
Subject: cannot ping internet from ve

Posted by [cesurasean](#) on Mon, 09 Jun 2008 08:59:57 GMT

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

I am having the same issues on two seperate machines.

One machine is Debian with OpenVZ installed, the other is RHEL.

Both CANNOT access the internet via the Container.

I don't understand what I am doing wrong here!

I set the parameters (ip address and name server) BEFORE I start the container,

and I also set the ip address and nameservers with the ifconfig command once the container is opened. Still, no success. Neither of those work.

I'm stumped. I'm thinking it has something to with virtualization of the ethernet card.

This is happening on two seperate machines (seperate distros also) !!!

```
server:~# ifconfig
```

```
venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
```

```
inet addr:192.168.1.101 P-t-P:192.168.1.101 Bcast:192.168.1.101 Mask:255.255.255.0
```

```
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

```
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
```

```
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
```

```
collisions:0 txqueuelen:0
```

```
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
server:~#
```

----this is taken from the host machine that runs openvz. it's able to ping the ve. ve can't access internet.

```
server:~# ping 192.168.1.101
```

```
PING 192.168.1.101 (192.168.1.101) 56(84) bytes of data.
```

```
64 bytes from 192.168.1.101: icmp_seq=1 ttl=64 time=0.151 ms
```

```
64 bytes from 192.168.1.101: icmp_seq=2 ttl=64 time=0.116 ms
```

```
64 bytes from 192.168.1.101: icmp_seq=3 ttl=64 time=0.128 ms
```

```
64 bytes from 192.168.1.101: icmp_seq=4 ttl=64 time=0.118 ms
```

64 bytes from 192.168.1.101: icmp_seq=5 ttl=64 time=0.122 ms
64 bytes from 192.168.1.101: icmp_seq=6 ttl=64 time=0.084 ms
64 bytes from 192.168.1.101: icmp_seq=7 ttl=64 time=0.111 ms
64 bytes from 192.168.1.101: icmp_seq=8 ttl=64 time=0.082 ms
64 bytes from 192.168.1.101: icmp_seq=9 ttl=64 time=0.100 ms
64 bytes from 192.168.1.101: icmp_seq=10 ttl=64 time=0.106 ms
64 bytes from 192.168.1.101: icmp_seq=11 ttl=64 time=0.130 ms
64 bytes from 192.168.1.101: icmp_seq=12 ttl=64 time=0.083 ms
64 bytes from 192.168.1.101: icmp_seq=13 ttl=64 time=0.132 ms
64 bytes from 192.168.1.101: icmp_seq=14 ttl=64 time=0.121 ms
64 bytes from 192.168.1.101: icmp_seq=15 ttl=64 time=0.125 ms

--- 192.168.1.101 ping statistics ---

15 packets transmitted, 15 received, 0% packet loss, time 13997ms

rtt min/avg/max/mdev = 0.082/0.113/0.151/0.024 ms

server:~#

For SOME reason I can actually ping the container and get valid responses back from another machine within the network. What's going on? I have checked, and there is no firewalls on the host machines, the routers allow for those ip addresses that are specified (otherwise there wouldn't be a response when I ping it from another machine).

Subject: Re: cannot ping internet from ve

Posted by [cesurasean](#) on Mon, 09 Jun 2008 09:39:25 GMT

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

Good ole' howtoforge. I was able to get it working via their instructions here:

<http://www.howtoforge.com/installing-and-using-openvz-on-debian-etch>

Also, maybe the nameservers might have been wrong? I added my router's address, and my loopback address to my nameservers this time.

Subject: Re: cannot ping internet from ve

Posted by [cesurasean](#) on Mon, 09 Jun 2008 10:26:55 GMT

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

It seems this is still happening on the machine running RHEL. The debian machine cleared up fine. Tried install Ubuntu, and Debian on the RHEL host. Neither will connect to the internet following the same instructions as given in the notes on those websites listed above.

Subject: Re: cannot ping internet from ve
Posted by [maratrus](#) on Mon, 09 Jun 2008 12:28:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

http://forum.openvz.org/index.php?t=tree&goto=27545&#msg_27545

Have you read?

This information will be useful:

- 1) "ip a l" (from HN and from VE)
- 2) "ip r l" (from HN and from VE)
- 3) sysctl parameters (/etc/sysctl.conf)
- 4) "arp -n" from HN

5) tcpdump output:

ping from inside the VE: ping x.x.x.x and at the same moment issue "tcpdump" on the HN and inside the VE.

Subject: Re: cannot ping internet from ve
Posted by [cesurasean](#) on Mon, 09 Jun 2008 21:44:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

#VE

root@test:/# ip rule list

0: from all lookup local

32766: from all lookup main

32767: from all lookup default

root@test:/# ip route list table all

192.0.2.1 dev venet0 scope link

default via 192.0.2.1 dev venet0

broadcast 127.255.255.255 dev lo table local proto kernel scope link src 127.0.0.1

local 75.125.210.115 dev venet0 table local proto kernel scope host src 75.125.210.115

broadcast 127.0.0.0 dev lo table local proto kernel scope link src 127.0.0.1

local 127.0.0.1 dev lo table local proto kernel scope host src 127.0.0.1

local 127.0.0.1 dev venet0 table local proto kernel scope host src 127.0.0.1

local 127.0.0.0/8 dev lo table local 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 local proto none metric 0 mtu 16436 advmss 16376 hoplimit 4294967295

unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255

