

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 6, 2009

V. Devarapalli
WiChorus
K. Weniger
Panasonic
February 2, 2009

Re-direct Mechanism for IKEv2
draft-ietf-ipsecme-ikev2-redirect-03.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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 August 6, 2009.

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 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.

Abstract

IKEv2 is a protocol for setting up VPN tunnels from a remote location

Internet-Draft

IKEv2 Re-direct

February 2009

to a gateway so that the VPN client can access services in the network behind the gateway. Currently there is no standard mechanism specified that allows an overloaded VPN gateway or a VPN gateway that is being shut down for maintenance to re-direct the VPN client to attach to another gateway. This document proposes a re-direct mechanism for IKEv2. The proposed mechanism can also be used in Mobile IPv6 to enable the home agent to re-direct the mobile node to another home agent.

Table of Contents

1.	Introduction	3
2.	Terminology	4
3.	IKEv2 Exchange with Redirect	4
4.	Use of Anycast Addresses with the Re-direct Mechanism	5
5.	Gateway Initiated Redirect	6
6.	Redirect Messages	7
6.1.	REDIRECT_SUPPORTED	7
6.2.	REDIRECT	7
6.3.	REDIRECTED_FROM	8
7.	Security Considerations	9
8.	IANA Considerations	10
9.	Acknowledgements	10
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	10
	Authors' Addresses	11

Internet-Draft

IKEv2 Re-direct

February 2009

1. Introduction

IKEv2 [2] is used for setting up IPsec-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 gateways 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 re-direct mechanism for IKEv2 that enables a VPN gateway to re-direct the VPN client to another VPN gateway, for example, based on the load condition. The re-direct is done during the IKE_SA_INIT exchange. The re-direct 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 re-direct can also happen because of administrative or optimal routing reasons. This document does not attempt to provide an exhaustive list of reasons for re-directing a VPN client to another VPN gateway.

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

There is a Home Agent Switch mechanism available for re-directing 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 re-directed to another home agent. The mechanism described in this document can be used for re-directing a mobile node before any state related to the Mobile IPv6 binding is created on the home agent.

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 Exchange with Redirect

To redirect an IKEv2 session to another VPN gateway, the VPN gateway that initially received the IKE_SA_INIT request selects another VPN gateway and responds to the VPN client with a REDIRECT Notification payload. The mechanism by which the initial VPN gateway selects another VPN gateway is out of scope for this document. The IP address of the selected VPN gateway is sent in the REDIRECT payload.

The VPN client indicates support for the IKEv2 redirect mechanism by including a REDIRECT_SUPPORTED notification message in the initial IKE_SA_INIT request. If the IKE_SA_INIT request did not include the REDIRECT_SUPPORTED payload, the responder MUST NOT send the REDIRECT payload to the VPN client.

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, IP_R)
```

When the VPN client receives the IKE_SA_INIT response with the REDIRECT payload, it initiates a new IKE_SA_INIT exchange with the VPN gateway listed in the REDIRECT payload. The VPN client includes the IP address of the original VPN gateway that re-directed the client. The IKEv2 exchange then proceeds as normal with the selected 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}
```

When this mechanism is used with Mobile IPv6, a mobile node's security associations with its home agent may expire while it still has a valid binding cache entry at the home agent. In this case, the mobile node MUST NOT use the original home agent address as the

destination address in the IKE_SA_INIT exchange to setup new security associations. It MUST try to setup security associations with its existing home agent.

4. Use of Anycast Addresses with the Re-direct 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 IKEv2 transaction between the VPN client and the VPN gateway is slightly modified. The VPN client sends the IKE_SA_INIT request 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

(IP_I:500 -> ANYCAST:500)
HDR(A,0), SAi1, KEi, Ni) -->
N(REDIRECT_SUPPORTED)

Responder (any VPN GW)

(ANYCAST:500 -> IP_I:500)
<-- HDR(A,0), N(REDIRECT, IP_R)

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 re-direct the VPN client to a unicast address of one of the VPN gateway. The VPN gateway that received the

IKE_SA_INIT request MAY re-direct 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. Gateway Initiated Redirect

