

Network Working Group	A. Ebalard	
Internet-Draft	EADS	
Intended status: Informational	July 26, 2010	
Expires: January 27, 2011		

[TOC](#)

Mobile IPv6 IPsec Route Optimization (IRO) draft-ebalard-mext-ipsec-ro-02

Abstract

This memo specifies an improved alternate route optimization procedure for Mobile IPv6 designed specifically for environments where IPsec is used between peers (most probably with IKE). The replacement of the complex Return Routability procedure for a simple mechanism and the removal of HAO and RH2 extensions from exchanged packets result in performance and security improvements.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 27, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1.](#) Disclaimer and conventions
 - [1.1.](#) Disclaimer
 - [1.2.](#) Conventions used in this document
- [2.](#) Introduction
 - [2.1.](#) Current situation
 - [2.2.](#) Characteristics of IRO
 - [2.3.](#) Motivation
 - [2.4.](#) Notes to the reader
- [3.](#) Overview
 - [3.1.](#) The big picture
 - [3.2.](#) Pre-binding steps
 - [3.3.](#) BU emission
 - [3.4.](#) Proof of CoA ownership
 - [3.5.](#) BA emission
 - [3.6.](#) Post-bindings steps
- [4.](#) Proof of CoA ownership
 - [4.1.](#) Position of the problem
 - [4.2.](#) Overview
 - [4.3.](#) Mobility Options
 - [4.3.1.](#) Nonce option
 - [4.4.](#) IRO Messages
 - [4.4.1.](#) Address Ownership Test Offer (AOTO)
 - [4.4.2.](#) Address Ownership Test Challenge (AOTC)
 - [4.4.3.](#) Address Ownership Test Response (AOTR)
 - [4.4.4.](#) Address Ownership Test Status (AOTS)
 - [4.5.](#) Concrete uses of AOT* Messages
 - [4.5.1.](#) Registration with a CN
 - [4.5.2.](#) Early test of CoA ownership
 - [4.5.3.](#) Test of HoA ownership
- [5.](#) Remapping rules
 - [5.1.](#) Overview
 - [5.1.1.](#) Remapping rules for outgoing traffic
 - [5.1.2.](#) Remapping rules for incoming traffic
 - [5.1.3.](#) On-wire addresses access from userland
 - [5.2.](#) Details of traffic processing
 - [5.2.1.](#) Non-MH traffic (data traffic)
 - [5.2.1.1.](#) Incoming traffic
 - [5.2.1.2.](#) Outgoing traffic
 - [5.2.2.](#) MH traffic
 - [5.2.2.1.](#) Incoming traffic
 - [5.2.2.2.](#) Outgoing traffic
 - [5.2.3.](#) Related traffic (ICMPv6 error messages ...)
- [6.](#) Extending advantages of IRO to the HA
 - [6.1.](#) Rationale and expected advantages
 - [6.2.](#) Changes to HA processing
 - [6.3.](#) Changes to MN processing
- [7.](#) Implementation Notes

7.1.	Nested SA
7.2.	Having IKE traffic flow via the IPsec tunnel to the HA
7.3.	Remapping rules and old IPsec architecture
7.4.	Userland and address remapping
8.	Security Considerations
8.1.	Proof of address ownership
8.1.1.	Position of the problem
8.1.2.	Home Address ownership
8.1.3.	Care-of Address ownership
8.2.	Remapping (comparison with explicit HAO/RH2 inclusion)
8.3.	Anonymity
8.4.	Limiting attack surface
9.	IANA Considerations
10.	Acknowledgements
11.	References
11.1.	Normative References
11.2.	Informative References
Appendix A.	Ability to send does not prove CoA ownership
Appendix B.	IKE exchanges use the HoA and the tunnel to the HA
Appendix C.	Arguments for no regular check of HoA ownership
Appendix D.	Lack of encryption between MN and HA
Appendix E.	What if I don't need protection?
Appendix F.	MTU Gains
Appendix G.	Compatibility with static keying
Appendix H.	Compatibility with the use of CoA in SP/SA
Appendix I.	Rationale for not specifying a new BU
Appendix J.	Anonymity
§	Author's Address

1. Disclaimer and conventions

[TOC](#)

1.1. Disclaimer

[TOC](#)

This memo covers MIPv6 Route Optimization in IPsec/IKE environments. For that reasons it is expected that the reader be familiar with the main reference documents associated with those topics.

This includes the main MIPv6 reference documents ([\[RFC3775\]](#) (Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.), [\[RFC3776\]](#) (Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents," June 2004.), [\[RFC4877\]](#) (Devarapalli, V., "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture," April 2007.), ...) and main

IPsec/IKE reference documents ([\[RFC4301\]](#) (Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.), [\[RFC4303\]](#) (Kent, S., "IP Encapsulating Security Payload (ESP)," December 2005.), [\[RFC4306\]](#) (Kaufman, C., "Internet Key Exchange (IKEv2) Protocol," December 2005.) and their previous versions).

For the discussions regarding the security of route optimization (proof of address ownership, mainly) [\[RFC4225\]](#) (Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background," December 2005.) is a must read and [\[RFC4651\]](#) (Vogt, C. and J. Arkko, "A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization," February 2007.) provides a good summary of the issues and previous work on possible solutions.

The Security Considerations section (section 6) of [\[RFC4866\]](#) (Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6," May 2007.) also provides a good security-oriented introduction to the address ownership problem.

1.2. Conventions used in this document

[TOC](#)

In this document, except otherwise specified:

*"IKE" is used as a placeholder for both [IKEv1](#) (Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)," November 1998.) [\[RFC2409\]](#) and [IKEv2](#) (Kaufman, C., "Internet Key Exchange (IKEv2) Protocol," December 2005.) [\[RFC4306\]](#).

*"Peer" is used as a placeholder for a MIPv6 end entity, i.e. a CN or a MN.

*"HAO" is used as a placeholder for Destination Options Header carrying a Home Address Option. When the address in a HAO is considered, it denotes the address found in the Home Address field of the Home Address option carried in the Destination Options Header.

*When "tunnel" is used to designate the IPv6-in-IPv6 path between the MN and its HA, IPsec in tunnel mode is assumed to be in place.

*When "MN-CN" is used it also obviously includes the MN-MN case where the second MN acts as a CN for the first MN.

*When "IPsec flow" applies to a MN-MN or MN-CN communication, the address of the MN considered for the associated SP is the HoA.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\] \(Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

2. Introduction

[TOC](#)

2.1. Current situation

[TOC](#)

[Mobile IPv6 specification \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) [RFC3775] mandates the use of IPsec for protection of communications (control and optionally data) between a Mobile Node and its Home Agent. Support for static keying is made mandatory, and dynamic keying optional. The protection is made possible by the trust relationship that preexists between the HA and the MN: they belong to a common trust domain (the same network, a PKI). Interactions between MIPv6 and IPsec/IKE for MN and HA exchanges are partly covered in [\[RFC3776\] \(Arkko, J., Devarapalli, V., and F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents," June 2004.\)](#) and [\[RFC4877\] \(Devarapalli, V., "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture," April 2007.\)](#). For implementation reasons outside the scope of previous reference documents, some additional changes to IPsec/IKE are required to support Mobile IPv6. [\[MIGRATE\] \(Ebalard, A. and S. Decugis, "PF_KEY Extension as an Interface between Mobile IPv6 and IPsec/IKE," August 2008.\)](#) specifies a way to implement those changes by extending PF_KEY framework.

[\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) also specifies a Route Optimization procedure which allows direct communications to occur between a Mobile Node (MN) and a Correspondent Node (CN), without suffering the delay associated with the routing through the MN's Home Agent. The setup of this optimized routing is based on a mechanism called Return Routability Procedure (RRP).

One of the main hypothesis behind the design of Return Routability Procedure is the lack of trust relationship between the MN and its CN. This results in a complete lack of security in terms of privacy and authentication of data: the procedure mainly provides a limited proof of MN's HoA and CoA addresses ownership to the CN.

In trust domains (networks with an underlying PKI infrastructure) where Mobile IPv6 gets deployed using dynamic keying (IKE) for negotiating Security Associations, Mobile Nodes are already provisioned with

credentials (X.509 certificates). In those environments, the initial hypothesis that led to the design of RRP and its associated limited security abilities does not hold anymore.

At the moment, [\[CNIPsec\] \(Dupont, F. and JM. Combes, "Using IPsec between Mobile and Correspondent IPv6 Nodes," August 2008.\)](#) only describes how IPsec can be used to protect signaling traffic between the Mobile Node and the Correspondent Node but only provides a limited coverage of the problem.

2.2. Characteristics of IRO

[TOC](#)

This document defines an extension of Mobile IPv6 protocol that aims at replacing common RRP and RO procedures by a mechanism called IPsec Route Optimization (IRO) in environments where IPsec and IKE are used. It allows MN to mount and maintain direct IPsec-protected communications with CN and other MN with which they share some trust relationship, in a completely transparent fashion for upper layer protocols.

IRO is not a detailed set of requirements for IPsec to work between MN and CN but a new mechanism resulting from the tight integration and joint efforts of MIPv6, IKE and IPsec to provide a secure and scalable mobility service.

The main functional and security advantages that best describe IRO are:

- *Protected MN-CN binding using IKE-negotiated IPsec SAs.
- *Complete transparency for IKE (negotiation, rekeying, movement, ...) and other upper layers, including layer 14 (the user).
- *Compatibility with both tunnel and transport mode IPsec protection between peers.
- *Compatibility with static keying (See [Appendix G \(Compatibility with static keying\)](#)).
- *In MN-CN case, non-disclosure of MN's HoA on its foreign link.
- *No additional changes to IPsec or IKE protocols and limited changes to MIPv6 via four simple messages and an option resulting in simple and generic integration within IPsec and Mobile IPv6 stacks.
- *Improved and more generic proof of address ownership mechanism.
- *Safe by default behavior avoiding direct unprotected traffic flows.

- *Complete removal of RH2 and HAO, resulting in simplified packet handling on both sides and possibly better compatibility with filtering implemented in the network.

- *Per packet MTU gains between 24 and 48 bytes in comparison with equivalent uses of IPsec in standard R0 context. Details are provided in [Appendix F \(MTU Gains\)](#).

The main prerequisites of IRO are:

- *Existence of a trust relationship between peers (i.e. shared secret or ability to use IKE).

- *Required protection of peers' exchanges (i.e. IPsec is used between peers). IRO does not apply to direct unprotected communications between peers. More precisely, IRO prevents them.

- *To fully benefit from IRO improvements, data traffic between the MN and its HA must be exchanged through an IKE-negotiated movement resistant IPsec tunnel. If this hypothesis is not fulfilled, IRO will still be usable but some security features listed previously will be lost ([Appendix D \(Lack of encryption between MN and HA\)](#)).

2.3. Motivation

[TOC](#)

The motivation behind this work is the direct need for both efficient and secure communications in Mobile IPv6 environments already benefiting from an underlying trust domain.

The first intended target of the mechanism described in this memo is the growing number of corporate networks where PKI are now widespread. This is generally due to the increasing number of services (802.1X, SSL/IPsec VPN, TLS Web portal, S/MIME, ...) that use them on a daily basis as the root of their security and to provide logical segregation. It is also suitable for other kinds of communities.

