
Subject: Two different IP blocks on OpenVZ = Unpingable help plz

Posted by [navigator](#) on Thu, 04 Dec 2008 16:16:29 GMT

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

Hi guys,

I have some OpenVZ nodes.

The server main IP (eth0) is 200.44.32.42, the allocation is 200.44.32.40/29 - 200.44.32.41 being the GATEWAY.

The ADDITIONAL IP allocation is 200.44.32.96/28, 200.44.32.97 being the GATEWAY.

Now, if I be specific on /etc/sysconfig/network-scripts/ifcfg-eth0 and declare everything like so:

```
DEVICE=eth0
BOOTPROTO=none
BROADCAST=10.20.47.255
HWADDR=
IPADDR=10.20.42.198
NETMASK=255.255.248.0
NETWORK=10.20.40.0
ONBOOT=yes
GATEWAY=10.20.40.1
TYPE=Ethernet
```

Please ignore all the above IPs config.

Now, using that VE nodes on the 200.44.32.96/28 range are unpingable. How can I fix that?

Please take note that I remove all settings from ifcfg, and leave IPADDR and GATEWAY, the internet can ping the nodes. But within the 200.44.32.0/24 no-one can ping the nodes on the 200.44.32.96/28 range (Which is my problem, that I came here to solve).

Here's the VE

```
root@NM [~]# ifconfig
```

```
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:138 errors:0 dropped:0 overruns:0 frame:0
      TX packets:138 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:13430 (13.1 KiB)  TX bytes:13430 (13.1 KiB)
```

```
venet0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      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:4355 errors:0 dropped:0 overruns:0 frame:0
TX packets:5095 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:378263 (369.3 KiB) TX bytes:2541022 (2.4 MiB)

venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:200.44.32.100 P-t-P:216.151.144.100 Bcast:200.44.32.100
Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1

```
root@NM [~]# ip route list table all
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
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 216.151.144.100 dev venet0 table 255 proto kernel scope host src 200.44.32.100
broadcast 200.44.32.100 dev venet0 table 255 proto kernel scope link src 200.44.32.100
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.1 dev venet0 table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
You have new mail in /var/spool/mail/root
```

```
root@NM [~]# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
iptables v1.3.5: can't initialize iptables table `nat': Table does not exist (do you need to insmod?)
Perhaps iptables or your kernel needs to be upgraded.
```

```
root@NM [~]# tcpdump -i venet0:0 -e host 200.44.32.2
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0:0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
0 packets captured
1 packets received by filter
0 packets dropped by kernel
```

```
root@NM [~]# tcpdump -i venet0 -e host 200.44.32.2
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes

0 packets captured
1 packets received by filter
0 packets dropped by kernel
```

Thanks!

Subject: Re: Two different IP blocks on OpenVZ = Unpingable help plz
Posted by [navigator](#) on Mon, 08 Dec 2008 23:32:34 GMT

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

any help?

Subject: Re: Two different IP blocks on OpenVZ = Unpingable help plz
Posted by [maratrus](#) on Tue, 09 Dec 2008 10:41:52 GMT

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

Hello,

Quote:

```
root@NM [~]# ifconfig
```

```
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:138 errors:0 dropped:0 overruns:0 frame:0
      TX packets:138 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:13430 (13.1 KiB)  TX bytes:13430 (13.1 KiB)
```

```
venet0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      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:4355 errors:0 dropped:0 overruns:0 frame:0
      TX packets:5095 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:378263 (369.3 KiB)  TX bytes:2541022 (2.4 MiB)
```

```
venet0:0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:200.44.32.100  P-t-P:216.151.144.100  Bcast:200.44.32.100
      Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

```
root@NM [~]# ip route list table all
```

```
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
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 216.151.144.100 dev venet0 table 255 proto kernel scope host src 200.44.32.100
broadcast 200.44.32.100 dev venet0 table 255 proto kernel scope link src 200.44.32.100
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
local 127.0.0.1 dev lo table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.1 dev venet0 table 255 proto kernel scope host src 127.0.0.1
local 127.0.0.0/8 dev lo table 255 proto kernel scope host src 127.0.0.1
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
local ::1 via :: dev lo table 255 proto none metric 0 mtu 16436 advmss 16376 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255
You have new mail in /var/spool/mail/root
```

seems ok

Quote:

```
root@NM [~]# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
iptables v1.3.5: can't initialize iptables table `nat': Table does not exist (do you need to insmod?)
Perhaps iptables or your kernel needs to be upgraded.
```

the first command fails and other haven't been executed.

Quote:

```
root@NM [~]# tcpdump -i venet0:0 -e host 200.44.32.2
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0:0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
0 packets captured
1 packets received by filter
0 packets dropped by kernel
```

```
root@NM [~]# tcpdump -i venet0 -e host 200.44.32.2
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on venet0, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
```

```
0 packets captured
1 packets received by filter
0 packets dropped by kernel
```

Did you try to ping your VE at that moment. If yes, you should run tcpdump on the HN too.

please show the same information from the HN.

And also

"arp -n", "ip rule list".
