strongSwan - Bug #1347

Route won't be created for passthrough with subnet other than /24

11.03.2016 16:22 - Arnaud G.

Status:	Closed	**Start date:**	11.03.2016
Priority:	Normal	**Due date:**	
Assignee:	Tobias Brunner	**Estimated time:**	0.00 hour
Category:	libcharon		
Target version:	5.5.0		
Affected version:	5.3.5	**Resolution:**	Fixed

Description

Hi,

My setup is (site to site, strongswan 5.3.5):
eth0 (internet) 10.10.10.5/24
eth1 (lan) 192.168.32.0/24

The passtrough connection in ipsec.conf:

```
conn local-net
        leftsubnet=192.168.32.0/24
        rightsubnet=192.168.32.0/24,172.17.72.0/22
        authby=never
        type=pass
        auto=route
```

When I start ipsec, the route for 192.168.32.0/24 is correctly created but **not** the route for 172.17.72.0/22:

```
# ip route show table 220
default via 10.10.10.1 dev eth0  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
```

After, I change the mask to /24:

```
conn local-net
        leftsubnet=192.168.32.0/24
        rightsubnet=192.168.32.0/24,172.17.72.0/24
        authby=never
        type=pass
        auto=route
```

When I start ipsec, the route for 192.168.32.0/24 **and** 172.17.72.0/24 are correctly created:

```
# ip route show table 220
default via 10.10.10.1 dev eth0  proto static  src 192.168.32.5  mtu 1446
172.17.72.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
```

If I add more /24 subnet, no problem... but if I try with /22, /16 or any other mask, the route wasn't created.

For trying to understand, I've added more loglevel for knl:

```
daemon {
                knl = 2
        }
```

The log for the first subnet (OK):

```
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 o
ut  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 i
n  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 f
wd  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] getting a local address in traffic selector 192.168
.32.0/24
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] using host 192.168.32.5
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] no nexthop found to reach 192.168.32.0/24
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] 192.168.32.5 is on interface eth1
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] installing route: 192.168.32.0/24 via (null) src 19
2.168.32.5 dev eth1
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] getting iface index for eth1
```

And the log for second subnet (NOT OK):

```
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 192.168.32.0/24 === 172.17.72.0/22 ou
t  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 172.17.72.0/22 === 192.168.32.0/24 in
  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] adding policy 172.17.72.0/22 === 192.168.32.0/24 fw
d  (mark 0/0x00000000)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] getting a local address in traffic selector 192.168
.32.0/24
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] using host 192.168.32.5
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] using 10.10.10.1 as nexthop to reach 172.17.72.0/22
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] 192.168.32.5 is on interface eth1
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] installing route: 172.17.72.0/22 via 10.10.10.1 src
 192.168.32.5 dev eth1
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] getting iface index for eth1
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] received netlink error: Network is unreachable (101
)
Mar 11 15:59:12 prodvpn2-nimes charon: 06[KNL] unable to install source route for 192.168.32.5
```

The main difference is that for this second subnet, he's finding a nexthop (10.10.10.1). But... it's wrong, and kernel don't accept to add a rule with a gateway on wrong interface.

Of course, I can split my /22 in 4 /24... but It's only a workaround.

Thanks in advance for your opinion on this issue.

Related issues:

Related to Bug #824: kernel_netlink plugin decides on wrong interface for route	**Closed**	**17.01.2015**

Associated revisions

Revision 0ed9430d - 21.03.2016 12:14 - Tobias Brunner

kernel-netlink: Fix lookup of next hops for destinations with prefix

References #1347.

Revision 96b1fab5 - 10.06.2016 18:15 - Tobias Brunner

Merge branch 'interface-for-routes'

Changes how the interface for routes installed with policies is
determined. In most cases we now use the interface over which we reach the
other peer, not the interface on which the local address (or the source IP) is
installed. However, that might be the same interface depending on the
configuration (i.e. in practice there will often not be a change).

Routes are not installed anymore for drop policies and for policies with
protocol/port selectors.

Fixes #809, #824, #1347.

History

#1 - 11.03.2016 16:28 - Tobias Brunner

- Status changed from New to Feedback

How does your default routing table look like? (Or any other tables that are in use).

#2 - 11.03.2016 16:32 - Arnaud G.

My routing table:

```
0.0.0.0         10.10.10.1      0.0.0.0         UG    0      0        0 eth0
192.168.32.0    0.0.0.0         255.255.255.0   U     0      0        0 eth1
10.10.10.0      0.0.0.0         255.255.255.248 U     0      0        0 eth0
```

#3 - 11.03.2016 16:46 - Tobias Brunner

> he's finding a nexthop (10.10.10.1). But... it's wrong, and kernel don't accept to add a rule with a gateway on wrong interface.

Why do you think that's wrong? Your default route has that next hop set and that's the route the kernel interface assumes is used to reach hosts in the 172.17.72.0/22 subnet. However, when configuring 172.17.72.0/24 the default route should get used too so I wonder what the log output looks like in that case. Could you post that.

Could you also add the output of ip route list table all.

#4 - 11.03.2016 17:10 - Arnaud G.