```
root@test:/# ip a l
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 75.125.210.115/32 scope global venet0:0
```

```
root@test:/# ip r l
192.0.2.1 dev venet0 scope link
default via 192.0.2.1 dev venet0
```

```
root@test:/# tcpdump
-bash: tcpdump: command not found
```

#HN

```
[root@dreamintsadecv cache]# ip rule list
```

```
0:    from all lookup 255
```

```
32766: from all lookup main
```

```
32767: from all lookup default
```

```
[root@dreamintsadecv cache]# ip route list table all
```

```
75.125.210.115 dev venet0 scope link
```

```
75.125.210.112/28 dev eth0 proto kernel scope link src 75.125.210.114
```

```
169.254.0.0/16 dev eth0 scope link
```

```
default via 75.125.210.113 dev eth0
```

```
broadcast 75.125.210.127 dev eth0 table 255 proto kernel scope link src 75.125.210.114
```

```
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
```

```
local 75.125.210.125 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.124 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.123 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.122 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.121 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.120 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.119 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.118 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.117 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.116 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.115 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
local 75.125.210.114 dev eth0 table 255 proto kernel scope host src 75.125.210.114
```

```
broadcast 127.0.0.0 dev lo table 255 proto kernel scope link src 127.0.0.1
```

```
broadcast 75.125.210.112 dev eth0 table 255 proto kernel scope link src 75.125.210.114
```

```
local 127.0.0.1 dev lo 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 ::/96 dev lo metric 1024 expires 21332063sec error -101 mtu 16436 advmss 16376
```

```
hoplimit 4294967295
```

```
unreachable ::ffff:0.0.0.0/96 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
```

```
advmss 16376 hoplimit 4294967295
```

```
unreachable 2002:a00::/24 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
```

```
advmss 16376 hoplimit 4294967295
```

```

unreachable 2002:7f00::/24 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
advmtss 16376 hoplimit 4294967295
unreachable 2002:a9fe::/32 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
advmtss 16376 hoplimit 4294967295
unreachable 2002:ac10::/28 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
advmtss 16376 hoplimit 4294967295
unreachable 2002:c0a8::/32 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
advmtss 16376 hoplimit 4294967295
unreachable 2002:e000::/19 dev lo metric 1024 expires 21332063sec error -101 mtu 16436
advmtss 16376 hoplimit 4294967295
unreachable 3ffe:ffff::/32 dev lo metric 1024 expires 21332063sec error -101 mtu 16436 advmtss
16376 hoplimit 4294967295
fe80::/64 dev eth0 metric 256 expires 21332060sec mtu 1500 advmtss 1440 hoplimit
4294967295
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 advmtss 16376 hoplimit
4294967295
local fe80::219:b9ff:fefe:a654 via :: dev lo table 255 proto none metric 0 mtu 16436 advmtss
16376 hoplimit 4294967295
ff00::/8 dev eth0 table 255 metric 256 expires 21332060sec mtu 1500 advmtss 1440 hoplimit
4294967295
unreachable default dev lo table unspec proto none metric -1 error -101 hoplimit 255

```

