

Network Working Group	Y. Nir	
Internet-Draft	Check Point	
Intended status: Informational	April 14, 2010	
Expires: October 16, 2010		

[TOC](#)

IPsec High Availability and Load Sharing Problem Statement draft-ietf-ipsecme-ipsec-ha-01

Abstract

This document describes a requirement from IKE and IPsec to allow for more scalable and available deployments for VPNs. It defines terminology for high availability and load sharing clusters implementing IKE and IPsec, and describes gaps in the existing standards.

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 October 16, 2010.

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.

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.

Table of Contents

- [1. Introduction](#)
 - [1.1. Conventions Used in This Document](#)
 - [2. Terminology](#)
 - [3. The Problem Statement](#)
 - [3.1. Lots of Long Lived State](#)
 - [3.2. IKE Counters](#)
 - [3.3. Outbound SA Counters](#)
 - [3.4. Inbound SA Counters](#)
 - [3.5. Missing Synch Messages](#)
 - [3.6. Simultaneous use of IKE and IPsec SAs by Different Members](#)
 - [3.6.1. Outbound SAs using counter modes](#)
 - [4. Security Considerations](#)
 - [5. Change Log](#)
 - [6. Informative References](#)
 - [§ Author's Address](#)
-

1. Introduction

[TOC](#)

IKEv2, as described in [\[RFC4306\]](#) (Kaufman, C., "Internet Key Exchange (IKEv2) Protocol," December 2005.) and [\[RFC4718\]](#) (Eronen, P. and P. Hoffman, "IKEv2 Clarifications and Implementation Guidelines," October 2006.), and IPsec, as described in [\[RFC4301\]](#) (Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.) and others, allows deployment of VPNs between different sites as well as from VPN clients to protected networks.

As VPNs become increasingly important to the organizations deploying them, there is a demand to make IPsec solutions more scalable and less prone to down time, by using more than one physical gateway to either share the load or back each other up. Similar demands have been made in the past for other critical pieces of an organizations's infrastructure, such as DHCP and DNS servers, web servers, databases and others.

IKE and IPsec are in particular less friendly to clustering than these other protocols, because they store more state, and that state is more

volatile. [Section 2 \(Terminology\)](#) defines terminology for use in this document, and in the envisioned solution documents.

In general, deploying IKE and IPsec in a cluster requires such a large amount of information to be synchronized among the members of the cluster, that it becomes impractical. Alternatively, if less information is synchronized, failover would mean a prolonged and intensive recovery phase, which negates the scalability and availability promises of using clusters. In [Section 3 \(The Problem Statement\)](#) we will describe this in more detail.

1.1. Conventions Used in This Document

[TOC](#)

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. Terminology

[TOC](#)

"Single Gateway" is an implementation of IKE and IPsec enforcing a certain policy, as described in [\[RFC4301\] \(Kent, S. and K. Seo, "Security Architecture for the Internet Protocol," December 2005.\)](#).

"Cluster" is a set of two or more gateways, implementing the same security policy, and protecting the same domain. Clusters exist to provide both high availability through redundancy, and scalability through load sharing.

"Member" is one gateway in a cluster.

"High Availability" is a condition of a system, not a configuration type. A system is said to have high availability if its expected down time is low. High availability can be achieved in various ways, one of which is clustering. All the clusters described in this document achieve high availability.

"Fault Tolerance" is a condition related to high availability, where a system maintains service availability, even when a specified set of fault conditions occur. In clusters, we expect the system to maintain service availability, when one or more of the cluster members fails.

"Completely Transparent Cluster" is a cluster where the occurrence of a fault is never visible to the peers.

"Partially Transparent Cluster" is a cluster where the occurrence of a fault may be visible to the peers.

"Hot Standby Cluster", or "HS Cluster" is a cluster where only one of the members is active at any one time. This member is also referred to as the "active", whereas the others are referred to as "stand-bys".

[\[VRRP\] \(Hinden, R., "Virtual Router Redundancy Protocol \(VRRP\)," April 2004.\)](#) is one method of building such a cluster.

