
Subject: Network problems (git kernel 2.4.24).

Posted by [Rafael Isturiz](#) on Wed, 02 Apr 2008 15:45:40 GMT

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

Hi,

I have network problem with openvz... I can't ping outside the HN with VE (i have stable openvz working with almost the same config in another server).

uname -a

Linux proxy 2.6.24-openvz1 #1 SMP Tue Apr 1 14:47:32 VET 2008 x86_64 GNU/Linux

ip route list table all

```
192.168.0.250 dev venet0 scope link src 192.168.0.200
192.168.0.201 dev venet0 scope link src 192.168.0.200
192.168.0.0/24 dev eth1 proto kernel scope link src 192.168.0.200
default via 192.168.0.1 dev eth1
broadcast 192.168.0.255 dev eth1 table 255 proto kernel scope link
src 192.168.0.200
broadcast 127.255.255.255 dev lo table 255 proto kernel scope link
src 127.0.0.1
local 192.168.0.200 dev eth1 table 255 proto kernel scope host src
192.168.0.200
broadcast 192.168.0.0 dev eth1 table 255 proto kernel scope link src
192.168.0.200
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 127.0.0.0/8 dev lo table 255 proto kernel scope host src
127.0.0.1
local ::1 via :: dev lo proto none metric 0 mtu 16436 advmss 16376
hoplimit 4294967295
local fe80::21a:64ff:fe98:656 via :: dev lo proto none metric 0 mtu
16436 advmss 16376 hoplimit 4294967295
fe80::/64 dev eth1 metric 256 expires 8567024sec mtu 1500 advmss 1440
hoplimit 4294967295
ff00::/8 dev eth1 metric 256 expires 8567024sec mtu 1500 advmss 1440
hoplimit 4294967295
unreachable default dev lo table unspec proto none metric -1 error
-101 hoplimit 255
```

iptables -L

Chain INPUT (policy ACCEPT)

| | | | |
|--------|----------|--------|-------------|
| target | prot opt | source | destination |
|--------|----------|--------|-------------|

Chain FORWARD (policy DROP)

```
target    prot opt source          destination
```

Chain OUTPUT (policy ACCEPT)

```
target    prot opt source          destination
```

```
tcpdump -i venet0 -e host 192.168.0.6
```

```
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
```

```
04:51:22.866438 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 1, length 64
```

```
04:51:23.868812 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 2, length 64
```

```
04:51:24.871478 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 3, length 64
```

```
04:51:25.874154 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 4, length 64
```

```
04:51:26.876810 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 5, length 64
```

```
04:51:27.879466 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 6, length 64
```

```
04:51:28.882174 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 7, length 64
```

```
04:51:29.884850 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 8, length 64
```

```
04:51:30.887506 In ethertype IPv4 (0x0800), length 100: 192.168.0.201 >  
192.168.0.6: ICMP echo request, id 31489, seq 9, length 64
```

```
9 packets captured
```

```
9 packets received by filter
```

```
0 packets dropped by kernel
```

(This when i ping inside the VE to 192.168.0.6)

```
ping 192.168.0.6
```

```
PING 192.168.0.6 (192.168.0.6) 56(84) bytes of data.
```

```
--- 192.168.0.6 ping statistics ---
```

```
12 packets transmitted, 0 received, 100% packet loss, time 11029ms
```

I can ping all VE's to HN... But not ping between them.

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

```
-----
```

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

```
net.ipv4.conf.default.proxy_arp = 0
net.ipv4.conf.all.rp_filter = 1
kernel.sysrq = 1
net.ipv4.conf.default.send_redirects = 1
net.ipv4.conf.all.send_redirects = 0
-----
```

```
/etc/network/interfaces
-----
```

```
auto lo
iface lo inet loopback
```

```
allow-hotplug eth0
iface eth0 inet dhcp
```

```
auto eth1
iface eth1 inet static
    address 192.168.0.200
    netmask 255.255.255.0
    broadcast 192.168.0.255
    gateway 192.168.0.1
    up sysctl -w net.ipv4.conf.eth1.proxy_arp=100
    pre-down sysctl -w net.ipv4.conf.eth1.proxy_arp=0
-----
```

Thanks for any help.

--

a.k.a. pelox (Debian Powered) | JID:risturiz@gusl.org.ve
Barquisimeto [VE] | LinuxUser [238188]
<http://pelox.gusl.org.ve>
