strongSwan - Feature #380

Allow /32 remote subnet with kernel-libipsec plugin

12.08.2013 16:13 - Mikael Magnusson

Status:	Closed	**Start date:**	12.08.2013
Priority:	Normal	**Due date:**	
Assignee:	Tobias Brunner	**Estimated time:**	0.00 hour
Category:	charon		
Target version:	5.1.1		
Resolution:	Fixed		

Description

When using the new kernel-libipsec plugin on a openvz vps I run into problems with /32 remote subnets being unsupported. It can be solved by first disabling the checks in kernel_libipsec_ipsec.c.

```
diff --git a/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c b/src
index 40f253d..5e2e1dc 100644
--- a/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c
+++ b/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c
@@ -464,6 +464,7 @@ static bool install_route(private_kernel_libipsec_ipsec_t *t
                policy->route = NULL;
        }

+#if 0
        if (dst_ts->is_host(dst_ts, dst))
        {
                DBG1(DBG_KNL, "can't install route for %R === %R %N, conflicts w
@@ -479,6 +480,7 @@ static bool install_route(private_kernel_libipsec_ipsec_t *t
                /* add exclude route for peer */
                add_exclude_route(this, route, src, dst);
        }
+#endif

        DBG2(DBG_KNL, "installing route: %R src %H dev %s",
                 dst_ts, route->src_ip, route->if_name);
```

Then I changed the ip rule to only have effect if a mark is set:

```
policy from all fwmark 0x4/0x4 lookup 220
```

On this server I only marked packets in the mangle OUTPUT chain since it isn't a VPN gateway/router. (A gateway also need to mark packets on PREROUTING.)

```
Chain OUTPUT (policy ACCEPT 309 packets, 38910 bytes)
 pkts bytes target     prot opt in     out     source               destination
  271 38404 MARK      !udp  --  *      *       0.0.0.0/0            0.0.0.0/0           MARK set
0x4
   81  5390 MARK       udp  --  *      *       0.0.0.0/0            0.0.0.0/0           udp dpt:!
4500 MARK set 0x4
```

Related issues:

Has duplicate Issue #403: v4 tunnel host-to-host "can't install route for 192...	**Rejected**	**07.09.2013**

Associated revisions

Revision 1ff63f15 - 11.10.2013 15:33 - Tobias Brunner

Merge branch 'fwmarks'

Allows setting a mark on outbound packets and the routing rule
installed by charon. With those settings it is possible to setup
tunnels with kernel-libipsec where the remote peer is part of the remote
traffic selector.

The following example settings in strongswan.conf show how this can be
configured:

```
charon {
plugins {
kernel-netlink {
fwmark = !0x42
}
socket-default {
fwmark = 0x42
}
kernel-libipsec {
allow_peer_ts = yes
}
}
}
```

To make it work it is necessary to set

```
net.ipv4.conf.all.rp_filter
```

appropriately, otherwise the kernel drops the packets.

References #380.

History

#1 - 12.08.2013 17:41 - Mikael Magnusson

Btw, the reverse path filter also needs to be disabled.

echo 0 > /proc/sys/net/ipv4/conf/all/rp_filter

#2 - 12.08.2013 18:48 - Tobias Brunner

- Description updated

- Status changed from New to Assigned

- Assignee set to Tobias Brunner

#3 - 14.08.2013 11:04 - Tobias Brunner

- Status changed from Assigned to Feedback

Interesting idea, thanks. I pushed some patches to the *fwmarks* branch to make this configurable without the need for additional *iptables* rules.

With the following entries in strongswan.conf they implement basically the same thing:

```
charon {
    plugins {
        kernel-netlink {
            fwmark = !0x4
        }
        socket-default {
            fwmark = 0x4
        }
        kernel-libipsec {
            allow_peer_ts = yes
        }
    }
}
```

The first setting configures the rule for table 220 to apply only if packets are **not** marked with *0x4*. The second applies a mark of *0x4* to all packets sent by charon (i.e. the routes in table 220 wont apply to them). Finally, the third setting disables the check you patched out.

When I tested it I actually had to set *rp_filter* to **2**, it didn't work when it was set to **0**.

#4 - 14.08.2013 23:27 - Mikael Magnusson

I have made a diff from the branch and applied to version 5.1.0 which I tested. Found two problems.

1. It didn't add any route for the /32 rightsubnet I had configured since it had an exclude route for the IKE peer. New patch below.
2. I also couldn't get it to work with the ip rule inserted by strongswan (containing a negative condition), it only resulted in “Destination Host Unreachable”.

```
diff --git a/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c b/src
index 686b755..af71a7f 100644
--- a/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c
+++ b/src/libcharon/plugins/kernel_libipsec/kernel_libipsec_ipsec.c
```

```
@@ -479,7 +479,7 @@ static bool install_route(private_kernel_libipsec_ipsec_t *t
                return FALSE;
        }
        /* if remote traffic selector covers the IKE peer, add an exclude route
-       if (dst_ts->includes(dst_ts, dst))
+       if (!this->allow_peer_ts && dst_ts->includes(dst_ts, dst))
        {
                /* add exclude route for peer */
                add_exclude_route(this, route, src, dst);
```

#5 - 15.08.2013 09:42 - Tobias Brunner

> It didn't add any route for the /32 rightsubnet I had configured since it had an exclude route for the IKE peer. New patch below.

I fixed the commit, thanks. It worked in my tests because I used a virtual IP, which caused the proper route to end up being installed.

> I also couldn't get it to work with the ip rule inserted by strongswan (containing a negative condition), it only resulted in “Destination Host Unreachable”.

I can't reproduce this, even for direct host-to-host tunnels (i.e. without virtual IPs). Do you get that for IKE or tunneled traffic? Have you remove the *iptables* rules you installed previously?

How does your setup look like exactly (config)? Is your OpenVZ host initiator or responder? What do you use on the other end? Do you use virtual IPs?

#6 - 17.08.2013 13:28 - Mikael Magnusson

It is working now with the fwmark settings in strongswan.conf, I don't know why it wasn't before.

(Setting marks in the mangle OUTPUT chain didn't work well because the output interface [from the rule based routing decision] wasn't updated before handled in the filter OUTPUT chain, which is the reason I tried the fwmark settings again.)

#7 - 18.08.2013 01:15 - Francesco Frassinelli

I have the same issue.

#8 - 19.08.2013 13:38 - Mikael Magnusson

I noticed another problem, for some reason the route to ipsec0 added by strongswan is removed. Maybe during renegotiations, I am not sure yet.

Update: Yes, the route in table 220 is in fact deleted during the rekeying, which of course means nothing will be sent over the connection any longer.

#9 - 11.10.2013 15:44 - Tobias Brunner

- Tracker changed from Issue to Feature

- Category set to charon

- Status changed from Feedback to Closed

- Target version set to 5.1.1

- Resolution set to Fixed

> Update: Yes, the route in table 220 is in fact deleted during the rekeying, which of course means nothing will be sent over the connection any longer.

Sorry for the delay. I merged the *fwmarks* branch together with a fix for the rekeying issue to master.