
Subject: *ISSUE* VE route problem

Posted by [stonenzj](#) on Fri, 08 Jun 2007 11:35:46 GMT

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

Hi, i reinstall the root server at yesterday. openvz working fine before reinstall.
Now,the VZ110 can't connect to outside , only connent to HN .

// Base Kernel

```
image = /boot/vmlinuz-2.6.9-55.ELsmp
  label = 2.6.9-55.ELsmp
  initrd = /boot/initrd-2.6.9-55.ELsmp.img
```

// VZ Kernel

```
Linux 84-16-233-207 2.6.9-023stab044.4-smp #1 SMP Mon May 28 23:18:17 MSD 2007 i686
i686 i386 GNU/Linux
```

// ip r l at HN

```
84.16.251.121 dev venet0 scope link
89.149.221.89 dev venet0 scope link
84.16.233.0/24 dev eth0 proto kernel scope link src 84.16.233.207
169.254.0.0/16 dev eth0 scope link
default via 84.16.233.1 dev eth0
```

// ip r l at VE110

```
192.0.2.0/24 dev venet0 scope host
169.254.0.0/16 dev venet0 scope link
default via 192.0.2.1 dev venet0
```

//ip a l at HN

```
2: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
4: eth0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
  link/ether 00:02:a5:0a:e5:30 brd ff:ff:ff:ff:ff:ff
  inet 84.16.233.207/24 brd 84.16.233.255 scope global eth0
11: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
  link/void
```

//ip a l at VE110

```
1: lo: <LOOPBACK,UP> 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 brd 127.255.255.255 scope host lo
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
  link/void
  inet 127.0.0.1/32 scope host venet0
  inet 89.149.221.89/32 brd 89.149.221.89 scope global venet0:0
  inet 84.16.251.121/32 brd 84.16.251.121 scope global venet0:1
```

//route at HN

```
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
84.16.251.121 0.0.0.0 255.255.255.255 UH 0 0 0 venet0
89.149.221.89 0.0.0.0 255.255.255.255 UH 0 0 0 venet0
84.16.233.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 0 0 0 eth0
0.0.0.0 84.16.233.1 0.0.0.0 UG 0 0 0 eth0
```

//route at ve110

```
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.0.2.0 0.0.0.0 255.255.255.0 U 0 0 0 venet0
169.254.0.0 0.0.0.0 255.255.0.0 U 0 0 0 venet0
0.0.0.0 192.0.2.1 0.0.0.0 UG 0 0 0 venet0
```

//ifconfig at Ve110

```
lo Link encap:Local Loopback
  inet addr:127.0.0.1 Mask:255.0.0.0
  UP LOOPBACK RUNNING MTU:16436 Metric:1
  RX packets:1230259009 errors:0 dropped:0 overruns:0 frame:0
  TX packets:1631407695 errors:0 dropped:0 overruns:0 carrier:0
  collisions:0 txqueuelen:0
  RX bytes:1140876388 (1.0 GiB) TX bytes:1095784005 (1.0 GiB)

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:8 errors:0 dropped:0 overruns:0 frame:0
TX packets:44 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:672 (672.0 b) TX bytes:3596 (3.5 KiB)

venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:89.149.221.89 P-t-P:89.149.221.89 Bcast:89.149.221.89 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

venet0:1 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:84.16.251.121 P-t-P:84.16.251.121 Bcast:84.16.251.121 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

//ifconfig at HN

eth0 Link encap:Ethernet HWaddr 00:02:A5:0A:E5:30
inet addr:84.16.233.207 Bcast:84.16.233.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:268804 errors:0 dropped:0 overruns:0 frame:0
TX packets:134942 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:150618261 (143.6 MiB) TX bytes:13052749 (12.4 MiB)

lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:34 errors:0 dropped:0 overruns:0 frame:0
TX packets:34 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:4353 (4.2 KiB) TX bytes:4353 (4.2 KiB)

venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:44 errors:0 dropped:0 overruns:0 frame:0
TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:3596 (3.5 KiB) TX bytes:672 (672.0 b)

//iptables all stop
//sysctl -w net.ipv4.ip_forward="1"
//Ping test

HN to ve110

PING 89.149.221.89 (89.149.221.89) 56(84) bytes of data.
64 bytes from 89.149.221.89: icmp_seq=0 ttl=64 time=0.197 ms
64 bytes from 89.149.221.89: icmp_seq=1 ttl=64 time=0.106 ms

64 bytes from 89.149.221.89: icmp_seq=2 ttl=64 time=0.129 ms
64 bytes from 89.149.221.89: icmp_seq=3 ttl=64 time=0.092 ms
64 bytes from 89.149.221.89: icmp_seq=4 ttl=64 time=0.124 ms

HN to DNS

PING 217.20.116.1 (217.20.116.1) 56(84) bytes of data.
64 bytes from 217.20.116.1: icmp_seq=0 ttl=61 time=0.332 ms
64 bytes from 217.20.116.1: icmp_seq=1 ttl=61 time=0.296 ms
64 bytes from 217.20.116.1: icmp_seq=2 ttl=61 time=0.283 ms
64 bytes from 217.20.116.1: icmp_seq=3 ttl=61 time=0.265 ms

VE110 TO HN

PING 84.16.233.207 (84.16.233.207) 56(84) bytes of data.
64 bytes from 84.16.233.207: icmp_seq=0 ttl=64 time=0.222 ms
64 bytes from 84.16.233.207: icmp_seq=1 ttl=64 time=0.110 ms
64 bytes from 84.16.233.207: icmp_seq=2 ttl=64 time=0.098 ms

VE110 to DNS

PING 217.20.116.1 (217.20.116.1) 56(84) bytes of data.

--- 217.20.116.1 ping statistics ---
32 packets transmitted, 0 received, 100% packet loss, time 31000ms

//traceroute test

HN TO DNS

traceroute to 217.20.116.1 (217.20.116.1), 30 hops max, 38 byte packets
1 84.16.224.1 0.531 ms 0.469 ms 0.451 ms
2 89.149.218.33 0.247 ms 0.232 ms 0.760 ms
3 89.149.218.14 0.492 ms 0.460 ms 0.506 ms
4 217.20.116.1 0.243 ms 0.252 ms 0.236 ms

VE TO DNS

traceroute to 217.20.116.1 (217.20.116.1), 30 hops max, 38 byte packets
1 84.16.233.207 (84.16.233.207) 0.285 ms 0.253 ms 0.084 ms
2 * * *
3 * * *

noting

Subject: Re: *ISSUE* VE route problem
Posted by [rickb](#) on Fri, 08 Jun 2007 15:11:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

can you see the VE->internet icmp packet when running tcpdump on the HN?

just to be sure- (people have configured sysctl.conf and not applied it in the past):

