
Subject: VENET, VETH + SNAT
Posted by [chinaman](#) on Wed, 18 Oct 2006 07:10:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

due to your help last time I've got the configuration using VETH devices running with multiple IPs inside VE.

Because some VE's running just several different services inside (different ports) I am thinking I could save some Public IPs using SNAT and private IPs inside some VEs.

I tried the setup as described in Howto, but i am running into the same problem as last time using Venet. The latency is much different inside VE and HW. Ping google around 200ms difference!

Is this a known problem or is there a solution about it?

Is it possible to use Veth and SNAT?

Sorry, but I am new with SNAT and IPTABLES.

Thanks. Best regards
Thomas

Subject: Re: VENET, VETH + SNAT
Posted by [Andrey Mirkin](#) on Wed, 18 Oct 2006 07:48:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you please describe here your configuration (ip a l, ip r l, iptables -nL).

Subject: Re: VENET, VETH + SNAT
Posted by [chinaman](#) on Thu, 19 Oct 2006 08:48:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
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 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 00:13:72:67:28:15 brd ff:ff:ff:ff:ff:ff
```

```

inet6 fe80::213:72ff:fe67:2815/64 scope link
  valid_lft forever preferred_lft forever
8: eth1: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc pfifo_fast qlen 1000
  link/ether 00:13:72:67:28:16 brd ff:ff:ff:ff:ff:ff
  inet6 fe80::213:72ff:fe67:2816/64 scope link
    valid_lft forever preferred_lft forever
1: veth0: <BROADCAST,POINTOPOINT,NOARP,UP> mtu 1500 qdisc noqueue
  link/void
10: vzbr0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:12:34:56:78:9a brd ff:ff:ff:ff:ff:ff
  inet 222.48.222.12/26 scope global vzbr0
  inet6 fe80::212:34ff:fe56:789a/64 scope link
    valid_lft forever preferred_lft forever
7: veth105.0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:12:34:56:78:9a brd ff:ff:ff:ff:ff:ff
  inet6 fe80::212:34ff:fe56:789a/64 scope link
    valid_lft forever preferred_lft forever
9: veth103.0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:ad:36:55:75:9a brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2ad:36ff:fe55:759a/64 scope link
    valid_lft forever preferred_lft forever
13: veth105.1: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:ad:56:55:45:9a brd ff:ff:ff:ff:ff:ff
  inet6 fe80::2ad:56ff:fe55:459a/64 scope link
    valid_lft forever preferred_lft forever
12: vzbr1: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:13:72:67:28:16 brd ff:ff:ff:ff:ff:ff
  inet6 fe80::213:72ff:fe67:2816/64 scope link
    valid_lft forever preferred_lft forever
15: veth104.0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:37:56:55:55:9a brd ff:ff:ff:ff:ff:ff
  inet6 fe80::237:56ff:fe55:559a/64 scope link
    valid_lft forever preferred_lft forever
21: veth106.0: <BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
  link/ether 00:1d:34:5c:78:9a brd ff:ff:ff:ff:ff:ff
  inet 192.168.10.5/32 scope global veth106.0
  inet6 fe80::21d:34ff:fe5c:789a/64 scope link
    valid_lft forever preferred_lft forever

```

```
ip r l
```

```

255.222.233.213 dev vzbr0 scope link
255.222.233.214 dev vzbr0 scope link
222.48.222.0/26 dev vzbr0 proto kernel scope link src 222.48.222.12
default via 222.48.222.1 dev vzbr0

```

