

Address Management for IKE version 2

[<draft-dupont-ikev2-addrmgmt-05.txt>](mailto:draft-dupont-ikev2-addrmgmt-05.txt)

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, or will be disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

Status of this Memo

This document is an Internet Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

Abstract

The current IKEv2 proposal [1] lacks an address management feature. As it is compatible with the NAT traversal capability, this document specifies a complete address management with support for multi-homing and mobility, and fulfill mobike working group [2] goals 1, 2, 3, 4, and 6 (for goal 5 look at [3]).

INTERNET-DRAFT

Address Management for IKEv2

Jun 2004

[1.](#) Introduction

In this document, the addresses used to transport IKE messages are named the "peer addresses" (term introduced by [\[4\]](#)). These peer addresses should no more be directly or indirectly included in identities ([\[5\]](#) and [\[6\]](#)) as it is commonly done for IKEv1.

The current IKEv2 draft [\[1\]](#) often makes the assumption that an address identifies a node when nodes behind a NAT can share the same address and a node can use many different addresses. This must be taken into account in implementations, for instance by reading this document before writing code...

This document describes the goals of an address management for IKEv2, including the requirements for multi-homing and mobility support (this part will be removed as soon as the mobility requirements document [\[12\]](#) is finalized), and finishes by a concrete proposal.

In this document, open questions are introduced by the word NOTE and will be refined in a dedicated section.

[2.](#) Goals

The goals of the address management proposed in the document can be divided in some general goals and in requirements for the three mechanisms which can change the peer addresses.

[2.1](#) Simplicity, Performance and Security

The address management should be as simple as possible, i.e., it should introduce minimal additions to the current IKEv2 draft [\[1\]](#) and each addition should be justified.

The performance is an important criterion. For instance, rekeying can update the peer addresses of an IKE SA or an IPsec SA pair, but rekeying is too expensive and a specific solution is needed.

As a security protocol, IKEv2 should get a high security level. Unfortunately we already showed that the NAT traversal feature comes with a security issue (the transient pseudo-NAT

attack [7]).

Such problems introduced by the peer address flexibility must be described in this document and at least be mitigated by options in configurations. For instance, the NAT traversal feature should never be enabled when one knows that there can not be a NAT as today in IPv6.

An other example of an insecure mechanism is to use the addresses in IP headers of CREATE_CHILD_SA messages as the endpoint addresses of the new IPsec SAs without further control on them: peer addresses must be managed.

[2.2](#) Terminology

The addresses of the two peers are named "peer addresses". With other words the peer addresses are the addresses IKE runs over but this document extends this basic definition. The primary peer address of a peer is initialized to the address used to transport messages of the initial exchanges, other addresses are "alternate peer addresses".

The proxy case is the setup of transport mode IPsec SAs on the behalf of another party, i.e., transport mode IPsec SAs where the traffic selectors do not match the primary peer addresses.

[2.3](#) Multi-homing requirements

In this document, the support of multi-homing is the support of nodes with several global addresses. Some of the addresses can be "better" than others, or "better" for some destinations. Some can, from time to time, be unavailable.

The main requirement for the support of multi-homing is the management of a set of peer addresses for each peer. The set can be partially ordered or some subset can be loosely associated with some destinations (i.e., some subset of the other peer address set).

For the communication between multi-addressed hosts, the support of the proxy case can be useful because it provides an easy way to

setup transport mode IPsec SAs with different addresses from one IKE SA. In such cases the other party is in fact the same host, this dramatically simplifies the authorization issue.

[2.4](#) Mobility requirements

In the context of Mobile IPv6 ([[8](#)] and for the special case of Home Agents [[9](#)]), the interaction of Mobility and IPsec was analyzed in another document [[10](#)]. This document assumes an IPv6 context as Mobile IPv6 is the most powerful mobility proposal available today.

An IPv6 mobile node is another type of multi-addressed node with:

- a care-of address in a prefix of the visited link.

The care-of address is used to route packets.

- the home address in a prefix of the home link.

The home address is used to identify the mobile node.

The care-of address is transient and usually the mobile node can not provide a proof that it is the node using it. So it has to be trusted and a return routability check (i.e., an enforced answer from this address) should be used if it is not.

With a common correspondent, the mobility is transparent and there is no reason to use another address than the home address. For optimized schemes, without an implementation of header compression in ESP tunnel mode (the goal 5 of mobike [[2](#)]) the choice between a transport mode using triangular routing (IPsec can be used to verify home address options) and a tunnel mode with routing optimization is not clear. But this case does not add new requirement, i.e., the home agent case includes them.

