

---

Subject: **\*SOLVED\*** Moving default NIC  
Posted by [jmeek](#) on Fri, 14 Jul 2006 16:18:45 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I just may be overly dense lately but I can't seem to figure out how to move the association of venet0 from eth0 to eth3. The host is dual homed and I want the virtuals to attach to the eth3 not eth0

Output:

```
[root@harpalyke ~]$ ip route show
10.30.102.101 dev venet0 scope link src 10.30.100.23
10.30.100.0/24 dev eth0 proto kernel scope link src 10.30.100.23
10.30.102.0/24 dev eth3 proto kernel scope link src 10.30.102.23
169.254.0.0/16 dev eth3 scope link
default via 10.30.100.2 dev eth0
```

Kind of important for an upcoming launch date, please help !

/The Dense One

---

---

Subject: Re: Moving default NIC  
Posted by [dev](#) on Sat, 15 Jul 2006 10:57:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

try this one in /etc/sysconfig/vz:

```
# The name of the device whose ip address will be used as source ip for VE.
# By default automatically assigned.
#VE_ROUTE_SRC_DEV="eth0"
```

---

---

Subject: Re: Moving default NIC  
Posted by [jmeek](#) on Mon, 17 Jul 2006 17:20:55 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Like I said, I was being dense. Later that day I found that change but had to destroy the VM and create another to effect the change. My current challenge is now (slightly) further down the line in that now that I'm on the right NIC I can see local network but I can't get the default g/w to stick. I have reviewed the forums about "ip route show" and related adjustments but just cannot see the easy way to ensure that the VM gets the default g/w of eth3 (I only get a default g/w of 191.<something> which is not on my network at all) and get the packets to leave the correct interface. As you can see below I am still leaving on the 100 network, not the 102.

(HARPALYKE is the host)

```
[root@harpalyke ~]$ ip route show
10.30.102.101 dev venet0 scope link src 10.30.100.23
10.30.100.0/24 dev eth0 proto kernel scope link src 10.30.100.23
10.30.102.0/24 dev eth3 proto kernel scope link src 10.30.102.23
169.254.0.0/16 dev eth3 scope link
default via 10.30.100.2 dev eth0
```

```
[root@harpalyke ~]$ vzctl exec 101 ip route show
191.255.255.0/24 dev venet0 scope host
169.254.0.0/16 dev venet0 scope link
default via 191.255.255.1 dev venet0
```

```
[root@harpalyke ~]$ vzctl exec 101 ping 10.30.102.5
PING 10.30.102.5 (10.30.102.5) 56(84) bytes of data.
64 bytes from 10.30.102.5: icmp_seq=0 ttl=127 time=550 ms
64 bytes from 10.30.102.5: icmp_seq=1 ttl=127 time=0.174 ms
```

```
[root@harpalyke ~]$ vzctl exec 101 ping 10.30.100.5
```

```
[root@harpalyke ~]$ vzctl exec 101 traceroute 10.30.100.5
traceroute to 10.30.100.5 (10.30.100.5), 30 hops max, 38 byte packets
 1 harpalyke.jupiter.private (10.30.100.23) 0.120 ms 0.048 ms 0.034 ms
```

```
[root@harpalyke ~]$ vzctl exec 101 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:84 errors:0 dropped:0 overruns:0 frame:0
        TX packets:84 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:12139 (11.8 KiB)  TX bytes:12139 (11.8 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:7 errors:0 dropped:0 overruns:0 frame:0
        TX packets:161 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:560 (560.0 b)  TX bytes:13327 (13.0 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:10.30.102.101  P-t-P:10.30.102.101  Bcast:10.30.102.101  Mask:255.255.255.255
        UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
```

Subject: Re: Moving default NIC  
Posted by [jmeek](#) on Thu, 20 Jul 2006 13:39:27 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I'm sorry for being a pest but once I can get past this issue I can go to Prod.

When I use VZ on a single nic machine everything is fantastic but when I am on a multihomed system I need to be able to control which NIC the traffic on the virtuals attach, especially the default gateway. I have been fiddling with this for two weeks and really need to help with what I'm sure is a easy solution.

So, said bluntly ... How do I "attach" a VM to a NIC on a multihomed system when the default gateway on the VM is not the default gateway of the host ?

---

Subject: Re: Moving default NIC  
Posted by [dev](#) on Thu, 20 Jul 2006 16:42:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

If you want a set of VEs to use one interface and gateway and another set of VEs to user another interface/gateway, you need to setup IP source based routing in host system.

Check [http://wiki.openvz.org/Source\\_based\\_routing](http://wiki.openvz.org/Source_based_routing)  
it is a draft I took from forums here.

In a few words: you need to make host system to route some of your packets from VEs to if1/gw1, and another set of IPs (VEs) to if2/gw2. This is done purely by routing.  
In developement OpenVZ version with 2.6.16 kernel it is also possible to do it with veth devices and bridges, but I wouldn't bother with it in your case.

---

Subject: Re: Moving default NIC  
Posted by [jmeek](#) on Thu, 20 Jul 2006 18:55:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Thank you so much. I knew it was something easy that I just wasn't getting !

---