

---

Subject: **\*SOLVED\*** disable venet?

Posted by [szocske](#) on Thu, 14 Sep 2006 21:00:38 GMT

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

---

Hi!

I'm new here, we are planning to set up openvz guests on debian (with debian inside) to create multiple independent test environments for our software product. As we attempt to bind to the interface, make vlans on it, and modify iptables rules on them, I assume we must use veth interfaces instead of the default venet.

So the question is, how do I disable the venet interfaces altogether to reduce the clutter?

Also, /etc/network/interfaces sais I should not edit it but only an interfaces.template file. Where is that file (host or guest?) and what is its format?

Is there any good guide I could read about such setups?

Oh, and when I try to add my veth<guestid>.0 interface to a bridge in the host system with "brctl add", I get "Invalid argument". What could I have done wrong?

Thanks for your time and this promising tool:

Gabor

---

---

Subject: Re: disable venet?

Posted by [szocske](#) on Fri, 15 Sep 2006 08:36:26 GMT

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

---

Here is the exact error message I get:

```
# brctl addbr br1
```

```
# brctl addif br1 veth101.0
```

```
can't add veth101.0 to bridge br1: Invalid argument
```

Any advice welcome!

---

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Fri, 15 Sep 2006 09:17:22 GMT

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

---

Can you please post here list of all interfaces in your system (ip a ls), you can change IP addresses to x.x.x.x

---

---

Subject: Re: disable venet?

---

On the host system:

```
ip a ls
```

```
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: eth0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:16:36:48:f8:40 brd ff:ff:ff:ff:ff:ff
    inet 192.168.100.131/24 brd 192.168.100.255 scope global eth0
    inet6 fe80::216:36ff:fe48:f840/64 scope link
        valid_lft forever preferred_lft forever
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
5: veth101.0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop
    link/ether 11:22:33:44:55:66 brd ff:ff:ff:ff:ff:ff
8: br1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop
    link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
```

```
dpkg -l | grep vz
```

```
ii kernel-image-2.6.8-stable-ovz-smp 022stab078.14      Linux kernel binary image for version
2.6.8-
ii kernel-patch-openvz              026test015-1      server virtualization solution - kernel patc
ii vzctl                            3.0.10-3          server virtualization solution - control too
ii vzquota                          3.0.8-2          server virtualization solution - quota tools
```

```
strace brctl addif br1 veth101.0
```

```
execve("/usr/sbin/brctl", ["brctl", "addif", "br1", "veth101.0"], [/* 15 vars */]) = 0
uname({sys="Linux", node="vservertest", ...}) = 0
brk(0) = 0x804f000
access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
mmap2(NULL, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7f66000
access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
open("/etc/ld.so.cache", O_RDONLY) = 3
fstat64(3, {st_mode=S_IFREG|0644, st_size=12704, ...}) = 0
mmap2(NULL, 12704, PROT_READ, MAP_PRIVATE, 3, 0) = 0xb7f62000
close(3) = 0
access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
open("/lib/libsysfs.so.2", O_RDONLY) = 3
read(3, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\300\35"..., 512) = 512
```

```

fstat64(3, {st_mode=S_IFREG|0644, st_size=36096, ...}) = 0
mmap2(NULL, 39108, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3, 0) =
0xb7f58000
mmap2(0xb7f61000, 4096, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x8) = 0xb7f61000
close(3) = 0
access("/etc/ld.so.nohwcap", F_OK) = -1 ENOENT (No such file or directory)
open("/lib/tls/libc.so.6", O_RDONLY) = 3
read(3, "\177ELF1\1\1\0\0\0\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\240O\1"... , 512) = 512
fstat64(3, {st_mode=S_IFREG|0755, st_size=1245676, ...}) = 0
mmap2(NULL, 1251484, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3, 0)
= 0xb7e26000
mmap2(0xb7f4e000, 28672, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x128) = 0xb7f4e000
mmap2(0xb7f55000, 10396, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0xb7f55000
close(3) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7e25000
mprotect(0xb7f4e000, 20480, PROT_READ) = 0
set_thread_area({entry_number:-1 -> 6, base_addr:0xb7e25940, limit:1048575, seg_32bit:1,
contents:0, read_exec_only:0, limit_in_pages:1, seg_not_present:0, useable:1}) = 0
munmap(0xb7f62000, 12704) = 0
socket(PF_INET, SOCK_STREAM, IPPROTO_IP) = 3
brk(0) = 0x804f000
brk(0x8070000) = 0x8070000
open("/proc/mounts", O_RDONLY) = 4
fstat64(4, {st_mode=S_IFREG|0444, st_size=0, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7f65000
read(4, "/dev/root / ext3 rw,data=ordered"... , 1024) = 410
close(4) = 0
munmap(0xb7f65000, 4096) = 0
lstat64("/sys/class/net", {st_mode=S_IFDIR|0755, st_size=0, ...}) = 0
access("/proc/net", R_OK) = 0
access("/proc/net/unix", R_OK) = 0
socket(PF_FILE, SOCK_DGRAM, 0) = 4
ioctl(4, SIOCGIFINDEX, {ifr_name="veth101.0", ifr_index=5}) = 0
close(4) = 0
ioctl(3, 0x89a2, 0xbff7a230) = -1 EINVAL (Invalid argument)
ioctl(3, SIOCDEVPRIVATE, 0xbff7a230) = -1 EINVAL (Invalid argument)
write(2, "can't add veth101.0 to bridge br"... , 52can't add veth101.0 to bridge br1: Invalid
argument
) = 52
exit_group(1) = ?
Process 6683 detached

```

guest 101 is an instance of

<http://download.openvz.org/template/precreated/debian-3.1-i386-minimal.tar.gz>

veth101.0 created like this:

```
vzctl set 101 --veth_add veth101.0,11:22:33:44:55:66,eth0,11:22:33:44:55:67 --save
```

Thank you for your kind attention.