"Load Sharing Cluster", or "LS Cluster" is a cluster where more than one of the members may be active at the same time. The term "load balancing" is also common, but it implies that the load is actually balanced between the members, and we don't want to even imply that this is a requirement.

"Failover" is the event where a one member takes over some load from some other member. In a hot standby cluster, this happens when a standby member becomes active due to a failure of the former active member, or because of an administrator command. In a load sharing cluster this usually happens because of a failure of one of the members, but certain load-balancing technologies may allow a particular load (an SA) to move from one member to another to even out the load, even without any failures.

"Tight Cluster" is a cluster where all the members share an IP address. This could be accomplished using configured interfaces with specialized protocols or hardware, such as VRRP, or through the use of multicast addresses, but in any case, peers need only be configured with one IP address in the PAD.

"Loose Cluster" is a cluster where each member has a different IP address. Peers find the correct member using some method such as DNS queries or [\[REDIRECT\] \(Devarapalli, V. and K. Weniger, "Redirect Mechanism for IKEv2," November 2009.\)](#).

"Synch Channel" is a communications channel among the cluster members, used to transfer state information. The synch channel may or may not be IP based, may or may not be encrypted, and may work over short or long distances. The security and physical characteristics of this channel are out of scope for this document, but it is a requirement that its use be minimized for scalability.

3. The Problem Statement

[TOC](#)

This document will make no attempt to describe the problems in setting up a cluster. The following subsections describe the problems related to the protocol itself.

We also ignore the problem of synchronizing the policy between cluster members, as this is an administrative issue that is not particular to either clusters or to IPsec.

Note that the interesting scenario here is VPN, whether tunneled site-to-site or remote access. host-to-host transport mode is not expected to benefit from this work.

[TOC](#)

3.1. Lots of Long Lived State

IKE and IPsec have a lot of long lived state:

- *IKE SAs last for minutes, hours, or days, and carry keys and other information. Some gateways may carry thousands to hundreds of thousands of IKE SAs.

- *IPsec SAs last for minutes or hours, and carry keys, selectors and other information. Some gateways may carry hundreds of thousands such IPsec SAs.

- *SPD Cache entries. While the SPD is unchanging, the SPD cache changes on the fly due to narrowing. Entries last at least as long as the SAD entries, but tend to last even longer than that.

A naive implementation of a high availability cluster would have no synchronized state, and a failover would produce an effect similar to that of a rebooted gateway. [\[resumption\] \(Sheffer, Y. and H. Tschofenig, "IKEv2 Session Resumption," January 2010.\)](#) describes how new IKE and IPsec SAs can be recreated in such a case.

3.2. IKE Counters

[TOC](#)

We can overcome the first problem described in [Section 3.1 \(Lots of Long Lived State\)](#), by synchronizing states - whenever an SA is created, we can synch this new state to all other members. However, those states are not only long-lived, they are also ever changing.

IKE has message counters. A peer may not process message *n* until after it has processed message *n-1*. Skipping message IDs is not allowed. So a newly-active member needs to know the last message IDs both received and transmitted.

Often, it is feasible to synchronize the IKE message counters for every IKE exchange. This way, the newly active member knows what messages it is allowed to process, and what message IDs to use on IKE requests, so that peers process them.

3.3. Outbound SA Counters

[TOC](#)

ESP and AH have an optional anti-replay feature, where every protected packet carries a counter number. Repeating counter numbers is considered an attack, so the newly-active member must not use a replay counter number that has already been used. The peer will drop those packets as duplicates and/or warn of an attack.

Though it may be feasible to synchronize the IKE message counters, it is almost never feasible to synchronize the IPsec packet counters for every IPsec packet transmitted. So we have to assume that at least for IPsec, the replay counter will not be up-to-date on the newly-active member, and the newly-active member may repeat a counter.

A possible solution is to synch replay counter information, not for each packet emitted, but only at regular intervals, say, every 10,000 packets or every 0.5 seconds. After a failover, the newly-active member advances the counters for outbound SAs by 10,000. To the peer this looks like up to 10,000 packets were lost, but this should be acceptable, as neither ESP nor AH guarantee reliable delivery.

