
Subject: NAT enabled but VE cannot communicate beyond HN

Posted by [cheuschober](#) on Sun, 13 May 2007 05:29:28 GMT

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

Hi. I have a HN with two NIC's (eth0=172.16.192.XX,eth1=172.16.192.YY). One of the NIC's (eth0) I'm reserving exclusively for a specific type of high-availability lan traffic. The other NIC (eth1) I want to handle all of my VE traffic.

In the case of my VE's I'm trying to run them on a different domain (172.16.91.---). As instructed by the wiki I set up a NAT route for traffic over eth1 for the VE (in this case a very generous route for debugging):

```
$ iptables -t nat -A POSTROUTING -o eth1 -j SNAT --to 172.16.192.YY
```

The HN and the VE can communicate with or without the route but I cannot get the LAN (or WAN) outside the HN to communicate to the VE or visa/versa.

At this point I'm not certain what I'm doing wrong so any guidance or help is appreciated. If I assign the VE on the same net as the HN (eg, 172.16.192.ZZ) I can communicate to the LAN and WAN with or without trouble (or a NAT entry for that matter) so I take that to mean this has something to do with NAT but I don't quite know what.

Again, any and all help is appreciated.

HN, \$ ip a l

```
2: 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
   inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
4: sit0: <NOARP> mtu 1480 qdisc noop
   link/sit 0.0.0.0 brd 0.0.0.0
6: eth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 1000
   link/ether 00:15:17:19:17:68 brd ff:ff:ff:ff:ff:ff
   inet 172.16.192.50/24 brd 172.16.192.255 scope global eth0
   inet6 fe80::215:17ff:fe19:1768/64 scope link
       valid_lft forever preferred_lft forever
8: eth1: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc pfifo_fast qlen 100
   link/ether 00:15:17:19:17:69 brd ff:ff:ff:ff:ff:ff
   inet 172.16.192.60/24 brd 172.16.192.255 scope global eth1
   inet6 fe80::215:17ff:fe19:1769/64 scope link
       valid_lft forever preferred_lft forever
7: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
   link/void
```

HN, /etc/sysctl.conf

```
net.ipv4.conf.all.rp_filter=1
```

```
net.ipv4.conf.eth1.forwarding=1
```

```
net.ipv4.ip_forward=1
net.ipv4.conf.eth1.proxy_arp=0
net.ipv4.tcp_ecn=0
net.ipv4.conf.eth1.send_redirects=1
net.ipv4.conf.all.send_redirects=0
net.ipv4.conf.all.arp_ignore=1
net.ipv4.conf.all.arp_announce=2
net.ipv4.conf.default.forwarding=1
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

Subject: Re: NAT enabled but VE cannot communicate beyond HN
Posted by [Vasily Tarasov](#) on Mon, 14 May 2007 10:17:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

Please, post here the output of `cat /etc/modprobe.conf` command and kernel version you're using. Also inform us, please, did you compile OpenVZ kernel yourself or did you use a binary image?

Thank you,
Vasily.

Subject: Re: NAT enabled but VE cannot communicate beyond HN
Posted by [morik](#) on Mon, 14 May 2007 12:28:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Please show output of

HW node:

```
# vzlist -a -o "veid,status,hostname,name,ip"
# ip r l
# lsmod | grep nat
```

from one of VE:

```
# traceroute -n 64.233.187.99
```
