

---

Subject: dual nic environment

Posted by [rudal](#) on Thu, 12 Oct 2006 21:37:28 GMT

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

---

Hello,

is there a way that you can pass eth1 traffic to the VE's and VE's out to eth1?  
Currently the host has two NICs, one eth0 and one eth1.

---

Subject: Re: dual nic environment


Posted by [Vasily Tarasov](#) on Fri, 13 Oct 2006 05:31:36 GMT

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

---

Hello,

If you're using venet device, than HN is used as a gateway for the VEs on these HN and you can picture that venet in VE is directly connected to venet on HN.  
After these assumptions it's simple to perform almost any configuration. For example in you case we have:



So the only thing you should do is proper routing on HN.

HTH,  
vass.

---

Subject: Re: dual nic environment

Posted by [rudal](#) on Fri, 13 Oct 2006 08:31:02 GMT

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

---

Yep that's what I thought too.  
So I tried this scenario:

```
eth0: 192.168.0.2 255.255.255.0
```

```
eth1: 10.10.10.2 255.255.252.0 (a big 10.x.x network)
```

So I did a route statement on the HN,  
route add -host 10.10.10.3 dev venet0

I did the dev on venet0 just like how the vzctl --ipadd command set the ip to venet0. I am not sure if that's the correct way to route that IP to one of the VE's.

After doing that, I did a vzctl set 100 --ipadd 10.10.10.3

I am able to ping from another machine, 10.10.10.10, but i can not ping from 10.10.10.3 to 10.10.10.10.

So what did I do wrong?

Thanks.

---

Subject: Re: dual nic environment  
Posted by [Andrey Mirkin](#) on Fri, 13 Oct 2006 09:49:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Please post here HN route table.  
And please post here result of command on HN:  
cat /proc/sys/net/ipv4/conf/eth1/\*

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Fri, 13 Oct 2006 17:50:24 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I forgot to add this warning message:

a = host node.

```
[root@a ~]# vzctl set 500 --ipadd 10.2.5.2
Adding IP address(es): 10.2.5.2
arp send: 10.2.5.2 is detected on another computer : 00:d0:00:58:e0:00
vps-net_add WARNING: arp send -c 1 -w 1 -D -e 10.2.5.2 eth0 FAILED
WARNING: Settings were not saved and will be resetted to original values on next start (use
--save flag)
```

```
[root@a ~]#
```

Inside the VE:

```
[root@a ~]# vzctl enter 500
```

```
entered into VPS 500
```

```
-bash-3.00#
```

```
-bash-3.00# ifconfig
```

```
...omitted...
```

```
venet0:1 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
```

```
inet addr:10.2.5.2 P-t-P:10.2.5.2 Bcast:10.2.5.2 Mask:255.255.255.255
```

```
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

On another box in the same switch:

```
[root@b ~]# arp -a | grep 10.2.5
```

```
? (10.2.5.2) at 00:0E:0C:7F:0E:D6 [ether] on eth0
```

```
[root@b ~]#
```

00:0E:0C:7F:0E:D6 -> the mac address of eth1 box a.

```
=====
```

on box a:

```
[root@a eth1]# more *
```

```
.....
```

```
accept_redirects
```

```
.....
```

```
1
```

```
.....
```

```
accept_source_route
```

```
.....
```

```
0
```

```
.....
```

```
arp_announce
```

```
.....
```

```
0
```

```
.....
```

```
arp_filter
```

```
.....
```

```
0
```

```
.....
```

```
arp_ignore
```

```
.....
```

```
0
```

```
.....
```

```
bootp_relay
```

```
.....
```

0  
.....  
disable\_policy  
.....  
0  
.....  
disable\_xfrm  
.....  
0  
.....  
force\_igmp\_version  
.....  
0  
.....  
forwarding  
.....  
1  
.....  
log\_martians  
.....  
0  
.....  
mc\_forwarding  
.....  
0  
.....  
medium\_id  
.....  
0  
.....  
proxy\_arp  
.....  
0  
.....  
rp\_filter  
.....  
1  
.....  
secure\_redirects  
.....  
1  
.....  
send\_redirects  
.....  
1  
.....  
shared\_media  
.....

1  
.....  
tag  
.....  
0

---

---

Subject: Re: dual nic environment  
Posted by [John Kelly](#) on Fri, 13 Oct 2006 19:39:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

rudal wrote on Fri, 13 October 2006 13:50# vzctl set 500 --ipadd 10.2.5.2

First you said 10.10.10, and now you say 10.2.5. And you did not post the route table, you only showed ifconfig. It's hard to help, given the conflicting and incomplete information you provided.

---

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Fri, 13 Oct 2006 19:48:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I apologize for that.  
My previous post about 10.10.x.x is just an example.  
Please disregards 10.10.x.x range as it is only an example.

The correct range for a 255.255.252.0 subnet mask is  
10.2.4.0 - 10.2.7.255

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
*.*.*.27	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
10.2.5.2	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
*.*.*.240	0.0.0.0	255.255.255.255	UH	0	0	0	venet0
*.*.*.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
10.2.4.0	0.0.0.0	255.255.252.0	U	0	0	0	eth1
169.254.0.0	0.0.0.0	255.255.0.0	U	0	0	0	eth1
0.0.0.0	*.*.*.1	0.0.0.0	UG	0	0	0	eth0

note: \*.\*.\* = is a public subnet.

---

---

Subject: Re: dual nic environment  
Posted by [Andrey Mirkin](#) on Mon, 16 Oct 2006 14:14:47 GMT

---

Please run  
tcpdump -nni <dev>  
for all interfaces (eth1, venet0 on HW, venet0 inside VPS) while pinging VPS from remote machine.  
Please post here tcpdump output.

---

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Mon, 16 Oct 2006 17:16:44 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hello Andrew,  
below is the requested information:

An external machine with 10.2.4.9 IP address is pinging the 10.2.5.2 (currently attached to the VE.)

The HW node has an IP of

