
Subject: can only ping locally - not network
Posted by [john32](#) on Tue, 01 Jul 2008 08:58:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
uname -a  
Linux deb-prq1 2.6.16-026test020.1-ovz-i386-smp #1 SMP Thu Oct 26 21:57:45 UTC 2006 i686  
GNU/Linux
```

```
vzctl --version  
vzctl version 3.0.22
```

I did an apt-get upgrade, and it installed new versions of openvz and other various system files. After the installation was complete I rebooted for good measure, and the VMs turned back on properly, but I was unable to access them from the network/internet as before the upgrade.

I can ping VMs from other VMs, and ping the local machine from the VMs and vice versa - but anything on the network can only access the physical machine, not any of the VMs.

Each VM has a unique public IP.

```
cat /proc/sys/net/ipv4/ip_forward  
1
```

Access to the physical machine seems unaffected, so I am hoping someone more familiar with OpenVZ can tell me what the problem is with my VMs...

Subject: Re: can only ping locally - not network
Posted by [maratrus](#) on Tue, 01 Jul 2008 11:52:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Quote:
2.6.16-026test020.1-ovz-i386-smp

From here <http://wiki.openvz.org/Download/kernel>

Quote:
These branches are not developed/supported anymore. They are here mostly for the historical reasons. Do not use it.

please show the following information:

1. "ip a l" from inside the VE and from inside the HN.
 2. "ip r l" from HN and from inside the VE
 3. "arp -n" from HN
 4. check your iptables rules
 5. if it is possible try to investigate with tcpdump utility whether your HN receive the packets.
- Briefly speaking, try to ping your VE from outside and at the same moment start tcpdump on the HN and inside VE simultaneously.
-

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Tue, 01 Jul 2008 16:42:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

I'm not surprised the kernel version is outdated - I will look into upgrading it once I figure out what the problem is with this.

in these examples XXX.XXX.XXX.130 is the VE
XXX.XXX.XXX.94 is the HN

From HN:

ip a l

```
2: lo: <LOOPBACK,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: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
6: eth3: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:18:f3:b6:35:09 brd ff:ff:ff:ff:ff:ff
    inet XXX.XXX.XXX.94/26 brd XXX.XXX.XXX.255 scope global eth3
    inet XXX.XXX.XXX.131/26 brd XXX.XXX.XXX.255 scope global eth3:0
    inet6 fe80::218:f3ff:feb6:3509/64 scope link
        valid_lft forever preferred_lft forever
8: eth1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:18:f3:b6:25:05 brd ff:ff:ff:ff:ff:ff
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
```

```
XXX.XXX.XXX.130 dev venet0 scope link
XXX.XXX.XXX.64/26 dev eth3 proto kernel scope link src XXX.XXX.XXX.94
XXX.XXX.XXX.128/26 dev eth3 proto kernel scope link src XXX.XXX.XXX.131
default via XXX.XXX.XXX.65 dev eth3 src XXX.XXX.XXX.131
default via XXX.XXX.XXX.65 dev eth3
```

```
arp -n
Address          HWtype HWaddress      Flags Mask    Iface
XXX.XXX.XXX.65   ether  00:04:23:AA:FF:DB  C           eth3
```

```
iptables -L
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
```

From VE:

```
ip a l
1: lo: <LOOPBACK,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
3: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
    inet 127.0.0.1/32 scope host venet0
    inet XXX.XXX.XXX.130/32 scope global venet0:0
```

```
ip r l
192.0.2.1 dev venet0 scope link
default via 192.0.2.1 dev venet0
```

```
iptables -L
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
```

I may be able to check the tcpdump results a little bit later, but in the meantime maybe someone can spot a problem with some of those results. Thanks!

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Thu, 03 Jul 2008 22:59:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Okay, I upgraded my kernel

```
uname -a
Linux deb-prq1 2.6.18-12-fza-686-bigmem #1 SMP Sun May 18 13:01:05 CEST 2008 i686
GNU/Linux
```

but I still have the same issue as before - nothing from the lan/internet is reaching the VE, but I can ping the VE from the HN just fine.

I ran tcpdump on the HN and pinged the VE a few times, here are the results:

```
tcpdump
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
```

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

Subject: Re: can only ping locally - not network
Posted by [maratrus](#) on Fri, 04 Jul 2008 07:16:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

sorry for delay

please try this (on the HN)

```
ip neigh add proxy VE_IP dev eth0
```

does this help?

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Fri, 04 Jul 2008 20:57:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

I tried both eth3 (active lan, I think that's the one you meant), and eth0 just to be safe. Neither changed any behavior, and I don't see any trace of anything different in this either:

