strongSwan - Issue #2964

Route to IKE Gateway Fails to Update Under Particular Configuration

11.03.2019 17:12 - Daniel Miess

Status:	Feedback		
Priority:	Normal		
Assignee:			
Category:	libipsec		
Affected version:	5.7.2	**Resolution:**	

Description

Hello,
I am currently running a strongSwan client connecting to a remote server. The connection is established for a full tunnel with a virtual IP address (10.10.10.10) assigned to the client. The tunnel is established using FIPS mode and is making use of the kernel libipsec plugin. I am also configured to use MOBIKE. Using this configuration I can successfully establish the tunnel and pass traffic.

The issue arises however when my active WAN interface (default route) changes. According to all the strongSwan logs strongSwan has detected the WAN switch and as far as it knows it is now operating on the new WAN interface. However there is a route in table 220 which directs traffic to the IKE gateway (x.y.z.129) which hasn't been updated so tunnel traffic continues to use the old WAN interface.

Before the WAN switch this what my table 220 looks like:

```
$ ip r show table 220
default dev ipsec0 proto static src 10.10.10.10
10.254.215.113 dev wan0 proto static src 10.254.215.113
192.168.1.0/24 dev br0 proto static src 192.168.1.1
x.y.z.129 via x.y.z.129 dev wan0 proto static src 10.254.215.113
```

In this case wan0 is the WAN interface as seen in the standard routing table:

```
$ ip r show
default dev wan0 scope link
```

Making the WAN switch wan1 is my new WAN interface:

```
$ ip r
default via 10.1.31.254 dev wan1
```

However my table 220 hasn't changed so all my VPN traffic continues using the wan0 interface

```
$ ip r show table 220
default dev ipsec0 proto static src 10.10.10.10
10.254.215.113 dev wan0 proto static src 10.254.215.113
192.168.1.0/24 dev br0 proto static src 192.168.1.1
x.y.z.129 via x.y.z.129 dev wan0 proto static src 10.254.215.113
```

From the logs I can tell that strongSwan was aware of the WAN switch:

```
Mar 11 16:00:42 info charon: 10[KNL] 10.1.31.111 appeared on wan1
Mar 11 16:00:42 info charon: 16[ENC] <tunnel1|1> generating INFORMATIONAL request 24 [ N(ADD_4_ADD
R) N(ADD_4_ADDR) N(ADD_4_ADDR) ]
Mar 11 16:00:42 info charon: 16[NET] <tunnel1|1> sending packet: from 10.254.215.113[4500] to x.y.
z.129[4500] (112 bytes)
Mar 11 16:00:42 info charon: 10[NET] <tunnel1|1> received packet: from x.y.z.129[4500] to 10.254.2
15.113[4500] (80 bytes)

...
```

```
Mar 11 16:00:44 info charon: 14[ENC] <tunnel1|1> generating INFORMATIONAL request 26 [ ]
Mar 11 16:00:44 info charon: 14[NET] <tunnel1|1> sending packet: from 10.1.31.111[4500] to x.y.z.1
29[4500] (80 bytes)
Mar 11 16:00:44 info charon: 10[NET] <tunnel1|1> received packet: from x.y.z.129[4500] to 10.1.31.
111[4500] (80 bytes)
```

From then on all log messages indicate that the IP address on wan1 is the intended source for VPN traffic. Running the command "swanctl --list-sas" also shows that the IP address on the client side of the connection has changed.

History

#1 - 11.03.2019 17:52 - Daniel Miess

In debuging this I've also found that the issue isn't specific to full tunnels but rather occurs when the IKE gateway is within the remote subnet for the tunnel.

#2 - 12.03.2019 09:48 - Tobias Brunner

- Status changed from New to Feedback

> is making use of the kernel libipsec plugin

Why?

> Before the WAN switch this what my table 220 looks like:
> [...]
>
> In this case wan0 is the WAN interface as seen in the standard routing table:
> [...]
>
> Making the WAN switch wan1 is my new WAN interface:
> [...]

What does that entail exactly? Is the wan0 interface completely gone? As long as a valid route to the remote with that interface is there, the daemon won't change anything (MOBIKE won't be triggered either to use a new source IP). Read the log for details.

> From the logs I can tell that strongSwan was aware of the WAN switch:
> [...]

No, it's only aware that a new interface appeared (at least that's what that part of the log shows, post the complete log if you need me explaining more of it to you). If wan0 is still there with all its routes, the daemon will not change anything. If you did some fiddling with metrics, there is an option (*charon.prefer_best_path*) that forces a switch to the "best" path even if the existing path is still valid.

> From then on all log messages indicate that the IP address on wan1 is the intended source for VPN traffic.

Show us those log messages. Using log level 2 for *knl* will also show more details about the route installation.

> Running the command "swanctl --list-sas" also shows that the IP address on the client side of the connection has changed.

So MOBIKE updated the address?

> In debuging this I've also found that the issue isn't specific to full tunnels but rather occurs when the IKE gateway is within the remote subnet for the tunnel.

What are "full tunnels"? And with *kernel-libipsec* there is an additional route if the peer is within the remote traffic selector to avoid that that traffic is routed to the TUN device or virtual IPs are used for it. Do you mean to say that's the route that's not updated? The other one is?