In environments where data confidentiality and privacy do matter (IPsec is used for the protection of data between the MN and its HA), current RRP and R0 between peers of the trust domain are usually deactivated:

- *to prevent direct unprotected communications between peers that would bypass protected tunneling through the Home Network.

- *because their implementation and setup with IPsec/IKE does not work out of the box and is not trivial even if [\[CNIPsec\] \(Dupont,](#)

[F. and JM. Combes, "Using IPsec between Mobile and Correspondent IPv6 Nodes," August 2008.](#)) helps for signaling.

This results in heavy constraints on the HA (which handles all the traffic to/from its MN) and the de-facto inability to get direct end-to-end IPsec-protected MN-CN and MN-MN communications.

The ability to reduce the number of communications performed by the Mobile Node that get tunneled through the HA is both an improvement in term of upload bandwidth consumption on the link to the HA, cryptographic processing requirements on the HA and also in term of latency for applications that directly benefit from end-to-end connections, like Chat, VoIP, Videoconferencing or direct file exchanges.

In a sense, there is a kind of vicious circle regarding the use of IPsec/IKE with various protocols, including MIPv6: because dynamic keying and IPsec are not considered the common case, they are not fully covered in specifications (static keying for simple modes). The net effect is that their implementation and deployment is then complicated, which results in limited use. In a sense, IRO tries to break that circle. Simply put, this specification considers IKE-enabled environments as the first target and then covers static keying cases.

2.4. Notes to the reader

[TOC](#)

The mechanism described in this memo is very simple from a design perspective. To keep this apparent simplicity and the reading of the document pleasant, all design decisions and main justifications are provided in the numerous appendices (around 10 pages). This allows to focus on the details of the mechanism in the body of the document (around 20 pages).

For previous reason, the reading of the document can be performed linearly. The not so curious reader can skip over the appendices which are only a must read for developers and security people to acquire a deep understanding of the mechanism and how security has been taken into account in its design.

Unlike many other IETF documents, this memo voluntarily provides a practical implementation feedback geared towards developers. Even if the associated section does not mandate an implementation design, it might be of interest anyway.

3. Overview

[TOC](#)

The whole document is geared towards improved security between MIPv6 nodes and also improved usability of IPsec/IKE with MIPv6. This section

provides to the reader a quick non-normative overview of how IRO works, before entering the details of the mechanism in next sections. The reader is expected to be familiar with the vocabulary used in [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#). We do not consider in this section the relationship between the MN and its HA, only the relationship between a MN and its CNs. In the whole document, IKE is considered as the default mechanism used for SA setup.

This section provides a quick and non-normative overview of IRO and introduces next sections that contain normative details. The first subsection provides a rough outline of IRO. It is followed by 5 small subsections that cover the steps of IRO processing, in the order they occur:

- *Pre-binding steps: installation of remapping rules, which IRO uses to prevent the use of RH2 and HAO in incoming and outgoing signaling packets (MH traffic).
- *BU emission: description of the steps that apply to the emission of the BU by the MN in the context of IRO.
- *Proof of CoA ownership: description of the steps that occur when a proof of CoA ownership is requested by the peer.
- *BA emission: description of the steps that apply to the emission of the BA to the MN in the context of IRO.
- *Post-binding steps: installation of additional remapping rules, which IRO uses to prevent the use of RH2 and HAO in incoming and outgoing data packets.

3.1. The big picture

[TOC](#)

This section simply lists the key ideas and design concepts behind IRO. When IPsec is used between two peers, every packet de facto contains a simple piece of information (the SPI) that gives access to many parameters. Among those parameters synchronized between the two IPsec stacks are address information for both peers.

Unlike IPsec, MIPv6 uses specific extensions (RH2 and HAO) to explicitly carry address information. When both protocols are used together and the IPsec SA/SP make use of HoA (i.e. not CoA), the RH2 and HAO extensions in packets carry the HoA. This is redundant information which could easily be provided by the IPsec stack. Based on previous observation, IRO removes RH2 and HAO extensions from packets and replaces them by simple additional steps on the sender and

the receiver: remapping rules for addresses based on the information maintained by the IPsec stack.

Previous removal of HAO and RH2 extensions from traffic between peers is also perfectly applicable to the traffic between a MN and its HA. This specification extend the remapping rules to the traffic between a MN and its HA. When IRO is used, RH2 and HAO extensions are simply not seen on the wire as depicted below:

Without IRO:

```
HA                                                    MN

      ESP(MH Type 5 (BU) w/ AltCoA Opt.)/
      DestOpt(HAO(HoA), nh=50)/
      IPv6(src=CoA, dst=HA, nh=60)
<-----
      ESP(MH Type 6 (BA))/
      RH2(HoA, nh=50)/
      IPv6(src=HA, dst=CoA, nh=43)
----->
```

With IRO:

```
HA                                                    MN

      ESP(MH Type 5 (BU) w/ AltCoA Opt.)/
      IPv6(src=CoA, dst=HA, nh=50)
<-----
      ESP(MH Type 6 (BA))/
      IPv6(src=HA, dst=CoA, nh=50)
----->
```

As stated previously, the hypothesis on which common R0 and RRP are based simply do not hold when peers are able to use IPsec/IKE between them. For that reason, even if some proof of address ownership is still required, a more suitable (read simple) mechanism is defined for that purpose.

To sum it up (simplistic vision):

- *IRO simplifies packets format by removing HAO and RH2 extensions.

- *IRO fully replaces RRP by a more suitable and simple mechanism taking into account the use of IPsec/IKE between peers.

- *IRO defines how this can be extended to MN-HA exchanges.

3.2. Pre-binding steps

[TOC](#)

Before any direct communication can take place between a MN and a CN, the CN must accept a binding between the CoA and the HoA of the MN. For that to happen, the CN must have acquired the proof of both HoA and CoA addresses ownership by the MN.

In RRP, the MN proves to the CN its ability to both send and receive traffic from and at those addresses by a four messages exchange combining both direct and HA-tunneled packets.

In the context of IRO, the binding registration between peers is IPsec-protected. It is expected that IKE be used for negotiating an initial pair of ESP transport mode IPsec SAs between the HoA of the MN and the address of the CN for protecting this registration (static keying is covered later in the document). IKE negotiation is performed between the HoA of the MN and the address of the CN: the use of the HoA implies that IKE packets are routed via the tunnel between the MN and its HA. This provides the CN the initial proof of HoA ownership by the MN. Using the IPsec-protected tunnel (expected routing path) between the MN and its HA for IKE negotiation is not straightforward; this is discussed in [Appendix B \(IKE exchanges use the HoA and the tunnel to the HA\)](#).

On both entities, the specific IPsec ESP transport mode SAs (protecting MH traffic) created between the peers are taken into account by IRO code in Mobile IPv6 stack. This triggers the setup of specific "remapping rules" on both entities, that will be applied to incoming and outgoing IPsec packets associated with the SAs:

1. On the MN, the outgoing IPsec traffic associated with the SA to the CN has its source address remapped to the address stored (by MIPv6 process) in the ancillary data of the packet (the CoA).
2. On the CN, the incoming IPsec traffic associated with the specific SA from the MN has its source address remapped to the source address in the SA (the HoA of the MN). The remapped address (CoA) is kept as an ancillary data in the local packet structure for further processing. The packet is then naturally handled by the IPsec stack.
3. On the CN, the outgoing IPsec traffic associated with the SA to the MN has its destination address remapped to the address stored in the ancillary data of the packet, if not null.

4. On the MN, the incoming IPsec traffic associated with the SA to the CN has its destination address compared with the CoA the MN is asking a binding for to the CN. On match, the destination address of the IPsec packet is remapped to the destination address in the SA (the HoA of the MN). The packet is then naturally handled by the IPsec stack.

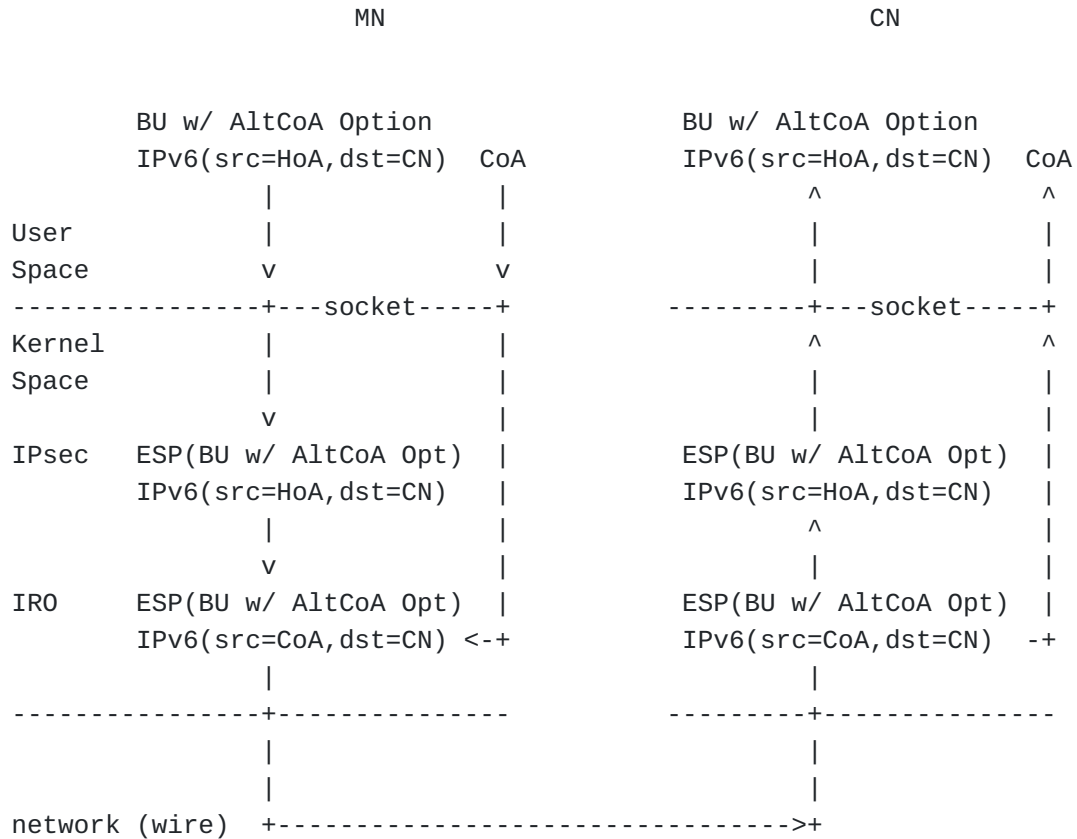
Simply stated, rules 1 and 3 will end up performing a remapping of HoA used in outgoing IPsec packets in their CoA counterparts and rules 2 and 4 will do the opposite on the other side for incoming IPsec packets. This is depicted in [Section 3.3 \(BU emission\)](#).

Note that these rules apply only to IPsec packets associated with SA that protect MH traffic. They are used before any data packet is received or sent by the entities using a direct path.

