

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 5, 2010

V. Devarapalli
WiChorus
K. Weniger
August 4, 2009

Redirect Mechanism for IKEv2
draft-ietf-ipsecme-ikev2-redirect-13.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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/lid-abstracts.txt>.

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

This Internet-Draft will expire on February 5, 2010.

Copyright Notice

Copyright (c) 2009 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

Internet-Draft

IKEv2 Redirect

August 2009

Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

IKEv2 is a protocol for setting up Virtual Private Network (VPN) tunnels from a remote location to a gateway so that the VPN client can access services in the network behind the gateway. This document defines an IKEv2 extension that allows an overloaded VPN gateway or a VPN gateway that is being shut down for maintenance to redirect the VPN client to attach to another gateway. The proposed mechanism can also be used in Mobile IPv6 to enable the home agent to redirect the mobile node to another home agent.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	IKEv2 Initial Exchange with Redirect	3
4.	Use of Anycast Addresses with the Redirect Mechanism	5
5.	Redirect During an Active Session	6
6.	Redirect During IKE_AUTH Exchange	7
7.	Handling Redirect Loops	8
8.	Using the Redirect Mechanism with Mobile IPv6	8
9.	Redirect Messages	9
9.1.	REDIRECT_SUPPORTED	9
9.2.	REDIRECT	10
9.3.	REDIRECTED_FROM	11
10.	Use of the Redirect Mechanism between IKEv2 Peers	12
11.	Security Considerations	12
12.	IANA Considerations	13
13.	Acknowledgements	13
14.	References	14
14.1.	Normative References	14
14.2.	Informative References	14
	Authors' Addresses	14

Internet-Draft

IKEv2 Redirect

August 2009

[1.](#) Introduction

IKEv2 [\[2\]](#) is used for setting up IPsec [\[8\]](#) based VPNs. The IP address of the VPN gateway can be configured on the VPN client. But this does not scale well, when the number of VPN gateways is large. Dynamic discovery of VPN gateways using DNS is quite widely used too. However, using DNS is not flexible when it comes to assigning a VPN gateway to the VPN client based on the load on the VPN gateways. The VPN client typically tries to connect to the IP address of the VPN gateway that appears first in the DNS response. If the VPN tunnel setup fails, then the VPN client tries to attach to the other VPN gateways returned in the DNS response.

This document proposes a redirect mechanism for IKEv2 that enables a VPN gateway to redirect the VPN client to another VPN gateway, for example, based on the load condition. The redirect can be done during the IKE_SA_INIT or the IKE_AUTH exchange. Gateway-initiated redirect in the middle of a session is also supported. The redirect mechanism can also be used in conjunction with anycast addresses. In this case, anycast address for the cluster of VPN gateways is stored in the DNS instead of a list of unicast IP addresses of the VPN gateways.

The redirect can also happen because of administrative or optimal routing reasons. This document does not attempt to provide an exhaustive list of reasons for redirecting a VPN client to another VPN gateway.

[2.](#) Terminology

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 [\[1\]](#).

3. IKEv2 Initial Exchange with Redirect

This section describes the use of Redirect mechanism during the IKE_SA_INIT exchange. Gateway-initiated redirect during an active session and the use of redirect during IKE_AUTH exchange are explained in subsequent sections.

The VPN client indicates support for the IKEv2 redirect mechanism and the willingness to be redirected by including a REDIRECT_SUPPORTED notification message in the initial IKE_SA_INIT request. (See [Section 9.1](#)). The gateway MUST keep track of those clients that indicated support for the redirect mechanism and those that didn't.

To redirect an IKEv2 session to another VPN gateway, the VPN gateway that initially received the IKE_SA_INIT request selects another VPN gateway (how the selection is made is beyond the scope of this document), and replies with an IKE_SA_INIT response containing a REDIRECT notification message. (See [Section 9.2](#)). The notification includes information about the selected VPN gateway, and the nonce data from the Ni payload in the IKE_SA_INIT request. If the IKE_SA_INIT request did not indicate support for the redirect mechanism, the responder MUST NOT send the REDIRECT payload to the VPN client. This is applicable to all REDIRECT scenarios described in this document.

