
Subject: IPVS/LVS OpenVZ

Posted by [neotic](#) on Tue, 06 Nov 2007 08:22:47 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi! I've been trying to set up IPVS/LVS (Linux Virtual Server) on OpenVZ without any success.

Let's take it from the beginning.

loadb1 is a separate machine on 192.168.30.139 which is the load balancer.

VPS is the VPS-HOST on 192.168.30.70

WWW1 is a virtual server on VPS-HOST with 192.168.30.242

WWW2 is a virtual server on VPS-HOST with 192.168.30.243

WWW3 is a virtual server on VPS-HOST with 192.168.30.244

192.168.30.245 is the VIRTUAL IP for the loadbalancer, which where ALL web-requests go.

(Every machine is Debian Etch)

First off iproute2 is installed on WWW1/WWW2/WWW3

Then we edit the /etc/sysctl.conf on those machines, adding:

```
net.ipv4.conf.all.arp_ignore = 1
net.ipv4.conf.eth0.arp_ignore = 1
net.ipv4.conf.all.arp_announce = 2
net.ipv4.conf.eth0.arp_announce = 2
and sysctl -p of course.
```

We add the VIRTUAL IP on WWW1/WWW2/WWW3

/etc/network/interfaces:

```
auto lo:0
iface lo:0 inet static
    address 192.168.30.245
    netmask 255.255.255.255
    pre-up sysctl -p > /dev/null
```

and bring it up by ifup lo:0

We can see it's set on the machines with

```
1: lo: <LOOPBACK,UP,10000> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet 192.168.30.245/32 scope global lo:0
```

And since Venet is kinda fubar we use VETH for these 3 virtual machines.

On the VPS-HOST:
ifconfig venet0 down

```
vzctl set <VEID-www1> --veth_add eth0,00:0C:29:11:80:9F,veth103,00:0C:29:3D:4E:32--save
vzctl set <VEID-www2> --veth_add eth0,00:0C:29:FE:16:D2,veth104,00:0C:29:15:D9:C7--save
vzctl set <VEID-www3> --veth_add eth0,00:0C:29:8F:1D:2A,veth105,00:0C:29:12:52:3A --save
```

```
vzctl set 103(www1) --veth_add veth103.0,00:0C:29:3D:4E:32,eth0,00:0C:29:11:80:9F --save
vzctl set 104(www2) --veth_add veth104.0,00:0C:29:15:D9:C7,eth0,00:0C:29:FE:16:D2 --save
vzctl set 104(www3) --veth_add veth105.0,00:0C:29:12:52:3A,eth0,00:0C:29:8F:1D:2A --save
```

```
ifconfig veth103.0 0
ifconfig veth104.0 0
ifconfig veth105.0 0
```

```
echo 1 > /proc/sys/net/ipv4/conf/veth103.0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/veth104.0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/veth105.0/forwarding
```

```
echo 1 > /proc/sys/net/ipv4/conf/veth103.0/proxy_arp
echo 1 > /proc/sys/net/ipv4/conf/veth104.0/proxy_arp
echo 1 > /proc/sys/net/ipv4/conf/veth105.0/proxy_arp
```

```
echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
```

Now we go into all WWWand issue
ifconfig eth0 0
ip addr add 192.168.30.242 dev eth0
ip route add default dev eth0

(just changing the ip for the different WWW1-3)

And now BACK to the VPS-HOST

```
ip route add 192.168.30.242 dev veth103.0
ip route add 192.168.30.243 dev veth104.0
ip route add 192.168.30.244 dev veth105.0
```

The problem NOW is that the VPS-HOST get's all the packets and won't send it to the WWW's.

On the loadbalancer we can see that our WWW's is up and running:

```
loadb1:~# ipvsadm -L -n
IP Virtual Server version 1.2.1 (size=4096)
Prot LocalAddress:Port Scheduler Flags
  -> RemoteAddress:Port      Forward Weight ActiveConn InActConn
TCP  192.168.30.245:80 rr
  -> 192.168.30.244:80        Route 1    0      0
  -> 192.168.30.243:80        Route 1    0      0
  -> 192.168.30.242:80        Route 1    0      0
```

But if we access 192.168.30.245 in our browser we end up at the VPS-HOST and NOT on the WWW1-3

Routes on VPS-HOST:

```
192.168.30.243 dev veth104.0 scope link
192.168.30.242 dev veth103.0 scope link
192.168.30.244 dev veth105.0 scope link
192.168.30.0/24 dev eth0 proto kernel scope link src 192.168.30.70
default via 192.168.30.1 dev eth0
```

Routes on the WWW:

```
default dev eth0 scope link
```

Any ideas? I'm really going crazy with this one.

Martin
