IPsec High Availability

Extensions to IKE & IPsec for Support of High Availability and Load Balancing Solutions

Yoav Nir

September 2009

Agenda

- Problem Statement
- Proposed Work Item
- Non Goals
- Goals
- Why this should be a WG item

The Problem

- IPsec services are required to be available all the time. Down time is not allowed
- But the computers used for IPsec cannot be available 100% of the time:
  - Electric and hardware outages
  - Scheduled and unscheduled maintenance
  - OS failures
  - Software bugs

The Problem

- Other services, such as DHCP, DNS and Web solve this by having load sharing and high availability solutions
  - Active / Stand-by solutions
- IKE and IPsec are not friendly towards high availability solutions
- A stand-by gateway going active would need to re-establish all tunnels:
  - Crash discovery
  - Session resumption

The Problem

- Both IKE and IPsec have sequence numbers. If the Stand-by implementation becomes active, it cannot continue, unless it has:
  - Exact IKE Counters, and approximate IPsec replay counters
  - SAD & SPD Cache synchronized
  - Some magical way of causing packets to be routed through it rather than the old gateway.

Proposed Work Item

- Assumptions:
  - Two or more gateways implement the same policy, and protect the same networks.
  - They have a way of passing state data to one another, but the use of this “synch channel” should be minimized.
  - Failover is detected “quickly”
  - Peers see them as a single gateway
    - May be implemented using a multicast address or a DNS name

Proposed Work Item

- Describe protocol extensions and/or best practices to allow implementations to be in high availability configurations
- Describe what data (in RFC 4301 terms) needs to be synchronized between members, and how often.
- Describe requirements from peers.

Non Goals

- We don't want to define protocols for interoperability between members from different vendors.
  - The synch protocol is explicitly out of scope
- We don't intend to describe how the failover is detected
  - Failure to send synch packets?
- How the clusters are set up is out of scope.

Goals

- Allow recovery of IPsec SA if replay counter gets un-synched
- Allow recovery of IKE SA if message counters are un-synched
  - Tolerate some strange DELETEs and INVALID_SPI notifications
  - Reset of message counters?
- Specify what needs to be synched
  - All IKE counters; ESP counters occasionally.

The Goal

Specify protocol extensions and behavior needed for interoperability between cluster implementations of IPsec and any other IPsec implementation.

Why this should be a WG item

- All implementations from “big vendors” have HA configurations
- Interoperability with other vendors has been iffy
  - They try too hard to seem like one gateway
- For acceptance of IPsec, HA solutions should work with no hiccups.
- HA solutions need help from the IKE peers.

Why this should be a WG Item

- The definition of “cluster” varies between vendors:
  - Different working distances
  - Different kinds of load balancers and discovery
  - Different capacities of the synch channel
- Different requirements from peers
- This shouldn't come from a single vendor
- Help wanted...

http://tools.ietf.org/html/draft-nir-ipsecme-ipsecha