
Subject: All routable IP addresses all on the same vlan

Posted by [thezeke](#) on Fri, 07 Dec 2007 05:20:20 GMT

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

Hello,

After creating a VE with an IP address in the same VLAN as the IP address of the HW Node I am unable to ping that VE's IP address and I'm unable to ping an outside IP address. If I stop the VE and add the VE's IP address to eth0 as an alias I can ping it from the outside as expected.

I want to use all routable IP addresses here. I guess I need to setup a bridge between the HW Node and each VE. How do I do this without too many custom scripts and special firewall rules.

```
# vzctl create 101 --ostemplate centos-4-i386-default
# vzctl set 101 --ipadd 209.160.28.232 --save
# vzctl set 101 --nameserver 66.36.226.110 --save
# vzctl set 101 --hostname vps101 --save
# vzctl start 101
# vzctl enter 101
[root@vps101 /]# ping -c 4 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 2999ms
```

HW NODE "ifconfig":

```
-----
[root@sls-ab2p14 ~]# ifconfig
eth0    Link encap:Ethernet HWaddr 00:0E:0C:7F:0E:C7
        inet addr:209.160.40.18 Bcast:209.160.47.255 Mask:255.255.248.0
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:356760 errors:0 dropped:0 overruns:0 frame:0
        TX packets:110181 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:250149540 (238.5 MiB) TX bytes:14660446 (13.9 MiB)
```

```
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:0 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:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
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:2366 errors:0 dropped:0 overruns:0 frame:0
        TX packets:13 errors:0 dropped:0 overruns:0 carrier:0
```

collisions:0 txqueuelen:0
RX bytes:198408 (193.7 KiB) TX bytes:1092 (1.0 KiB)

VE "ifconfig":

```
-----  
[root@vps101 /]# ifconfig  
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:3517831808 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:32766 (31.9 KiB) TX bytes:510 (510.0 b)  
  
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:2271 errors:0 dropped:0 overruns:0 carrier:0  
      collisions:0 txqueuelen:0  
      RX bytes:0 (0.0 b) TX bytes:190764 (186.2 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:209.160.28.232  P-t-P:209.160.28.232  Bcast:209.160.28.232  
      Mask:255.255.255.255  
      UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500  Metric:1
```

SYSCTL.CONF (sysctl -p recently executed)

```
-----  
[root@sls-ab2p14 ~]# cat /etc/sysctl.conf  
net.ipv4.ip_forward = 1  
net.ipv4.conf.default.proxy_arp = 1  
net.ipv4.conf.default.rp_filter = 1  
net.ipv4.conf.all.rp_filter = 1  
kernel.sysrq = 1  
net.ipv4.conf.default.send_redirects = 1  
net.ipv4.conf.all.send_redirects = 1  
net.ipv4.conf.default.accept_source_route = 0
```

ROUTING TABLES:

```
[root@sls-ab2p14 ~]# netstat -r  
Kernel IP routing table  
Destination  Gateway      Genmask      Flags  MSS  Window  irtt  Iface  
209.160.40.0 *            255.255.248.0  U      0 0      0 eth0  
169.254.0.0  *            255.255.0.0   U      0 0      0 eth0  
default      vl102-gw.acc.se 0.0.0.0      UG     0 0      0 eth0
```

```
[root@vps101 /]# netstat -r
Kernel IP routing table
Destination  Gateway      Genmask      Flags  MSS Window  irtt Iface
192.0.2.0    *           255.255.255.0  U      0 0      0 venet0
169.254.0.0 *           255.255.0.0   U      0 0      0 venet0
default     192.0.2.1   0.0.0.0       UG     0 0      0 venet0
```

```
[root@sls-ab2p14 ~]# iptables -nL
Chain INPUT (policy ACCEPT)
target     prot opt source                destination

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination
ACCEPT    all  --  0.0.0.0/0             255.255.255.255

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

Subject: Re: All routable IP addresses all on the same vlan
Posted by [maratrus](#) on Fri, 07 Dec 2007 18:05:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

If I correctly understand you, you want to use VLAN inside VE?
Please look at this page, I hope it will help you
<http://wiki.openvz.org/VLAN>

Thank You!

Subject: Re: All routable IP addresses all on the same vlan
Posted by [thezeke](#) on Wed, 12 Dec 2007 00:25:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hrmmm, I don't think that's what I want to do but I'm not entirely sure.

Both 209.160.40.18 and 209.160.28.232 are already on the same VLAN configured on the switch port that the HW Node is physically connected to. The HW node can use either or both of them successfully. I guess you could treat them as being in the same subnet although that's not technically true. I want to assign 209.160.28.232 to the VE but when I do I cannot ping it from the outside and I cannot ping out. I can however ping the HW Node's IP address.

Subject: Re: All routable IP addresses all on the same vlan

Posted by [maratrus](#) on Wed, 12 Dec 2007 21:45:49 GMT

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

Hi,

1. Can you please show an output of the "ip a l" from HN?
2. Are you sure that you have shown your routing table from HN when your VPS was started?
"ip ro l" on HN (when VPS started)
3. "arp -n" on HN (when VPS started)
4. You can use "tcpdump" to determine the packages behaviour.

When you ping VPS from the outside please provide us with such information:

- Is arp request comes to HN?
- Is it replies?
- Is ping request comes to HN?
- Can we observe it on venet on HN?
- Can we observe it inside VPS?
- Is VPS replies?
- Can we observe this reply on HN?
- Is it goes outside?

Thank You!

Subject: Re: All routable IP addresses all on the same vlan

Posted by [thezeke](#) on Thu, 13 Dec 2007 18:13:31 GMT

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

Quote:Hi,

1. Can you please show an output of the "ip a l" from HN?