3.3. BU emission

[TOC](#)

The MIPv6 stack on the MN emits a Binding Update packet containing a Mobility Header AltCoA option which carries the CoA it is proposing a binding for to the CN. This packet is sent from the HoA of the MN to the address of the peer. The CoA is put as an ancillary data in the local packet structure for further processing. As it matches the IPsec SA put in place between the MN and the peer, it gets handled by the IPsec stack to be ESP protected. Before leaving the MN, it passes the set of MIPv6 rules for the MN; a match is found against rule 1, so that the source address of the packet is remapped to the address available as an ancillary data in the packet, the CoA of the MN. This is depicted below (note that the picture is also valid when the CN is the MN's HA):



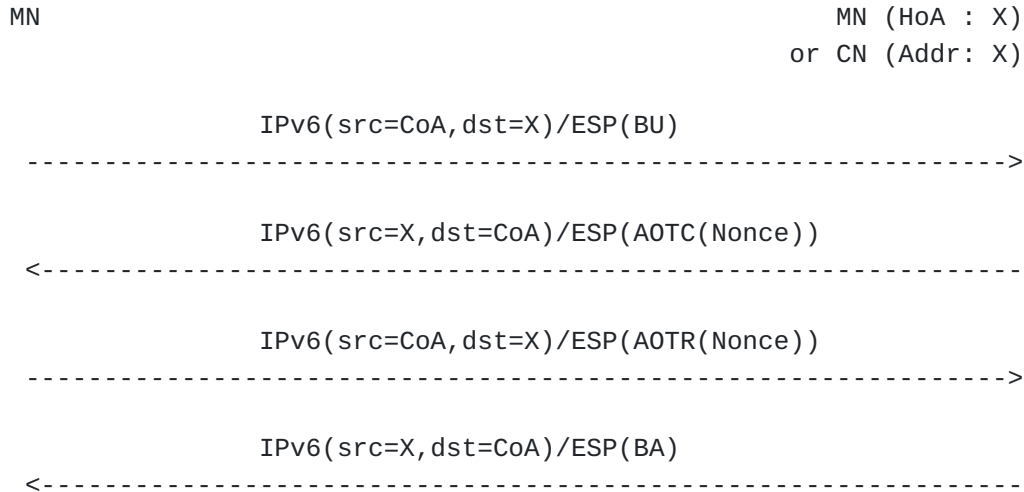
When the IPsec protected BU hits the remote MN, it passes the set of MIPv6 rules for the CN. It matches rule 2 so that its source address is remapped to the source address of the SA (the HoA of the MN). The source address found in the packet is stored as ancillary data. The packet is handled by the IPsec module, matches the SA, is decrypted and passed to the upper layer, the MIPv6 process. During parsing, the CN compares the content of the AltCoA option with the address previously stored as ancillary data. This is depicted on previous picture.

3.4. Proof of CoA ownership

[TOC](#)

At that point, before accepting the binding and replying with a BA, the CN must have the proof of CoA ownership from the MN. If one is already available, it simply goes on and sends a BA, as described in 3.5. Otherwise, it first performs following steps. It sends a newly defined AOTC message (Address Ownership Test Challenge) to the MN, providing the CoA as ancillary data, so that the remapping rules will make the packet use the CoA for the address found in the IPv6 header destination field. This packets carries a freshly generated nonce. It is recorder locally for a future check.

On the MN, the packet follows the reverse remapping process, the CoA being remapped to the HoA and passed as ancillary data. The MIPv6 stack replies with an IPsec protected AOTR message (Address Ownership Test Request) to the CN, using the HoA as local source but providing the CoA as ancillary data. The remapping rule makes the CoA the on-wire address of the packet. This packets carries the nonce sent by the CN. The CN receives the packet and after having checked the source address and the nonce against the one previously sent, the MIPv6 stack records the address ownership of the CoA for that MN, and continues with the steps described in [Section 3.5 \(BA emission\)](#)



3.5. BA emission

[TOC](#)

The CN constructs a Binding Acknowledgement packet to be sent to the HoA of the MN. The CoA of the MN is put as an ancillary data in the local packet structure for further processing. Now, as the BA matches the SA, it is ESP-protected and passes the set of MIPv6 rules for the CN. It matches rule 3 and the HoA is replaced with the address available in the ancillary data of the packet (MN's CoA). The packet is then sent. At that moment, if the status code in the BA is 0 (Binding Update Accepted), the binding is effective on the CN.

On the MN, the IPsec protected BA is received, it is passed through the set of MIPv6 rules for the MN and matches rule 4. The destination address is changed to the destination address of the SA (HoA of the MN). It is then handled by the IPsec module, and then by the MIPv6 process. If the status code in the BA is 0, the binding is effective on the MN.

For the rest of this section, we consider the binding is effective on both sides. Other scenarios are covered in details in [Section 3 \(Overview\)](#).

3.6. Post-bindings steps

[TOC](#)

In [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#), when the RRP has completed successfully, routing of traffic between the MN and the CN is automatically modified to follow a direct path. With IRO, on the contrary, a successful binding between a MN and a CN does not trigger any change in routing of `_regular_` traffic between the MN and the CN. It still flows using the IPsec tunnel through the MN's HA. Only IPsec traffic is optimized. This design decision provides a "safe by default" behavior and avoids a successful binding to lead to unprotected direct communications. Furthermore, only IPsec flows will be able to take advantage of the direct path between the MN and the CN. Arguments for this design are provided in [Appendix E \(What if I don't need protection?\)](#).

So, if existing IPsec SAs protecting non-signaling traffic (data) are already available on both sides, that have the HoA and the address of the MN as address selectors, remapping rules are put in place to perform the same kind of address changes presented in four previous rules. Those rules are static ones (they do not use any ancillary data) that remap CoA to HoA and HoA to CoA (after some checks), respectively before and after IPsec processing.

They simply replace the HAO and RH2 headers inclusion and parsing on both sides by using address information provided by the IPsec stack.

4. Proof of CoA ownership

[TOC](#)

4.1. Position of the problem

[TOC](#)

A CN accepting a binding for the CoA of a peer is not something harmless. In the context of IRO, this decision is based on:

- *a proof of HoA ownership by the MN at the time the SA is negotiated.

- *a proof of CoA ownership by the MN.

The existence of a strong trust relationship between the two (pairs of SA) and an easy proof of emission capability from the CoA are unfortunately insufficient proofs of CoA ownership. As covered in

[Appendix A \(Ability to send does not prove CoA ownership\)](#), a proof of the ability for the MN to receive traffic at its asserted CoA is required to workaround the lack of ingress-filtering at the scale of Internet: it avoids the CN to involuntarily take part in a DoS against the provided CoA.

4.2. Overview

[TOC](#)

As the proof of HoA ownership is only required to occur once in the context of IRO, the mechanism focuses on the proof of CoA ownership. Instead of reusing the complicated RRP, IRO directly benefits from the available IPsec protection between the MN and its CN to simplify things.

Furthermore, in the context of IRO, the lifetime of the provided proof is no longer limited and generally de-correlated from registration steps. This already reduces the amount of transferred data and leaves room for further optimizations (nodes with multiple simultaneous connections, nodes with limited numbers of foreign networks, ...). As CoTI and CoT messages have some associated requirements, options and semantic, and also lacks some expressiveness, they are not reused for IRO proof of address ownership. It is based on the four new extremely simple messages introduced previously:

*AOTO (Address Ownership Test Offer): Sent by a MN to a CN to offer to prove address ownership.

*AOTC (Address Ownership Test Challenge): Sent by a CN to the MN at the address to be tested, with a Nonce option that will be returned in an AOTR message.

*AOTR (Address Ownership Test Response): Sent by the MN as a response to an AOTC, with the received Nonce option.

*AOTS (Address Ownership Test Status): Sent by the CN to the MN to provide a status on ongoing address ownership test.

The format of those messages is provided in the next following sections.

4.3. Mobility Options

[TOC](#)

Reception of this message can trigger the emission of either:

- *an AOTC containing a Nonce option, sent back to the source address of the AOTO.
- *an AOTS with status 0, indicating that the peer does not allow the peer to pre-register CoA ownership information.
- *an AOTS with status 1, indicating to the peer that the proof of Address ownership is still valid.
- *nothing if it is invalid or sent by an unregistered MN.

The format of the message is as follows:

[illegible]

Reserved field is not used yet but might be for future need. It currently serves padding requirements. It should be set to null on emission and ignored on reception by peers complying with this specification.

AOT0 messages do not carry options. MH Type field in Mobility Header takes value XX when carrying an AOT0 message.

XXX For test purposes, MH Type field should use value 30

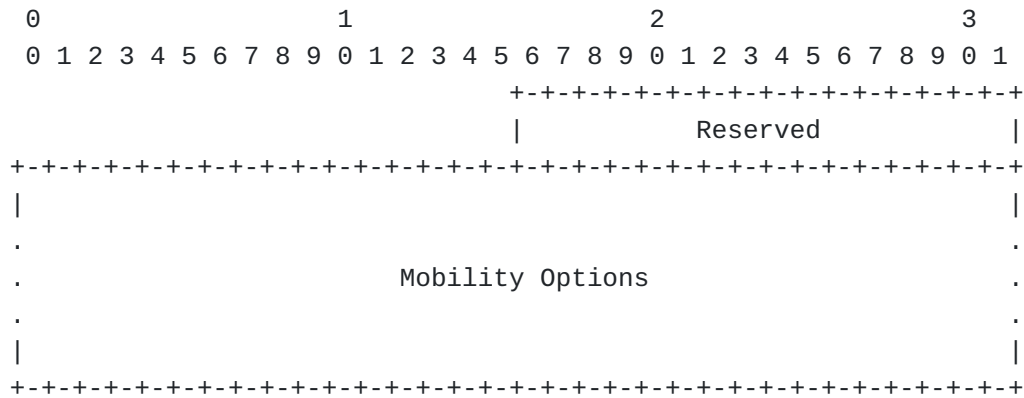
4.4.2. Address Ownership Test Challenge (AOTC)

TOC

The purpose of this message is to provide a nonce to a MN at the address the MN wants to provide proof of ownership for. The ability for the MN to return the nonce to the CN (in an AOTR) provides a live proof of its ability to receive traffic at that address. This message is possibly sent by a CN to a MN in two situations:

- *After receiving a Binding Update message that the CN is willing to accept but for which it does not already have a proof of address ownership for the originating CoA the packet was sent from (correlated with the content of the AltCoA option).
- *After receiving an AOTO from a MN that wants to perform a proof of address ownership for the source address of the packet.

The format of the message is as follows:



The Mobility Options field must always contain a Nonce option. The nonce must be stored locally by the CN for that MN, along with the address being tested. The nonce will be compared with the content of the nonce option found in the AOTR message.

MH Type field in Mobility Header takes value XX when carrying an AOTC message.

XXX For test purposes, MH Type field should be set to 31

4.4.3. Address Ownership Test Response (AOTR)

[TOC](#)

This message is sent by the MN as a result of receiving an AOTC (resulting from an initial action, BU or AOT0). It contains the same nonce, in a Nonce option, the peer had included in its AOTC. The AOTR message is sent from the address to be tested (the on-wire destination address of the AOTC).

When received by the CN, on-wire source address is used to access the stored nonce previously sent in an AOTC message. It is compared with the one in the Nonce option found in the message. On match, the address ownership by the peer is considered proved. It is dropped otherwise. The format of the message is the same as the AOTC message except for MH Type field in Mobility Header which takes value XX when carrying an AOTR message.


XXX For test purposes, MH Type field should be set to 32

4.4.4. Address Ownership Test Status (AOTS)

[TOC](#)

This message is sent by the CN with a status regarding a proof of address ownership. The status can be generic (not associated to an address whose ownership is being proved), for instance if this CN does not allow MN-initiated Address Ownership Tests to occur. It can also be

specific to an ongoing or already performed Test of Address Ownership, for instance to explicitly acknowledge the result of the test.



