
Subject: network issue

Posted by [soptom](#) on Mon, 29 Oct 2007 13:35:26 GMT

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

Dear Friends,

i created a VE, but i have network issues due to my poor networking knowledge.
I provided to VE an external ip and the proper subnet but still have no access to internet world.
Here is my ifconfig:

```
[root@master2 /]# 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:204 errors:0 dropped:0 overruns:0 frame:0
        TX packets:204 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:18520 (18.0 KiB)  TX bytes:18520 (18.0 KiB)
```

```
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:83 errors:0 dropped:0 overruns:0 frame:0
        TX packets:58 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:6136 (5.9 KiB)  TX bytes:9528 (9.3 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:91.194.90.126  P-t-P:91.194.90.126  Bcast:91.194.90.126  Mask:255.255.255.128
        UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

Subject: Re: network issue

Posted by [Valmont](#) on Mon, 29 Oct 2007 14:31:32 GMT

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

Please provide next info:

```
ip r s
ip a s
```

```
cat /etc/sysctl.conf
```

```
on hardware node.
```

Subject: Re: network issue

Posted by [soptom](#) on Mon, 29 Oct 2007 17:05:42 GMT

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

```
[root@master2 ~]# ip r s
91.194.90.126 dev venet0 scope link
91.194.90.0/25 dev eth0 proto kernel scope link src 91.194.90.25
169.254.0.0/16 dev eth0 scope link
default via 91.194.90.1 dev eth0
```

```
[root@master2 ~]# ip a s
2: lo: <LOOPBACK,UP,LOWER_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 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
4: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:02:b3:b7:29:75 brd ff:ff:ff:ff:ff:ff
    inet 91.194.90.25/25 brd 91.194.90.127 scope global eth0
    inet6 fe80::202:b3ff:feb7:2975/64 scope link
        valid_lft forever preferred_lft forever
6: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/void
```

```
[root@master2 ~]# cat /etc/sysctl.conf
# Kernel sysctl configuration file for Red Hat Linux
#
# For binary values, 0 is disabled, 1 is enabled. See sysctl( and
# sysctl.conf(5) for more details.
```

```
# Disables packet forwarding
net.ipv4.ip_forward=1
```

```
# Disables IP source routing
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.lo.accept_source_route = 0
net.ipv4.conf.eth0.accept_source_route = 0
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.conf.default.proxy_arp = 0
```

```
# Enable IP spoofing protection, turn on source route verification
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.lo.rp_filter = 1
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
```

```
# Disable ICMP Redirect Acceptance
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.lo.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
net.ipv4.conf.default.accept_redirects = 0

# Enable Log Spoofed Packets, Source Routed Packets, Redirect Packets
net.ipv4.conf.all.log_martians = 0
net.ipv4.conf.lo.log_martians = 0
net.ipv4.conf.eth0.log_martians = 0

# Disables IP source routing
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.lo.accept_source_route = 0
net.ipv4.conf.eth0.accept_source_route = 0
net.ipv4.conf.default.accept_source_route = 0

# Enable IP spoofing protection, turn on source route verification
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.lo.rp_filter = 1
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1

# Disable ICMP Redirect Acceptance
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.lo.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0

# Disables the magic-sysrq key
kernel.sysrq = 1

# Decrease the time default value for tcp_fin_timeout connection
net.ipv4.tcp_fin_timeout = 15

# Decrease the time default value for tcp_keepalive_time connection
net.ipv4.tcp_keepalive_time = 1800

# Turn off the tcp_window_scaling
net.ipv4.tcp_window_scaling = 0

# Turn off the tcp_sack
net.ipv4.tcp_sack = 0

# Turn off the tcp_timestamps
net.ipv4.tcp_timestamps = 0
```

```
# Enable TCP SYN Cookie Protection
net.ipv4.tcp_syncookies = 1

# Enable ignoring broadcasts request
net.ipv4.icmp_echo_ignore_broadcasts = 1

# Enable bad error message Protection
net.ipv4.icmp_ignore_bogus_error_responses = 1

# Log Spoofed Packets, Source Routed Packets, Redirect Packets
net.ipv4.conf.all.log_martians = 1

# Increases the size of the socket queue (effectively, q0).
net.ipv4.tcp_max_syn_backlog = 1024

# Increase the tcp-time-wait buckets pool size
net.ipv4.tcp_max_tw_buckets = 1440000

# Allowed local port range
net.ipv4.ip_local_port_range = 16384 65536
```

Subject: Re: network issue
Posted by [Valmont](#) on Mon, 29 Oct 2007 17:36:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

And what about proxy_arp?