Note that when the IKE_SA_INIT response includes the REDIRECT notification, the exchange does not result in the creation of an IKE_SA and the responder SPI will be zero.

Initiator	Responder (initial VPN GW)
-----	-----
(IP_I:500 -> Initial_IP_R:500)	
HDR(A,0), SAi1, KEi, Ni, -->	
N(REDIRECT_SUPPORTED)	
	(Initial_IP_R:500 -> IP_I:500)
	<-- HDR(A,0), N(REDIRECT, New_GW_ID, Ni_data)

When the client receives the IKE_SA_INIT response, it MUST verify that the nonce data matches the value sent in the IKE_SA_INIT

request. If the values do not match, the client MUST silently discard the response (and keep waiting for another response). This prevents certain Denial-of-Service attacks on the initiator that could be caused by an attacker injecting IKE_SA_INIT responses with the REDIRECT payloads.

After verifying the nonce data, the client initiates a new IKE_SA_INIT exchange with the VPN gateway listed in the REDIRECT payload provided this is allowed by its PAD entries. In the IKE_SA_INIT exchange with the new VPN gateway, the client MUST include the REDIRECTED_FROM payload. (See [Section 9.3](#)). The VPN client includes the IP address of the original VPN gateway that redirected the client in the REDIRECTED_FROM notification. The IKEv2 exchange then proceeds as it would have proceeded with the original VPN gateway.

Initiator -----	Responder (Selected VPN GW) -----
(IP_I:500 -> IP_R:500) HDR(A,0), SAi1, KEi, Ni, --> N(REDIRECTED_FROM, Initial_IP_R)	
	(IP_R:500 -> IP_I:500) <-- HDR(A,B), SAR1, KEr, Nr, [CERTREQ]
(IP_I:500 -> IP_R:500) HDR(A,B), SK {IDi, [CERT,] [CERTREQ,] [IDr,]AUTH, SAi2, TSi, TSr} -->	
	(IP_R:500 -> IP_I:500) <-- HDR(A,B), SK {IDr, [CERT,] AUTH, SAR2, TSi, TSr}

The client MAY get redirected again by the new VPN gateway if the new VPN gateway cannot also serve the client. The client does not have to include the REDIRECT_SUPPORTED payload again in the IKE_SA_INIT exchange with the new gateway after a redirect. The presence of the

REDIRECT_FROM payload in the IKE_SA_INIT exchange with the new gateway indicates to the new gateway that the client supports the redirect mechanism.

When the client gets redirected, it MUST use the same Peer Authorization Database (PAD) and Security Policy Database (SPD) entries as it would have used with the original gateway. Receiving a redirect notification MUST NOT result in the modification of any PAD or SPD entries. In practice, this means the new gateway either has to use the same responder identity (IDr) as the original gateway, or both should be part of a group of responders that are authorized by the same PAD entry. See section 4.4.3.1 of [8] on using DNS names to represent a group of peers in a PAD entry.

[4.](#) Use of Anycast Addresses with the Redirect Mechanism

The use of anycast addresses will avoid having to configure a particular VPN gateway's IP address in the DNS. Instead, the anycast address that represents the group of VPN gateways is stored in the DNS. When the VPN client performs a DNS lookup for the VPN gateway, it receives the anycast address of the VPN gateway in the DNS response.

If an anycast address is returned in response to DNS resolution of an FQDN, the VPN client sends the IKE_SA_INIT request to the anycast

address. The REDIRECT_SUPPORTED payload is included in the IKE_SA_INIT request sent to the anycast address. The IKE_SA_INIT request is routed to one of the VPN gateways that is part of the anycast group. The VPN gateway that receives the IKE_SA_INIT request responds with an IKE_SA_INIT reply from the anycast address.