3.4. Inbound SA Counters

[TOC](#)

An even tougher issue, is the synchronization of packet counters for inbound SAs. If a packet arrives at a newly-active member, there is no way to determine whether this packet is a replay or not. The periodic synch does not solve the problem at all, because suppose we synchronize every 10,000 packets, and the last synch before the failover had the counter at 170,000. It is probable, though not certain, that packet number 180,000 has not yet been processed, but if packet 175,000 arrives at the newly-active member, it has no way of determining whether or not that packet has or has not already been processed. The synchronization does prevent the processing of really old packets, such as those with counter number 165,000. Ignoring all counters below 180,000 won't work either, because that's up to 10,000 dropped packets, which may be very noticeable.

The easiest solution is to learn the replay counter from the incoming traffic. This is allowed by the standards, because replay counter verification is an optional feature. The case can even be made that it is relatively secure, because non-attack traffic will reset the counters to what they should be, so an attacker faces the dual challenge of a very narrow window for attack, and the need to time the attack to a failover event. Unless the attacker can actually cause the failover, this would be very difficult. It should be noted, though, that although this solution is acceptable as far as RFC 4301 goes, it is a matter of policy whether this is acceptable.

Another possible solution to the inbound SA problem is to rekey all child SAs following a failover. This may or may not be feasible depending on the implementation and the configuration.

[TOC](#)

3.5. Missing Synch Messages

The synch channel is very likely not to be infallible. Before failover is detected, some synchronization messages may have been missed. For example, the active member may have created a new Child SA using message n. The new information (entry in the SAD and update to counters of the IKE SA) is sent on the synch channel. Still, with every possible technology, the update may be missed before the failover. This is a bad situation, because the IKE SA is doomed. the newly-active member has two problems:

- *It does not have the new IPsec SA pair. It will drop all incoming packets protected with such an SA. This could be fixed by sending some DELETES and INVALID_SPI notifications, if it wasn't for the other problem...

- *The counters for the IKE SA show that only request n-1 has been sent. The next request will get the message ID n, but that will be rejected by the peer. After a sufficient number of retransmissions and rejections, the whole IKE SA with all associated IPsec SAs will get dropped.

The above scenario may be rare enough that it is acceptable that on a configuration with thousands of IKE SAs, a few will need to be recreated from scratch or using session resumption techniques. However, detecting this may take a long time (several minutes) and this negates the goal of creating a high availability cluster in the first place.

3.6. Simultaneous use of IKE and IPsec SAs by Different Members

[TOC](#)

For load sharing clusters, all active members may need to use the same SAs, both IKE and IPsec. This is an even greater problem than in the case of HA, because consecutive packets may need to be sent by different members to the same peer gateway.

The solution to the IKE SA issue is up to the application. It's possible to create some locking mechanism over the synch channel, or else have one member "own" the IKE SA and manage the child SAs for all other members. For IPsec, solutions fall into two broad categories. The first is the "sticky" category, where all communications with a single peer, or all communications involving a certain SPD cache entry go through a single peer. In this case, all packets that match any particular SA go through the same member, so no synchronization of the replay counter needs to be done. Inbound processing is a "sticky" issue, because the packets have to be processed by the correct member based on peer and SPI. Another issue is that commodity load balancers will not be able to match the SPIs of the encrypted side to the clear

traffic, and so the wrong member may get the the other half of the flow.

The other way, is to duplicate the child SAs, and have a pair of IPsec SAs for each active member. Different packets for the same peer go through different members, and get protected using different SAs with the same selectors and matching the same entries in the SPD cache. This has some shortcomings:

- *It requires multiple parallel SAs, which the peer has no use for. Section 2.8 or [\[RFC4306\] \(Kaufman, C., "Internet Key Exchange \(IKEv2\) Protocol," December 2005.\)](#) specifically allows this, but some implementation might have a policy against long term maintenance of redundant SAs.

