pfSense - Feature #8544

Routed IPsec using FreeBSD if_ipsec(4) VTI

05/30/2018 03:53 PM - Jim Pingle

Status:	Resolved	**Start date:**	05/30/2018
Priority:	Normal	**Due date:**	
Assignee:	Jim Pingle	**% Done:**	100%
Category:	IPsec	**Estimated time:**	0.00 hour
Target version:	2.4.4		

Description

Add routed IPsec using if_ipsec(4) VTI (Virtual Tunnel Interfaces) from FreeBSD 11.1 and later with strongSwan.

- Add code to create and manage the interfaces like other interfaces (can assign, setup static routes, specific rules, packet capture, NAT, etc)
- Add a new Phase2 mode for VTI which defines the local and remote ipsec interface endpoints (like gif or tun)
- Add input validation to restrict usage of VTI to supported cases

To me, I have a patch to commit.

Associated revisions

Revision bd4c337c - 05/30/2018 03:53 PM - Jim Pingle

Please welcome routed IPsec using if_ipsec VTI interfaces. Implements #8544

To use, create a P1/P2 and set P2 to VTI using local/remote network as tunnel endpoint addresses, then assign the interface (enable, but IP type = none), and use like any other interface for routing.

Revision 50c4282d - 05/31/2018 08:15 AM - Jim Pingle

Add vpn.inc changes for IPsec VTI that missed the previous commit. Ticket #8544

Revision e8f7e051 - 05/31/2018 08:53 AM - Jim Pingle

A couple vpn.inc refinements for VTI. Ticket #8544

Revision 65767828 - 06/04/2018 01:21 PM - Jim Pingle

IPsec VTI interface refinements/fixes. Ticket #8544

Revision 235c051f - 06/05/2018 04:00 PM - Jim Pingle

Rework how IPsec VTI interfaces and reqid specifications for same are formed. Ticket #8544

Revision aea2a0c3 - 06/06/2018 09:19 AM - Jim Pingle

Fix IPsec VTI gateway generation to match interface changes. Fixes #8544

History

#1 - 05/30/2018 04:10 PM - Jim Pingle

- Status changed from Assigned to Feedback

- % Done changed from 0 to 100

Applied in changeset bd4c337c061f989c4be1bbeaf207447cd8af4989.

#2 - 05/31/2018 09:50 AM - Jim Pingle

Once a new snapshot is up with the later two commits it should be OK for testing.

#3 - 06/03/2018 11:10 AM - Michael OBrien

Jim Pingle wrote:

> Once a new snapshot is up with the later two commits it should be OK for testing.

Just tested with 6/2 9:59:05 CDT snapshot - I'm not sure I 100% understand the instructions, but I did the following on each side:

1. Added P1 IKEV2 with the usual config I'd use for a transport P2 that I was throwing GREs over
2. Added P2, set local network as 192.168.240.1, remote 192.168.240.2 on partner 1, reversed on partner 2
3. Assigned OPT interfaces to the new ipsecX int on both sides, enabled but did nothing else in config (like OpenVPN)
4. Created allow-all firewall rules on each OPT interface for testing

IPSec status shows P1 up, but no traffic over P2 (out or in).

Both gateways show down, and when I try to ping from each box to the other on the 192.168.240 addresses, I get "ping: sendto: Network is down".

Route table looks right (not trying to do any "real" routing yet):

```
192.168.240.1    link#9        UHS      lo0
192.168.240.2    link#9        UH     ipsec1
```

Interfaces are odd, they show netmask of /24... When I tested with 10.200.0.1/2 addresses, the interfaces both showed netmasks of /8. I suspect it should be a /30, but the netmask was greyed out in P2 so I'm not sure how this gets set.

tcpdump doesn't show any traffic on either interface.

Let me know if you'd like further testing or info!

#4 - 06/04/2018 01:10 PM - Jim Pingle

- Status changed from Feedback to Assigned

Reopening as there are some issues with how the tunnel addresses are applied to the interface (local and remote should be in the same subnet, but it may work with both sides as address or /32 as well.