Initiator	Responder (any VPN GW)
-----	-----
(IP_I:500 -> ANYCAST:500)	
HDR(A,0), SAi1, KEi, Ni) -->	
N(REDIRECT_SUPPORTED)	
	(ANYCAST:500 -> IP_I:500)
	<-- HDR(A,0), N(REDIRECT, New_GW_ID, Ni_data)

If the destination address on the IKE_SA_INIT request is an anycast address, the VPN gateway that received the IKE_SA_INIT request MUST include the REDIRECT payload to redirect the VPN client to a unicast address of one of the VPN gateway. The VPN gateway that received the IKE_SA_INIT request MAY redirect the client to its own unicast address, if it is not overloaded.

The rest of the IKEv2 exchange is the same as described in [Section 3](#).

[5](#). Redirect During an Active Session

The redirect mechanism may also be used by a VPN gateway to redirect the client to another VPN gateway in middle of a session. To redirect a client, the gateway should send an INFORMATIONAL message with the REDIRECT Notify payload. The REDIRECT payload MUST carry information about the new VPN gateway. The gateway MUST NOT include any nonce data in the REDIRECT payload, since it is a gateway-initiated redirect and is protected by the IKEv2 security association. When the client receives this message, it sends a response (usually empty) to the gateway. The gateway retransmits the redirect INFORMATIONAL message as described in [\[2\]](#), until it gets a response. The following illustrates the INFORMATIONAL message exchange for gateway-initiated redirect.

Initiator (VPN client)	Responder (VPN GW)
-----	-----
	<-- HDR, SK {N(REDIRECT, New_GW_ID)}
HDR, SK {} -->	

The INFORMATIONAL message exchange described above is protected by the existing IKEv2 SA between the client and the gateway.

Once the client sends an acknowledgement to the gateway, it SHOULD delete the existing security associations with the old gateway by sending an INFORMATIONAL message with a DELETE payload. The gateway MAY also decide to delete the security associations without any signaling from the client, again by sending an INFORMATIONAL message

with a DELETE payload. However, it should allow sufficient time for the client to setup the required security associations with the new security gateway. This time period should be configurable on the gateway.

6. Redirect During IKE_AUTH Exchange

If the gateway decides to redirect the client during the IKE_AUTH exchange, based on the identity presented by the client in the IKE_AUTH request message, it prevents the creation of a CHILD SA and sends the REDIRECT payload in the IKE_AUTH response. The gateway MUST verify the client's AUTH payload before sending the Redirect payload, and the client MUST verify the gateway's AUTH payload before acting on the Redirect payload. Since the AUTH payloads were exchanged and successfully verified, the IKEv2 security association is valid. When the client receives the IKE_AUTH response with the REDIRECT payload, it SHOULD delete the IKEv2 security association with the gateway by sending an INFORMATIONAL message with a DELETE payload.

Initiator -----	Responder (VPN GW) -----
(IP_I:500 -> IP_R:500) HDR(A,0), SAi1, KEi, Ni, --> N(REDIRECTED_SUPPORTED)	
	(IP_R:500 -> IP_I:500) <-- HDR(A,B), SAR1, KEr, Nr, [CERTREQ]
(IP_I:500 -> IP_R:500) HDR(A,B), SK {IDi, [CERT,] [CERTREQ,] [IDr,]AUTH, SAi2, TSi, TSr} -->	
	(IP_R:500 -> IP_I:500) <-- HDR(A,B), SK {IDr, [CERT,] AUTH, N(REDIRECT, New_GW_ID)}

In case the IKE_AUTH exchange involves EAP authentication as

methods as described in [RFC 4739](#) [6], the gateway may decide to redirect the client based on the interaction with the AAA server or the external authentication server. In this case, the gateway MUST send the REDIRECT Notification payload in either the first or the last IKE_AUTH response. The client and the gateway MUST verify the AUTH payloads as described above.

When EAP is used, the gateway MAY also redirect the client based on the unauthenticated identity presented by the client in the first IKE_AUTH exchange itself. Since EAP is used as the authentication mechanism, the client does not include AUTH payload to authenticate his identity, but the server still MUST include his own AUTH payload, and client MUST verify it. Note that the IKEv2 SA is not created in this case and the client does not have to explicitly delete the IKEv2 SA.

In all of the cases above, the client MUST accept the REDIRECT notification only in the first IKE_AUTH response or the last IKE_AUTH response. It MUST NOT accept the REDIRECT notification in an intermediate IKE_AUTH response.