MH Type field in Mobility Header takes value XX when carrying an AOTS message. Code field provides the status code the message carries. The list of status codes is provided below:

AOTS status codes:

```
*0 AOTO not allowed
```

*1 Proof of Address Ownership is valid

XXX For test purposes, MH Type field should be set to value 33

4.5. Concrete uses of AOT* Messages

TOC

4.5.1. Registration with a CN

TOC

The registration process between a MN and its HA is simple and efficient, being made of a simple BU/BA exchange. This is because the proof of CoA ownership is not required by the HA from the MN.

Like with other route optimization procedures, with IRO, the CN is required to have a proof of CoA ownership available for the MN before accepting a binding and replying with a Binding Ack. More precisely, the proof is needed only before sending traffic to the CoA of the MN but does not impact the reception of traffic from the CoA. This is of particular importance in the rest of the discussion.

Unlike in other more common environments where the proof has to be made at every binding, or "renewed", IRO uses proofs with unlimited lifetimes. This does not mean that once the ownership has been proved to a CN the CoA will indefinitely belong to a MN. The decision is always left to the CN, with the expectation that some sufficient temporary storage will make it capable to keep the binding for a while. This means that if a proof of CoA ownership for a MN is available locally on its CN, no live proof is required and a simple BU/BA exchange is sufficient for the registration to occur. This also means

that inside small or medium communities, where MN move between few locations, the number of potential CoA remains quite low and stable, and can be kept locally on nodes acting as CN.

For instance, without limiting the possible uses, a typical scenario for a daily use includes an address at home (wifi, ...), another on a mobile network (3G, ...) and another at work (wifi, Ethernet, ...).

With IRO, when a MN sends a BU to a CN for registration or reregistration purposes, it starts directing its traffic instantly after the emission of the BU to the address of the CN. Then, the CN will either ask for proof of CoA ownership if it has none available from that MN for that CoA or send a BA to the peer. In all cases, it puts in place the remapping rules for accepting traffic from the CoA (and not the one for emission). That way, there is no disruption of traffic from the MN to the CN.

If the CN replies with an AOTC message sent to the CoA of the MN, the MN replies with an AOTR, proving its complete ownership. The CN then replies with the expected BA and puts in place the required remapping rules for the traffic to flow to the MN at its CoA.

Regarding re-emission, if the MN has no reply from the CN (i.e. no BA or AOTC), common re-emission rules apply. Then, if the CN has sent an AOTC, but receives no reply, it can keep things that way or garbage collect the remapping rule (i.e. remove it after some time). If the MN receives no BA from the CN, it performs re-emission of the AOTR (This implies that the Nonce must be kept locally on the CN even after the emission of the BA).

4.5.2. Early test of CoA ownership

[TOC](#)

There are cases where a MN will be willing to perform early proof of address ownership, allowing it to avoid the delay during movement. In that case, the MN sends an AOTO message to the CN, and receives either an AOTC or an AOTS. If the received message is an AOTS, the exchange is over. If the message is an AOTC, it replies with an AOTR and waits for an AOTS.

MN
Y Z
(3G) (Wifi)

MN (HoA : X)
or CN (Addr: X)

Reg. of Z
w/ proof of
ownership

```
IPv6(src=Z,dst=X)/ESP(BU)
|----->

IPv6(src=X,dst=Z)/ESP(A0TC(Nonce))
<-----|

IPv6(src=Z,dst=X)/ESP(A0TR(Nonce))
|----->

IPv6(src=X,dst=Z)/ESP(BA)
<-----|

...
```

Early proof of
ownership of Y

```
IPv6(src=Y,dst=X)/ESP(A0T0)
|----->

IPv6(src=X,dst=Y)/ESP(A0TC(Nonce))
<-----|

IPv6(src=Y,dst=X)/ESP(A0TR(Nonce))
|----->

IPv6(src=X,dst=Y)/ESP(BA)
<-----|

...
```

Switch to Y
Reg w/o proof

```
IPv6(src=Y,dst=X)/ESP(BU)
|----->

IPv6(src=X,dst=Y)/ESP(BA)
```

A possible use of that early test of CoA ownership is by multi-homed nodes that already have a list of possible CoAs they will switch to if they lose their primary connectivity mean. Note that:

- *this is only a possible optimization allowed by the AOT* framework introduced in this document, not a requirement.

- *it still requires the MN to be registered with the CN.

- *it requires the MN to exchange AOT* messages using the address whose ownership is to be proved.

4.5.3. Test of HoA ownership

[TOC](#)

IRO does not mandate regular proofs of HoA ownership, for the reasons covered in [Appendix C \(Arguments for no regular check of HoA ownership\)](#). For those who have the need, and can afford to lose the associated bits on a regular basis, the AOT* messages can be used for that purpose. The following helps supporting that.

If a CN wants to get a live proof of HoA ownership from a MN, it simply emits an AOTC message (with a fresh Nonce option) to the HoA of the MN for which it has already accepted a registration. The MN MUST reply with an AOTR message containing the received Nonce option. The exchange occurs using the HoA as on-wire destination and source address respectively. This implies that the packets are tunneled through MN's HA. In the MN-MN case, this mainly results in packets never following a direct path.

Note that this specification does not define the action taken by a CN if it does not receive AOTR messages as response to its AOTC messages sent to the HoA.

5. Remapping rules

[TOC](#)

This section covers the heart of IRO processing, the remapping rules that are applied to incoming and outgoing IPsec protected traffic.

[TOC](#)

5.1. Overview

As introduced in the beginning of this memo, in order to remove RH2 and HAO from traffic exchanged between MIPv6 peers, IRO defines the concept of remapping rules.

A remapping rule is associated with a specific type of outgoing or incoming (IPsec-protected) traffic (typically the traffic associated with a given SP). A remapping rule modifies either the source or destination address of matching traffic. IRO remapping rules only apply to IPsec-protected traffic. Unprotected traffic is not modified by IRO remapping process.

5.1.1. Remapping rules for outgoing traffic

[TOC](#)

An IRO remapping rule for outgoing traffic is used to change either the on-wire source or destination address of IPsec-protected emitted packets. For a matching outgoing packet, this remapping process is performed after the processing by the IPsec stack:

- *If such a rule remaps the destination address of a packet, it is used to turn the HoA of the remote peer into its CoA, so that the traffic be routed to its current location.

- *If such a rule remaps the source address of a packet, it is used to turn the MN's HoA into its current CoA, so that the source address of the packet is compatible with the network the MN is currently in.

The address to be used (CoA) for the remapping may be found in the rule or may be provided via an ancillary path with the packet. This second approach is described in [Section 5.1.3 \(On-wire addresses access from userland\)](#) below.

5.1.2. Remapping rules for incoming traffic

[TOC](#)

An IRO remapping rule for incoming traffic is used to change either the on-wire source or destination address of IPsec-protected received packets. For each incoming packet, this remapping process is performed before the processing by the IPsec stack::

- *If such a rule remaps the destination address of a packet, it is used to turn the CoA of the MN into its HoA, so that the IPsec

stack gets the right address referenced in the SA associated with the packet.

*If such a rule remaps the source address of a packet, it is used to turn the peer's CoA into its current HoA, so that the IPsec gets the right address referenced in the SA associated with the packet.

The address to be used (CoA) for the remapping may be found in the rule, in which case it is checked against the on-wire address found in the packet before the remapping is performed. If the address found in the rule is the unspecified address (::), then the on-wire address found in the packet is not verified. Additionally, as described in [Section 5.1.3 \(On-wire addresses access from userland\)](#) below, the on-wire address found in the packet may be passed to the userland if it requests it (e.g. using socket options): this may be used for additional userland checks.

5.1.3. On-wire addresses access from userland

[TOC](#)

With IRO, there is a need for the MIPv6 processing engine to both pass and get on-wire source and destination addresses of received and emitted IPsec protected MH packets. This need is mainly associated with the proof of address ownership and binding exchanges. The need is similar to the one associated with the ability to set and get HAO/RH2 for a common MIPv6 process. Instead of having explicit information in the packet, an ancillary path is required.

This requirement is limited only to MH traffic in general and some specific MH types in particular.

For incoming IPsec protected MH packets, this means that during the handling by remapping rules, the remapped on-wire address must be kept in the local packet structure as an ancillary data that the MIPv6 process will be able to access.

For outgoing MH packets, this means that the addresses MUST be made available as ancillary data in the local packet structures by the MIPv6 process and then be used, if available, by the remapping rules.

For all incoming IPsec packets associated with a coarse or fine grained SA for MH traffic, if a remapping rule is applied to the traffic, the on-wire source and destination addresses MUST be made available as ancillary data to the userland process that will process the packet (i.e. at socket level). In all cases (remapping rule being applied or not), if an on-wire source or destination address is not changed, the associated ancillary data MUST contain the unspecified address (::).

[TOC](#)

5.2. Details of traffic processing

5.2.1. Non-MH traffic (data traffic)

[TOC](#)

Data traffic exchanged between MN and CN using IRO has simple requirements in term of remapping. We consider here only IPsec packets that are not associated with a transport mode IPsec SA protecting MH traffic.

5.2.1.1. Incoming traffic

[TOC](#)

When an incoming IPsec packet is handled by IRO, as the last step before being processed by the IPsec module (or at the beginning of the processing by the IPsec module), source and destination addresses remapping rules for the traffic are applied (at most one for each). Conceptually, the SPI found in the packet can be considered as the main key to access remapping rules but using it directly in practice has associated side effects (e.g. tracking and update upon changes). For that reason, the specific way to bind IRO remapping rules to incoming IPsec-protected traffic is left implementation dependent. Interested reader can consult [Section 7 \(Implementation Notes\)](#) ("Implementation notes") for possible solutions.

Each rule has an expected on-wire address, which is checked against the on-wire one found in the packet. If it matches, the remapping occurs. Note that the remapping rules for source and destination addresses are applied in an independent fashion.

5.2.1.2. Outgoing traffic

[TOC](#)

When an outgoing IPsec protected packet is handled by IRO (just after or at the end of the processing by the IPsec stack), source and destination addresses remapping rules for that traffic (at most one for each) are applied. Conceptually, the SPI found in the packet can be considered as the main key to access remapping rules but (as discussed previously for incoming traffic) using it in practice has associated side effects. For that reason, the specific way to bind IRO remapping rules to outgoing IPsec-protected traffic is left implementation dependent. Interested reader can consult [Section 7 \(Implementation Notes\)](#) ("Implementation notes") for possible solutions.

The expected address is checked against the one found in the IPv6 header of the packet. If it matches, the remapping occurs. Note that the remapping rules for source and destination addresses are applied in an independent fashion.

5.2.2. MH traffic

[TOC](#)

MH traffic emitted and received by a MIPv6 entity using IRO has specific additional requirements compared to common data traffic exchanged between those MIPv6 entities.

Basically, the checks and settings on source and destination addresses are relaxed to allow IPsec-protected traffic sent from a new non-registered CoA to pass through (nonetheless, the associated SA protecting traffic are already in place). In the MIPv6 stack of the CN, checks are done using an ancillary path that allows the on-wire address to be passed for verification.

Here, we only consider IPsec-protected traffic associated with transport mode SAs whose selectors provide protection of MH traffic. Granularity considerations are covered below.