With the home agent, there are three main cases (c.f. [[9](#)]):

- The mobility signaling which is mandatory protected and raises a specific issue in its initial phase: the IKE SA must be setup using the care-of address as the peer address but this IKE SA is used to build an IPsec SA pair with the home address as traffic selector. This IPsec SA will protect the home registration which will make the home address available. This can be considered as a specialized proxy case.

- Other genuine communications between the home agent and the mobile node can be covered by the proxy case support too. Note this is the only case at the exception of signaling where mobility behaves in a different way than a mobile IPsec VPN (so we proposed to relax the corresponding rule in a future version of [8] and [9]).

- The traffic relayed by the home agent through a tunnel with the mobile node can be partially or fully protected by IPsec SA pair(s). Encapsulation should be performed only once, including for degenerated (but not for free) encapsulation like the home address option or the mobility routing header.

In all cases of interaction with the home agent, the mobile node peer address should be a care-of address. When the mobile node moves, another care-of address is used and some SAs, including the IKE SA, must be updated to use the new address.

Usually the previous peer address is no more usable. In order to avoid a trivial denial of services, a strong sequencing of updates is required with a way to cancel possible pending updates when fast multiple handoff happen.

The IPsec pair which protects the mobility signaling uses the home address as its traffic selector for the mobile node. It must not be updated at each handoff. The update mechanism must provide a fine grain (i.e., per SA) update.

[3.](#) Proposal

The proposal for an address management in IKEv2 is spawn from the NAT traversal mechanisms, mainly with a new peer address update payload. But there are some points that have to be kept or clarify as they already are in the current IKEv2 draft [\[1\]](#).

[3.1](#) Kept points from/clarification to the IKEv2 draft 14

The peer addresses MUST be stable during the initial exchanges, i.e., the IKE_SA_INIT and IKE_AUTH exchanges MUST be transported using the same peer address pair.

The peer addresses are used to transport messages. The reply to a request MUST be sent to the source of the request from the destination request, i.e., addresses and ports are reversed between the request and its reply. There is no exception to this rule.

For tunnel mode IPsec SAs, the endpoint addresses are the primary peer addresses. We don't propose an alternate way to specify them. The same requirement applies to transport mode IPsec SAs at the exception of the proxy case.

[3.2](#) Minor points

In retransmission of requests or responses, copies of messages do not include peer addresses. So a peer MAY retransmit an IKE message from or to a different address.

The primary peer addresses are IKE SA parameters and are specified by the IKE_SA_INIT exchange. Note that when NAT traversal is not active, they are implicitly protected by the NAT_DETECTION notifications.

All the text below applies only to the case where NAT traversal is not active.

In the proxy case, the initiator is acting as a client negotiator on the behalf of another party. The address of this other party is sent in the initiator traffic selector and will become the address of this end of the transport mode IPsec SA pair. A proper authorization in the local policy of the responder is REQUIRED, the default rule SHOULD be:

- using an alternate peer address is permitted
- other cases are denied.

Return routability checks are done using an informational exchange carrying a nonce in order to get a proof the probed peer really receives the request. Of course the replay MUST contain the same nonce payload than the request.

[3.3](#) Peer address notifications

The peer address notifications are copied from the current NAT-DETECTION-SOURCE-IP and NAT-DETECTION-DESTINATION-IP notifications. They includes the peer source or destination address with its family and an operation code. They MUST be in an encrypted payload. Operations are PRIMARY, ADD and DELETE (last two for alternate addresses, see open issue section for the empty set one and the delete operation on the primary peer address).

All messages after the first exchange involving an alternate peer address MUST include at least one peer address notification for each peer, i.e., at least one for the source and at least one for the destination.

Such messages belong to IKE_AUTH or CREATE_CHILD_SA exchanges, or carry the peer address update payload defined below, or are pure peer address set management (add/delete).

They provide a cryptographically proof of no en-route alteration of the peer addresses and enable operations on the sets of peer addresses, i.e., change of the primary peer address of a peer, addition to and deletion from the peer address set of a peer.

When the peer address notifications are not supported, the capability to use an alternate peer address, and only this, is lost.

[3.4](#) Explicit peer address update payload

A new payload has to be defined for an explicit peer address update mechanism. We propose to copy it from the delete payload, see Annex B.