```
#cat /proc/sys/net/ipv4/ip_forward
```

Subject: Re: *ISSUE* VE route problem
Posted by [stonenzj](#) on Fri, 08 Jun 2007 15:36:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for the help!Yes,i catch some ICMP info on HN,when run "ping 217.20.116.1" at VE110 .
And i sure net.ipv4.ip_forward="1"

```
// TCPDUMP ICMP ON HN  
16:28:37.828222 IP (tos 0x0, ttl 64, id 0, offset 0, flags [DF], proto 1, length: 84) 89.149.221.89 >  
217.20.116.1: icmp 64: echo request seq 256
```

```
//sysctl -a
```

```
net.ipv4.ip_nonlocal_bind = 0  
net.ipv4.ip_no_pmtu_disc = 0  
net.ipv4.ip_autoconfig = 0  
net.ipv4.ip_default_ttl = 64  
net.ipv4.ip_forward = 1  
net.ipv4.tcp_retrans_collapse = 1  
net.ipv4.tcp_sack = 1  
net.ipv4.tcp_window_scaling = 1  
net.ipv4.tcp_timestamps = 1
```

```
//
```

Subject: Re: *ISSUE* VE route problem
Posted by [rickb](#) on Fri, 08 Jun 2007 15:42:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Do you have access to 217.20.116.1? I would say that assuming there are no postrouting rules and no filters between you and 217.20.116.1, 217.20.116.1 has received your packet.

Subject: Re: *ISSUE* VE route problem
Posted by [stonenzj](#) on Fri, 08 Jun 2007 15:54:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

IP's (217.20.116.1) is DNS server, I have not access to this server.
Just catch 89.149.221.89 to 217.20.116.1 ICMP info, have not received any packet form
217.20.116.1 to 89.149.221.89

Subject: Re: *ISSUE* VE route problem
Posted by [rickb](#) on Fri, 08 Jun 2007 15:56:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

on HN:

```
#iptables -L -n  
#iptables -L -n -t nat
```

Subject: Re: *ISSUE* VE route problem
Posted by [stonenzj](#) on Fri, 08 Jun 2007 15:59:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
//iptables -L -n
```

```
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
```

```
//iptables -L -n -t nat
```

```
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
```

Subject: Re: *ISSUE* VE route problem
Posted by [rickb](#) on Fri, 08 Jun 2007 16:04:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

ok. so your HN is on a public IP, VE on a private, and you want your VE to communicate with the internet. How can it do this when you have not configured any NATing, masquerading, or otherwise any way for the 192.168.* packet to reach the net? you need to rewrite it somehow, somewhere before a router that isn't configured to route 192.168.* receives it.

check out this wiki page, http://wiki.openvz.org/Using_NAT_for_VE_with_private_IPs

this scenario describes your setup exactly:

How to provide access for VE to Internet

To enable the VEs, which have only internal IP addresses, to access the Internet, SNAT (Source Network Address Translation, also known as IP masquerading) should be configured on the Hardware Node. This is ensured by the standard Linux iptables utility. To perform a simple SNAT setup, execute the following command on the Hardware Node:

```
# iptables -t nat -A POSTROUTING -s src_net -o eth0 -j SNAT --to ip_address
```

Subject: Re: *ISSUE* VE route problem
Posted by [stonenzj](#) on Fri, 08 Jun 2007 16:25:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

VE 89.149.221.89 and HN 84.16.233.207 all public IPs, assign from the provide.
And i just re-install the root-server at yesterday, same as OS images , IP's and config. Only different is openvz kernel version.
Old kernel 2.6.8-022stab078.10 on HN, And VE110 IP's 89.149.221.89 can working, install about one year ago. The root-server and VE can access to outside before re-install.

Subject: Re: *ISSUE* VE route problem
Posted by [rickb](#) on Fri, 08 Jun 2007 16:30:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

oops, sorry. Ok. the last thing I would do, is reassign the IP which you have in the VE to eth0:10 on the HN, and try "ping openvz.org -I IP", specify that IP as the source. If that works, its an openvz problem and someone else will have to throw some ideas at this, as I am empty.

If doesn't work, reboot your box into the non openvz kernel, and do the same thing. If you can't ping with that source from the HN (non vz kernel) again, its not a openvz problem.

Subject: Re: *ISSUE* VE route problem

Posted by [stonenzj](#) on Fri, 08 Jun 2007 17:20:25 GMT

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

Okay, I try add IP to virtual interfaces with eth0:6 and ping -I 89.149.221.89 google.com respond timeouts. both kernel also.

I contact my provide first, Thanks for your time.