[7.](#) Handling Redirect Loops

The client could end up getting redirected multiple times in a sequence, either because of wrong configuration or a DoS attack. The client could even end up in a loop with two or more gateways redirecting the client to each other. This could deny service to the client. To prevent this, the client SHOULD be configured not to accept more than a certain number of redirects (MAX_REDIRECTS) within a short time period (REDIRECT_LOOP_DETECT_PERIOD) for a particular IKEv2 SA setup. The default value for MAX_REDIRECTS configuration variable is 5. The default value for REDIRECT_LOOP_DETECT_PERIOD configuration variable is 300 seconds. Client implementations may allow these variables to be configured depending on a specific deployment or system configuration.

[8.](#) Using the Redirect Mechanism with Mobile IPv6

Mobile IPv6 [3] may use IKEv2 for mutual authentication between the mobile node and the home agent, for home address configuration and for setting up security associations for protecting Mobile IPv6 signaling messages [4]. The IKEv2 exchange, if IKEv2 is used, precedes the exchange of Mobile IPv6 signaling messages. Therefore, the mechanism described in this document can also be used by a Mobile IPv6 home agent to redirect a mobile node to another home agent.

There is a Home Agent Switch mechanism available for redirecting a mobile node to another home agent, described in [5]. The Home Agent Switch mechanism can only be used after the binding cache had been created at the home agent for the mobile node. The disadvantage with this is that quite a bit of state is created on the home agent before the mobile node can be redirected to another home agent. The mechanism described in this document can be used for redirecting a mobile node before any state related to the Mobile IPv6 binding is created on the home agent.

When running IKEv2 between a Mobile IPv6 Mobile Node (MN) and Home Agent (HA), redirecting the IKEv2 exchange to another HA is not enough; the Mobile IPv6 signalling also needs to be sent to the new HA address. The MN MAY treat the information received in the IKE_SA_INIT response in similar way as it would treat HA discovery information received from other unauthenticated (and potentially untrustworthy) sources (such as DNS lookups not protected with DNSSEC). However, if the MN has authenticated information about its Home Agent, it MUST NOT be updated based on the IKE_SA_INIT response.

If the REDIRECT notification is received during the IKE_AUTH exchange (after the HA has been authenticated; see [Section 6](#)), the MN MAY pass the new address to Mobile IPv6 and treat it in similar fashion as information from the Home Agent Switch Message [5].

Gateway-initiated REDIRECT notifications exchanged in INFORMATIONAL exchanges (see [Section 5](#)) MUST NOT result in updating any Mobile IPv6 state. In such cases, the Home Agent Switch Message specified in [5] is used instead.

[9.](#) Redirect Messages

[9.1.](#) REDIRECT_SUPPORTED

The REDIRECT_SUPPORTED payload is included in the initial IKE_SA_INIT request by the initiator to indicate support for the IKEv2 redirect mechanism described in this document.

1										2										3											
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(=0)										SPI Size (=0)										Notify Message Type											