```
eth1:0  Link encap:Ethernet  HWaddr 00:0E:0C:7F:0E:D6
        inet addr:10.2.4.124  Bcast:10.2.7.255  Mask:255.255.252.0
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
```

From inside a VPS:

```
-bash-3.00# tcpdump -nni venet0
09:47:50.704721 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 5
09:47:50.704745 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 5
09:47:51.704670 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 6
09:47:51.704696 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 6
09:47:52.704606 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 7
09:47:52.704630 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 7
09:47:53.704574 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 8
09:47:53.704601 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 8
```

From inside the HN on eth1:

```
09:50:45.372464 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 5
09:50:45.372490 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 5
09:50:46.373410 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 6
09:50:46.373439 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 6
09:50:47.373368 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 7
09:50:47.373408 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 7
09:50:48.374317 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 8
09:50:48.374342 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 8
```

From inside the HN on venet0:

```
09:56:26.916456 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 5
09:56:26.916533 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 5
09:56:27.917379 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 6
09:56:27.917411 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 6
09:56:28.918339 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 7
09:56:28.918367 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 7
09:56:29.919429 IP 10.2.4.9 > 10.2.5.2: icmp 64: echo request seq 8
09:56:29.919454 IP 10.2.5.2 > 10.2.4.9: icmp 64: echo reply seq 8
```

So now I tried to ping from inside the VE to the external machine:

```
-bash-3.00# ping 10.2.4.9
PING 10.2.4.9 (10.2.4.9) 56(84) bytes of data.
```

and at the same time on the HW node: (xx.xxx.x.27 is a public IP on the eth0 of the HW node)

```
tcpdump -nni venet0
```

```
09:58:55.591271 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 10
09:58:56.591045 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 11
09:58:57.590816 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 12
09:58:58.590591 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 13
09:58:59.590362 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 14
09:59:00.590136 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 15
09:59:01.589909 IP xx.xxx.x.27 > 10.2.4.9: icmp 64: echo request seq 16
```

So apparently, the VE uses eth0 as the outgoing interface? But pinging from the outside to the VE, seems to have worked and the VE knew which interface to use, but not the other way?

My understanding will be since the VE outgoing interface (if initiated from inside the VE) will be defaulted (bounded) to eth0 (packet header will have a source interface of eth0)? But if it's initiated from the outside, and the packet header has a source interface eth1, then the VE will use that source interface to return the echo back to the sender, is that correct?

Hope the information above helps.

Thanks.

---

Subject: Re: dual nic environment

Posted by [John Kelly](#) on Mon, 16 Oct 2006 18:42:26 GMT

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

---

rudal wrote on Fri, 13 October 2006 13:50:[root@a ~]# vzctl set 500 --ipadd 10.2.5.2  
Adding IP address(es): 10.2.5.2  
arp send: 10.2.5.2 is detected on another computer : 00:d0:00:58:e0:00

vps-net\_add WARNING: arpsend -c 1 -w 1 -D -e 10.2.5.2 eth0 FAILED

It says 10.2.5.2 is "detected on another computer." That's the problem you need to fix first.

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Mon, 16 Oct 2006 18:47:42 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

John Kelly wrote on Mon, 16 October 2006 14:42  
It says 10.2.5.2 is "detected on another computer." That's the problem you need to fix first.

vps-net is sending the arp to eth0 (not eth1) and 10.2.5.2 lives on eth1 in the same physical server.  
that's why it said "detected on another computer" but in fact it's only on eth1 of the same server.

---

Subject: Re: dual nic environment  
Posted by [John Kelly](#) on Mon, 16 Oct 2006 19:01:29 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

[rudal](#) wrote on Mon, 16 October 2006 14:47  
[John Kelly](#) wrote on Mon, 16 October 2006 14:42  
It says 10.2.5.2 is "detected on another computer." That's the problem you need to fix first.

vps-net is sending the arp to eth0 (not eth1) and 10.2.5.2 lives on eth1 in the same physical server. that's why it said "detected on another computer" but in fact it's only on eth1 of the same server.

Well don't do that. Trying to use the same IP address on two different interfaces does not work.

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Mon, 16 Oct 2006 19:12:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

[John Kelly](#) wrote on Mon, 16 October 2006 15:01  
Well don't do that. Trying to use the same IP address on two different interfaces does not work.



nope2 heheh.. I didnt use the same ip address on two different interfaces.  
We have a server with dual NICs  
one NIC has the public address  
the other NIC has the private address.

Hehe

---

---

Subject: Re: dual nic environment  
Posted by [John Kelly](#) on Mon, 16 Oct 2006 19:25:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

rudal wrote on Mon, 16 October 2006 15:12John Kelly wrote on Mon, 16 October 2006 15:01Well don't do that. Trying to use the same IP address on two different interfaces does not work.

nope2 heheh.. I didnt use the same ip address on two different interfaces.

Yes you did. venet0 \*is\* an interface, just like eth1. Trying to use the same IP address on venet0 and eth1, will not work.

I'm starting to wonder if you're serious, or just having some fun posting to the forum.

---

---

Subject: Re: dual nic environment  
Posted by [rudal](#) on Mon, 16 Oct 2006 22:55:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

John Kelly wrote on Mon, 16 October 2006 15:25  
Yes you did. venet0 \*is\* an interface, just like eth1. Trying to use the same IP address on venet0 and eth1, will not work.

I'm starting to wonder if you're serious, or just having some fun posting to the forum.

John, first of all, I would ask you if you understand the question that I initially posted in this forum.

Secondly, I am sorry but I dont do fun posting.

Thirdly, you are not respecting my posting by making that statement.

If you think this topic is some sort of a joke to you, please do not even view this posting.

Anyway, just in case you are still interested:

Venet0 inside the VE has two IP attached to it.

Venet0 -> the public IP (accessible from the outside)

Venet0:0 -> 10.2.5.2