The new peer address update payload has strong sequencing requirements. IKEv2 messages have a protected sequence number so the only sequencing issues are the window of processing and pending exchanges. Any messages with a peer address update payload MUST be processed in order.

When the receiver of an update request has to check the validity of the new primary peer address, it MAY use a return routability

check sending an informational request carrying a nonce payload at the new address and waiting for an answer. As informational exchanges are protected no more is needed.

Example of a return routability check:

```
I ----- address update request -----> R
I <-- informational RR [Ni] request - R
I --- informational RR {Ni} reply --> R
    now the responder knows the initiator should be where it
    claimed to be.
I <----- address update reply ----- R
```

As for the delete payload, the peer address update payload specifies the SPIs of the IPsec and IKE SAs it applies to. But a simple way to specify all SAs (i.e., the IKE SA and all the tunnel mode IPsec SAs it negotiated) is needed so is provided.

[3.5](#) Open issues

Notification/payload/exchange: the current choice is a notification for peer addresses (copied on NAT detection notifications) and a payload for peer address update (copied on SA delete payload).

Interaction with NAT-T: the current choice is to avoid the case where one peer is behind a NAT then uses NAT-T and the other peer uses MOBIKE: in this situation NAT-T usage by both peers MUST be enforced.

Path failure detection: the proposal does not provide any dedicated mechanism, the generic mobility or multi-homing control SHOULD be used instead, including for simultaneous changes.

When to perform a return routability check?: this is a policy issue, some answers are in the next section.

Can a peer address set be empty?: still open. Mechanisms permit this...

New error notification for address problems: likely to be necessary.

Peer address addition request from an unknown address: (here

unknown means not in the peer address set even after processing of the message) this is the only circumstance where return routability checks are clearly REQUIRED:

- if it succeeds for the peer address the message MUST be accepted
- if it fails for the peer address but succeeds for the unknown source address the peer has moved behind a NAT

[4. Security Considerations](#)

Great care was used to avoid to introduce security threats.

The NAT traversal feature comes with a security flaw (the transient pseudo-NAT attack [\[7\]](#)) which can not be easily avoid. IMHO the NAT traversal feature should be enabled only when the presence of NATs is likely/possible.

When the NAT traversal feature is disabled, the address of the other peer can not be changed en-route by an attacker but the proofs the peer is really at its address are:

- the trust in the peer
- the source address is topologically plausible
- the proof that the peer can receive messages sent to its address.

[draft-dupont-ikev2-adrrmgmt-05.txt](#)

[Page 8]

INTERNET-DRAFT

Address Management for IKEv2

Jun 2004

The second (a.k.a. the return routability check) works only with at least two of three not-trivial messages, i.e., for the initial exchange (with the address stability requirement) and for explicit checks. IMHO these checks SHOULD be required for a new alternate peer address as soon as there is no proof of the address validity, for instance when:

- the message does not come from this address (ingress filtering [\[13\]](#) can drop a message with a fake source address),
- and there is no authorization for this address.

[5. Acknowledgments](#)

The rare people in the Mobility world with IPsec interests, or in the IPsec world with Mobility interest, should receive all thanks because without them we (me and all the future co-authors) have given up for a long time.

Tero Kivinen helped to improve the NAT traversal part of this

proposal. Tero and Jari Arkko proposed another form of peer address update based on the IKE SA addresses.

7. Normative References

[1] C. Kaufman, ed., "Proposal for the IKEv2 Protocol", [draft-ietf-ipsec-ikev2-14.txt](#), May 2004.

8. Informative References

[2] IKEv2 Mobility and Multihoming (mobike), "charter", <http://www.ietf.org/html.charters/mobike-charter.html>.

[3] J. Vilhuber, "IP header compression in IPsec ESP", [draft-vilhuber-hcoesp-00.txt](#), January 2003.

[4] B. Korver, "The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX", [draft-ietf-pki4ipsec-ikecert-profile-00.txt](#), May 2004.

[5] P. Hoffman, "Adding revised identities to IKEv2", <http://www.vpnc.org/ietf-ipsec/>,
Message-Id: <p05200f06b9edf48ac57b@[165.227.249.18]>,
November 2002.

[draft-dupont-ikev2-addrmgmt-05.txt](#)

[Page 9]

INTERNET-DRAFT

Address Management for IKEv2

Jun 2004

[6] M. Kaat, "Overview of 1999 IAB Network Layer Workshop", [RFC 2956](#), October 2000.