---

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Fri, 15 Sep 2006 09:59:20 GMT

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

---

You can't use such MAC address (11:22:33:44:55:66) for ethernet device.

The thing is in first bit of first byte of MAC address - this bit has special meaning (broadcast).

So you can't set it to 1.

If you change MAC to 10:22:33:44:55:66 or to 00:22:33:44:55:66 everything will work perfect.

---

---

Subject: Re: disable venet?

Posted by [szocske](#) on Fri, 15 Sep 2006 10:57:41 GMT

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

---

Thank you!

Now I can't cause lockups with "vzctl stop 101" anymore either

Sooo, where do I put the /etc/network/interfaces templates and how does their format differ from the regular interfaces file?

---

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Fri, 15 Sep 2006 11:43:32 GMT

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

---

szocske wrote on Fri, 15 September 2006 06:57: Now I can't cause lockups with "vzctl stop 101" anymore either

You can't do what? I don't understand you...

veth device is removed from system when you stop VPS.

szocske wrote on Fri, 15 September 2006 06:57

Sooo, where do I put the /etc/network/interfaces templates and how does their format differ from the regular interfaces file?

You want to configure veth device automatically on VPS start?

Unfortunately we do not have in vzctl functionality which will configure veth device in VE0 on VPS start, so you will need to create script which will configure veth device and run it manually after VPS is started.

---

---

Subject: Re: disable venet?

Posted by [szocske](#) on Fri, 15 Sep 2006 11:58:27 GMT

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

---

Andrey Mirkin wrote on Fri, 15 September 2006 07:43szocske wrote on Fri, 15 September 2006 06:57Now I can't cause lockups with "vzctl stop 101" anymore either  
You can't do what?

All is fine now.

The multicast MAC address I used had horrible side effects on the VPS shutdown process, even preventing the clean shutdown of the host OS.

Please don't make me re-create it to get you the exact error message, it was something like waiting for veth interface to close or something, repeating every few seconds, infinitely.

Andrey Mirkin wrote on Fri, 15 September 2006 07:43

You want to configure veth device automatically on VPS start?

Unfortunately we do not have in vzctl functionality which will configure veth device in VE0 on VPS start, so you will need to create script which will configure veth device and run it manually after VPS is started.

I just wanted to edit the guest /etc/network/interfaces file as on a normal system, but it starts with a comment reminding me to use an interfaces.template file instead, which I can't find any information on. I don't need any vzctl support, I just want to specify the content of my /etc/networks/interfaces file

---

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Fri, 15 Sep 2006 13:10:34 GMT

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

---

szocske wrote on Fri, 15 September 2006 07:58Please don't make me re-create it to get you the exact error message, it was something like waiting for veth interface to close or something, repeating every few seconds, infinitely.

No, you don't need to repeat this

szocske wrote on Fri, 15 September 2006 07:58I just wanted to edit the guest

/etc/network/interfaces file as on a normal system, but it starts with a comment reminding me to use an interfaces.template file instead, which I can't find any information on. I don't need any vzctl support, I just want to specify the content of my /etc/networks/interfaces file

You can use VEID.start script which is executed inside VE on start to configure veth interface inside VE. Something like this:

```
#!/bin/bash
```

```
ifconfig eth0 0
```

```
ip a add a.b.c.d dev eth0
```

```
ip ro add default dev eth0
```

---

---

Subject: Re: disable venet?

Posted by [szocske](#) on Fri, 15 Sep 2006 13:33:46 GMT

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

---

Thank you again. (who marks the title \*SOLVED\*?)

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Fri, 15 Sep 2006 14:57:01 GMT

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

---

szocske wrote on Fri, 15 September 2006 07:58 The multicast MAC address I used had horrible side effects on the VPS shutdown process, even preventing the clean shutdown of the host OS. Can you please specify what kernel version you use? What OS distributive is installed on HW node?

I want to investigate this issue more deeply.

Thanks!

---

Subject: Re: disable venet?

Posted by [szocske](#) on Mon, 18 Sep 2006 08:46:50 GMT

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

---

HW node:

i386 debian testing

kernel: 2.6.16 from kernel.org + patch-026test015-combined,

config: kernel-2.6.16-026test015-i686-smp.config.ovz

Everything else left at default.

I added the veth to a debian-3.1 guest, ifconfig-ed the resulting eth0 up inside, (maybe the vethxxx outside, too?) then tried to "vzctl stop" the guest.

When we are done experimenting on this computer and move to the real target servers, I'll get back to re-create the scenario, but without clean shutdown, I do not want to risk file system corruption before that.

---

Subject: Re: disable venet?

Posted by [Andrey Mirkin](#) on Mon, 18 Sep 2006 08:59:19 GMT

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

---

Thanks!

I'll try to reproduce and investigate this issue.

---