```
ip neigh show
XX.XX.X.65 dev eth3 lladdr xx:xx:xx:xx:xx:xx DELAY
```

So basically the host is not routing the traffic to the venet0 interface properly, is that right?

Subject: Re: can only ping locally - not network
Posted by [maratrus](#) on Wed, 09 Jul 2008 15:47:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

please sorry for delay
You've only shown "iptables -L"
Please show also "iptables -t nat -L", "iptables -t mangle -L".

P.S. Could you give me an access to your HN via private message? And to the node from which I'll be able to ping your VE to test the connection between external node and your VE/HN.

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Wed, 09 Jul 2008 19:00:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
iptables -t nat -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
```

```
iptables -t mangle -L
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

I can give you any information you need, but I don't feel comfortable giving access to a HN that I don't have physical access to myself.

Subject: Re: can only ping locally - not network

Posted by [maratrus](#) on Tue, 15 Jul 2008 10:22:41 GMT

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

Hello,

did you manage to solve the problem?

If yes, could you possibly put here the answer?

If no, could you possibly conduct the following experiments:

1. ping your VE from external node. (please, show what IP address does it have). At the same moment run

- tcpdump -n -i eth0 (on the external node) (if it has eth0 active interface)
- tcpdump -n -i eth3 (on the HN)
- tcpdump -n -i venet0 (on the HN)
- tcpdump -n -i venet0 (inside the VE)

2. Please show "arp -n" from the HN

then do "ip neigh add proxy VE_IP dev eth3"

and show "arp -n" again. The output should contain the following string.

VE_IP	*	*	MP	eth3
-------	---	---	----	------

Subject: Re: can only ping locally - not network

Posted by [john32](#) on Thu, 17 Jul 2008 22:30:08 GMT

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

Please note, I used a different VE on the same HN for these tests, that is why the IP is 132 instead of 130.

Here is the arp after the neighbor add - before the add it only contained the first line.

```
arp -n
```

Address	HWtype	HWaddress	Flags	Mask	Iface
XX.XX.XX.65	ether	00:04:23:AA:FF:DB	C		eth3
XX.XX.XX.132	*	<from_interface>	MP		eth3

tcpdump on the external node shows the ping requests outgoing and host unreachablees.

tcpdump on the HN shows one line like this for each ping:

```
arp who-has XX.XX.XX.132 tell XX.XX.XX.65
```

There were never any arp replies in the logs, which seems like the problem, though I am not sure how to fix this.

tcpdump on venet0 on HN does not receive any packets

I do not have tcpdump installed on the VE, so I cannot test from inside, but since the HN venet0 does not see any packets I doubt the VE would either.

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Fri, 18 Jul 2008 02:00:40 GMT

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

after a bit of messing about I found a temporary solution - though I still don't know why it doesn't work by default as it did before.

if I issue the following command I am then able to ping the VE both internally and externally. However, the VE still is not able to ping or access the net from inside it - so there must be something slightly different I need to do.

```
ip route add local XXX.XXX.XXX.132 dev eth3
```

Subject: Re: can only ping locally - not network
Posted by [maratrus](#) on Fri, 18 Jul 2008 07:41:07 GMT

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

Hi,

"ip rule list" from HN
"ip route list table all"?

Subject: Re: can only ping locally - not network
Posted by [john32](#) on Fri, 18 Jul 2008 18:09:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

ip rule list

0: from all lookup local

32766: from all lookup main

32767: from all lookup default

ip route list table all

XXX.XXX.XXX.132 dev venet0 scope link

XXX.XXX.XXX.131 dev venet0 scope link

XXX.XXX.XXX.130 dev venet0 scope link

XXX.XXX.XXX.64/26 dev eth3 proto kernel scope link src XXX.XXX.XXX.94

XXX.XXX.XXX.128/26 dev eth3 proto kernel scope link src XXX.XXX.XXX.131

default via XXX.XXX.XXX.65 dev eth3 src XXX.XXX.XXX.130

default via XXX.XXX.XXX.65 dev eth3 src XXX.XXX.XXX.131

default via XXX.XXX.XXX.65 dev eth3

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

broadcast XXX.XXX.XXX.64 dev eth3 table local proto kernel scope link src XXX.XXX.XXX.94

local XXX.XXX.XXX.131 dev eth3 table local proto kernel scope host src XXX.XXX.XXX.131

local XXX.XXX.XXX.130 dev eth3 table local proto kernel scope host src XXX.XXX.XXX.131

broadcast XXX.XXX.XXX.255 dev eth3 table local proto kernel scope link src XXX.XXX.XXX.94

broadcast XXX.XXX.XXX.255 dev eth3 table local proto kernel scope link src

XXX.XXX.XXX.131

broadcast XXX.XXX.XXX.127 dev eth3 table local proto kernel scope link src XXX.XXX.XXX.94

local XXX.XXX.XXX.94 dev eth3 table local proto kernel scope host src XXX.XXX.XXX.94

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 127.0.0.0/8 dev lo table local proto kernel scope host src 127.0.0.1

local ::1 via :: dev lo proto none metric 0 mtu 16436 rtt 9ms rttvar 15ms cwnd 99 advmss 16376

hoplimit 4294967295

local fe80::218:f3ff:feb6:3509 via :: dev lo proto none metric 0 mtu 16436 advmss 16376

hoplimit 4294967295

fe80::/64 dev eth3 metric 256 expires 2017154sec mtu 1500 advmss 1440 hoplimit 4294967295

ff00::/8 dev eth3 metric 256 expires 2017154sec mtu 1500 advmss 1440 hoplimit 4294967295

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

132 is the only real VE on that list currently - 130 and 131 are added in the interfaces file on the HN as secondary IPs for eth3.

Also I got a warning (which I've never seen before) when I restarted 132 that proxy_arp was set to 0, so I set all instances of proxy_arp to 1 but that had no effect.

EDIT: somehow since posting on here the ip_forward got turned off (maybe the kernel upgrade you suggested changed my sysctl.conf?)

anyway, after adding these lines to sysctl.conf and rebooting, the VEs work again. Thanks for your help.

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

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

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