```
[root@sls-ab2p14 ~]# 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 brd 127.255.255.255 scope host lo
```

```
4: eth0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:0e:0c:7f:0e:c7 brd ff:ff:ff:ff:ff:ff
    inet 209.160.40.18/21 brd 209.160.47.255 scope global eth0
```

```
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
    link/void
```

Quote:2. Are you sure that you have shown your routing table from HN when your VPS was started?

"ip ro l" on HN (when VPS started)

```
[root@sls-ab2p14 ~]# vzlist
```

```
  VEID  NPROC STATUS IP_ADDR  HOSTNAME
   101   16 running 209.160.28.232 vps101
```

```
[root@sls-ab2p14 ~]# ip ro l
```

```
209.160.40.0/21 dev eth0 proto kernel scope link src 209.160.40.18
```

169.254.0.0/16 dev eth0 scope link
default via 209.160.40.1 dev eth0
Quote:3. "arp -n" on HN (when VPS started)

```
[root@sls-ab2p14 ~]# arp -n
```

Address	HWtype	HWaddress	Flags	Mask	Iface
209.160.40.1	ether	00:00:0C:07:AC:0F	C		eth0
209.160.28.232	*	*	MP		eth0

Quote:4. You can use "tcpdump" to determine the packages behaviour.

When you ping VPS from the outside please provide us with such information:

- Is arp request comes to HN?

Using "tcpdump -nn -i eth0 icmp or arp" for these tests.

I see an arp request

Quote:- Is it replies?

There is no reply, every ping packet triggers and arp request but there is no reply.

Quote:- Is ping request comes to HN?

After sending some pings from outside I see no ICMP traffic on the HN whatsoever. I do see the arp requests as mentioned above however.

Quote:- Can we observe it on venet on HN?

I see absolutely no traffic on venet0.

Quote:- Can we observe it inside VPS?

There is nothing, I tried tcpdump after entering the VE but saw absolutely no traffic whatsoever.

Quote:- Is VPS replies?

I cannot ping the HN from within the VE and I cannot ping the VE from the HW

(HN-209.160.40.18, VE-209.160.28.232)

Quote:- Can we observe this reply on HN?

No echo request from within the VE and no reply on the HN (of course)

Quote:- Is it goes outside?

Nothing goes outside from the VE.

Quote:Thank You!

Thank you.

Subject: Re: All routable IP addresses all on the same vlan

Posted by [maratrus](#) on Fri, 14 Dec 2007 09:59:33 GMT

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

Hi,

1. Quote:[root@sls-ab2p14 ~]# ip ro l

```
209.160.40.0/21 dev eth0 proto kernel scope link src 209.160.40.18
```

```
169.254.0.0/16 dev eth0 scope link
```

```
default via 209.160.40.1 dev eth0
```

When your VPS starts vzctl add additional route to your route table (to get your VPS) but I can't see it in your case. Let's conduct such experiment:

a) we add it manually:

