Subject: Bridgin thru vethxxx.0
Posted by alvakoldo10 on Fri, 30 Jul 2010 11:25:34 GMT
View Forum Message <> Reply to Message

Hi.
I have followed instructions at http://wiki.openvz.org/Using_veth_and_brctl_for_protecting_HN_and_saving_IP_addresses but I just have one phisical NIC eth0. So I issue next commands in order to have my 2 containers connected between them and to the network (see commands below).
Problem is I neither get one of my goals, CTs do not see each other and even worse ! whenever I issue the final brctl addif br0 eth0 my phisical network falls down !

See what I have done to get nothing working

---------- COMMANDS ----------

```
ifconfig venet0:0 down
ifconfig venet0 down
#============================   CONTAINER 1
echo "ACTIONS ON THE CONTAINER 201 ....."
vzctl set 201 --netif_add eth0,00:0C:29:FA:54:FF,veth201.0,00:0C:29:FA:F3:D3 --save
ifconfig veth201.0 0
echo 0 > /proc/sys/net/ipv4/conf/veth201.0/forwarding
echo 0 > /proc/sys/net/ipv4/conf/veth201.0/proxy_arp
#
vzctl exec 201 ifconfig venet0:0 down
vzctl exec 201 ifconfig venet0 down
vzctl exec 201 ifconfig eth0 0
vzctl exec 201 ip addr add 192.168.2.201 dev eth0
vzctl exec 201 ip route add default dev eth0
vzctl exec 201 route add default gw 192.168.1.35
#
#============================   CONTAINER 2
echo "ACTIONS ON THE CONTAINER 202 ....."
vzctl set 202 --netif_add eth0,00:0C:29:FA:54:FF,veth202.0,00:0C:29:FA:F3:D3 --save
ifconfig veth202.0 0
echo 0 > /proc/sys/net/ipv4/conf/veth202.0/forwarding
echo 0 > /proc/sys/net/ipv4/conf/veth202.0/proxy_arp
#
vzctl exec 202 ifconfig venet0:0 down
vzctl exec 202 ifconfig venet0 down
vzctl exec 202 ifconfig eth0 0
vzctl exec 202 ip addr add 192.168.2.202 dev eth0
vzctl exec 202 ip route add default dev eth0
vzctl exec 202 route add default gw 192.168.1.35
#
#============================   BRIDGE
ifconfig eth0 0
```

```
echo 0 > /proc/sys/net/ipv4/conf/eth0/forwarding
echo 0 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
#
brctl addbr br0
brctl addif br0 eth0
brctl addif br0 veth201.0
brctl addif br0 veth202.0
ifconfig br0 0
#
echo 0 > /proc/sys/net/ipv4/conf/br0/forwarding
echo 0 > /proc/sys/net/ipv4/conf/br0/proxy_arp
# ============ MY HARVEST
dhclient eth0
route add -net 192.168.2.0 netmask 255.255.255.0 dev br0
```

I have tried different iptable's rules and even stop it at all.

Subject: Re: Bridgin thru vethxxx.0
Posted by alvakoldo10 on Mon, 02 Aug 2010 06:39:38 GMT
View Forum Message <> Reply to Message

Sorry to insist but this missconfiguration because of my lack of knowledge i guess is getting me into troubles.

Has anyone faced a situation like this? I have seen some other posts related and followed some of their recommendations, but no one seemed to fit at last.

Subject: Re: Simple Architecure 1 HN with 2 CTs
Posted by alvakoldo10 on Mon, 02 Aug 2010 10:22:55 GMT
View Forum Message <> Reply to Message

Ok, I have gone a bit forward, now my CT's see each other and apparently good communication with the HN. However, CT's keep on failing to reach the external Internet.

Some one could help please ? whaqt am I missing ?

I'll show my config:

****** NETwork on the HN ******
** ifconfig

```
br0       Link encap:Ethernet  HWaddr 00:0c:29:42:f4:91
          inet addr:192.168.1.35  Bcast:192.168.1.255  Mask:255.255.255.0
```

```
        inet6 addr: fe80::212:3fff:fe64:99c4/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        ...

eth0      Link encap:Ethernet  HWaddr 00:12:3f:64:99:c4
        inet6 addr: fe80::212:3fff:fe64:99c4/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        ...

veth201.0 Link encap:Ethernet  HWaddr 00:0c:29:42:f4:91
        inet6 addr: fe80::20c:29ff:fe42:f491/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        ...
veth202.0 Link encap:Ethernet  HWaddr 00:0c:29:a4:5c:29
        inet6 addr: fe80::20c:29ff:fea4:5c29/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        ...

** route -n

Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.168.2.0     0.0.0.0         255.255.255.0   U     0      0        0 br0
192.168.1.0     0.0.0.0         255.255.255.0   U     0      0        0 br0
0.0.0.0         192.168.1.1     0.0.0.0         UG    0      0        0 br0


** IPTABLES rules

-A FORWARD -i br0 -o eth0 -j ACCEPT
-A FORWARD -s 192.168.1.0/24 -i eth0 -o br0 -j ACCEPT
-A FORWARD -i br0 -o br0 -j ACCEPT
-A POSTROUTING -s 192.168.2.201/32 -o eth0 -j SNAT --to-source 192.168.1.35
-A POSTROUTING -s 192.168.2.202/32 -o eth0 -j SNAT --to-source 192.168.1.35


** brctl show

bridge name     bridge id               STP enabled     interfaces
br0             8000.000c2942f491       no              eth0
                                                        veth201.0
                                                        veth202.0


****** NETwork on the CT 201 (202 is analogue) ******
*** ifconfig

eth0      Link encap:Ethernet  HWaddr 00:0c:29:1c:ff:3f
        inet addr:192.168.2.201  Bcast:192.168.2.255  Mask:255.255.255.0
```

inet6 addr: fe80::20c:29ff:fe1c:ff3f/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:178 errors:0 dropped:0 overruns:0 frame:0
TX packets:203 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:11554 (11.2 KiB) TX bytes:18618 (18.1 KiB)

*** route -n

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.168.2.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
0.0.0.0	192.168.1.35	0.0.0.0	UG	0	0	0	eth0
0.0.0.0	0.0.0.0	0.0.0.0	U	0	0	0	eth0

Subject: Re: Simple Architecure 1 HN with 2 CTs
Posted by alvakoldo10 on Tue, 03 Aug 2010 10:06:12 GMT
View Forum Message <> Reply to Message

It seems like crazy ! (not just because of no one seems to be reading this topic), the command lines shown above, i got them in a script which I use to start up my container's configuration. Well, TODAY CTs unable to ping each other again !?!?!?! ..... these network settings will drive me into a sanitarium.

Subject: Re: Simple Architecure 1 HN with 2 CTs
Posted by alvakoldo10 on Thu, 05 Aug 2010 15:04:05 GMT
View Forum Message <> Reply to Message

Thanks no one for nothing .... nice forum this one.

For whom may be interesting :

It was this missing

```
iptables -t nat -A POSTROUTING -s 192.168.2.0/24 -o br0\
        -j SNAT --to-source 192.168.1.35
```

Instead of this

```
iptables -t nat -A POSTROUTING -s 192.168.2.0/24 -o eth0\
        -j SNAT --to-source 192.168.1.35
```

conception failure like communication.