The search for remapping rules is done in the same fashion as previously described for data traffic. Only checks and application of the rules are changed as described below.

5.2.2.1. Incoming traffic

[TOC](#)

If a remapping rule is found for source address, which contains the unspecified address as check, the remapping is performed without checking the source address of the packet. The unspecified address is used as a wildcard.

In source rule case, the on-wire address found in the packet is stored as an ancillary data for further processing and decision by the MIPv6 stack (commonly in userland).

Note that this specification does not explicitly mandate when the unspecified address should be used in the source remapping rule, and leave that to implementors, as it is highly dependent of following facts:

- *if the system does not support fine-grained SP/SA or simply does not use them for MH traffic with a peer, then the use of the unspecified address will be required.

- *if fine-grained SA are used, the MIPv6 stack will use the unspecified address if the traffic received protected by that SA

can't be reliably mapped to a specific CoA for the peer, i.e. if it is expected and authorized that the peer sends traffic from another [possibly to be registered] CoA. This is the case for BU, AOTO, AOTR traffic for instance.

*if the MIPv6 stack supports extensions to [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#)-defined MH messages, the use of IRO will still remain possible with those extensions.

Note that this decision is not expected to create interoperability issues, as the use of the unspecified address is based on non-ambiguous criteria defined in the documents specifying the purpose of MH traffic. Also note that the use of the unspecified address for checks and the passing of the on-wire address to the MIPv6 stack for further processing is equivalent from a security standpoint to the decision that occurs in common MIPv6 processing of HAO extension.

5.2.2.2. Outgoing traffic

[TOC](#)

Emission of IPsec-protected MH traffic by an IRO-enabled node is conceptually similar to the steps described in [Section 5.2.1.2 \(Outgoing traffic\)](#) regarding the application of the remapping rules for data traffic.

Nonetheless, unlike remapping rules applying to data traffic (for which the address used for the remapping is expected to be available statically in the rule), it is possible that a userland MIPv6 implementation uses an ancillary path to provide the address used for the remapping. This may for instance provide more control and granularity. This is partly discussed in [Section 7 \(Implementation Notes\)](#) ("Implementation notes").

5.2.3. Related traffic (ICMPv6 error messages ...)

[TOC](#)

When an ICMPv6 Error message (.e.g. for PMTUD) is received by a node implementing IRO, carrying some IPsec traffic in its citation, the IPsec stack processes the message as usual (probably based on what is described in section 6 to 8 of [\[RFC4301\] \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#)). Obviously, because the citation in the ICMPv6 Error message may contain remapped addresses, an additional step is needed before (or at the beginning of) the processing by the IPsec stack, in order to check if a the reverse remapping of the addresses in the packet is required. Conceptually, this additional step is similar to the processing

performed on citations containing RH2 and HAO elements, except that the real source and destination addresses are not explicitly available in the packet, but need to be provided by the IPsec stack, using the SPI available in the citation.

6. Extending advantages of IRO to the HA

[TOC](#)

6.1. Rationale and expected advantages

[TOC](#)

IRO's primary purpose is to improve security and efficiency of MIPv6 communications in IPsec environments. Because most of them are expected to occur directly between peers, IRO is oriented towards MN-CN and MN-MN flows.

But the flows between a MN and its HA can also benefit from the improvements: using the SPI information available on both sides to perform the remapping of incoming and outgoing IPsec traffic, the need of RH2 and HAO extensions between the MN and its HA simply disappears. This provides anonymity (See [Appendix J \(Anonymity\)](#)) of the MN on a foreign link by hiding its HoA to eavesdropper on the path (if IKE does not leak that information). It also makes the MN fully capable in networks where only IPsec is allowed to flow (500/udp is required for the initial negotiation of SA and infrequently for rekeying).

6.2. Changes to HA processing

[TOC](#)

IRO does not mandate a detection mechanism ([Appendix I \(Rationale for not specifying a new BU\)](#)) and transparently reuses most of [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#)-defined messages. For that reason, MN and HA must be explicitly configured to use IRO.

The changes to HA processing for the peer, required by the use of IRO, are simply the use of remapping rules instead of HAO and RH2 extensions.

With IRO, the relationship between a MN and one of its CN is basically the same as the relationship between the MN and its HA, with the following simple differences:

HA and MN never exchange AOT messages as the MN is always trusted regarding the CoA information it provides in its BU.

*HA and MN have a larger set of possible messages they can exchange. Remapping rules should be able to handle that.

*Chances are high that an IPsec tunnel be used for protecting traffic relayed through the HA. Remapping rules do not interfere with that as the associated SA is not tracked. This is due to the fact that the SA references the CoA of the MN as the source (on-wire) address of the communication and not the HoA.

6.3. Changes to MN processing

[TOC](#)

The changes to MN processing for IRO to be used with its HA are quite comparable to the one previously described for the MN, i.e. they are naturally deduced from the basic requirement that RH2 and HAO must be replaced by the use of remapping rules.

7. Implementation Notes

[TOC](#)

The content of this section is not meant to be normative but only informative.

This section provides some explicit feedback associated with the implementation of IRO on Linux (Linux kernel, UMIP mobility daemon and racoon IKE daemon). Based on the specific targeted system, some of the points discussed in this section may be completely irrelevant.

7.1. Nested SA

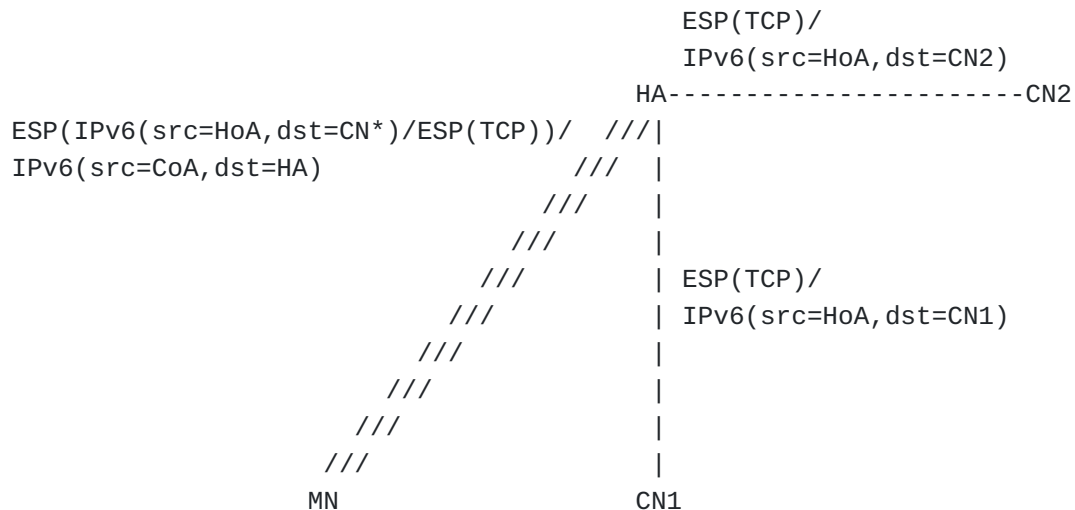
[TOC](#)

7.1.1. Problem

[TOC](#)

The main purpose of IRO is to route optimize IPsec traffic exchanged between the MN (from its HoA) and its CN. Before that optimization takes places (i.e. before the AOT* exchanges), that traffic is expected to follow the natural path, i.e. be routed via the tunnel to the HA. In IPsec environments, chances are high the data path to the HA will be IPsec-protected, using IPsec in tunnel mode as (optionally) expected by MIPv6 specifications. In that case, before the optimization is effective, traffic sent by the MN to the remote IPsec peer will have to

undergo both the protection of the specific SA protecting the end-to-end traffic with the peer and the tunnel one protecting the data traffic to the HA. This is required (at least for topological reasons) because the HoA is only valid as an inner address for tunneled traffic. This is depicted below with a MN having E2E IPsec-protected communications with two peers (CN1 and CN2) via its HA. The layout of traffic is given for packets sent by the MN; the format of traffic sent by the (the format of return traffic). :



```
/// IPsec in tunnel mode
--- IPsec in transport mode
```

This is basically the kind of setup described in Appendix E of [\[RFC4301\] \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#). Supporting that kind of nested IPsec configuration, even temporarily before the route optimization is in place, is not straightforward. Obviously, this usually does not require anything specific on the HA or the CN, but the MN case may be trickier. In the specific context of IRO, there may be a need for both previously described IPsec SP/SA to exist separately. The main reason is that, even if the traffic to the CN initially needs to undergo both SP (and in the end protection provided by associated SAs), this need disappears when the optimization is in place. At that moment, the traffic only undergo the end-to-end SP and the protection of associated SA. This is depicted below with the setup of IPsec Route Optimization between the MN and CN1:

```

                                ESP(TCP)/
                                IPv6(src=HoA,dst=CN2)
                                HA-----CN2
ESP(IPv6(src=HoA,dst=CN2)/ESP(TCP))/  ///
IPv6(src=CoA,dst=HA)                    ///
                                     ///
                                     ///
                                     ///
                                     ///
                                     ///
                                     ///
                                     ///
                                     ///
MN-----CN1
    ESP(TCP)/
    IPv6(src=CoA,dst=CN)

/// IPsec in tunnel mode
--- IPsec in transport mode

```

7.1.2. Selected solution

[TOC](#)

XXX FIXME

7.2. Having IKE traffic flow via the IPsec tunnel to the HA

[TOC](#)

7.2.1. Problem

[TOC](#)

Traffic generated by an IKE daemon needs to bypass the system's security policies in order to avoid chicken and eggs issues. Unix IKE daemons implementation usually achieve that using a specific IPsec bypass `setsockopt()` call on their sockets.

In common situations, previous method works just fine. But when an IPsec data tunnel already exists on the system as it is usually the case for corporate VPN clients, this method prevents the use of the inner tunnel address for IKE negotiation with remote peers accessible only via the remote security gateway (e.g. hosts in the corporate

network). Stated differently, this prevents the setup of end-to-end IPsec protection via the IPsec tunnel.

More precisely, if the IKE daemon tries and use the inner address for negotiation with a peer, the IPsec bypass `setsockopt()` call on associated socket prevents associated IKE traffic to flow correctly. This is because, from a topological standpoint, the only valid path for traffic originating from that address is via the IPsec tunnel.

When MIPv6 data traffic between the MN and its HA is required by policy to undergo IPsec tunnel protection, the same limitation as described above exists. For the specific needs of IRO, this is an issue because the HoA is used for the IKE negotiation with the CN (whose side effect is also to prove HoA ownership to the peer).

7.2.2. Selected solution

[TOC](#)

For the sake of the discussion, we consider the existence of some wide IPsec SP requiring protection of all traffic from the HoA to a given peer behind the HA. To make things clear, IKE traffic (500/udp) between the HoA and the address of the peer matches this SP's selectors. Obviously, the simple removal of the IPsec bypass `setsockopt()` on the socket associated with the HoA is not sufficient to make things work. This would have initially been sufficient to have associated IKE traffic undergo the IPsec tunnel mode SP (protecting data traffic between the MN and its HA). But the existence of the additional SP creates a chicken and eggs situation and prevents things to work that easily.