```
#ip ro add <VE_IP> dev venet0 src <HN_IP>
```

<VE_IP> - your VPS ip

<HN_IP> - your HN ip

- b) after this experiment can you ping your VPS from HN and HN from inside the VPS?
 - c) can you ping your VE from outside?
 - d) if no. Can you please observe once again is there "arp reply" from you HN when you ping your VPS from the outside or ping your HN from the outside.
2. What version of vzctl do you use?
 3. What kernel do you use?
 4. What kind of OS template do you use for your VE?

Thank You!

Subject: Re: All routable IP addresses all on the same vlan

Posted by [thezeke](#) on Fri, 14 Dec 2007 19:49:38 GMT

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

maratrus wrote on Fri, 14 December 2007 04:59Hi,

When your VPS starts vzctl add additional route to your route table (to get your VPS) but I can't see it in your case. Let's conduct such experiment:

a) we add it manually:

```
#ip ro add <VE_IP> dev venet0 src <HN_IP>
```

<VE_IP> - your VPS ip

<HN_IP> - your HN ip

This returns an error...

```
[root@sls-ab2p14 ~]# ip ro add 209.160.28.232 dev venet0 src 209.160.40.18
```

```
RTNETLINK answers: No such device
```

```
[root@sls-ab2p14 ~]# ifconfig
```

```
eth0    Link encap:Ethernet  HWaddr 00:0E:0C:7F:0E:C7
        inet addr:209.160.40.18  Bcast:209.160.47.255  Mask:255.255.248.0
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:7017484 errors:0 dropped:0 overruns:0 frame:0
        TX packets:150131 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:886948228 (845.8 MiB)  TX bytes:20420162 (19.4 MiB)
```

```
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:216 errors:0 dropped:0 overruns:0 frame:0
        TX packets:216 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:284460 (277.7 KiB)  TX bytes:284460 (277.7 KiB)
```

```
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:2833 errors:0 dropped:0 overruns:0 frame:0
        TX packets:13 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
```

RX bytes:237604 (232.0 KiB) TX bytes:1092 (1.0 KiB)

Quote:b) after this experiment can you ping your VPS from HN and HN from inside the VPS?
c) can you ping your VE from outside?
d) if no. Can you please observe once again is there "arp reply" from you HN when you ping your VPS from the outside or ping your HN from the outside.

I can't do the above tests.

Quote:2. What version of vzctl do you use?

```
[root@sls-ab2p14 ~]# vzctl | head -1
```

```
vzctl version 3.0.18
```

Quote:3. What kernel do you use?

```
[root@sls-ab2p14 ~]# uname -a
```

```
Linux sls-ab2p14.sea2.superbrowsers.com 2.6.9-023stab044.11 #1 Sun Sep 30 13:30:10 MSD  
2007 i686 i686 i386 GNU/Linux
```

```
[root@sls-ab2p14 ~]# cat /etc/redhat-release
```

```
CentOS release 4.5 (Final)
```

Quote:4. What kind of OS template do you use for your VE?

```
centos-4-i386-default
```

Subject: Re: All routable IP addresses all on the same vlan

Posted by [maratrus](#) on Fri, 14 Dec 2007 21:24:19 GMT

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

Hi,

1. Where have you got this kernel? Have you compiled it yourself (if yes can you please show your config) or use RPM?
2. /etc/init.d/vz restart
3. Is there any error messages when you are starting VPS?
4. It is not a decision but can you try before VPS start:
 - a) down venet0
 - b) ifconfig venet0 up
 - c) or ifconfig venet0 0

Subject: Re: All routable IP addresses all on the same vlan

Posted by [thezeke](#) on Fri, 14 Dec 2007 21:31:14 GMT

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

I think restarting vz did it. Argh... Thank you for ALL of your help with this. I only wish I had tried restarting vz or even rebooting earlier.

I restarted vz using 'service vz restart' and then executed:
ifconfig venet0 down

```
ifconfig venet0 up
ifconfig venet0 0
```

I know those were probably unnecessary. By the way, the kernel I am using is from the ovz yum mirrors.

```
[openvz-kernel-rhel4]
name=OpenVZ RHEL4-based kernel
#baseurl= http://download.openvz.org/kernel/branches/rhel4-2.6.9/current/
mirrorlist=http://download.openvz.org/kernel/mirrors-rhel4-2.6.9
enabled=1
gpgcheck=1
gpgkey=http://download.openvz.org/RPM-GPG-Key-OpenVZ
```

```
[openvz-kernel-2.6.8]
name=OpenVZ 2.6.8 kernel
#baseurl=http://download.openvz.org/kernel/branches/2.6.8/current
mirrorlist=http://download.openvz.org/kernel/mirrors-2.6.8
enabled=1
gpgcheck=1
gpgkey=http://download.openvz.org/RPM-GPG-Key-OpenVZ
```
