurance that MN is where it claims to be. RRv3 will in this way be secure against attacks from elsewhere the HA-CN link, but it is still a possibility for an attacker at this link to launch a redirecting attack [9]. This is though a huge improvement, reducing redirecting attacks from the entire Internet to the HA-CN link.

The ROM protocol also uses two different routes, but this is not done as a security cause as in RRv3, but for the sake of faster handovers. One route is used for verification of the belonging home link and for sending of hash values. The other route is used for managing the actual handover when necessary as quickly as possible.

As in RRv3, redirecting attacks where Eve is not at the HA-CN link is also avoided in our scheme. ROM as a matter of fact narrows the attackers range to launch its attack from the HA-CN link as in RRv3, to the necessity of being within the home link of the attacked MN. In message 2b, figure 4, CN checks if the hash is real and is sent from the given home link by challenging the HA.

If the HA suddenly is not able to act as MN's HA, e.g. due to down time, MN's session with a CN when changing location is not interrupted, assuming CN has a valid hash value. MN's negotiation with the home network getting hold of a node willing to act as its HA does not lead to any further delays in the connection re-establishment with the CNs in the general case. This is a great advantage compared to RRv3. Seamless handovers are not dependent of the immediate cooperation of the HA.

IV. Attacks on ROM

Constructing a secure protocol it is important considering different possible attacks. As in RRv3 we will in ROM not consider attacks also possible in fixed IPv6, e.g. man-in-the-middle attacks. We only

attempt to avoid attacks possible due to the mobility feature.

A. Redirecting a mobile node's session

In a redirecting attack Eve (the attacker) must generate her own nonce and belonging hash value as a part of the attempt redirecting a MN's incoming traffic. This is the only feasible solution as long as the MN uses a one-way hash function.

Eve sending a message to CN, impersonating another MN's HA, is not possible having the 2b and 2c messages in figure 4 verifying the location of the HA and the fact that HA actually sent the hash originating from its MN. The message 2b contains the received hash value and a challenge, and in return HA will answer the challenge and indicate to CN whether CN's received hash is legitimate.

Eve is not able of redirecting a MN's incoming data by sending the CN a false hash value for which Eve has a valid nonce value, unless Eve is a node at the attacked MN's home network. This is because the home address (HoA) of the attacked MN has to be equal in prefix to Eve's prefix when Eve is acting as the HA. The HA's address and the MN's HoA has to originate from the same network, otherwise the CN will not update its table with the new hash value. The nonce value could now be used for sending a false BU message as if coming from MN, resulting in a redirecting attack on MN.

B. Flooding attacks

The CNs in this protocol must have a table as shown in figure 6. If e.g. a PDA initiates a session with a MN, and the MN is currently not at home, the PDA must store the location information it receives in its table. This may sound as a drawback, but managing route optimisation, a node has to allocate memory for a location table. The restrictions and properties of this table will be subject to further research.

HoA	h(n)	CoA
⋮		

Figure 6. Location table at CN

This table consists of a MN's HoA, hashed value and CoA. CN records this information for each of the MNs it currently has established route optimization for. There is a possibility of flooding attacks against this table at CN, possibly resulting in denial of service (DoS). We may think of an Eve sending a movement indication message (2a, figure 4), resulting in CN storing the HoA and h(n) values. This only occurs if

Eve is responding to the 2b with the 2c message. Route optimization is though an optimization, and may be turned off by a CN under attack.

1) If Eve is impersonating another HA and sends a CN new information (HoA and h(n), message 2a in figure 4), CN will not get the expected 2c message in return from the impersonated HA, but possibly an error indication, and will in turn reject the information sent by Eve.

2) A flooding attack against CN's buffer may be launched by an Eve who uses her own address acting as a HA and sends a 2c message in reply to the challenge from CN as mentioned. Eve is now able of sending many different HoAs and h(n)s to this CN. CN will store these values in its table as long as the prefixes of the HoAs are similar to the HA prefix. This can be considered a problem, but may be minimised by restricting the number of MNs a HA can have in its domain. Another possibility is CN restricting the number of allowed nodes from a given network performing route optimization. A distributed attack is still possible, even though these rows will be deleted after a short time when not used. Anyway, this seems to be the most critical part of our protocol. It may lead to DoS attacks, and should most definitely be topic of further research. It must be added; the existence of a table at CNs for managing route optimization can't easily be omitted.

3) Eve may also flood a CN by sending lots of BU messages from different random generated addresses. CN should remember these BUs in case it receives an authentic 2a message, figure 4, shortly after from a MN establishing route optimization. In this case the MN sends the BU message before the h(n) value has reached the CN. This may only occur in some specific cases, which we will not elaborate in this paper. Preventing an attack, CN must drop the BUs that it can't yet authenticate after e.g. 3 seconds. In this case, if a non-fraudulent node's BU message was dropped, because of delays in the MN-HA-CN route, it has to retransmit its BU message. MN is almost certain of CN's availability when not receiving an error message from HA within a few seconds, indicating that CN is unavailable/off-line, message 2d in figure 4.