But once the IPsec bypass `setsockopt()` call is removed for the IKE socket associated with the HoA, associated traffic basically undergo the system security policies. Then, the addition of a high priority SP with selectors specifically suited for IKE traffic from the HoA to the address of the peer is sufficient to have the IKE traffic be IPsec tunneled using the existing SA already protecting MIPv6 data traffic. From an implementation perspective, the removal of the IPsec bypass `setsockopt()` call has been implemented by adding a simple option ('no_bypass') to the racoon IKE daemon allowing the user to specify an address for which associated socket should not undergo the `setsockopt()` call.

With regard to the addition of the high priority IPsec security policy matching IKE traffic between the HoA and the address of the peer, it was easier to have that job done inside UMIP. This is basically because UMIP is already the one handling the installation of other security policies for MIPv6 and data traffic. Another reason is that UMIP will need to update the tunnel mode SP after a handover (via [\[MIGRATE\]](#) (Ebalard, A. and S. Decugis, "PF_KEY Extension as an Interface between Mobile IPv6 and IPsec/IKE," August 2008.)).

7.3. Remapping rules and old IPsec architecture

[TOC](#)

7.3.1. Problem

[TOC](#)

In the new IPsec architecture [\[RFC4301\] \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#), "for an SA used to carry unicast traffic, the Security Parameters Index (SPI) by itself suffices to specify an SA". Section 4.1 of [\[RFC4301\] \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#) provides additional guidance on the topic.

In the old IPsec architecture [\[RFC2401\] \(Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol," November 1998.\)](#), a SA "is uniquely identified by a triple consisting of a Security Parameter Index (SPI), an IP Destination Address and a security protocol (AH or ESP) identifier".

If an IPsec stack only supports the behavior mandated by the old IPsec architecture, SAD lookup on inbound packets require the use of both the SPI and the destination address of the SA.

For inbound IPsec traffic, IRO remapping rules may exist on the MN to remap the destination address (CoA) into the HoA. In that case, by design, the address found in the destination address field of the packet (CoA) does not match the one in the SA (HoA).

If an IRO implementation relies on the information maintained in the SADB to access the HoA, it needs to be able to perform a SA lookup using only the SPI. This is possible if the IPsec stack is compliant with the new IPsec architecture but additional changes are required if it is compliant with the old IPsec architecture.

7.3.2. Selected solution

[TOC](#)

Linux IPsec stack currently follows the model defined in the old IPsec Architecture; it makes use of the destination address, the SPI and the IPsec protocol for lookups of SA in the SADB upon reception of an IPsec-protected packet. Those parameters are hashed together for efficient access to the associated SA via a hashtable.

To support IRO needs, this behaviour has been changed in order for the IPsec stack to perform lookup of SA for inbound packets without considering the destination address, i.e. using mainly the SPI.

For traffic matching an IRO remapping rule mandating the remapping of either source or destination address, the one available in the found SA

(after the lookup by SPI) is used after a check that received on-wire address matches the expected address in the remapping rule. For common inbound IPsec packets not undergoing any IRO processing, mismatch between the destination address in the packet and the one in the SA simply results in the packet being dropped. This simple modification has allowed the implementation of IRO remapping based on address information stored in the IPsec stack (HoA to be used for remapping), even for inbound IPsec traffic for which the on-wire destination address has been set to the CoA by the emitter (remote CN).

7.4. Userland and address remapping

[TOC](#)

7.4.1. Problem

[TOC](#)

[\[RFC3542\]](#) (Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface (API) for IPv6," [ay 2003.](#)) specifies an Advanced Socket API for IPv6, which allows (among other things) applications running on systems supporting this interface to:

- *control the inclusion of RH2 or HAO in outgoing packets (using IPV6_RTHDR and IPV6_DSTOPTS socket options).

- *access the content of the RH2 or HAO in incoming packets (using IPV6_RECVRTHDR and IPV6_RECVDSTOPTS socket options).

IRO remapping process is expected to be done in kernel space, from information (states or ancillary data) passed by userland. For that purpose, a similar interface is between kernel and userland, even if none is currently available at the time of writing on most systems.

7.4.2. Selected solution

[TOC](#)

Linux supports the Advanced Socket API for IPv6 [\[RFC3542\]](#) (Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface (API) for IPv6," [ay 2003.](#)). It is used by UMIP for previously described tasks. To support IRO needs, a similar set of

socket options have been defined. The straightforward mapping is the following:

- *IPV6_RECVIROSRC: once enabled on a raw socket, the on-wire source address found in the incoming packet (before IRO and IPsec processing) is passed as ancillary data. This is the IRO equivalent of [\[RFC3542\] \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface \(API\) for IPv6," ay 2003.\)](#)-defined IPV6_RECVDSTOPTS.
- *IPV6_IROSRC: allows passing as ancillary data of the address to be used by IRO as on-wire source address after IPsec processing. This is the IRO equivalent of [\[RFC3542\] \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface \(API\) for IPv6," ay 2003.\)](#)-defined IPV6_DSTOPTS.
- *IPV6_RECVIRODST: once enabled on a raw socket, the on-wire destination address found in the incoming packet (before IRO and IPsec and processing) is passed as ancillary data. This is the IRO equivalent of [\[RFC3542\] \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface \(API\) for IPv6," ay 2003.\)](#)-defined IPV6_RECVRTHDR.
- *IPV6_IRODST: allows passing as ancillary data of the address to be used by IRO as on-wire destination address after IPsec processing. This is the IRO equivalent of [\[RFC3542\] \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface \(API\) for IPv6," ay 2003.\)](#)-defined IPV6_RTHDR.

Unlike [\[RFC3542\] \(Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, "Advanced Sockets Application Program Interface \(API\) for IPv6," ay 2003.\)](#)-defined socket options for which specific structures are defined (to support passing of IPv6 extensions headers), IRO ones only require the passing of IPv6 addresses when needed. The details of the interface and its implementation are not provided here for brevity but they are freely available.

8. Security Considerations

[TOC](#)

[TOC](#)

8.1. Proof of address ownership

8.1.1. Position of the problem

[TOC](#)

As a CN, registering a binding between a CoA and a HoA is not something harmless. This can be seen as a modification of local routing table, like an order from a peer to direct traffic to a specific address. For that reason, the CN needs some proof regarding this binding. In MIPv6, RRP has been designed with the hypothesis that there is no initial trust relationship between a MN and its CN. The solution to provide confidence to the CN in the HoA and CoA binding has consisted in showing the ability for the MN to send _and_ receive traffic both at the HoA and CoA.

With IRO, there is an initial trust relationship between a MN and the CN it will contact. This is expected to take the form of cryptographic credentials (X.509 certificates, ...) that will allow an IKE negotiation to occur to setup SAs to protect the binding. Static keying case is covered in [Appendix G \(Compatibility with static keying\)](#). Those SA only references the HoA of the MN and not at all its CoA.

The point here is that the existence of SAs does not directly provide to the CN any _live_ proof of address ownership as it occurs with RRP. Furthermore, as summarized in section 6.2 of [\[RFC4866\] \(Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6," May 2007.\)](#) paragraph 4, the trust relationship between a HA and its MN is very different from the one between a MN and a CN, even if both use IPsec/IKE to authenticate.

8.1.2. Home Address ownership

[TOC](#)

The proof of HoA ownership to the CN is one of the reason behind the design decision to have MN and CN perform the IKE negotiation via the tunnel to the HA (i.e. using the HoA). That way, the existence of the SAs gets bound to a successful initial exchange between the CN and the MN. This proves to the CN the ability for the MN to send/receive traffic from/at that address.

Nonetheless, as IKE basically allows negotiation to be performed from a different address than the one the SA contain ([MIGRATE] has such an use), this behavior MUST be prevented on the CN for the purpose of negotiations of the initial SAs that will protect MH traffic for IRO's binding between the MN and the CN.

While MN and CN are capable of performing an IKE exchange between them using a set of credentials, there are many possible reasons for which those credentials might in fact be invalid at the time the negotiation

occurs. This might for instance be the case if the CN has not up-to-date revocation information. This can also result from the use by the MN of a different set of credentials for the purpose of protecting its HA registration and the registration with its CN.

Mandating the IKE negotiation to be routed through the tunnel to the HA provides the proof that the MN is still granted ownership of the address by the network it belongs to at the time of negotiation. It should be noted that the proof of HoA ownership occurs at SA setup time and remains valid till the SA is rekeyed, i.e. each rekeying providing a new live proof. This specification does not mandate regular check of HoA ownership between rekeying. [Appendix C \(Arguments for no regular check of HoA ownership\)](#) provides arguments on that topic.

The case of static keying is covered in [Appendix G \(Compatibility with static keying\)](#).

8.1.3. Care-of Address ownership

[TOC](#)

The proof of CoA ownership by the CN is an especially important point in the security of large scale deployments of IRO. As stated in the introduction of this section, the acceptance of a binding by a CN for a CoA is a modification of local routing table to send current and future traffic to that address when it is destined to the HoA.

The proof by the MN that it is able to both send and receive traffic at this address is a primary concern in the security of the protocol.

[Appendix A \(Ability to send does not prove CoA ownership\)](#) covers the reasons why the only ability to send is an insufficient proof of CoA ownership, even in the context of IRO.

8.2. Remapping (comparison with explicit HAO/RH2 inclusion)

[TOC](#)

For every remapping, the practical impact is the same as the explicit one resulting from the inclusion of a RH2 or HAO.

8.3. Anonymity

[TOC](#)

At the moment, this section is empty. See [Appendix J \(Anonymity\)](#).

[TOC](#)

8.4. Limiting attack surface

IRO can provide the ability to have port 500/udp open for remote negotiations on the HoA for the purpose of the inbound contacts and not on the CoA. CoA is only used for the discussion between the MN and its HoA, which allow to put some specific firewalling rules in place for that purpose.

9. IANA Considerations

[TOC](#)

The values for following mobility header messages MUST be assigned by IANA:

*Address Ownership Test Offer message (AOTO, see [Section 4.4.1 \(Address Ownership Test Offer \(AOTO\)\)](#))

*Address Ownership Test Challenge message (AOTC, see [Section 4.4.2 \(Address Ownership Test Challenge \(AOTC\)\)](#))

*Address Ownership Test Response message (AOTR, see [Section 4.4.3 \(Address Ownership Test Response \(AOTR\)\)](#))

*Address Ownership Test Status message (AOTS, see [Section 4.4.4 \(Address Ownership Test Status \(AOTS\)\)](#))

The values for following mobility option MUST be assigned by IANA:

*Nonce Option (see [Section 4.3.1 \(Nonce option\)](#))

10. Acknowledgements

[TOC](#)

The author acknowledge the comments and correction of Guillaume Valadon on the initial version of the document.
This document was generated by xml2rfc.