The re-direct mechanism may also be used by a VPN gateway to re-direct the client to another VPN gateway in middle of a session. To re-direct 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. When the client receives this message, it MUST send an empty message as an acknowledgement. Until the client responds with an acknowledgement, the gateway SHOULD re-transmit the re-direct INFORMATIONAL message as described in [2]. The following illustrates the INFORMATIONAL message exchange for gateway-initiated redirect.

```
Initiator (VPN client)           Responder (VPN GW)
-----
                                     <-- HDR, SK {N[REDIRECT, IP_R/FQDN_R]}
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. 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 Messages

6.1. REDIRECT_SUPPORTED

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

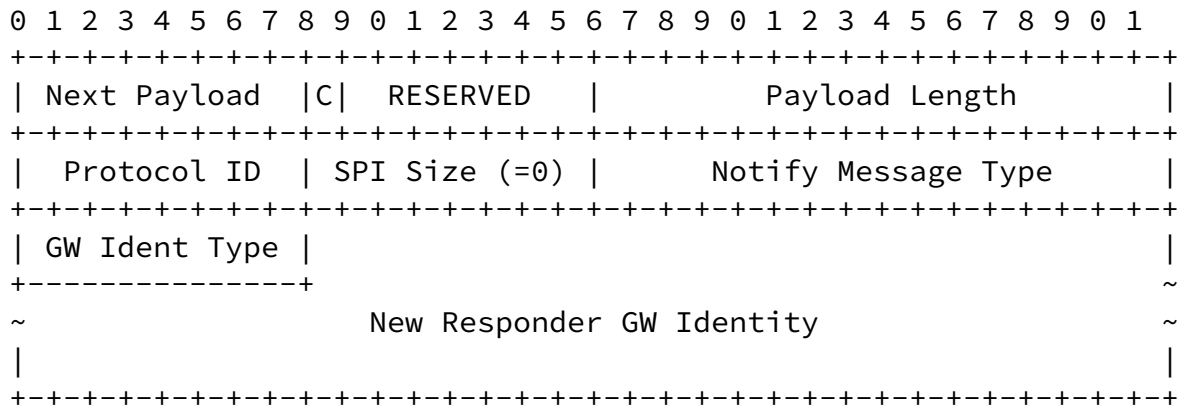


The 'Next Payload', 'Payload Length', 'Protocol ID', 'SPI Size' and the 'Notify Message Type' fields are the same as described in [Section 3.10](#) of [2]. The 'SPI Size' field MUST be set to 0 to indicate that the SPI is not present in this message.

The 'Payload Length' field MUST be set to '8'. The 'Notify Message Type' field is set to indicate the REDIRECT_SUPPORTED payload <value to be assigned by IANA>.

6.2. REDIRECT

The REDIRECT payload is included in an IKE_SA_INIT response from the responder when the responder wants to re-direct the initiator to another VPN gateway. The message includes the new responder's IP address.



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.

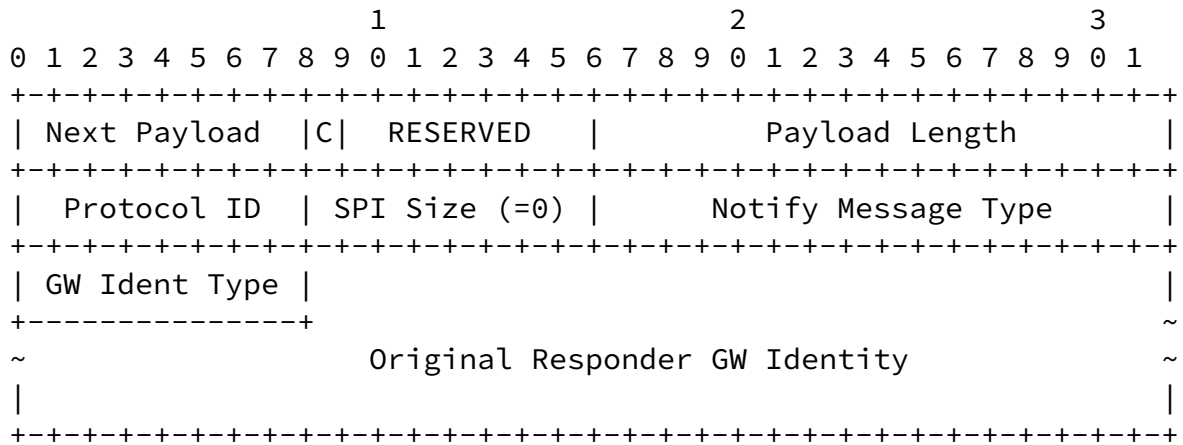
If the IP address of the new VPN gateway is sent, the 'Payload Length' field MUST be set to either '13' or '25' depending on whether an IPv4 or IPv6 address is sent in the message. If the FQDN of the new VPN gateway is sent, the 'Payload Length' field is set to the length of the FQDN plus the fixed fields in the message. The 'Notify Message Type' field is set to indicate the REDIRECT 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 reserved by this document.

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

All other values for this field are reserved and MUST NOT be used. The identity of the new VPN gateway is carried in the 'New Responder GW Identity' field.

6.3. REDIRECTED_FROM

The REDIRECTED_FROM message type is included in the IKE_SA_INIT request from the initiator to the new VPN gateway to indicate the IP address of the original VPN gateway that re-directed the initiator. The original VPN gateway's IP address is included in the message.



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 'Payload Length' field MUST be set to either '13' or '25' depending on whether an IPv4 or IPv6 address of the original VPN gateway is sent in the message. 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 reserved by this document.

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

All other values for this field are reserved and MUST NOT be used. The identity of the original VPN gateway is carried in the 'Original Responder GW Identity' field.

7. 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

Internet-Draft

IKEv2 Re-direct

February 2009

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 re-directing a large number of VPN clients to a particular VPN gateway. It does not prevent such a DoS attack.

8. IANA Considerations

This document defines four new IKEv2 Notification Message types as described in [Section 6](#). The four Notify Message Types must be assigned values between 16396 and 40959.

- o REDIRECT_SUPPORTED
- o REDIRECT
- o REDIRECTED_FROM

9. Acknowledgements

The use of anycast address with IKEv2 was first described in [\[6\]](#). It was then added to an early draft version of [RFC 5026](#) and later removed before the RFC was published. Therefore the authors of [\[6\]](#) 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. The authors would also like to thank Yaron Sheffer, Sunil Kumar, Fan Zhao, Yoav Nir and Arnaud Ebalard for their reviews and comments.

10. References

10.1. 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.

[10.2](#). Informative References

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

Devarapalli & Weniger

Expires August 6, 2009

[Page 10]

Internet-Draft

IKEv2 Re-direct

February 2009

- [4] Giaretta, 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] 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.

Authors' Addresses

Vijay Devarapalli
WiChorus
3590 North First St
San Jose, CA 95134
USA

Email: vijay@wichorus.com

Killian Weniger
Panasonic R&D Center Germany
Monzastr. 4c
Langen 63225
Germany

Email: kilian.weniger@eu.panasonic.com