The 'Next Payload', 'Payload Length', 'Protocol ID', 'SPI Size' and

Internet-Draft

IKEv2 Redirect

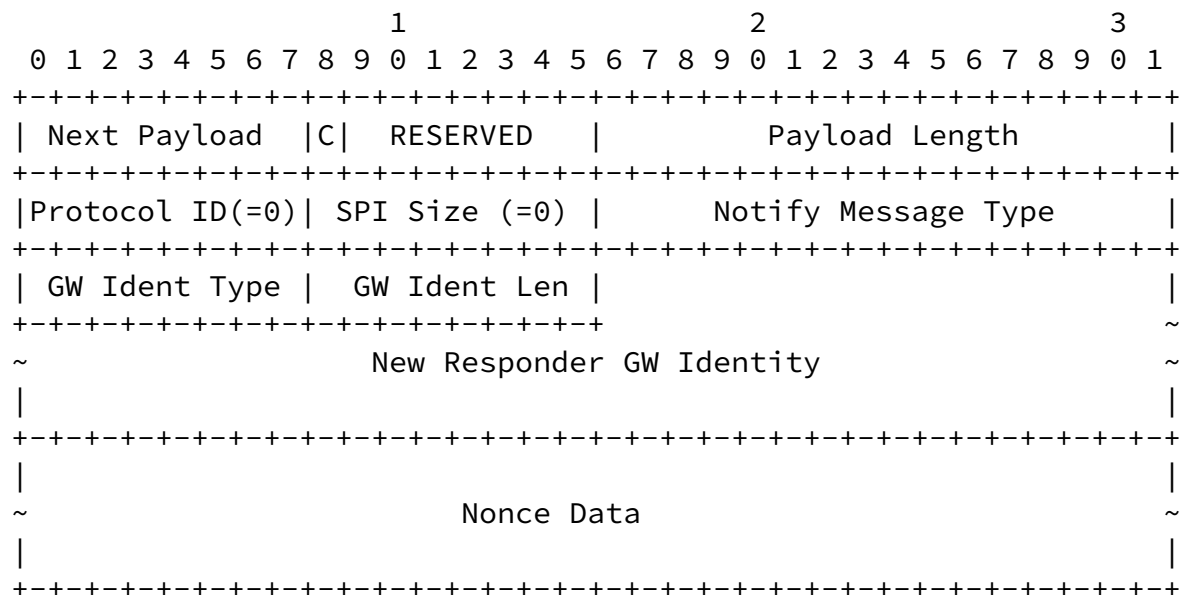
August 2009

the 'Notify Message Type' fields are the same as described in [Section 3.10](#) of [2]. The 'SPI Size' field MUST be set to 0 to indicate that the SPI is not present in this message. The 'Protocol ID' MUST be set to 0, since the notification is not specific to a particular security association.

The 'Payload Length' field is set to the length in octets of the entire payload, including the generic payload header. The 'Notify Message Type' field is set to indicate the REDIRECT_SUPPORTED payload <value to be assigned by IANA>.

9.2. REDIRECT

The REDIRECT payload is included in an IKE_SA_INIT response from the responder or an INFORMATIONAL message from the responder, when the responder wants to redirect the initiator to another VPN gateway. The message includes the new responder's IP address or DNS name.