[7] F. Dupont, J.-J. Bernard, "Transient pseudo-NAT attacks or how NATs are even more evil than you believed", [draft-dupont-transient-pseudonat-03.txt](#), February 2004.

[8] D. Johnson, C. Perkins, J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.

[9] J. Arkko, V. Devarapalli, F. Dupont, "Using IPsec to Protect Mobile IPv6 Signaling between Mobile Nodes and Home Agents", [RFC 3776](#), June 2004.

[10] F. Dupont, W. Haddad, "How to make IPsec more mobile IPv6

friendly", [draft-dupont-ipsec-mip6-05.txt](#), February 2004.

[11] D. McDonald, C. Metz, B. Phan, "PF_KEY Key Management API, Version 2", [RFC 2367](#), July 1998.

[12] T. Kivinen, "Design of the MOBIKE protocol", [draft-kivinen-mobike-design-00.txt](#), February 2004.

[13] P. Ferguson, D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [RFC 2827](#) / [BCP 38](#), May 2000.

9. Author's Address

Francis Dupont
ENST Bretagne
Campus de Rennes
2, rue de la Chataigneraie
CS 17607
35576 Cesson-Sevigne Cedex
FRANCE
Fax: +33 2 99 12 70 30
EMail: Francis.Dupont@enst-bretagne.fr

[draft-dupont-ikev2-adrrgmt-05.txt](#)

[Page 10]

INTERNET-DRAFT

Address Management for IKEv2

Jun 2004

Annex A. Peer Address Notification Format.

The following diagram illustrates the content of the Peer Address Notification:

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
! Next Payload !C!  RESERVED   !           Payload Length           !
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
! Protocol-ID  !  SPI Size    !           Notify Message Type       !
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
!           Address Family           !           Operation           !
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                                     Address                           ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The notification header is for IKE SA (Protocol-ID 0, SPI Size 0

updated. In this case the update is for the IKE-SA (Protocol-ID 0, SPI size 0, no SPI and number of SPIs 0). MUST be set to zero when an individual SA is updated.

- o Protocol_Id (7 bits) - MUST be zero for an IKE_SA, 1 for ESP, or 2 for AH.
- o SPI Size (1 octet) - Length in octets of the SPI as defined by the Protocol-Id. Zero for IKE (the SPI is got from the message header) or four for AH or ESP.
- o # of SPIs (2 octets) - The number of SPIs contained in the Peer Address Update Notification. The size of each SPI is defined by the SPI Size field.
- o Security Parameter Index(es) (variable length) - Identifies the specific security association(s) to delete. The lengths of these fields are determined by the SPI Size and the number of SPIs fields.

The C[ritical] bit MUST be set to one even a peer which does not support Peer Address Update Payloads does not support Peer Address Notifications too.

ESP and AH SAs always exist in pairs, with one SA in each direction. When an SA is updated for a peer address, both members of the pair MUST be updated. When SAs are nested, as when data (and IP headers if in tunnel mode) are encapsulated first with IPcomp, then with ESP, and finally with AH between the same pair of endpoints, all of the SAs MUST be updated together. Each endpoint MUST update the SAs it receives on and allow the other endpoint to update the other SA in each pair.

To update a peer address of an SA, an Informational Exchange with one or more peer address update payloads is sent listing the SPIs (as they would be placed in the headers of inbound packets) of the SAs to be updated, and with a peer address notification setting the primary peer address. The recipient MUST update the designated SAs. Normally, the reply in the Informational Exchange will contain peer address update payloads for the paired SAs going in the other direction. Note there is no special case for update collision.

The proposed name is the Update (U) payload.

Annex C. PF_KEY version 2 SADB_X_ADDUPD

This annex describes an extension to PF_KEYv2 [11] which provides a way to ask a peer address update of an IPsec SA and all its siblings (i.e., an update with the All flag set to one).

The format of the message is:

<base, SA(*), address(SD), new_address(SD)>

and is sent the registered socket listeners by or via the kernel. No answer is needed because if it fails it will be done again.

New values are needed for SADB_X_ADDUPD and for SADB_X_EXT_NEW_ADDRESS_SRC and SADB_X_EXT_NEW_ADDRESS_DST which should have the same layout than SADB_EXT_ADDRESS_*, i.e., sadb_address structure.

NOTE: IKE itself needs a PF_KEYv2 extension for individual updating of an IPsec SA.