```

[root@dreamintsadecv cache]# iptables -t nat -L && iptables -t filter -L && iptables -t mangle -L
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

```

```

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

```

```

Chain PREROUTING (policy ACCEPT)

```

```

target    prot opt source                destination

```

```

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

Chain POSTROUTING (policy ACCEPT)
target prot opt source destination

```
[root@dreamintsadecv cache]# iptables -F  
[root@dreamintsadecv cache]#
```

```
[root@dreamintsadecv cache]# ip a l  
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 00:19:b9:fe:a6:54 brd ff:ff:ff:ff:ff:ff  
    inet 75.125.210.114/28 brd 75.125.210.127 scope global eth0  
    inet 75.125.210.115/28 brd 75.125.210.127 scope global secondary eth0:0  
    inet 75.125.210.116/28 brd 75.125.210.127 scope global secondary eth0:1  
    inet 75.125.210.117/28 brd 75.125.210.127 scope global secondary eth0:2  
    inet 75.125.210.118/28 brd 75.125.210.127 scope global secondary eth0:3  
    inet 75.125.210.119/28 brd 75.125.210.127 scope global secondary eth0:4  
    inet 75.125.210.120/28 brd 75.125.210.127 scope global secondary eth0:5  
    inet 75.125.210.121/28 brd 75.125.210.127 scope global secondary eth0:6  
    inet 75.125.210.122/28 brd 75.125.210.127 scope global secondary eth0:7  
    inet 75.125.210.123/28 brd 75.125.210.127 scope global secondary eth0:8  
    inet 75.125.210.124/28 brd 75.125.210.127 scope global secondary eth0:9  
    inet 75.125.210.125/28 brd 75.125.210.127 scope global secondary eth0:10  
    inet6 fe80::219:b9ff:fefe:a654/64 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@dreamintsadecv cache]# ip r l
75.125.210.115 dev venet0 scope link
75.125.210.112/28 dev eth0 proto kernel scope link src 75.125.210.114
169.254.0.0/16 dev eth0 scope link
default via 75.125.210.113 dev eth0
```

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.

# Controls IP packet forwarding
#net.ipv4.ip_forward = 0

# Controls source route verification
net.ipv4.conf.default.rp_filter = 1

# Do not accept source routing
net.ipv4.conf.default.accept_source_route = 0

# Controls the System Request debugging functionality of the kernel
#kernel.sysrq = 0

# Controls whether core dumps will append the PID to the core filename
# Useful for debugging multi-threaded applications
kernel.core_uses_pid = 1

# Controls the use of TCP syncookies
net.ipv4.tcp_syncookies = 1

# Controls the maximum size of a message, in bytes
kernel.msgmnb = 65536

# Controls the default maximum size of a message queue
kernel.msgmax = 65536

# Controls the maximum shared segment size, in bytes
kernel.shmmax = 4294967295

# Controls the maximum number of shared memory segments, in pages
kernel.shmall = 268435456
```



```
# On Hardware Node we generally need
# packet forwarding enabled and proxy arp disabled
net.ipv4.ip_forward = 1
net.ipv4.conf.default.proxy_arp = 0
# Enables source route verification
net.ipv4.conf.all.rp_filter = 1
net.ipv4.icmp_echo_ignore_broadcasts = 1
net.ipv4.conf.default.forwarding = 1
# Enables the magic-sysrq key
kernel.sysrq = 1
# TCP Explicit Congestion Notification
#net.ipv4.tcp_ecn = 0
# we do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

```
[root@dreamintsadecv cache]# arp -n
```