```
# sysctl -a | grep proxy_arp
net.ipv4.conf.venet0.proxy_arp = 0
net.ipv4.conf.eth2.proxy_arp = 1
net.ipv4.conf.eth1.proxy_arp = 1
net.ipv4.conf.eth0.proxy_arp = 0
net.ipv4.conf.lo.proxy_arp = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.proxy_arp = 0
```

You should have proxy_arp turned on at your gateway interface.

Subject: Re: network issue
Posted by [soptom](#) on Mon, 29 Oct 2007 18:13:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Valmont wrote on Mon, 29 October 2007 19:36 And what about proxy_arp?

```
# sysctl -a | grep proxy_arp
net.ipv4.conf.venet0.proxy_arp = 0
net.ipv4.conf.eth2.proxy_arp = 1
net.ipv4.conf.eth1.proxy_arp = 1
net.ipv4.conf.eth0.proxy_arp = 0
net.ipv4.conf.lo.proxy_arp = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.proxy_arp = 0
```

You should have proxy_arp turned on at your gateway interface.

```
This is my sysctl -a | grep proxy_arp
net.ipv4.conf.venet0.proxy_arp = 0
net.ipv4.conf.eth0.proxy_arp = 0
net.ipv4.conf.lo.proxy_arp = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.proxy_arp = 0
```

Please tell me Exact what to do.
Really thank you in advance

Subject: Re: network issue
Posted by [Valmont](#) on Mon, 29 Oct 2007 20:48:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
echo 'net.ipv4.conf.eth0.proxy_arp = 1' >> /etc/sysctl.conf
sysctl -p
```

Subject: Re: network issue
Posted by [soptom](#) on Mon, 29 Oct 2007 21:43:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
Valmont wrote on Mon, 29 October 2007 22:48echo 'net.ipv4.conf.eth0.proxy_arp = 1' >>
/etc/sysctl.conf
sysctl -p
```

```
[root@master2 ~]# echo 'net.ipv4.conf.eth0.proxy_arp = 1' >> /etc/sysctl.conf
[root@master2 ~]# sysctl -p
net.ipv4.ip_forward = 1
net.ipv4.conf.all.accept_source_route = 0
```

```
net.ipv4.conf.lo.accept_source_route = 0
net.ipv4.conf.eth0.accept_source_route = 0
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.lo.rp_filter = 1
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.lo.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.all.log_martians = 0
net.ipv4.conf.lo.log_martians = 0
net.ipv4.conf.eth0.log_martians = 0
net.ipv4.conf.all.accept_source_route = 0
net.ipv4.conf.lo.accept_source_route = 0
net.ipv4.conf.eth0.accept_source_route = 0
net.ipv4.conf.default.accept_source_route = 0
net.ipv4.conf.all.rp_filter = 1
net.ipv4.conf.lo.rp_filter = 1
net.ipv4.conf.eth0.rp_filter = 1
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.lo.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
kernel.sysrq = 1
net.ipv4.tcp_fin_timeout = 15
net.ipv4.tcp_keepalive_time = 1800
net.ipv4.tcp_window_scaling = 0
net.ipv4.tcp_sack = 0
net.ipv4.tcp_timestamps = 0
error: "Operation not permitted" setting key "net.ipv4.tcp_syncookies"
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.icmp_ignore_bogus_error_responses = 1
net.ipv4.conf.all.log_martians = 1
net.ipv4.tcp_max_syn_backlog = 1024
net.ipv4.tcp_max_tw_buckets = 1440000
net.ipv4.ip_local_port_range = 16384 65536
net.ipv4.conf.eth0.proxy_arp = 1
```

Now what?

Thank you

Subject: Re: network issue
Posted by [Valmont](#) on Mon, 29 Oct 2007 22:50:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sorry, I was wrong about proxy_arp. It's help me in one unusual conf, but it's not right answer

Please, check

iptables -nvL FORWARD

also.

Subject: Re: network issue
Posted by [soptom](#) on Mon, 29 Oct 2007 23:22:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Valmont wrote on Tue, 30 October 2007 00:50Sorry, I was wrong about proxy_arp. It's help me in one unusual conf, but it's not right answer

Please, check

iptables -nvL FORWARD

also.

