pfSense Packages - Bug #10503

Flapping any GW in multi-WAN influences restating all IPsec tunnels in FRR which leads to dropping all IPsec VTI static routes and related BGP issues

04/28/2020 08:24 AM - Constantine Kormashev

Status:	New	**Start date:**	04/28/2020
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	0%
Category:	FRR	**Estimated time:**	0.00 hour
Target version:			
Affected Version:		**Affected Architecture:**	

Description

There are 2 nodes with a multi-WAN setup: 2 WANs, 2 Gateways. The are 2 IPsec VTI tunnel every working through its own Gateway.
There is a FRR BGP setup with sessions via IPsec VTI tunnels. But both sessions sends and receives updates using loopback interfaces and static routes via IPsec VTI.

```
      +->loopback1-->IPsec VTI1-->WANGW1--v                v--WANGW3<--IPsec VTI3<--loopback3<-+
Node1 |                                   +->the internet<-+                                   | N
ode2
      +->loopback2-->IPsec VTY2-->WANGW2--^                ^--WANGW4<--IPsec VTI4<--loopback4<-+
```

FRR recursively finds Next-Hop for BGP routes via static routes via IPsec. So Node1 can reach routes that are behind Node2 via Node2 loopbacks (loopback3 and loopback4) and vice versa, Node2 can reach Node1 routes via loopback1 and loopback2.
If one of Gateway flapping, even if it is not default Gateway, it seems leading to remove static routes for all IPsec tunnel, due event /rc.newipsecdns and ipsec_reload_package_hook() which executes

```
`function frr_ipsec_reload() {
        require_once('interfaces.inc');
        $vti_ifs = array_keys(interface_ipsec_vti_list_all());
        foreach ($vti_ifs as $vif) {
                mwexec('/usr/local/bin/frrctl cycleinterface ' . escapeshellarg($vif));
        }
}`
```

The interesting thing here is that, existing BGP routes and BGP table entries are not removed from FRR routing table and BGP table, probably because BGP large session timeout. But at the same time these BGP routes are removed from system routing table. And the more interesting, is that, even if static routes via IPsec returned to system routing table and FRR routing table, these BGP routes are not exported back to system routing table by FRR.
On system it looks like:

Static routes through IPsec in FRR table

```
K>* 25.0.0.1/32 [0/0] via 66.0.0.1, 1d01h00m
K>* 26.0.0.1/32 [0/0] via 66.0.1.1, 1d01h00m
```

BGP routes in FRR table

```
B>  10.16.0.0/16 [20/0] via 25.0.0.1 (recursive), 2d05h00m
```

```
  *                         via 66.0.0.1, 2d05h00m
```

FRR BGP entries

```
*  10.16.0.0/16   25.0.0.1              0             50 65501 i
*>                26.0.0.1              0    150     100 65501 i
```

System route table has static routes through IPsec

```
25.0.0.1/32  66.0.0.1  UGS 3750    1400    ipsec3000
26.0.0.1/32  66.0.1.1  UGS 3752    1400    ipsec1000
```

But there are not BGP routes even if they, as we can see, exist in FRR routing table and BGP table. Pay attention on routes uptime. BGP session uptime is the same as BGP routes uptime.

History

#1 - 04/28/2020 08:26 AM - Jim Pingle

- Category set to FRR

#2 - 05/08/2020 07:51 PM - Alhusein Zawi

Working around the issue by splitting FRR from Vti

- Add new VIPs to Local host. (one to each side , do not use the same subnet).

- Use VTI interface to route VIPs between your sites by using static routes. System>Routing>Static Routes.

- Use remote VIP as BGP Neighbors IP Name/Address.

- Use the local VIP as Update Source Services>FRR> BGP> Edit>Neighbors.