
Subject: Networking issue with FC4

Posted by [wasrv](#) on Fri, 03 Nov 2006 13:11:38 GMT

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

Hello,

The installation of OpenVZ on the server was a breeze, both the kernel replacement and the daemon installation:

```
[07:59:05][~]> rpm -qa | grep vz
```

```
vzquota-3.0.9-1
```

```
vzpkg-2.7.0-18
```

```
vzctl-lib-3.0.12-1
```

```
vzctl-3.0.12-1
```

```
vzrpm44-4.4.1-22.5
```

```
vzrpm44-python-4.4.1-22.5
```

```
vztmpl-fedora-core-5-2.0-2
```

```
ovzkernel-2.6.9-023stab032.1
```

```
vzyum-2.4.0-11
```

```
vzrpm43-python-4.3.3-7_nonptl.6
```

```
[07:59:12][~]> uname -a
```

```
Linux X.Y.Z 2.6.9-023stab032.1 #1 Fri Oct 20 02:48:36 MSD 2006 i686 athlon i386 GNU/Linux
```

I modified the kernel tuning parameters as mentioned in the guide and networking parameters are all set:

```
[08:03:26][~]> cat /proc/sys/net/ipv4/ip_forward
```

```
1
```

```
[08:03:33][~]> iptables -L
```

```
Chain FORWARD (policy ACCEPT)
```

```
target    prot opt source                destination
```

```
Chain INPUT (policy ACCEPT)
```

```
target    prot opt source                destination
```

```
Chain OUTPUT (policy ACCEPT)
```

```
target    prot opt source                destination
```

I have created a FC5 VE and added it an IP as described in your guide, however networking doesn't work. Ifconfig & route on the VE lists the following info:

```
-bash-3.1# ifconfig -a
```

```
lo        Link encap:Local Loopback
```

```
inet addr:127.0.0.1 Mask:255.0.0.0
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:7 errors:0 dropped:0 overruns:0 frame:0
TX packets:7 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:1084 (1.0 KiB) TX bytes:1084 (1.0 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:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:164 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 b) TX bytes:23688 (23.1 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:11.22.33.44 P-t-P:11.22.33.44 Bcast:11.22.33.44 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

```
-bash-3.1# route -n
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
191.255.255.0	0.0.0.0	255.255.255.0	U 0 0	0	0	0	venet0
169.254.0.0	0.0.0.0	255.255.0.0	U 0 0	0	0	0	venet0
0.0.0.0	191.255.255.1	0.0.0.0	UG 0 0	0	0	0	venet0

Where 11.22.33.44 is the IP I added for the VE. Now here's the output for the HN:

```
[08:09:05][~]> ifconfig -a
```

```
eth0 Link encap:Ethernet HWaddr 00:00:21:F4:53:1F
inet addr:11.22.33.99 Bcast:11.22.33.255 Mask:255.255.255.128
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:1081802 errors:0 dropped:0 overruns:0 frame:0
TX packets:300108 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:88613125 (84.5 MiB) TX bytes:185321665 (176.7 MiB)
Interrupt:17 Base address:0xd000
```

```
lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:33065 errors:0 dropped:0 overruns:0 frame:0
TX packets:33065 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:3462234 (3.3 MiB) TX bytes:3462234 (3.3 MiB)
```

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

UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:164 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:23688 (23.1 KiB) TX bytes:0 (0.0 b)

Where 11.22.33.99 is my main IP address.

After browsing these forums, I tried several tricks but none worked (the preserve arguments for /bin/cf and the early load of the ip_contrack module).

Any suggestions for my problem?
Thank you and for your great product.

Subject: Re: Networking issue with FC4
Posted by [Vasily Tarasov](#) on Fri, 03 Nov 2006 13:28:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you ping HW from VE?

Thanks.

Subject: Re: Networking issue with FC4
Posted by [wasrv](#) on Fri, 03 Nov 2006 13:48:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

Nope, the only successful ping I can do from the VE is pinging to its assigned IP.