```
iptables -nvL FORWARD
Chain FORWARD (policy ACCEPT 90 packets, 5506 bytes)
pkts bytes target    prot opt in     out    source          destination
 45 2700 LOG      all  --  *      eth0    0.0.0.0/0        0.0.0.0/0        LOG flags 0 level 7 prefix
`BANDWIDTH_OUT:'
 45 2806 LOG      all  --  eth0    *      0.0.0.0/0        0.0.0.0/0        LOG flags 0 level 7 prefix
`BANDWIDTH_IN:'
[root@master2 ~]#
```

Subject: Re: network issue
Posted by [Valmont](#) on Mon, 29 Oct 2007 23:33:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

In VPS

```
ping -c5 -q 91.194.90.25
ping -c5 -q 91.194.90.1
traceroute 91.194.90.1
```

On HN (just to be assured)

```
ping -c5 -q 91.194.90.126
ping -c5 -q 91.194.90.1
```

Subject: Re: network issue
Posted by [soptom](#) on Mon, 29 Oct 2007 23:57:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

Valmont wrote on Tue, 30 October 2007 01:33In VPS

```
ping -c5 -q 91.194.90.25
ping -c5 -q 91.194.90.1
traceroute 91.194.90.1
```

```
[root@master2 ~]# vzctl enter 001
entered into VE 1
[root@master2 /]# ping -c5 -q 91.194.90.25
PING 91.194.90.25 (91.194.90.25) 56(84) bytes of data.
```

```
--- 91.194.90.25 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000ms
rtt min/avg/max/mdev = 0.044/0.067/0.117/0.028 ms
[root@master2 /]# ping -c5 -q 91.194.90.1
PING 91.194.90.1 (91.194.90.1) 56(84) bytes of data.
```

```
--- 91.194.90.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4003ms
rtt min/avg/max/mdev = 0.360/0.650/0.724/0.145 ms
[root@master2 /]# traceroute 91.194.90.1
traceroute to 91.194.90.1 (91.194.90.1), 30 hops max, 40 byte packets
 1 91.194.90.25 (91.194.90.25) 0.053 ms 0.026 ms 0.025 ms
 2 91.194.90.1 (91.194.90.1) 0.201 ms 0.169 ms 0.144 ms
[root@master2 /]#
```

On HN (just to be assured)

```
[B]ping -c5 -q 91.194.90.126[/B]
```

```
[root@master2 ~]# ping -c5 -q 91.194.90.126
PING 91.194.90.126 (91.194.90.126) 56(84) bytes of data.
```

--- 91.194.90.126 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4000ms
rtt min/avg/max/mdev = 0.029/0.036/0.057/0.012 ms

[B]ping -c5 -q 91.194.90.1[/B]

[root@master2 ~]# ping -c5 -q 91.194.90.1
PING 91.194.90.1 (91.194.90.1) 56(84) bytes of data.

--- 91.194.90.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4002ms
rtt min/avg/max/mdev = 0.219/0.727/1.163/0.365 ms

Subject: Re: network issue
Posted by [Valmont](#) on Tue, 30 Oct 2007 00:24:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

charmant

I didn't see problems with network in your case

hard day

At last point you can check also

cat /etc/resolv.conf

To be assured in proper dns configuration.

Subject: Re: network issue
Posted by [soptom](#) on Tue, 30 Oct 2007 00:35:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dear Friend,

i want to THANK you from deep from my heart for your help.
Finally, i have network to my VE! I can't find more words to thank you for it. Thank you, thank you, thank you

Subject: Re: network issue
Posted by [Valmont](#) on Tue, 30 Oct 2007 00:37:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

You are welcome)

Subject: Re: network issue
Posted by [soptom](#) on Thu, 01 Nov 2007 02:23:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Valmont wrote on Tue, 30 October 2007 00:50 Sorry, I was wrong about proxy_arp. It's help me in one unusual conf, but it's not right answer

Please, check

iptables -nvL FORWARD

also.

If you are still there dear friend, i ahve one question: I created another VE and i followed all those steps, everything seems to be ok , but iptables -nvL FORWARD produces nothing, and i have no network in new VE.

Any suggestions?

Thank you

Subject: Re: network issue
Posted by [Valmont](#) on Thu, 01 Nov 2007 07:28:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you do

ip a s
ping and traceroute as in last example.

from vps?

Subject: Re: network issue

Posted by [soptom](#) on Thu, 01 Nov 2007 11:16:45 GMT

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