With 172.17.72.0/24 it's not wrong:

```
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 out  (mark 0/
0x00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 in  (mark 0/0
x00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 192.168.32.0/24 === 192.168.32.0/24 fwd  (mark 0/
0x00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] getting a local address in traffic selector 192.168.32.0/24
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] using host 192.168.32.5
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] no nexthop found to reach 192.168.32.0/24
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] 192.168.32.5 is on interface eth1
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] installing route: 192.168.32.0/24 via (null) src 192.168.32.5 d
ev eth1
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] getting iface index for eth1
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 192.168.32.0/24 === 172.17.72.0/24 out  (mark 0/0
x00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 172.17.72.0/24 === 192.168.32.0/24 in  (mark 0/0x
00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] adding policy 172.17.72.0/24 === 192.168.32.0/24 fwd  (mark 0/0
x00000000)
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] getting a local address in traffic selector 192.168.32.0/24
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] using host 192.168.32.5
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] no nexthop found to reach 172.17.72.0/24
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] 192.168.32.5 is on interface eth1
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] installing route: 172.17.72.0/24 via (null) src 192.168.32.5 de
v eth1
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] getting iface index for eth1
```

And of course, route is fine:

```
# ip route show table 220
default via 10.10.10.1 dev eth0  proto static  src 192.168.32.5  mtu 1446
172.17.72.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
```

The only think changed in this configuration and the configuration before is the mask /22 -> /24.

```
# ip route list table all
default via 10.10.10.1 dev eth0  table 220  proto static  src 192.168.32.5  mtu 1446
172.17.72.0/24 dev eth1  table 220  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  table 220  proto static  src 192.168.32.5  mtu 1446
```

```
default via 10.10.10.1 dev eth0
192.168.32.0/24 dev eth1  proto kernel  scope link  src 192.168.32.5
10.10.10.0/24 dev eth0  proto kernel  scope link  src 10.10.10.5
broadcast 127.0.0.0 dev lo  table local  proto kernel  scope link  src 127.0.0.1
local 127.0.0.0/8 dev lo  table local  proto kernel  scope host  src 127.0.0.1
local 127.0.0.1 dev lo  table local  proto kernel  scope host  src 127.0.0.1
broadcast 127.255.255.255 dev lo  table local  proto kernel  scope link  src 127.0.0.1
broadcast 192.168.32.0 dev eth1  table local  proto kernel  scope link  src 192.168.32.5
local 192.168.32.5 dev eth1  table local  proto kernel  scope host  src 192.168.32.5
broadcast 192.168.32.255 dev eth1  table local  proto kernel  scope link  src 192.168.32.5
broadcast 10.10.10.0 dev eth0  table local  proto kernel  scope link  src 10.10.10.5
local 10.10.10.5 dev eth0  table local  proto kernel  scope host  src 10.10.10.5
broadcast 10.10.10.255 dev eth0  table local  proto kernel  scope link  src 10.10.10.5
```

If I had a third /24 subnet, it's OK:

```
conn local-net
        leftsubnet=192.168.32.0/24
        rightsubnet=192.168.32.0/24,172.17.72.0/24,172.18.72.0/24
        authby=never
        type=pass
        auto=route
```

All the needed routes are here:

```
# ip route show table 220
default via 10.10.10.1 dev eth0  proto static  src 192.168.32.5  mtu 1446
172.17.72.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
172.18.72.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
```

But if this third subnet mask was set to /22, it's not OK:

```
conn local-net
        leftsubnet=192.168.32.0/24
        rightsubnet=192.168.32.0/24,172.17.72.0/24,172.18.72.0/22
        authby=never
        type=pass
        auto=route
```

The route for 172.18.72.0/22 is missing:

```
# ip route show table 220
default via 10.10.10.1 dev eth0  proto static  src 192.168.32.5  mtu 1446
172.17.72.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
192.168.32.0/24 dev eth1  proto static  src 192.168.32.5  mtu 1446
```

#5 - 11.03.2016 19:31 - Tobias Brunner

- Tracker changed from Issue to Bug

- Category set to libcharon

- Assignee set to Tobias Brunner

```
Mar 11 16:51:37 prodvpn2-nimes charon: 08[KNL] no nexthop found to reach 172.17.72.0/24
```

That's strange. The same default route should have been found that was found while attempting to install a route for the /22 subnet.

I found the reason for the above anomaly, which was a small bug in the *kernel-netlink* plugin (it only worked correctly if the source address matched the same route as the destination subnet for which the lookup was done). I pushed a fix for it to the *1347-nexthop-prefix* branch.

However, this fix now breaks the installation of the /24 route too. The problem with the route is not the next hop but the outbound interface. It should be eth0, not eth1 (which the code currently determines via the found source address, i.e. 192.168.32.5 in your case, which is installed on eth1). Similar problems have been reported in #824 and #809.

I tried to solve this with the idea proposed in #824-1, that is, return the outbound interface in the lookup for the next hop and use that instead of the interface on which the local address is installed. I pushed some additional commits that implement this to the mentioned branch. They seem to fix the issue in my tests.

#6 - 11.03.2016 19:31 - Tobias Brunner

- Related to Bug #824: kernel_netlink plugin decides on wrong interface for route added

#7 - 14.03.2016 14:23 - Arnaud G.

Thanks you for your research.
Is it possible to test the patch ?
Have you planned to integrate this in future 5.4.0 ?

#8 - 15.03.2016 10:31 - Tobias Brunner

> Is it possible to test the patch ?

Sure, use the code (or only the patches) from the said branch.

> Have you planned to integrate this in future 5.4.0 ?

Not sure, we already released the RC and plan to do the release next week.

#9 - 10.06.2016 11:07 - Arnaud G.

Hi,
Sorry to ask the same question, have you planned to integrate the patch in version 5.4.1 ?
This bug is annoying for us and using a patched version is not possible.

Thanks in advance.

#10 - 10.06.2016 18:37 - Tobias Brunner

- Status changed from Feedback to Closed

- Target version set to 5.5.0

- Resolution set to Fixed