Subject: VE not reachable via ipsec-Tunnel using openswan and bridge-devices
Posted by dtombeil on Mon, 28 Jun 2010 13:24:21 GMT
View Forum Message <> Reply to Message

Hi,

I have problems with the following setup:

The scenario:

- I have a VZ-Server based on Debian Squeeze AMD64 using the latest OpenVZ Kernel from Debian unstable 2.6.32-15. The system's other packages are up2date.

- The system has two ethernet-devices eth0 (external) and eth1, which is a phys. interface for a .1q-Trunk.

- I have two bridge-devices br0 and br1 with own IP-subnet and no attached phys. ethernet device to form "virtual dmz" on the host. The traffic is routed between the networks.

- There is a VE attached to every bridge device. It uses veth as network subsystem.

- I use Openswan 1:2.6.26+dfsg-1 for ipsec tunnels

The Problem:

I can access the VE from the LAN attached to e.g. eth0 or eth1.100 (VLAN) without any problem. I can also ping from one VE to the other or to hosts on the LAN.

I can use the VPN-Tunnel to ping hosts on the phys. LAN and I can also ping the host's IP-addresses on the bridge-device. But I _cannot_ ping the VE's IP itself using the ipsec tunnel. I can see the packages travelling to br1 in tcpdump, but the VE does not answer. I can also see the packages _inside_ the VE using tcpdump on eth0 but the VE does not answer.

The most strange thing is, if I ping back from the VE to the VPN-Client IP-Address I can see bidirectional traffic on br1 using tcpdump but the ping-command inside the VE does not get any packet back. The VPN-Client is 192.168.10.1 the VE has 172.16.231.129. This is what I see in tcpdump:

15:06:30.496483 IP 172.16.231.129 > 192.168.10.1: ICMP echo request, id 318, seq 10, length 64
15:06:30.498103 IP 192.168.10.1 > 172.16.231.129: ICMP echo reply, id 318, seq 10, length 64
15:06:31.504440 IP 172.16.231.129 > 192.168.10.1: ICMP echo request, id 318, seq 11, length 64
15:06:31.507335 IP 192.168.10.1 > 172.16.231.129: ICMP echo reply, id 318, seq 11, length 64
15:06:32.512414 IP 172.16.231.129 > 192.168.10.1: ICMP echo request, id 318, seq 12, length 64
15:06:32.532765 IP 192.168.10.1 > 172.16.231.129: ICMP echo reply, id 318, seq 12, length 64
15:06:33.520455 IP 172.16.231.129 > 192.168.10.1: ICMP echo request, id 318, seq 13, length

64
15:06:33.524663 IP 192.168.10.1 > 172.16.231.129: ICMP echo reply, id 318, seq 13, length 64
15:06:34.528431 IP 172.16.231.129 > 192.168.10.1: ICMP echo request, id 318, seq 14, length 64
15:06:34.530911 IP 192.168.10.1 > 172.16.231.129: ICMP echo reply, id 318, seq 14, length 64

And this what the ping shows if stopped after a while:

root@proxy:/# ping 192.168.10.1
PING 192.168.10.1 (192.168.10.1) 56(84) bytes of data.
^C
--- 192.168.10.1 ping statistics ---
84 packets transmitted, 0 received, 100% packet loss, time 83663ms

Here comes the network-debug-output:

Routing-Table inside the proxy-VE:

root@proxy:/# ip route list table all
172.16.231.128/25 dev eth0 proto kernel scope link src 172.16.231.129
default via 172.16.231.254 dev eth0
local 172.16.231.129 dev eth0 table local proto kernel scope host src 172.16.231.129
broadcast 172.16.231.128 dev eth0 table local proto kernel scope link src 172.16.231.129
broadcast 127.255.255.255 dev lo table local proto kernel scope link src 127.0.0.1
broadcast 172.16.231.255 dev eth0 table local proto kernel scope link src 172.16.231.129
broadcast 127.0.0.0 dev lo table local proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table local proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table local proto kernel scope host src 127.0.0.1
fe80::/64 dev eth0 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::218:51ff:febd:fe1d via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
ff00::/8 dev eth0 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255