```
[root@fc5test /]# ip a s
```

```
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
```

```
inet6 ::1/128 scope host
```

```
valid_lft forever preferred_lft forever
```

```
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue
```

```
link/void
```

```
inet 127.0.0.1/32 scope host venet0
```

```
inet 91.194.90.126/25 brd 91.194.90.127 scope global venet0:0
```

```
[root@fc5test /]# ping -c5 -q 91.194.90.25
```

```
PING 91.194.90.25 (91.194.90.25) 56(84) bytes of data.
```

```
--- 91.194.90.25 ping statistics ---
```

```
5 packets transmitted, 5 received, 0% packet loss, time 3999ms
```

```
rtt min/avg/max/mdev = 0.044/0.054/0.083/0.014 ms
```

```
[root@fc5test /]# ping -c5 -q 91.194.90.1
```

```
PING 91.194.90.1 (91.194.90.1) 56(84) bytes of data.
```

```
--- 91.194.90.1 ping statistics ---
```

```
5 packets transmitted, 0 received, 100% packet loss, time 3999ms
```

```
[root@fc5test /]# traceroute 91.194.90.1
```

```
traceroute to 91.194.90.1 (91.194.90.1), 30 hops max, 40 byte packets
```

```
1 * * *
2 * * *
3 * * *
4 * * *
5 * * *
6 * * *
7 * * *
8 * * *
9 * * *
10 * * *
11 * * *
12 * * *
13 * * *
14 * * *
15 * * *
16 * * *
17 * * *
18 * * *
19 * * *
20 * * *
21 * * *
```

22 * * *
23 * * *
24 * * *
25 * * *
26 * * *
27 * * *
28 * * *
29 * * *
30 * * *

Subject: Re: network issue

Posted by [Valmont](#) on Thu, 01 Nov 2007 11:42:51 GMT

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

looks like problem with ip forwarding on HN, _but_.

this seems to be strange for me:

```
inet 91.194.90.126/25 brd 91.194.90.127 scope global venet0:0
```

why venet0, which should be point-to-point interface has such configuration.

for example, one of my vpses:

```
# vzlist | grep 109
```

```
109      7 running 10.0.5.9   ns2
```

```
# vzctl enter 109
```

```
entered into VE 109
```

```
# ip a s
```

```
1: lo: <LOOPBACK,UP,LOWER_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 scope host lo
```

```
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
```

```
link/void
```

```
inet 127.0.0.1/32 scope host venet0
```

```
inet 10.0.5.9/32 brd 10.0.5.9 scope global venet0:0
```

You see, the network interface venet0 has mask /32 and ip's is identical.

Did you modify network configuration in vps?

//you should not worry about iptables -nvL FORWARD if default policy is not "DROP".

Subject: Re: network issue

Posted by [soptom](#) on Thu, 01 Nov 2007 11:54:45 GMT

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

Valmont wrote on Thu, 01 November 2007 13:42 looks like problem with ip forwarding on HN, _but_.

this seems to be strange for me:

```
inet 91.194.90.126/25 brd 91.194.90.127 scope global venet0:0
```

why venet0, which should be point-to-point interface has such configuration.

for example, one of my vpses:

```
# vzlist | grep 109
 109      7 running 10.0.5.9   ns2
# vzctl enter 109
entered into VE 109
# ip a s
1: lo: <LOOPBACK,UP,LOWER_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 scope host lo
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/void
    inet 127.0.0.1/32 scope host venet0
    inet 10.0.5.9/32 brd 10.0.5.9 scope global venet0:0
```

You see, the network interface venet0 has mask /32 and ip's is identical.

Did you modify network configuration in vps?

I had to change the netmask (ifconfig venet0:0 netmask 255.255.255.128)

//you should not worry about iptables -nvL FORWARD if default policy is not "DROP".

Subject: Re: network issue

Posted by [Valmont](#) on Thu, 01 Nov 2007 12:14:13 GMT

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

You shouldn't do it.

OpenVZ startup scripts doing all necessary work for vps's network.

Subject: Re: network issue

Posted by [soptom](#) on Thu, 01 Nov 2007 12:52:37 GMT

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

I will change those setting (ip nameserver) with vzctl again.

But i had to do that due to netmask issue which had to be 255.255.255.128

brb

Subject: Re: network issue

Posted by [Valmont](#) on Thu, 01 Nov 2007 13:21:10 GMT

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

There is enough if mask is set on HN.

Subject: Re: network issue

Posted by [soptom](#) on Thu, 01 Nov 2007 16:03:22 GMT

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

[root@fc5test /]# ip a s

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

inet6 ::1/128 scope host

valid_lft forever preferred_lft forever

3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,10000> mtu 1500 qdisc noqueue

link/void

inet 127.0.0.1/32 scope host venet0

inet 91.194.90.126/32 brd 91.194.90.126 scope global venet0:0

Seems to be OK, but:

[root@fc5test /]# ping -c5 -q 91.194.90.1

PING 91.194.90.1 (91.194.90.1) 56(84) bytes of data.

--- 91.194.90.1 ping statistics ---

5 packets transmitted, 0 received, 100% packet loss, time 3999ms

Is that no connection to gateway?

P.S By the way i did not want to make any ifconfig changes but the vzctl --ipaddr x.x.x.x --nameserver x.x.x.x -save did not provide any external network to VE...