Also interface changes are not applied when IPsec changes are applied, only when the interface itself is saved/applied which is counterintuitive. Should be easy to fix that while I'm at it.

#5 - 06/04/2018 02:28 PM - Jim Pingle

- Status changed from Assigned to Feedback

Changes pushed, next snapshots should be better for testing.

#6 - 06/05/2018 09:11 AM - Jim Pingle

- Status changed from Feedback to Assigned

There is a problem with how the interfaces are numbered, since with more tunnels and phase 2 entries around the ID used for the interface does not match the reqid needed to properly hook into the P2, so some changes are needed in how the interfaces are determined.

#7 - 06/05/2018 09:24 PM - Jim Pingle

Interface numbering is fixed, VTI reqids work as expected and line up between strongswan and ipsecX numbering and uses, and I can pass traffic over interfaces that did not work previously.

The gateways are not being generated properly, and the interface address is not being added to the tonatsubnets macro, so there is still some work to do.

#8 - 06/06/2018 09:30 AM - Jim Pingle

- Status changed from Assigned to Feedback

Applied in changeset aea2a0c333407c0d8b74a51a9dec0829dc78db72.

#9 - 06/08/2018 10:27 AM - Jim Pingle

Another fix in d4b43c48ed1636d3fcd6e47d73ba721bd63d883a

#10 - 07/02/2018 05:09 PM - Aidan Mountford

Howdy,

Similar to Michael Obrien, I tested this on snapshot from 1st of July.

192.168.90.1/30 (Junos) to 192.168.90.2(pfsense).

TCPDUMP on freebsd reports packets from junos being received in good order.

Ping from the pfsense box toward juniper reports "Network is Down"

IKEv2 (P1) is Up.   ESP (P2) is up

Configuration is a described (tunnel endpoints in IPSEC P2 config, interface ipsec2000 configured end up with no address associated)

Any assistance I can provide in helping resolve this, please shout and thanks for providing a much needed feature!

ifconfig:

```
ipsec2000: flags=8111<UP,POINTOPOINT,PROMISC,MULTICAST> metric 0 mtu 1400
inet6 fe80::1c1c:e745:6b76:4db8%ipsec2000 prefixlen 64 tentative scopeid 0x8
inet 192.168.90.2 --> 192.168.90.1 netmask 0xfffffffc
nd6 options=21<PERFORMNUD,AUTO_LINKLOCAL>
reqid: 2000
groups: ipsec
```

Route table:

```
192.168.90.1       link#8            UH     ipsec200
192.168.90.2       link#8            UHS         lo0
```

Traffic on ecn0 from remote side (bgp opens):

```
07:48:57.313859 (authentic,confidential): SPI 0xc92d0d56: IP 192.168.90.1.57271 > 192.168.90.2.179: Flags [S], seq 3717540734, win 16384,
options [mss 9152,sackOK,eol], length 0
07:49:00.514212 (authentic,confidential): SPI 0xc92d0d56: IP 192.168.90.1.57271 > 192.168.90.2.179: Flags [S], seq 3717540734, win 16384,
options [mss 9152,sackOK,eol], length 0
07:49:03.714669 (authentic,confidential): SPI 0xc92d0d56: IP 192.168.90.1.57271 > 192.168.90.2.179: Flags [S], seq 3717540734, win 16384,
options [mss 9152,sackOK,eol], length 0
```

#11 - 07/02/2018 05:10 PM - Aidan Mountford

I should note that ipsec2000 is clipped to ipsec200 in the above

Its also worth noting that on the ifconfig the interface is not identified as "RUNNNING" in the flags.

Many thanks

A

#12 - 07/02/2018 05:13 PM - Jim Pingle

This site isn't good for discussion and diagnosis of that nature, please post on the forum and we can talk about it there. Please include any other related configuration info such as the assigned interface settings.

For the route table output, use netstat -rnW for wide output so the interface name is not cut off.

#13 - 08/02/2018 03:32 PM - Jim Pingle

- Status changed from Feedback to Resolved

The core of this is solid. Any other issues that come up can be handled as separate tickets.