
Subject: Strange networking problem

Posted by [Boden](#) on Thu, 18 Sep 2008 16:35:36 GMT

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

New to OpenVZ.

I have an OpenVZ server with two physical NICs. One NIC is public facing and I can't reach it without hopping out to my ISP and back. The other NIC is inward facing, but I'm not doing any sort of routing. I use the internal NIC to service the machine. This works fine.

The problem is that I would like one of the VEs on the machine to also have two NICs, one public and one internal. The reason is that I have to transfer large files to and from an application in this VE, and would prefer a more direct route to the system.

I got this "working" by creating a veth device for the VE, such that it has both venet0:0 with a public IP and eth1 with a private IP. eth1 corresponds to veth101.1 on the HN. After a bit of configuration, it seemed to work, except the connection is intermittent.

The VE's "internal" ip is 10.0.2.10. If I ping this interface from the hardware node, sometimes I'll get a "destination unreachable" error on the first ping, and then all subsequent pings will succeed. When I connect to 10.0.2.10 from another machine on the network, the connection very often will fail, but occasionally work. If I let a ping to 10.0.2.10 from the hardware node run continuously, connections from elsewhere on the network always seem to work. It's like the interface falls asleep until the hardware node pokes it with a stick and says "hey buddy, wake up".

The configuration is this:

HN:

```
eth0, physical: public facing  
eth1: physical: private, IP 10.0.2.6
```

The VE .conf file declares:

```
NETIF=" ifname=eth1,mac=00:18:51:A4:CD:89,host_ifname=veth101.1,host  
_mac=00:18:51:6F:65:14 "
```

```
#ip route
```

```
10.0.4.20 dev venet0 scope link src 10.0.4.6  
10.0.4.10 dev venet0 scope link src 10.0.4.6  
10.0.2.10 dev veth101.1 scope link src 10.0.2.6  
10.0.4.0/24 dev eth0 proto kernel scope link src 10.0.4.6  
10.0.2.0/24 dev eth1 proto kernel scope link src 10.0.2.6  
default via 10.0.4.1 dev eth0 metric 100
```

```
echo 1 > /proc/sys/net/ipv4/conf/veth101.1/forwarding  
echo 1 > /proc/sys/net/ipv4/conf/veth101.1/proxy_arp  
echo 1 > /proc/sys/net/ipv4/conf/eth1/forwarding 1  
echo 1 > /proc/sys/net/ipv4/conf/eth1/proxy_arp 1
```

VE

```
# ifconfig
```

```
eth1  Link encap:Ethernet HWaddr 00:18:51:a4:cd:89
      inet addr:10.0.2.10 Bcast:10.0.2.255 Mask:255.255.255.0
      inet6 addr: fe80::218:51ff:fea4:cd89/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:1460 errors:0 dropped:0 overruns:0 frame:0
      TX packets:1582 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:189047 (184.6 KB) TX bytes:1308848 (1.2 MB)
```

```
lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      inet6 addr: ::1/128 Scope:Host
      UP LOOPBACK RUNNING MTU:16436 Metric:1
      RX packets:128046 errors:0 dropped:0 overruns:0 frame:0
      TX packets:128046 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:16350942 (15.5 MB) TX bytes:16350942 (15.5 MB)
```

```
venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
      RX packets:1115 errors:0 dropped:0 overruns:0 frame:0
      TX packets:974 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:1262900 (1.2 MB) TX bytes:349670 (341.4 KB)
```

```
venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:10.0.4.10 P-t-P:10.0.4.10 Bcast:0.0.0.0 Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

Subject: Re: Strange networking problem
Posted by [Boden](#) on Thu, 18 Sep 2008 18:04:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Nevermind. I didn't have forwarding enabled on eth1 on the hardware node when I first started playing with getting this running. Enabling forwarding and proxy_arp on eth1 and then simply assigning multiple IP addresses to the VE (without creating a new veth interface) worked just fine.