
Subject: Re: networking Problem venet / veth
Posted by [dasicebaer](#) on Thu, 15 Feb 2007 10:46:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Vasily Tarasov wrote on Tue, 13 February 2007 10:42Hello,

First of all I want to say, that veth is less secure only in untrusted environments. The thing is that many people use OpenVZ for hosting: they sell VEs to various people. In such case this is rather insecure to give veth device to VE, because it will be able to create any ethernet packets, can be sniffed and etc.

Well, this is why I don't want to use veth for the dmzserv. From a security aspect, it's not a good idea to have dmzserv and lanserv running on the same machine at all, since the dmz's only purpose is to confine an attacker in it's own lansegment, where he can't do any harm to the internal network. I hope to achieve this confinement through some firewallrules, but if the attacker could simply change his ip to move the dmzserv from the dmz to the inner lan, that would render the dmz pretty useless.

Quote:Secondly, why venet0 in lanserv bothers you? Let it be there, but not configured. It will not prevent veth somehow!

Okay. First of all, I recognized the debianpackage of vzutils from etch did not contain /usr/sbin/vznetcfg. After reinstalling from source, I did not experience any more lags where I had to wait up to five minutes after a "vzctl enter VEID".

Secondly, I don't know how to specify that the ipaddress will be bound to the vethdevice and not to the venet. My configuration is like this:

```
NAME="lanserv"  
VETH="veth2.0,00:00:00:00:00:02,eth0,00:00:00:00:00:22"  
IP_ADDRESS="172.17.1.242"
```

After starting the VE, vzlist will show me the correct ipaddress (but not the hostname btw), but neither VE nor VE0 have the correct routes added:

```
VE0:/etc/vz/conf# ip route ls  
172.17.1.242 dev venet0 scope link [...]
```

```
lanserv:/# ip route ls  
lanserv:/# ifconfig -a  
eth0    Link encap:Ethernet HWaddr 00:00:00:00:00:22  
        BROADCAST MULTICAST MTU:1500 Metric:1  
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0  
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0  
        collisions:0 txqueuelen:0  
        RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
lo      Link encap:Local Loopback
```

```
LOOPBACK MTU:16436 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
BROADCAST POINTOPOINT NOARP MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

Basically, I have to go through the complete procedure in http://wiki.openvz.org/Virtual_Ethernet_device#Simple_configuration_with_virtual_ethernet_device again in VE and VE0 to be able to use networking of the VE. Is there any option I forgot to give in the VEID.conf-file or anything?

Quote:As concerns appropriate routes for veth. There is a file /usr/sbin/vznetcfg for it. It is a bash script, which is invoked by vzctl on VE start. Look at its source and you'll understand what to do in order to add appropriate routes.

My vznetcfg doesn't do anything beside an "ifconfig \${dev} up". So I added veth2.0 to VE0's /etc/network/interfaces, but it wouldn't help - still the same as quoted above:

```
VE0:/etc/network/# cat interfaces
[...]iface veth2.0 inet static
    up sysctl -w net.ipv4.conf.veth2.0.forwarding=1
    up sysctl -w net.ipv4.conf.veth2.0.proxy_arp=1
    up route add 172.17.1.242 dev veth2.0
```

Any more ideas? If not, I would start writing some bashscripts to run via vznetcfg for adding the appropriate rules to VE0 / some init.d-script for the VEs.

Thanks for helping so far!