Address	HWtype	HWaddress	Flags	Mask	Iface
75.125.210.113	ether	00:1C:58:31:53:7F	C		eth0
75.125.210.115	*	*	MP	eth0	

```
[root@dreamintsadecv cache]# tcpdump -i venet0 -e host 204.152.191.5
tcpdump: WARNING: arptype 65535 not supported by libpcap - falling back to cooked socket
tcpdump: WARNING: venet0: no IPv4 address assigned
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
16:22:57.397488 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org:
ICMP echo request, id 21807, seq 1, length 64
16:22:58.398286 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org:
ICMP echo request, id 21807, seq 2, length 64
```

16:22:59.397486 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 3, length 64
16:23:00.397673 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 4, length 64
16:23:01.397894 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 5, length 64
16:23:02.398032 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 6, length 64
16:23:03.397255 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 7, length 64
16:23:04.397433 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 8, length 64
16:23:05.397644 In ethertype IPv4 (0x0800), length 100: test.domain.com > pub1.kernel.org: ICMP echo request, id 21807, seq 9, length 64

I think this is the best I can do this time. If I can run anymore other tests, please tell me. Can someone PLEASE tell me why VE can't access the internet?

Subject: Re: cannot ping internet from ve
Posted by [curx](#) on Mon, 09 Jun 2008 22:35:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

have you read:

http://wiki.openvz.org/Unrouteable_private_ip_address
http://wiki.openvz.org/Using_NAT_for_container_with_private_IPs

Subject: Re: cannot ping internet from ve
Posted by [cesurasean](#) on Mon, 09 Jun 2008 22:56:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Why would I want to use a private IP instead of the public IPs I have? I have plenty of public ips to use...

and like I said, I'm able to ping that other ip address.

should I just use a local ip, and then forward the main ip to that address or something?

Subject: Re: cannot ping internet from ve
Posted by [cesurasean](#) on Mon, 09 Jun 2008 23:05:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

I setup another container with ip address 192.168.1.100, and it's still unable to access the internet.

Subject: Re: cannot ping internet from ve
Posted by [curx](#) on Mon, 09 Jun 2008 23:49:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

###

container : (Debian ?)

root@test:/# ip route list table all

local 75.125.210.115 dev venet0 table local proto kernel scope host src 75.125.210.115

###

hostnode : (RHEL)

[root@dreamintsadecv cache]# ip route list table all

75.125.210.115 dev venet0 scope link

(...)

local 75.125.210.115 dev eth0 table 255 proto kernel scope host src 75.125.210.114

(...)

@cesurasean:

Have you configured a eth0:x alias on your hostnode with the same ip like your container ???

Subject: Re: cannot ping internet from ve
Posted by [curx](#) on Mon, 09 Jun 2008 23:58:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

cesurasean wrote on Tue, 10 June 2008 01:05I setup another container with ip address 192.168.1.100, and it's still unable to access the internet.

And have you configured SNAT (Source Network Address Translation, also known as IP masquerading) to access a private IP to external NET like:

On your Hostnode "dreamintsadecv":

```
# iptables -t nat -A POSTROUTING -s 192.168.1.0/24 -o eth0 -j SNAT --to  
<configured_ip_on_your_hostnode>
```

Subject: Re: cannot ping internet from ve
Posted by [cesurasean](#) on Tue, 10 Jun 2008 00:02:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

I have now added:

```
iptables -t nat -A POSTROUTING -s 192.168.1.0/24 -o eth0 -j SNAT --to  
<configured_ip_on_your_hostnode>
```

Still no success on the ve accessing internet.

Subject: Re: cannot ping internet from ve
Posted by [cesurasean](#) on Tue, 10 Jun 2008 00:19:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

problem has been solved. the public ip address in question was being used by the host machine on eth0:0. I had to turn this off, and it worked like a charm.

Thanks to dgym on irc.freenode.net #openvz
