
Subject: Private virtual network with multiple VPN and NATPosted by [jmslkn](#) on Fri, 08 Jun 2007 17:39:50 GMT[View Forum Message](#) <> [Reply to Message](#)

I am trying to build a new server with OpenVZ but I have some problem with the networking. Right now each VEs can communicate each other in bridged mode other but the host and internet access still does not work. The machine has a physical network interface (eth0) and a public IP.

My Goals:

- 1) Multiple VE on a virtual private network, with VE<->VE communication
- 2) Each VE can access the host machine (HM) and HM can access the VEs
- 3) Each VE can access the internet (NAT)
- 4) Each VE accessible from VPN (OpenVPN)
- 5) Each VE is accessible from VPN (PPTP)
- [..]
- n) Multiple HM support on the same virtual private network

The settings:

```
ip r
123.123.2.0/25 dev eth0 proto kernel scope link src 123.123.2.123
169.254.0.0/16 dev eth0 scope link
192.168.0.0/16 dev vzbr0 proto kernel scope link src 192.168.0.254
default via 123.123.2.1 dev eth0

ip a
2: lo: <LOOPBACK,UP,LOWER_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: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen 1000
    link/ether 00:15:17:1b:46:58 brd ff:ff:ff:ff:ff:ff
    inet 123.123.2.0/25 brd 123.123.2.127 scope global eth0
    inet6 fe80::215:17ff:fe1b:4658/64 scope link
        valid_lft forever preferred_lft forever
6: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:15:17:1b:46:59 brd ff:ff:ff:ff:ff:ff
8: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
1: venet0: <BROADCAST,POINTOPOINT,NOARP,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/void
3: veth101.0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:18:51:14:9f:43 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::218:51ff:fe14:9f43/64 scope link
        valid_lft forever preferred_lft forever
5: veth102.0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:18:51:46:22:0c brd ff:ff:ff:ff:ff:ff
```

```

inet6 fe80::218:51ff:fe46:220c/64 scope link
    valid_lft forever preferred_lft forever
10: dummy0: <BROADCAST,NOARP,PROMISC,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether de:ad:be:ef:00:00 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::dcad:beff:feef:0/64 scope link
        valid_lft forever preferred_lft forever
7: tap0: <BROADCAST,MULTICAST,PROMISC,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether d6:98:42:d0:f4:cc brd ff:ff:ff:ff:ff:ff
    inet6 fe80::d498:42ff:fed0:f4cc/64 scope link
        valid_lft forever preferred_lft forever
12: vzbr0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue
    link/ether 00:18:51:14:9f:43 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.254/16 brd 192.168.0.255 scope global vzbr0
    inet6 fe80::218:51ff:fe14:9f43/64 scope link
        valid_lft forever preferred_lft forever

```

iptables:

```

*nat
:PREROUTING ACCEPT [1968:350481]
:POSTROUTING ACCEPT [5:1718]
:OUTPUT ACCEPT [5:1718]
-A POSTROUTING -s 192.168.0.0/255.255.0.0 -o eth0 -j SNAT --to-source 123.123.2.123
COMMIT
# Completed on Sat Jun 9 03:20:09 2007
# Generated by iptables-save v1.3.5 on Sat Jun 9 03:20:09 2007
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [394:150903]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp -m icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -p esp -j ACCEPT
-A RH-Firewall-1-INPUT -p ah -j ACCEPT
-A RH-Firewall-1-INPUT -d 224.0.0.251 -p udp -m udp --dport 5353 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp -m udp --dport 631 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m tcp --dport 631 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 25 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 80 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 443 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp -m state --state NEW -m udp --dport 137 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp -m state --state NEW -m udp --dport 138 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 139 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m state --state NEW -m tcp --dport 445 -j ACCEPT

```

```
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
# Completed on Sat Jun 9 03:20:09 2007
# Generated by iptables-save v1.3.5 on Sat Jun 9 03:20:09 2007
*mangle
:PREROUTING ACCEPT [7198:1207164]
:INPUT ACCEPT [7198:1207164]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [394:150903]
:POSTROUTING ACCEPT [394:150903]
COMMIT
# Completed on Sat Jun 9 03:20:09 2007
```

on VE101:

```
ip r
169.254.0.0/16 dev veth0 scope link
192.168.0.0/16 dev veth0 scope host
default dev veth0 scope link
```

```
ip a
1: lo: <LOOPBACK,UP,10000> 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
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
3: venet0: <BROADCAST,POINTOPOINT,NOARP> mtu 1500 qdisc noop
    link/void
5: veth0: <BROADCAST,MULTICAST,UP,10000> mtu 1500 qdisc noqueue
    link/ether 00:18:51:31:dc:03 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.101/32 brd 192.168.0.255 scope global veth0
    inet6 fe80::218:51ff:fe31:dc03/64 scope link
        valid_lft forever preferred_lft forever
```

I have started to ping to an internet address on VE101, and the tcpdump revealed that the private network is full with these kind of messages:

```
03:26:35.033522 arp who-has x.y.z tell 192.168.0.101
[.]
```

Any Idea, recommendation? Thanks. L.

Subject: Re: Private virtual network with multiple VPN and NAT
Posted by [jmslkn](#) on Mon, 11 Jun 2007 06:26:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Any idea?

Subject: Re: Private virtual network with multiple VPN and NAT

Posted by [ovzuser](#) on Mon, 11 Jun 2007 12:25:05 GMT

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

Hi,

did you add HN's eth0 to the bridge? I don't understand your iptable stuff so just this idea...

Take care: eth0 will become a member of the bridge! Don't configure it as network device, configure the bridge instead. Means your HN will have a network device "vzbr0" which is used to communicate to outside world (including VEs). eth0 will not have a ip-address any more.

I have done what you want to do with opensuse and it works. SuSE has a ifup-bridge script which does the details of setting up the brigde (but you can do it per hand of course). The veth-devices are added later during VE startup.

Tip: After adding a new device to a bridge, it takes 30s before data transmission. Look at brctl documentation. If you ping too early it can not work.

Markus

Subject: Re: Private virtual network with multiple VPN and NAT

Posted by [jmslkn](#) on Mon, 11 Jun 2007 15:35:35 GMT

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

Thanks for the answer. I'll try it asap (in fact I tried it but without reconfiguration - that was the last command that I saw on the ssh:)

However I am not sure how the private VE network will remain private with this solution. What is the proper way to provide NAT-ed internet access for the private bridged VE-s?