- *Different packets that belong to the same flow may be protected by different SAs, which may seem "weird" to the peer gateway, especially if it is integrated with some deep inspection middleware such as a firewall. It is not known whether this will cause problems with current gateways. It is also impossible to mandate against this, because the definition of "flow" varies from one implementation to another.

- *Reply packets may arrive with an IPsec SA that is not "matched" to the one used for the outgoing packets. Also, they might arrive at a different member. This problem is beyond the scope of this document and should be solved by the application, perhaps by forwarding misdirected packets to the correct gateway for deep inspection.

3.6.1. Outbound SAs using counter modes

[TOC](#)

For SAs involving counter mode ciphers such as [\[CTR\] \(Housley, R., "Using Advanced Encryption Standard \(AES\) Counter Mode," January 2009.\)](#) or [\[GCM\] \(Viega, J. and D. McGrew, "The Use of Galois/Counter Mode \(GCM\) in IPsec Encapsulating Security Payload \(ESP\)," June 2005.\)](#) there is yet another complication. The initial vector for such modes must never be repeated, and senders use methods such as counters or LFSRs to ensure this. An SA shared between more than one active member, or even failing over from one member to another need to make sure that they do not generate the same initial vector. See [\[COUNTER MODES\] \(McGrew, D. and B. Weis, "Using Counter Modes with Encapsulating Security Payload \(ESP\) and Authentication Header \(AH\) to Protect Group Traffic," March 2010.\)](#) for a discussion of this problem in another context.

4. Security Considerations

[TOC](#)

Implementations running on clusters MUST be as secure as implementations running on single gateways. In other words, no extension or interpretation used to allow operation in a cluster may facilitate attacks that are not possible for single gateways. Moreover, thought must be given to the synching requirements of any protocol extension, to make sure that it does not create an opportunity for denial of service attacks on the cluster.

As mentioned in [Section 3.4 \(Inbound SA Counters\)](#), allowing an inbound child SA to fail over to another member has the effect of disabling replay counter protection for a short time. Though the threat is arguably low, it is a policy decision whether this is acceptable.

5. Change Log

[TOC](#)

This is the first version, re-spun as an WG document

6. Informative References

[TOC](#)

[COUNTER_MODES]	McGrew, D. and B. Weis, " Using Counter Modes with Encapsulating Security Payload (ESP) and Authentication Header (AH) to Protect Group Traffic ," draft-ietf-msec-ipsec-group-counter-modes (work in progress), March 2010 (TXT , HTML).
[CTR]	Housley, R., " Using Advanced Encryption Standard (AES) Counter Mode ," RFC 3686, January 2009 (TXT , HTML , XML).
[GCM]	Viega, J. and D. McGrew, " The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP) ," RFC 4106, June 2005 (TXT , HTML , XML).
[REDIRECT]	Devarapalli, V. and K. Weniger, " Redirect Mechanism for IKEv2 ," RFC 5685, November 2009 (TXT , HTML , XML).
[RFC2119]	Bradner, S., " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[RFC4301]	Kent, S. and K. Seo, " Security Architecture for the Internet Protocol ," RFC 4301, December 2005 (TXT , HTML , XML).
[RFC4306]	Kaufman, C., " Internet Key Exchange (IKEv2) Protocol ," RFC 4306, December 2005 (TXT , HTML , XML).

[RFC4718]	Eronen, P. and P. Hoffman, " IKEv2 Clarifications and Implementation Guidelines ," RFC 4718, October 2006 (TXT , HTML , XML).
[VRRP]	Hinden, R., " Virtual Router Redundancy Protocol (VRRP) ," RFC 3768, April 2004 (TXT , HTML , XML).
[resumption]	Sheffer, Y. and H. Tschofenig, " IKEv2 Session Resumption ," RFC 5723, January 2010 (TXT , HTML , XML).

Author's Address

[TOC](#)

	Yoav Nir
	Check Point Software Technologies Ltd.
	5 Hasolelim st.
	Tel Aviv 67897
	Israel
Email:	ynir@checkpoint.com