Subject: Re: Networking issue with FC4
Posted by [wasrv](#) on Fri, 03 Nov 2006 13:54:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

In addition to my previous post, here's what tcpdump captured while pinging google:

HN:

```
[08:51:31][~]> tcpdump -n -i venet0
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
08:51:48.180973 IP 11.22.33.44 > 64.233.167.99: icmp 64: echo request seq 0
08:51:49.180063 IP 11.22.33.44 > 64.233.167.99: icmp 64: echo request seq 1
```

```
08:51:50.179916 IP 11.22.33.44 > 64.233.167.99: icmp 64: echo request seq 2
08:51:51.179758 IP 11.22.33.44 > 64.233.167.99: icmp 64: echo request seq 3
```

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

VE:

```
[08:50:59][~]> vzctl enter 102
entered into VE 102
-bash-3.00# ping 64.233.167.99
PING 64.233.167.99 (64.233.167.99) 56(84) bytes of data.
```

```
--- 64.233.167.99 ping statistics ---
4 packets transmitted, 0 received, 100% packet loss, time 2998ms
```

Subject: Re: Networking issue with FC4
Posted by [Vasily Tarasov](#) on Fri, 03 Nov 2006 14:06:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sorry if I misunderstood, you but:

As far as I understand your HN IP address is 11.22.33.99, not 64.233.167.99. So, please, ping you HN (11.22.33.99) from VE(11.22.33.44), and show us the tcpdump.

Thanks.

BTW, can you provide access to the node?

Subject: Re: Networking issue with FC4
Posted by [wasrv](#) on Fri, 03 Nov 2006 14:12:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Sorry, I misunderstood Here's the output:

VE:

```
-bash-3.00# ping 11.22.33.99
PING 11.22.33.99 (11.22.33.99) 56(84) bytes of data.
```

```
--- 11.22.33.99 ping statistics ---
4 packets transmitted, 0 received, 100% packet loss, time 2999ms
```

HN:

```
[09:09:48][~]> tcpdump -n -i venet0
```

```
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
```

```
09:10:36.381161 IP 11.22.33.44 > 11.22.33.99: icmp 64: echo request seq 0
```

```
09:10:37.380548 IP 11.22.33.44 > 11.22.33.99: icmp 64: echo request seq 1
```

```
09:10:38.380406 IP 11.22.33.44 > 11.22.33.99: icmp 64: echo request seq 2
```

```
09:10:39.380247 IP 11.22.33.44 > 11.22.33.99: icmp 64: echo request seq 3
```

```
4 packets captured
```

```
4 packets received by filter
```

```
0 packets dropped by kernel
```

11.22.33.44 - IP assigned to the VE

11.22.33.99 - main IP of the HN

About the access, it's a production server, I'm not allowed to

Subject: Re: Networking issue with FC4

Posted by [Vasily Tarasov](#) on Fri, 03 Nov 2006 14:27:32 GMT

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

In the subject of your post FC4 VE is pointed at, but in the actual post you mentioned FC5, so what template do you actually use? Are there other VEs (perhaps with other templates), that work with network well?

Thanks.

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Fri, 03 Nov 2006 14:36:08 GMT

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

HN runs FC4, while the I tested more templates, including FC5 I mentioned. CentOS fails as well.

Subject: Re: Networking issue with FC4

Posted by [Vasily Tarasov](#) on Fri, 03 Nov 2006 14:59:12 GMT

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

Hello,

can you post us output of the following command:

```
# ip r get 11.22.33.44
# route -an
```

Thanks.

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Fri, 03 Nov 2006 15:06:51 GMT

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

Here's the output after executing the commands on the HN. Please note that the "a" argument failed for my /sbin/route, and the ip command output is identical on the VE too. The route table of the VE id displayed in the first post of the thread.

```
[10:03:45][~]> ip r get 11.22.33.44
```

```
local 11.22.33.44 dev lo src 11.22.33.44
```

```
cache <local> mtu 16436 advmss 16396 metric 10 64
```

```
[10:04:11][~]> route -n
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
11.22.33.44	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
11.22.33.128	0.0.0.0	255.255.255.128	U	0	0	0	eth0
169.254.0.0	0.0.0.0	255.255.0.0	U	0	0	0	eth0
0.0.0.0	11.22.33.129	0.0.0.0	UG	0	0	0	eth0

Thank you

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Sat, 04 Nov 2006 11:38:28 GMT

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

Any suggestions for my problem?

I'm now 100% positive that the issue is not template dependent, I've tested all published distro templates and all get to be unconnected.

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Sun, 05 Nov 2006 14:17:45 GMT

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

I think I found something. While testing the suggested trick by using mod ip_conntrack I found that it is loaded AFTER ip_tables for the HN - a excerpt from dmesg follows:

```
Nov  1 16:40:24 master kernel: ip_tables: (C) 2000-2002 Netfilter core team
```