11. References

[TOC](#)

11.1. Normative References

[TOC](#)

[RFC2119]	Bradner, S., " Key Words for Use in RFCs to Indicate Requirement Levels ," RFC 2119, March 1997 (TXT).
[RFC2367]	McDonald, D., Metz, C., and B. Phan, " PF_KEY Key Management API, Version 2 ," RFC 2367, July 1998 (TXT).
[RFC2401]	Kent, S. and R. Atkinson, " Security Architecture for the Internet Protocol ," RFC 2401, November 1998 (TXT).
[RFC2409]	Harkins, D. and D. Carrel, " The Internet Key Exchange (IKE) ," RFC 2409, November 1998 (TXT).
[RFC3542]	Stevens, W., Thomas, M., Nordmark, E., and T. Jinmey, " Advanced Sockets Application Program Interface (API) for IPv6 ," RFC 3542, ay 2003 (TXT).
[RFC3775]	Johnson, D., Perkins, C., and J. Arkko, " Mobility Support in IPv6 ," RFC 3775, June 2004 (TXT).
[RFC3776]	Arkko, J., Devarapalli, V., and F. Dupont, " Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents ," RFC 3776, June 2004 (TXT).
[RFC4301]	Kent, S. and K. Seo, " Security Architecture for the Internet Protocol ," RFC 4301, December 2005 (TXT).
[RFC4303]	Kent, S., " IP Encapsulating Security Payload (ESP) ," RFC 4303, December 2005 (TXT).
[RFC4306]	Kaufman, C., " Internet Key Exchange (IKEv2) Protocol ," RFC 4306, December 2005 (TXT).
[RFC4835]	Manral, V., " Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH) ," RFC 4835, April 2007 (TXT).
[RFC4877]	Devarapalli, V., " Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture ," RFC 4877, April 2007 (TXT).

11.2. Informative References

[TOC](#)

[CNIPsec]	Dupont, F. and JM. Combes, " Using IPsec between Mobile and Correspondent IPv6 Nodes ," draft-ietf-mip6-cn-ipsec-08 (work in progress), August 2008 (TXT).
[MIGRATE]	Ebalard, A. and S. Decugis, " PF_KEY Extension as an Interface between Mobile IPv6 and IPsec/IKE ," draft-ebalard-mext-pfkey-enhanced-migrate-00 (work in progress), August 2008 (TXT).
[RFC4225]	Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, " Mobile IP Version 6 Route Optimization Security Design Background ," RFC 4225, December 2005 (TXT).
[RFC4651]	

	Vogt, C. and J. Arkko, " A Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization ," RFC 4651, February 2007 (TXT).
[RFC4866]	Arkko, J., Vogt, C., and W. Haddad, " Enhanced Route Optimization for Mobile IPv6 ," RFC 4866, May 2007 (TXT).
[RFC4882]	Koodli, R., " IP Address Location Privacy and Mobile IPv6: Problem Statement ," RFC 4882, May 2007 (TXT).

Appendix A. Ability to send does not prove CoA ownership

[TOC](#)

With IRO, negotiation of the protection (via IKE) for the registration traffic between the peers is done using the tunnel to the Home Agent (i.e. the HoA). Then, the protected Binding Update message is emitted by the MN. It locally uses the HoA but a remapping process makes the CoA the address appearing on the wire for this packet. When the CN receives that packet, the address is remapped to the HoA of the MN by the MIPv6 stack. The remapped address appearing as source of the packet is passed as ancillary data to the MIPv6 process where a check will be performed during parsing against the content of the required Alternate Care-of Address option.

This process proves the CN that the MN is able to `_send_` traffic using the CoA.

Then, IRO requires additional steps for the MN to prove its ability to receive traffic at that address. This appendix covers the threats prevented by the addition of this proof of reception capability in IRO. The trust model between the MN and its HA is based on the existence of the IPsec protection between the two. The only requirement for the update of the tunnel endpoint when a movement occur is the reception by the HA of a protected Binding Update message containing the new CoA in the Alternate CoA option. The use of the new CoA as source of the packet is not even mandatory.

One would argue that the same kind of trust relationship exists between the MN and its CN as they already have an established trust relationship, materialized by the pair of SA protecting MH traffic.

Nonetheless there are many difference between both situations.

The first and main one is that all the traffic emitted by the HA to the CoA provided by the MN has a traceable source: the address of the HA, which belongs to the home network. It allows to track the source of the traffic emitted to the CoA back to the Home Network. In the context of the flow from the CN to the MN, the source address might possibly be that of a foreign network (if the CN is also acting as a MN) and the destination is the one that would be provided by the MN.

Now consider the following, still in the context of a proof of address ownership based solely on the ability of the MN to emit traffic from the CoA. A MN has been compromised by an attacker that has the ability to emit traffic with the address of a target (no ingress-filtering by

its provider). The attacker would now be able to mount connections with the HA and then, using IRO, with all the peers that trust the MN. At the moment, it still uses some valid address where it can emit/receive traffic from/to. After having some traffic intensive connection running with a peer, it simply warns the peer of a change of CoA by advertising the address of the target. As the CN does not require a proof of reception capability, all the IPsec traffic gets redirected to the target. This might not be a problem with a single peer and some connected protocol but it is expected that the protocol be used in vast trust domains where the number of peers is not directly limited. In the end, requiring that the proof of CoA ownership includes a proof of reception capability by the MN at the CoA prevents that compromise of a MN by an attacker provides her with a potentially unlimited number of anonymous and unwilling "bots" to DoS a target other than herself. In the design of IRO, to maintain the efficiency of the protocol in term of latency associated with movement, the proof of reception capability is not required to occur before the CN can emit traffic to the CoA.

Appendix B. IKE exchanges use the HoA and the tunnel to the HA

[TOC](#)

Remainder: in this appendix, we still consider that the data tunnel between the HA and the MN is IPsec protected. Some security arguments in this appendix should be modulated if this hypothesis is known to be invalid.

We provide here some arguments regarding the use of the HoA for performing the IKE exchanges with the peers, using the tunnel through the HA.

The first simple rule which always applies with IRO is that no connection happens directly if it is not IPsec-protected. No difference is made for IKE exchanges even if those flows have intrinsic protection mechanisms.

The need for performing IRO to get direct routing between the peers is motivated by the net performance impact in terms of bandwidth, delay and jitter by avoiding triangular routing and the bottleneck of HA. This is of interest for specific flows like VoIP, direct file exchanges, ... but are mainly useless for infrequent flows like IKE negotiations. When a MN performs IKE negotiation with a peer, having IKE_SA (or ISAKMP SA for IKEv1) set up is only a matter of few packets (IKEv1 Main mode exchange uses six). Then, all next CHILD_SA (or IPsec SA for IKEv1) will reuse the same IKE_SA and generally complete after a three packet exchange. As rekeying is supposed to occur extremely infrequently and does not need the advantage of direct routing, this is unneeded. The argument regarding the loss associated with the routing through the tunnel gets the same answer: the impact is very limited

given the amount of traffic. Furthermore, when certificates are used, IKE packets already get fragmented even with a full 1500 bytes PMTU. In fact, the advantage of using the HoA and the IPsec tunnel to the HA for performing IKE negotiation with peers is the stability guaranteed by the migration process when movement occurs. MIPv6 simply makes things transparent for all IKE daemon connections from the HoA. To conclude and after all previous functional arguments, there are also some security advantages in performing IKE negotiations with peers using the protected IPsec tunnel to the HA. The most important one is anonymity. A positive side effect of having the negotiation performed through the IPsec tunnel to the HA (ESP with meaningful encryption is assumed) is that it hides everything to people in MN's foreign network. IKE traffic is only accessible on the path between the HA and the peer. In fact, in the MN-MN situation eavesdroppers on both foreign networks are unable to get the HoA of the peer on the other network. It requires being on the path between the two HA. The same is also true for the identity information that might appear during the IKE negotiation depending on the modes peers use. Another security advantage with that policy is that a peer is able to statefully filter 500/udp traffic received on its CoA and allow only outbound initiated connections addressed to the HoA. This policy simply allows reducing the network attack surface of the node in the foreign network.

Appendix C. Arguments for no regular check of HoA ownership

[TOC](#)

As presented in [Section 8.1.2 \(Home Address ownership\)](#), when dynamic keying is used, the initial IKE negotiation protecting the registration traffic between the MN and the CN provides to the CN the proof of HoA ownership by the MN. This proof remains valid till this SA is rekeyed. This is also true for further SA negotiated between the MN and the CN. The initial proof of HoA ownership is easily obtained as it results from a positive information: packets exchanged with the MN at this address. Note that the inability for the MN or the CN to get traffic routed to the HA at that moment results in the inability to get direct connectivity as the IKE negotiation cannot be performed. In the same fashion, after the initial proof, there is no defined way to track a loss of HoA ownership through a positive event: the CN is simply not warned that the MN has been removed ownership of its HoA by its home network (resulting from a compromise, change of network prefix, ...). Discovery of a loss of HoA ownership cannot be tracked by a negative event either, such as the inability to exchange traffic with the MN at a specific moment. In fact, a crash of the HA, the loss of connectivity between the MN and its HA, or between the CN and the HA are to be expected. In that context, such a mechanism would simply amplify the existence of points of failure or allow DoS to occur. Avoiding that

provides resilience and allows direct communications to survive previous failure conditions related to the HA.

Another reason to prevent regular proof of HoA ownership is the use of the HoA in IRO. It acts as a local identifier on both peers. It allows the MN to acquire movement independence and can be seen as a convenience in the relationship between the peers, to find themselves initially, no matter where they are located. With IRO, the HoA never appears anymore in packet exchanged directly between the peers (due to removal of HAO and RH2). It is only understood locally in the context of ongoing IPsec communications between the peers.

The last argument for not including this requirement (capability is provided, see [Section 4.5.3 \(Test of HoA ownership\)](#)) in the protocol is that different CN or MN might have different more efficient methods for performing that tracking. For instance, inside a home network, instead of using a constant regular polling by all MNs, an administrator revoking the credentials of a MN will easily be able to request all MNs to update their revocation information, before shutting down communications with associated MN (i.e. replacing polling by push). In the context of IRO, no mechanism to perform regular checks of HoA ownership is included. This capability is outside the scope of this specification.

Appendix D. Lack of encryption between MN and HA

[TOC](#)

In this specification, the use of IPsec tunnel protection of data traffic is expected. Note that section 5.5 of [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) only specifies that:

For traffic tunneled via the home agent, additional IPsec ESP encapsulation MAY be supported and used. If multicast group membership control protocols or stateful address autoconfiguration protocols are supported, payload data protection MUST be supported.

The logic behind previous expectation is associated with the availability of credentials between the MN and HA, and also the kind of environment in which it will get deployed.

However, the lack of IPsec protection of tunneled data does not prevent IRO; it only removes some security advantages of this protection. This loss is covered in this appendix.

When negotiating IRO, the MN uses the tunnel to its HA for routing IKE negotiation with the peer. As IKE is designed for robustness, the advantage of the privacy when IPsec is used for protecting the data tunnel (i.e. non NULL encryption) is the insurance that the address of the peer or its cryptographic credentials are not disclosed on MN's

network. Note that MN's HoA and associated identity are expected to be disclosed to eavesdroppers during registration of the MN to its HA (if IRO is not extended to HA-MN exchanges to remove HAO and RH2 from BU/BA).

As a conclusion, removing the hypothesis of privacy for data tunneled to the HA removes the anonymity provided to peer's identity (HoA or CoA, and possibly cryptographic identity appearing during IKE exchange).

Appendix E. What if I don't need protection?

[TOC](#)

IRO mandates the use of IPsec for all direct communications between MIPv6 peers. As IPsec is only a framework, the level of protection might vary, along with the additional requirements, environments and capabilities of end devices.

There will certainly be some very specific and limited cases where people will see a need in downgrading the security for performance or other reasons. To be fair, except in some very specific conditions, achieving performance while still keeping security is possible. For instance, if authentication is a real requirement but privacy is not (but it is still activated by default), and CPU limits the throughput, keeping only authentication services of IPsec as provided by ESP with NULL encryption or by AH will clearly boost performance.