```
iptables -nL
```

Chain INPUT (policy ACCEPT)

target prot opt source destination

Chain FORWARD (policy ACCEPT)

target prot opt source destination

Chain OUTPUT (policy ACCEPT)

target prot opt source destination

brctl show

bridge name	bridge id	STP enabled	interfaces
vzbr0	8000.00123456789a	no	eth0 veth103.0 veth105.0 veth105.1
vzbr1	8000.001372672816	no	eth1 veth104.0

That's the current configuration (IPs are changed). There are 2 physical eth devices in HW. Each is bridged to different veth. each veth device has got it's own public IP inside VE.

What I want is to setup on top of this configuration venet (or veth) devices with SNAT. So some more VEs can share one public IP.

Thomas

Subject: Re: VENET, VETH + SNAT

Posted by [Andrey Mirkin](#) on Thu, 19 Oct 2006 10:59:26 GMT

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

SNAT can be performed only with routing, i.e. you will need to remove eth0 and eth1 interfaces from bridges and configure correct route table and SNAT on HW node.

Subject: Re: VENET, VETH + SNAT

Posted by [chinaman](#) on Thu, 19 Oct 2006 12:48:23 GMT

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

That means it's impossible to use veth with some VEs inside a bridge and some VEs using SNAT. I thought using the bridge interface i.e. vzbr0 instead of eth0 together with SNAT, but that's not working?

What about the latency difference using venet and veth? If i could solve this it might be possible to use just venet.

Thomas

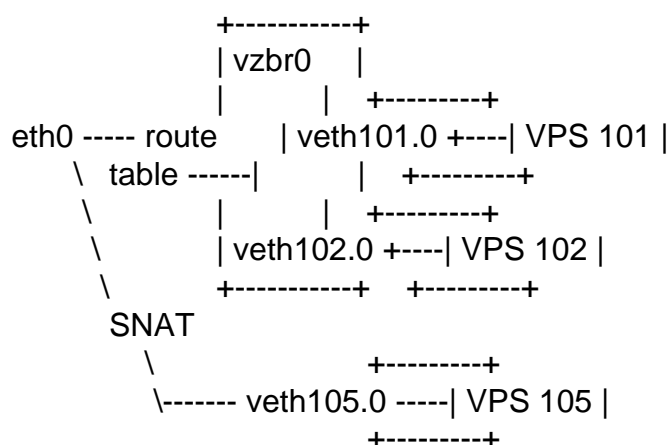
Subject: Re: VENET, VETH + SNAT

Posted by [Andrey Mirkin](#) on Thu, 19 Oct 2006 15:03:49 GMT

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

chinaman wrote on Thu, 19 October 2006 08:48 That means it's impossible to use veth with some VEs inside a bridge and some VEs using SNAT.

It is possible - just do not add eth0 interface to the bridge and keep some veth devices in bridge:



Quote:I thought using the brige interface i.e. vzbr0 instead of eth0 together with SNAT, but that's not working?

To use SNAT you will need "route stage", but when all interfaces are added to bridge routing is not performed.

Quote:What about the latency difference using venet and veth? If i could solve this it might be possible to use just venet. We didn't measure latency in such network configurations, so it will be interesting for us to see your results.

Subject: Re: VENET, VETH + SNAT

Posted by [chinaman](#) on Fri, 20 Oct 2006 07:07:47 GMT

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

Ok, I took out of the bridge the eth1 interface for testing purpose. But now I got a bit confused with routing and SNAT. Here is what I did.

+-----+

```
eth1 ----- SNAT ---- veth104.0 ----| VPS 104 |
(222.140.33.246)                +-----+
                                (192.168.10.104)
```

I tried to setup SNAT as explained in Howto:

```
# iptables -t nat -A POSTROUTING -s 192.168.10.0/24 -o eth1 -j SNAT --to 222.140.33.246
```

I setup a route from HW to VPS 104

```
# ip r a 192.168.10.104 dev veth104.0
```

and inside VPS 104

```
ip r a default via 0.0.0.0 dev eth0
```

In this point I am not sure what should be the gateway in HW for VPS?

If I ping from VPS to HW external IP i get

```
From 192.168.10.104 icmp_seq=48 Destination Host Unreachable
```

Tcpdump inside HW shows

```
# tcpdump -i veth104.0
```

```
tcpdump: WARNING: veth104.0: no IPv4 address assigned
```

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
```

```
listening on veth104.0, link-type EN10MB (Ethernet), capture size 96 bytes
```

```
14:50:31.240156 arp who-has 222.140.33.246 tell 192.168.10.104
```

What am I missing? Do I need to assign an IP to veth104.0 in HW?

Do I need to change something else? In Howto I read something about problems with 2.6.8 kernel, and other problems with Debian. Maybe its important to say that I am running Debian with 2.6.16 kernel.

Subject: Re: VENET, VETH + SNAT

Posted by [Andrey Mirkin](#) on Fri, 20 Oct 2006 09:52:24 GMT

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

Can you give me an access to node to check what is wrong?

Can you post here your route table.

Do you have correct route for 222.140.33.0 network via eth1?

Subject: Re: VENET, VETH + SNAT

Posted by [chinaman](#) on Fri, 20 Oct 2006 10:26:15 GMT

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

After a whole day searching around I just figured out one of the problems: Due to multiple external IP i've had source based routing enabled with multiple routing tables. Now I set the route into the correct table and I am able to ping from HW to VPS and back After another while i found out that

forwarding hasn't been enabled for eth1 - stupid me.

Now it's working.

Thanks for ur help.

Would it be possible to use venet instead of veth with SNAT and privat IPs? What would be the setup to do this? Routing? I would like to try the difference in speed of both ways.

Thomas

Subject: Re: VENET, VETH + SNAT

Posted by [Andrey Mirkin](#) on Fri, 20 Oct 2006 10:58:55 GMT

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

chinaman wrote on Fri, 20 October 2006 06:26 Would it be possible to use venet instead of veth with SNAT and privat IPs? What would be the setup to do this? Routing? I would like to try the difference in speed of both ways.

Yes, it is possible to configure SNAT using venet, just add private IP to your VPS with help of vzctl (route to this VPS via venet0 will be added by vzctl) and add the same SNAT rule. That's all.

It would be interesting to see results of your testing. Please post them and test procedure you used to obtain them.