Routing-Table on the Host-System:

root@vzhost01:~# ip route list table all
213.178.168.248/29 dev eth1.100 proto kernel scope link src 213.178.168.253
212.9.191.0/25 dev eth0 proto kernel scope link src 212.9.191.121
172.16.231.128/25 dev br1 proto kernel scope link src 172.16.231.254
172.16.231.0/25 dev br0 proto kernel scope link src 172.16.231.126
default via 212.9.191.1 dev eth0
broadcast 212.9.191.127 dev eth0 table local proto kernel scope link src 212.9.191.121
broadcast 172.16.231.128 dev br1 table local proto kernel scope link src 172.16.231.254
broadcast 127.255.255.255 dev lo table local proto kernel scope link src 127.0.0.1

broadcast 213.178.168.248 dev eth1.100 table local proto kernel scope link src 213.178.168.253
broadcast 172.16.231.0 dev br0 table local proto kernel scope link src 172.16.231.126
broadcast 213.178.168.255 dev eth1.100 table local proto kernel scope link src 213.178.168.253
local 212.9.191.121 dev eth0 table local proto kernel scope host src 212.9.191.121
local 213.178.168.253 dev eth1.100 table local proto kernel scope host src 213.178.168.253
local 172.16.231.126 dev br0 table local proto kernel scope host src 172.16.231.126
broadcast 172.16.231.127 dev br0 table local proto kernel scope link src 172.16.231.126
local 172.16.231.254 dev br1 table local proto kernel scope host src 172.16.231.254
broadcast 212.9.191.0 dev eth0 table local proto kernel scope link src 212.9.191.121
broadcast 172.16.231.255 dev br1 table local proto kernel scope link src 172.16.231.254
broadcast 127.0.0.0 dev lo table local proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table local proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table local proto kernel scope host src 127.0.0.1
fe80::/64 dev eth1 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev eth1.100 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev br1 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev br0 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev eth0 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev veth100.0 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
fe80::/64 dev veth200.0 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::218:51ff:fe7f:1f38 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::218:51ff:fe86:1506 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::5054:ff:fe91:f85d via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::5054:ff:fe91:f85d via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::5054:ff:fedd:5e72 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::ac33:a8ff:fe5b:a9e9 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
local fe80::cc5e:d0ff:fe76:5956 via :: dev lo table local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295
ff00::/8 dev eth1 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev eth1.100 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev br1 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev br0 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev eth0 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev veth100.0 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
ff00::/8 dev veth200.0 table local metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
unreachable default dev lo table unspec proto kernel metric -1 error -101 hoplimit 255

I am sure, that there are no iptables-filters active. Here comes the dump:

```
root@vzhost01:~# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination
Chain INPUT (policy ACCEPT)
target     prot opt source               destination

Chain FORWARD (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination
Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

Chain INPUT (policy ACCEPT)
target     prot opt source               destination

Chain FORWARD (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source               destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source               destination
```

Can anybody help?

Subject: Re: VE not reachable via ipsec-Tunnel using openswan and bridge-devices
Posted by buzzer on Mon, 18 Apr 2011 17:26:30 GMT
View Forum Message <> Reply to Message

Hi,

I'm experiencing the same problem with strongswan instead openswan on debian squeeze. (all up to date)
I would like to know if you find a solution to this problem ?

thanks in advance.

I can produce config snapshot if needed.

Subject: Re: VE not reachable via ipsec-Tunnel using openswan and bridge-devices
Posted by JohnDoe on Mon, 04 Jul 2011 14:28:30 GMT
View Forum Message <> Reply to Message

Hi all,

same problem on racoon in VETH-OpenVZ-VM.

No solutions yet?

Subject: Re: VE not reachable via ipsec-Tunnel using openswan and bridge-devices
Posted by JohnDoe on Mon, 04 Jul 2011 14:58:35 GMT
View Forum Message <> Reply to Message

Hi @all,

I got the same problem with IPSec on Debian using RACOON.

I can establish the tunnel with racoon. the VZ host can ping the VPN client and VPN client the host's private IP, but the VPSs "behind" the host get the ICMP packets on their interface, but it is ignored or dropped ?!

e.g., when I ping from VPS 10.0.2.123 to the VPN client 192.168.100.101, I see the packet on the VPS eth0, then on the host's bridge vmbr1, then on the VPN client with Wireshark. Then there is a ICMP reply generated. I can trace the reply on the host vmbr1 and on/in the VPS 10.0.2.123 eth0 interface, but then it got ignored or dropped.

No suggestions or solutions?

Best Regards,
JD.