The 'Next Payload', 'Payload Length', 'Protocol ID', 'SPI Size' and the 'Notify Message Type' fields are the same as described in [Section 3.10](#) of [2]. The 'SPI Size' field MUST be set to 0 to indicate that the SPI is not present in this message. The 'Protocol ID' MUST be


```

| Next Payload |C| RESERVED | Payload Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Protocol ID(=0) | SPI Size (=0) | Notify Message Type |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| GW Ident Type | GW Ident Len |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
~ Original Responder GW Identity ~
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The 'Next Payload', 'Payload Length', 'Protocol ID', 'SPI Size' and the 'Notify Message Type' fields are the same as described in [Section 3.10](#) of [2]. The 'SPI Size' field MUST be set to 0 to indicate that the SPI is not present in this message. The 'Protocol ID' MUST be set to 0, since the notification is not specific to a particular security association.

The 'Payload Length' field is set to the length in octets of the entire payload, including the generic payload header. The 'Notify Message Type' field is set to indicate the REDIRECTED_FROM payload <value to be assigned by IANA>. The 'GW Identity Type' field indicates the type of information that is sent to identify the new VPN gateway. The following values are valid in the REDIRECTED_FROM payload.

- 1 - IPv4 address of the original VPN gateway
- 2 - IPv6 address of the original VPN gateway

The 'GW Ident Len' field is set to the length of the gateway identity information. The identity of the original VPN gateway is carried in the 'Original Responder GW Identity' field.

[10.](#) Use of the Redirect Mechanism between IKEv2 Peers

The Redirect mechanism described in this document is mainly intended for use in client-gateway scenarios. However, the mechanism can also be used between any two IKEv2 peers. But this protocol is asymmetric, meaning that only the original responder can redirect the original initiator to another server.

11. Security Considerations

An eavesdropper on the path between VPN client and server may send a redirect to the client upon receiving an IKE_SA_INIT message from this client. This is no problem regarding DoS attacks for the VPN connection, since an on-path-attacker can as well drop the IKE_SA_INIT requests to prevent VPN access for the client. But an eavesdropper on the path between VPN client and server can redirect a large number of clients to a victim, which is then flooded with IKE_SA_INIT requests. Flooding only happens if many clients initiate IKEv2 exchange at almost the same time, which is considered a rare event. However, this may happen if a Home Agent/VPN server is shutdown for maintenance and all clients need to re-establish VPN connections with another Home Agent/VPN server or if the on-path attacker forces all IPsec security associations to expire by dropping all received IKEv2 messages.

The use of REDIRECTED_FROM payload is intended to discourage a rogue VPN gateway from redirecting a large number of VPN clients to a particular VPN gateway. It does not prevent such a DoS attack.

The redirect mechanism MUST NOT update any state on the client apart from the VPN gateway information. When used with Mobile IPv6, care

must be taken to ensure that the home agent information that the mobile node has configured is not modified wrongly by the redirect message.

Redirecting based on the unauthenticated identities from the client might leak out information about the user when an active attacker, pretending to be a VPN client can get information to which gateway the real user was redirected to. If redirection is based on some internal information of the user, it might leak information to attacker about the user which might not be available otherwise. To prevent these kind of attacks, redirection based on unauthenticated ID should be avoided and should be done only after the client has also authenticated itself.

12. IANA Considerations

This document defines three new IKEv2 Notification Message types as

described in [Section 9](#). The three Notify Message Types must be assigned values between 16396 and 40959.

- o REDIRECT_SUPPORTED
- o REDIRECT
- o REDIRECTED_FROM

This document creates a new namespace called the "Gateway Identity Type". This is used to indicate the type of information regarding the VPN gateway that is carried in the REDIRECT ([Section 9.2](#)) and REDIRECTED_FROM ([Section 9.3](#)) Notification payloads. The following values are assigned.

- 1 - IPv4 address of the new VPN gateway
- 2 - IPv6 address of the new VPN gateway
- 3 - FQDN of the new VPN gateway

Values '0', and 4-240 are reserved. New values can be allocated by Expert Review [\[9\]](#). Values 241-255 are set aside for private use. A specification that extends this registry MUST also mention which of the new values are valid in which Notification payload.

[13.](#) Acknowledgements

The use of anycast address with IKEv2 was first described in [\[7\]](#). It was then added to an early draft version of [RFC 5026](#) and later removed before the RFC was published. Therefore the authors of [\[7\]](#) and [RFC 5026](#) are acknowledged.

Thanks to Pasi Eronen, with whom the solution described in this document was extensively discussed. Thanks to Tero Kivinen for suggesting the use of REDIRECTED_FROM payload and other comments which helped improve the document. The authors would also like to thank Yaron Sheffer, Sunil Kumar, Fan Zhao, Yoav Nir, Richard Graveman, Kanagavel Rajan, Srini Addepalli, Raj Singh, and Arnaud Ebalard for their reviews and comments.

[14.](#) References

14.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", [RFC 4306](#), December 2005.

14.2. Informative References

- [3] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [4] Giarretta, G., Kempf, J., and V. Devarapalli, "Mobile IPv6 Bootstrapping in Split Scenario", [RFC 5026](#), October 2007.
- [5] Haley, B., Devarapalli, V., Deng, H., and J. Kempf, "Mobility Header Home Agent Switch Message", [RFC 5142](#), January 2008.
- [6] Eronen, P. and J. Korhonen, "Multiple Authentication Exchanges in the Internet Key Exchange (IKEv2) Protocol", [RFC 4739](#), November 2006.
- [7] Weniger, K. and F. Dupont, "IKEv2-based Home Agent Assignment in Mobile IPv6/NEMO Bootstrapping", [draft-dupont-ikev2-haassign-02](#) (work in progress), January 2007.
- [8] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [9] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

Authors' Addresses

Vijay Devarapalli
WiChorus

3590 North First St
San Jose, CA 95134
USA

Email: vijay@wichorus.com

Kilian Weniger

Email: kilian.weniger@gmail.com