C. Bombing attacks

We may consider a scenario where an attacker is not impersonating anyone, but connects to a CN offering video streaming services. The attack is to redirect this stream of data to a victim node anywhere in the entire Internet. This attack is feasible since the attacker may generate a nonce value and a belonging hash value. Eve will initiate the sequence getting CN ready for her spurious movement. The CN is now provided with Eve's correct HoA and a h(n) value waiting to be verified by a BU message sent from Eve from her new location. Forging the BU message using the HoA of Eve and inserting the verifiable nonce

value, Eve uses the victim's address as the CoA. CN will now redirect its streaming service to the specified address in the source address part of the BU message's IPv6 header, the victim MN. See [1] for the contents of the BU message.

This bombing is only efficient if the attacker at some point has redirected the stream from itself. By this strategy gaining information about the sequence numbers in the incoming TCP packets. The attacker has to send false TCP acknowledgements back to the streaming provider once in a while, inserting the address of the victim as the source address in the packets, to make the attack continual. The attacker only has to send one acknowledgement message per TCP window to realize its intention.

Mobile IPv6 defines a new routing variant, type 2 routing header, used for sending packets optimized to a MN from a CN. The bombing, as described, seems efficient, and can't easily be stopped by the attacked node. In this attack, the attacker must use its own HoA as a home address option and the CoA of the attacked node as the source address part of the BU message to fulfil its bombing intension. The CN will now redirect the stream to its victim. When receiving the stream, the attacked node will drop any packet at IP level when detecting that wrong HoA is used as the final destination. This is due to the routing header type 2 used for route optimization, see [1] for explanation of the use of routing headers. If all packets of the stream are silently discarded, it is not possible for the MN to process the attack at a higher level generating a message telling CN to halt the stream.

If possible somehow, informing the CN to stop the bombing, the solution depicted in figure 7 might be a solution.




Figure 7. Authentic BU message

The CN could piggyback the received nonce value to the MN in the mobility options of the returned BUAck message (figure 8). By IETF the reason for having the mobility option as a part of the BUAck message is to allow future extensions to be defined. If the MN didn't send the nonce, an error message to

CN should be sent. The added time penalty is limited using piggybacking.

	Status	Reserved
Sequence #	Lifetime	
Mobility Options		

Figure 8. BUAck message format

MN may also overcome an attack where Eve located at MN's home link updates MN's CN with a false $h(n)$ value. This attack can be handled if CN sends the MN the currently stored $h(n)$ as a part of the mobility options in the BUAck message. This should only be done if the stored $h(n)$ at CN is different from the hash value generated at CN from the received nonce value in the BU message from MN. Sending the stored $h(n)$ value back to MN, MN knows if it has been attacked by a node at its home link having changed the $h(n)$ value, which is possible in our scheme, but only from within the home link. The MN must now send a new hash value via HA to CN, making another attempt binding with CN.

D. Amplification attacks

In our protocol an attacker is not able to send one message and thereby making the receiver send more than one message. HA and MN has a pair of security associations, and neither HA nor MN can be used for launching this attack against one of the other nodes, while sending one message to them only results in at most one message sent from them. Also; considering CN (figure 4, 5 and 7) we can see that sending one message to it only results in one message sent from it. The attacker can thereby not amplify its attack.

V. Concluding remarks

Our ROM scheme is intended to make the handover in Mobile IPv6 more seamless than RRv3 manage, i.e. to speed up the actual handover and also provide similar security characteristics. This is attempted done by using some info sent to the CNs before changing location. Then, when changing location, only one verifiable BU message is sent to each of the currently used CNs.

The main drawback of the ROM protocol as we see it is the allocation of a table at CNs, which may be any node. A CN will store states in certain columns of the table after executing the three-way handshake with a HA. This is a potential for launching DoS attacks filling it.

The dangerous bombing attack in section IV.C might be solved omitting the use of a routing header

type 2 in the BUAck message. This will be a topic of further research.

Redirecting attacks as mentioned in [9] at the RRv3 protocol from the HA-CN link is in ROM limited to the home link of the attacked node as mentioned in section IV.A.

The home agent in our scheme is also a less crucial node, and handovers can be carried out seamlessly even if the HA is down. This is a great improvement compared to the RRv3 protocol.

As we all know, bandwidth in mobile and wireless networks is very unpredictable and often low. This is the main benefit of the ROM protocol; the reduction of the amount of messaging necessary to re-establish the connection when arriving at a new network.

It is also very important that we understand all the threats the new technology creates before a possible deployment.

VI. References

- [1] D. Johnson, C. Percins, J. Arkko, RFC 3775, "Mobility Support in IPv6", Jun. 2004.
- [2] S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Spec.", Dec. 1998.
- [3] T. Aura, "Cryptographically Generated Addresses (CGA)", Feb. 2004.
- [4] G. O'Shea, M. Roe, "Child-proof Authentication for MIPv6 (CAM)", 2001.
- [5] P. Nikander, J. Ylitalo, J. Wall, "Integrating Security, Mobility and Multi-homing in a HIP Way", 2003.
- [6] S. Okazaki, A. Densai, C. Gentry, J. Kempf, A. Silverberg, Y. L. Yin, "Securing MIPv6 Binding Updates Using Address Based Keys (ABKs)", Oct. 2002.
- [7] M. Roe, T. Aura, G. 'Shea, J. Arkko, "Authentication of Mobile IPv6 Binding Updates and Acknowledgments", 2002.
- [8] IETF, "IP security Protocol (IPsec)".
- [9] R. H. Deng, J. Zhou, F. Bao, "Defending Against Redirect Attacks in Mobile IP".