
Subject: Networking Issue from inside VE
Posted by [fredtt](#) on Tue, 09 Aug 2011 16:47:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello everyone,

I've been trying to get network access from within my VE Working for the past 48 hours. I've searched this forum and google, tried many different solutions, and even started from scratch twice. Here is what I got...

From within the VE:
[root@test /]# ping google.com
ping: unknown host google.com

```
tracert 8.8.8.8
traceroute to 8.8.8.8 (8.8.8.8), 30 hops max, 40 byte packets
 1 * * *
 2 * * *
 3 * * *
 4 * * *
```

```
[root@test /]# ip route list table all
192.0.2.0/24 dev venet0 scope host
169.254.0.0/16 dev venet0 scope link
default dev venet0 scope link
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link src 127.0.0.1
local 67.227.246.10 dev venet0 table 255 proto kernel scope host src 67.227.246.10
broadcast 67.227.246.10 dev venet0 table 255 proto kernel scope link src 67.227.246.10
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
```

Running a TCPDUMP on the HN while the VE is trying to ping 8.8.8.8

```
[root@cloud2 ~]# tcpdump -i venet0 -e host 8.8.8.8
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
12:34:41.436241 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 1, length 64
12:34:42.435563 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 2, length 64
```

```
12:34:43.435571 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 3, length 64
12:34:44.435581 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 4, length 64
12:34:45.435590 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 5, length 64
12:34:46.435599 In ethertype IPv4 (0x0800), length 100: 67.227.246.10 >
google-public-dns-a.google.com: ICMP echo request, id 55301, seq 6, length 64
```

On the HN

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

```
# Kernel sysctl configuration file for Red Hat Linux
```

```
#
```

```
# For binary values, 0 is disabled, 1 is enabled. See sysctl(8) and
```

```
# sysctl.conf(5) for more details.
```

```
# Controls IP packet forwarding
```

```
net.ipv4.ip_forward = 1
```

```
net.ipv4.conf.default.proxy_arp = 0
```

```
# Controls source route verification
```

```
net.ipv4.conf.default.rp_filter = 1
```

```
net.ipv4.conf.all.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 = 1
```

```
# 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
```

```
# We do not want all our interfaces to send redirects
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
```

On the VE

```
[root@test /]# cat /etc/resolv.conf
nameserver 8.8.8.8
nameserver 8.8.4.4
```

On the HN

```
[root@cloud2 ~]# service iptables restart
Flushing firewall rules: [ OK ]
Setting chains to policy ACCEPT: filter [ OK ]
Unloading iptables modules: [ OK ]
Applying iptables firewall rules: [ OK ]
Loading additional iptables modules: ipt_REJECT ipt_tos ipt_TOS ipt_LOG ip_conntrack ipt_limit
ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_length ipt_state
iptables_nat ip_nat_ftp ipt_owner ipt_REDIRECT ip_conntrack_ftp [ OK ]
[root@cloud2 ~]# service network restart
Shutting down interface eth0: [ OK ]
Shutting down loopback interface: [ OK ]
Disabling IPv4 packet forwarding: net.ipv4.ip_forward = 0
[ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0: [ OK ]
Bringing up interface venet0: Bringing up interface venet0:
Configuring interface venet0:
net.ipv4.conf.venet0.send_redirects = 0
[ OK ]
```

Any ideas? I've tried disabling iptables on the HN and still no luck.

Subject: Re: Networking Issue from inside VE
Posted by [curx](#) on Tue, 09 Aug 2011 22:03:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

---8<...
(...)

```
[root@cloud2 ~]# service network restart
Shutting down interface eth0:          [ OK ]
Shutting down loopback interface:      [ OK ]
Disabling IPv4 packet forwarding: net.ipv4.ip_forward = 0
(...)
---8<...
```

- please check the sysctl ip_forwarding, is set to 1 after all ?
- any ct iptables rule set ?
- on your hostnode, any outgoing traffic to any host are ok ?

Bye,
Thorsten

Subject: Re: Networking Issue from inside VE
Posted by [fredtt](#) on Tue, 09 Aug 2011 22:35:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quote:- please check the sysctl ip_forwarding, is set to 1 after all ?

Yes it is. That notice is during the "shutdown" portion of the network service, not during the boot up.

Quote:- any ct iptables rule set ?

Nope.

Quote:- on your hostnode, any outgoing traffic to any host are ok ?

Everything on the hostnode works just fine.

~Fred