```
Nov  1 16:40:24 master kernel: NET: Unregistered protocol family 17
```

```
Nov  1 16:40:24 master kernel: NET: Registered protocol family 17
```

```
Nov  1 16:40:24 master kernel: ip_conntrack : parameter ip_conntrack_enable_ve0 is obsolete. In kernel >= 2.6.15 connection tracking on hardware node is enabled by default, use ip_conntrack_disable_ve0=1 parameter to disable.
```

```
Nov  1 16:40:24 master kernel: ip_conntrack version 2.1 (4095 buckets, 32760 max) - 328 bytes per conntrack
```

I just placed the options line at the beginning of /etc/modprobe.conf, what should I do to load it before ip_tables?

Thank you.

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Sun, 05 Nov 2006 14:56:35 GMT

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

I loaded ip_conntrack before the other ip_ modules but didn't work anyway

Subject: Re: Networking issue with FC4

Posted by [hbidad](#) on Mon, 06 Nov 2006 23:35:20 GMT

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

Seems I am having the exact same problem...any insight?

Subject: Re: Networking issue with FC4

Posted by [hbidad](#) on Mon, 06 Nov 2006 23:42:59 GMT

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

Although, if I remove the VE IP from the host, I can connect directly to the VE from a remote workstation. I thought all would be good after this, but I can not get the VE to resolve any domain (i.e. install cpanel)

I have setup my nameservers and verified them in resolv.conf. my nameservers run off a vlan (intranet IP's) The host resolves everything fine...but beyond this.. I am stumped

I am at SoftLayer datacenter..

Subject: Re: Networking issue with FC4

Posted by [Vasily Tarasov](#) on Tue, 07 Nov 2006 06:27:58 GMT

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

Hello,

hbidad, It's better to start new thread about your problem.

Thank,
Vasily.

Subject: Re: Networking issue with FC4

Posted by [Vasily Tarasov](#) on Tue, 07 Nov 2006 07:21:25 GMT

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

Hello,

Sorry for delay, we have some holidays here.

As you see the problem is in routing:

Quote:[10:03:45][~]> ip r get 11.22.33.44

local 11.22.33.44 dev lo src 11.22.33.44

cache <local> mtu 16436 advmss 16396 metric 10 64

[10:04:11][~]> route -n

Every packet that goes to VE (11.22.33.44) thanks to routing goes to lo interface, but should go to venet0.

The question is why this happens?

I'll think about it.

Subject: Re: Networking issue with FC4

Posted by [hbidad](#) on Tue, 07 Nov 2006 07:35:20 GMT

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

Vasily, I am sure you were talking to me about starting a new thread. However, our problems are exactly alike. I get the same output as the OP. Down to the routing problem as you described.

I was just trying to if eth0 being vlan and eth1 being public would have something to with it.

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Tue, 07 Nov 2006 10:59:27 GMT

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

Vasily, good find, I couldn't figure that myself in years. Unfortunately I'm not very good with these kind of tasks related to routing, so if you have any suggestions for me please let me know.

Also, I don't know if this helps, but in the HN route table (which I listed in a previous post) the network masks seem to be different, could this be a problem?

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
11.22.33.44	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
11.22.33.128	0.0.0.0	255.255.255.128	U	0	0	0	eth0

I mean 255.255.255.255 vs 255.255.255.128 for the actual internet connection ...

Thanks a lot!

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Sun, 26 Nov 2006 20:46:41 GMT

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

Any suggestions for my issue?

Thanks!

Subject: Re: Networking issue with FC4

Posted by [hbidad](#) on Mon, 27 Nov 2006 01:07:00 GMT

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

Hell, I had the exact problem as you (see above) .. I resolved my issue..

go to /etc/sysconfig/network-scripts/ and tell us the filenames that start with ifcfg-eth1* (e.g. ifcfg-eth1-range*)

Subject: Re: Networking issue with FC4

Posted by [wasrv](#) on Mon, 27 Nov 2006 01:12:21 GMT

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

[20:10:55][~/etc/sysconfig/network-scripts]> ls ifcfg-eth*

```
-rw-r--r-- 1 root root 142 Sep 28 2005 ifcfg-eth0  
-rw-r--r-- 1 root root 166 Sep 27 2005 ifcfg-eth1
```

Thanks for the relpy!

Both files configure real NICs, but only eth0 is cabled. However no *range* files there

Thanks again.