#3 - 12.03.2019 17:36 - Daniel Miess

- File charon.log added

> What does that entail exactly? Is the wan0 interface completely gone? As long as a valid route to the remote with that interface is there, the

> daemon won't change anything (MOBIKE won't be triggered either to use a new source IP). Read the log for details.

At the beginning there is one WAN interface that all remote traffic exits over. When the "WAN Switches" occurs another interface receives an IP address while the existing one remains up. I change the default route when this new interface comes up so that WAN traffic is now routed out the new interface although the old interface is still there. For example if I had a WAN connection over Wi-Fi but then plugged in an Ethernet cable. I want WAN over Ethernet to take precedence so I make this the new default route.

> Show us those log messages. Using log level 2 for knl will also show more details about the route installation.

I'm attaching full logs with this reply. The WAN switch occurs about half way through.

> So MOBIKE updated the address?

Yes, here is the output of list-sas prior to the WAN switch:

```
$ swanctl --list-sas
tunnel1: #1, ESTABLISHED, IKEv2, 4c386211da4a8b34_i* ddd480a583ead941_r
  local  'XXXXXXXX' @ 10.254.24.125[4500] [10.10.10.10]
  remote 'x.y.z.129' @ x.y.z.129[4500]
  AES_CBC-256/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_2048
  established 40s ago, rekeying in 7090s
  tunnel1: #1, reqid 1, INSTALLED, TUNNEL-in-UDP, ESP:AES_CBC-256/HMAC_SHA2_256_128
    installed 40s ago, rekeying in 6601s, expires in 7880s
    in  28735fe6,      0 bytes,     0 packets
    out c5d55f2e,    630 bytes,    10 packets,    14s ago
    local  10.10.10.10/32
    remote 0.0.0.0/0
```

And here is the output of list-sas after the WAN switch:

```
$ swanctl --list-sas
tunnel1: #1, ESTABLISHED, IKEv2, 4c386211da4a8b34_i* ddd480a583ead941_r
  local  'XXXXXXXX' @ 10.1.31.111[4500] [10.10.10.10]
  remote 'x.y.z.129' @ x.y.z.129[4500]
  AES_CBC-256/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_2048
  established 316s ago, rekeying in 6814s
  tunnel1: #2, reqid 1, INSTALLED, TUNNEL-in-UDP, ESP:AES_CBC-256/HMAC_SHA2_256_128/MODP_2048
    installed 15s ago, rekeying in 6626s, expires in 7906s
    in  f0812520,      0 bytes,     0 packets
    out c224176e,    126 bytes,     2 packets,     0s ago
    local  10.10.10.10/32
    remote 0.0.0.0/0
```

> What are "full tunnels"? And with kernel-libipsec there is an additional route if the peer is within the remote traffic selector to avoid that that traffic is routed to the TUN device or virtual IPs are used for it. Do you mean to say that's the route that's not updated? The other one is?

When I say full tunnel I'm referring to a tunnel with a remote subnet of 0.0.0.0/0. And yes, the route in table 220 that directs traffic for the IKE gateway over the WAN interface rather than the ipsec0 interface is set correctly on the first connection but when the new interface comes up it still continues to point to the "old" WAN interface. The default route in the main routing table is updated because that's something that I do myself.

#4 - 12.03.2019 18:03 - Tobias Brunner

> > is making use of the kernel libipsec plugin
>
> Why?

You didn't answer this. Because, really, if there is no absolutely compelling reason to use it, don't!

> And yes, the route in table 220 that directs traffic for the IKE gateway over the WAN interface rather than the ipsec0 interface is set correctly on the first connection but when the new interface comes up it still continues to point to the "old" WAN interface.

OK, it looks like this is not supported. As long as the route for IPsec traffic doesn't change, the route that excludes IKE traffic is not touched. I guess that could be improved, then again, you can avoid that route by configuring what's described here (the *allow_peer_ts* option disables the installation of the exclusion route).

#5 - 12.03.2019 22:34 - Daniel Miess

> You didn't answer this. Because, really, if there is no absolutely compelling reason to use it, don't!

Unfortunately this has to do with the way I'm using FIPS and at least short term I can't switch away from using it this way.

> OK, it looks like this is not supported. As long as the route for IPsec traffic doesn't change, the route that excludes IKE traffic is not touched. I guess that could be improved, then again, you can avoid that route by configuring what's described here (the allow_peer_ts option disables the installation of the exclusion route).

Thanks for the link. I tried the steps described under Configuration and Host-to-Host tunnels and was able to get my use case working. This ticket can be closed now but I would ask that you consider refreshing of the IKE gateway route on a MOBIKE change for future releases.

#6 - 13.03.2019 10:38 - Tobias Brunner

> > You didn't answer this. Because, really, if there is no absolutely compelling reason to use it, don't!
>
> Unfortunately this has to do with the way I'm using FIPS and at least short term I can't switch away from using it this way.

You can't use the Linux kernel in FIPS mode?

Files

charon.log	81.4 KB	12.03.2019	Daniel Miess