Now, if there is a desperate need to suppress security services between MIPv6 peers for some reason, the best thing is to use another route optimization like common R0 as specified in [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#), instead of trying to circumvent them.

For those with some imagination, who still think the author is wrong and think about simply negotiating NULL authentication and NULL encryption, next paragraph might be worth reading.

[\[RFC4835\] \(Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload \(ESP\) and Authentication Header \(AH\)," April 2007.\)](#) defines mandatory-to-implement cryptographic algorithms for use with ESP (and AH). NULL encryption algorithm and NULL authentication algorithm must both be implemented. In section 3.2 of [\[RFC4303\] \(Kent, S., "IP Encapsulating Security Payload \(ESP\)," December 2005.\)](#), some requirements are specified on those algorithms, preventing their simultaneous uses:

Note that although both confidentiality and integrity are optional, at least one of these services MUST be selected, hence both algorithms MUST NOT be simultaneously NULL.

Standard R0 is based on the use of RH2 and HAO to explicitly carry the HoA of the MN, respectively as real destination or final source of the packet sent directly between the nodes:

*From MN to CN: HAO

*From CN to MN: RH2

*From MN to MN: HAO and RH2, simultaneously.

The inclusion of these explicit containers generates a loss of MTU. In common case, where no other specific extensions or options are used (to remove padding considerations), HAO and RH2 each consume 24 bytes. The loss of MTU associated with those MIPv6 extension for direct MN-CN communications is 24 bytes. For direct MN-MN communications, it is 48 bytes. As an initial comparison, unprotected routing via the HA through an IPv6-in-IPv6 tunnel consumes 40 bytes. When an IPsec tunnel is used, the loss of MTU depends on the authentication and encryption algorithm, negotiation of ESN, padding requirement.

As IRO has been designed to provide secure IPsec-protected direct communications between MIPv6 peers, it is difficult (and does not make that much sense) to compare the loss of MTU associated with IRO and the one of standard unprotected R0. In term of header inclusion, IRO neither use RH2 nor HAO but require AH or ESP. Depending on the size of ESP or AH header(s) and the specific type of communication (MN-MN or MN-CN), one route optimization type might consume more bandwidth than the other:

*ESP in transport mode using HMAC-SHA1-96 (required in all implementation) as authentication algorithm, NULL for encryption and no ESN consumes 24 bytes (with 2 bytes of padding). Adding a meaningful encryption algorithm will make that higher.

*AH in transport mode also using HMAC-SHA1-96 and no ESN will give a similar value.

To make things simple, using IRO with some ESP with NULL encryption or with AH for MN-CN communications provides similar MTU loss as standard R0. Having a meaningful encryption algorithm (expected) with ESP give a little advantage to standard R0 regarding MTU loss.

When considering MN-MN, IRO will clearly consumes less bandwidth than standard R0 in all possible combinations of algorithms for AH or ESP. Now, considering the same level of protection, i.e. by using IPsec for standard R0 carried packets (we do not take into account padding variations), IRO simply gets a direct advantage: 24 bytes for MN-CN

communications and 48 bytes for MN-MN communications. It is due to the complete removal of HAO and RH2 from packets exchanged directly between peers.

In fact, regarding MTU considerations, IRO provides a zero cost mobility service to IPsec protected connections between end nodes.

Appendix G. Compatibility with static keying

[TOC](#)

IRO has been designed for enabling direct secure communications between MIPv6 entities belonging to a common trust domain. Scalability was a primary concern; This is the reason why the specification covers SA negotiation under the hypothesis that IKE is used for that purpose. But IRO is also fully compatible with static keying.

In fact, the specification is not specifically bound to the use of either static or dynamic keying for SA setup; it is left as a local configuration decision to domain administrators.

This appendix quickly covers the differences regarding the use of static keying with IRO.

One great difference between static and dynamic keying is the removal of the IKE negotiation. For IRO, the first negotiation performed with the peer provides an additional information to the CN: a live proof of address (HoA) ownership by the MN. The removal of this step also removes the live check. This is a fact administrators should be aware of.

There should be no specific reason preventing the simultaneous use of static and dynamic keying with IRO.

Appendix H. Compatibility with the use of CoA in SP/SA

[TOC](#)

SP and SA are not changed at any moment by MIPv6 stack when IRO is used. Only incoming and outgoing IPsec packets can undergo source or destination address remapping.

When using IRO, MIPv6 stack will only act upon IPsec traffic associated with SA that involve the HoA of the MN or of a remote IRO peers's (MN) HoA as source or destination addresses (endpoints for tunnel mode SA). Outgoing IPsec packets are only possibly modified to change the HoA into the CoA. CoA of outgoing IPsec packets are never modified by MIPv6 stack, when IRO is used.

Incoming IPsec packets will have their source modified (from CoA to peer's MN HoA) if their are associated with a SA that expects the HoA of an IRO peer. This implies that no incoming packet with a CoA source will be modified if the associated SA references that CoA (and not peer's HoA). Regarding destination address of an incoming IPsec packet, remapping of a CoA will occur if the SA expects an HoA. This implies

that no incoming packet with a CoA destination will be modified if IRO remapping rules references that CoA (and not our HoA).

As a conclusion, the work of IRO is compatible with the use of CoA as destination or source address (endpoint addresses for tunnel mode) of any SP/SA.

Appendix I. Rationale for not specifying a new BU

[TOC](#)

IRO is not designed as a fallback mode for IPsec communications between MIPv6 entities but as an improved alternative.

You cannot use IRO and common mode with the same peer. You either need the security advantages of IRO for communications with a peer or you can afford unprotected direct communications with it, for which common RO has been developed. Parallel uses of common mode and IRO mode with different MIPv6 entities (including its HA) is not forbidden but strongly discouraged as it suppresses the anonymity of the MN on its foreign link.

For that reasons, IRO does not come with some detection algorithm against peers that do not have IRO activated to perform a fallback to common mode. Considering the setup associated with the protection mechanisms required by IRO and the kind of environments it is expected to be used in, requiring that entities be configured to specifically use IRO for a peer (or by default, preventing the common mode) is required.

This has many positive impacts both on development costs, deployment and debugging. This notably provides the ability to reuse messages without creating parallel versions where needed. As only a few things changes when IRO is activated between two entities, most of the code remains usable. In fact, the two main changes introduced by IRO are:

- *the removal of HAO/RH2 and the replacement by the remapping process on selected IPsec traffic.

- *a simplified registration procedure with CN (AOT* framework)

Let's go a little further. One can think that it would have been possible to create specific mobility options to discriminate IRO mode from the common mode. This was impossible for multiple reason. First, from a specification perspective, section 6.1.7 of [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) requires that "The receiver MUST ignore and skip any options which it does not understand", which prevents the reuse of MIPv6 messages with a slightly modified semantic if peers are not aware of that. For options to have an interest, you have to be aware that the peer support it (not necessarily that it is activated).

Anyway, there is a better reason that makes the use of common mode and IRO mode incompatible between peers: IRO remapping process must be activated on the receiver for the packets to be valid. If a MN that uses IRO sends an IPsec protected Binding Update message to a peer that is not using IRO, no remapping will occur and the checksum will end up being invalid (if it passes the IPsec stack). Section 9.2 of [\[RFC3775\] \(Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6," June 2004.\)](#) requires the following rule to be applied to such packet: "The checksum must be verified as per Section 6.1. Otherwise, the node MUST silently discard the message".

Appendix J. Anonymity

[TOC](#)

There are mainly 2 kinds of identifiers that can appear during an IPsec protected communication in IKE/MIPv6 environments: addresses (HoA, CoA, CN addresses) and cryptographic identifiers (IKE credentials, i.e. X.509 certificates).

In MN-CN communications, the addresses of the CN and the CoA of the MN will obviously be disclosed, but they should not be meaningful. We see later that in MN-MN case, the address of the CN that appears on the wire might temporarily be the HoA, before registration has been performed by the second MN.

In MIPv6, the use of explicit containers (HA0 and RH2) makes the Home Address of the MN available in all cases. With IRO, the complete removal of this extensions prevents the disclosure of the HoA during direct MN-CN communications and MN-HA communications.

The removal applies to:

- *the IPsec protected signaling traffic exchanged directly between the two peers (BU/BA for instance)

- *the IPsec protected data traffic exchanged in MN-MN and MN-CN cases (when tunnel mode is not already used to avoid RH2/HA0).

From the perspective of an eavesdropper on the FL of the MN, when IRO is used the visible exchanges that occur are (in order, for MN-MN case, with registrations performed on that link, i.e. worst case scenario):

- *IKE negotiation between the CoA of the MN and its HA

- *IPsec protected BU/BA exchanges using CoA and HA@ as on-wire addresses (remapping rules applied on both ends)

- *IPsec protected (tunnel mode this time) traffic with peers

*IKE exchange from the HoA, IPsec tunneled to the HA, with a CN

*Direct IPsec-protected BU/BA exchanges using the CoA and the address of the the CN for on-wire addresses.

*Direct IPsec-protected data traffic exchange with the CN, between the CoA and the address of the CN.

In all those exchanges, the only addresses that are disclosed to an eavesdropper on the FL of the MN (if ESP with a meaningful encryption is used for all IPsec exchanges) are the CoA, the address of MN's HA and the address of the CN. The HoA of the MN never appears in those exchanges.

For IKE case, even if it is used as an ID in Phase 2 for bootstrapping as described in [\[MIGRATE\] \(Ebalard, A. and S. Decugis, "PF_KEY Extension as an Interface between Mobile IPv6 and IPsec/IKE," August 2008.\)](#), the exchanges are encrypted and the HoA does not appear. For the negotiation with the CN, because the HoA is used for the exchanges, the IPsec tunnel to the HA protects traffic to/from the Home Network.

Regarding cryptographic identifiers, the certificate of the MN is not expected to appear on the wire. In all cases, the only information one can get are associated with the MN's Home Network (HA address and possibly certificate), but nothing more specific should be disclosed. Now, considering the specific case of MN-MN communications, on the network of the initiating MN1 (the first to register with its peer), after the IKE negotiation as been performed relayed by both Home Agents, the IPsec protected Binding Update packets is emitted with the HoA of MN2 as destination (the address of the CN in previous list is the HoA of MN2). Let's consider associated SPI is 42. The packet is sent directly with the CoA of MN1 on the wire and is routed to the Home Network of the peer, before it is tunneled to it. The BA follows a reverse path but with a different SPI (say 43). After the second registration is over, the MH traffic using those SPI values (42 and 43) flows directly (remapping rules are now in place on both ends). From an eavesdropper perspective on the FL of MN1, this provides "a clue" about the association between the HoA and the CoA of the second MN2. This is introduced in [\[RFC4882\] \(Koodli, R., "IP Address Location Privacy and Mobile IPv6: Problem Statement," May 2007.\)](#).

Note that this is the only "leaking" that happens and only on the FL of the first MN. It is no more the case on FL that are visited later. Anyway, from the perspective of an eavesdropper monitoring that, the information will be that a Mobile Node from a known Home Network (HA@) has performed IPsec communications with a MN having a known HoA (no credentials).

Author's Address[TOC](#)

	Arnaud Ebalard
	EADS Innovation Works
	12, rue Pasteur - BP76
	Suresnes 92152
	France
Phone:	+33 1 46 97 30 28
Email:	arnaud.ebalard@eads.net