```
-bash-3.00# ifconfig
```

```
venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:x.x.x.27 P-t-P:x.x.x.27 Bcast:x.x.x.27 Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

```
venet0:1 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
      inet addr:10.2.5.2 P-t-P:10.2.5.2 Bcast:10.2.5.2 Mask:255.255.255.255
      UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

and the HW node eth1 IP is:

```
eth1      Link encap:Ethernet HWaddr 00:0E:0C:7F:0E:D6
      inet addr:10.2.4.124 Bcast:10.2.7.255 Mask:255.255.252.0
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

Thanks.

---

Subject: Re: dual nic environment

Posted by [John Kelly](#) on Tue, 17 Oct 2006 02:57:49 GMT

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

---

rudal wrote on Mon, 16 October 2006 18:55: John, first of all, I would ask you if you understand the question that I initially posted in this forum. Secondly, I am sorry but I don't do fun posting. Thirdly, you are not respecting my posting by making that statement. If you think this topic is some sort of a joke to you, please do not even view this posting.

Earlier you said "10.2.5.2 lives on eth1 in the same physical server." And now you say:

```
eth1 Link encap:Ethernet HWaddr 00:0E:0C:7F:0E:D6
inet addr:10.2.4.124 Bcast:10.2.7.255 Mask:255.255.252.0
```

You keep changing the information and contradicting yourself. Sorry, but I have no more time for that.

---

Subject: Re: dual nic environment

Posted by [rudal](#) on Tue, 17 Oct 2006 07:02:28 GMT

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

---

John Kelly wrote on Mon, 16 October 2006 22:57

Earlier you said "10.2.5.2 lives on eth1 in the same physical server." And now you say:

```
eth1 Link encap:Ethernet HWaddr 00:0E:0C:7F:0E:D6  
inet addr:10.2.4.124 Bcast:10.2.7.255 Mask:255.255.252.0
```

You keep changing the information and contradicting yourself. Sorry, but I have no more time for that.

You are not understanding it again, if you have no time for this, that's fine john. I respect that and I apologize also if my statement is incomprehensible to you. Anyway, I do not want to prolong this debate and I rather move on with anything that I found out so far.

---

Subject: Re: dual nic environment

Posted by [rudal](#) on Tue, 17 Oct 2006 07:15:52 GMT

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

---

Sorry everyone, about my previous posts with John.

Anyway, this is what I found out theoritically. \*correct me if I'm wrong\*

This is just popped into my discussion with my other colleague:

If we set the default gateway to be some private IP lives on eth1's local area network which falls on the same subnet, we may be able to ping the rest of the other hosts on the local area network inside the VE. Of course, we have not tried it yet.

But that will not solve the problem of the external interface (eth0).

I wonder if the venet0 depends on the default gateway if the packet header's source interface field is empty or maybe something else?

Maybe a better question will be how does venet0 traverse out of the VE to the host and to the outside network?

My other choices will be source routing or NAT-ing.

Thank you.

---

Subject: Re: dual nic environment

Posted by [Andrey Mirkin](#) on Tue, 17 Oct 2006 08:36:37 GMT

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

---

try this from VPS:

```
-bash-3.00# ping 10.2.4.9 -I 10.2.5.2
```

Ping command takes first IP address assigned to interface by default. That is why you have to specify IP address (by -I option).

---