Subject: Routing blues
Posted by laotse on Sun, 14 Dec 2008 22:22:53 GMT
View Forum Message <> Reply to Message

Well, probably I'm simply missing something and it's not related to OpenVZ ...

I try to set-up a container, which will hold the end-points for my VPNs, but currently I'm still stuck at basic connectivity, i.e. no VPN in place so far.

The problem is to ping eth3 attached to the container from the node. I can ping eth3 = 172.16.2.1 from the container=172.16.6.7, and I can ping the container from the node 172.16.1.66 and vice versa. I can also ping the WLAN router on the 172.16.2.0 network from the container. But ping 172.16.2.1 from the node fails.

I used route add -net 172.16.2.0/24 gw 172.16.6.7 venet0 for announcing the target network to the node.

There is a basic iptables firewall on the node, but as yet it does only filter for ppp0 and other interfaces. In particular, it does not filter for any LAN addresses. And of course ipv4 forwarding is enabled.

Thanks for any ideas,
 - lars

Container:

```
dergon:/# ifconfig
eth3      Link encap:Ethernet  HWaddr 00:1e:58:df:a4:4a
          inet addr:172.16.2.1  Bcast:172.16.2.255  Mask:255.255.255.0
          inet6 addr: fe80::21e:58ff:fedf:a44a/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:366 errors:0 dropped:0 overruns:0 frame:0
          TX packets:12 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:120254 (117.4 KiB)  TX bytes:980 (980.0 B)
          Interrupt:17 Base address:0xe800

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:336 (336.0 B)  TX bytes:336 (336.0 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:342 errors:0 dropped:0 overruns:0 frame:0
        TX packets:503 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:440975 (430.6 KiB)  TX bytes:122500 (119.6 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:172.16.6.7  P-t-P:172.16.6.7  Bcast:0.0.0.0  Mask:255.255.255.255
        UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1

dergon:/# cat /proc/sys/net/ipv4/ip_forward
1
dergon:/# route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
192.0.2.1       0.0.0.0         255.255.255.255 UH    0      0        0 venet0
172.16.2.0      0.0.0.0         255.255.255.0   U     0      0        0 eth3
0.0.0.0         192.0.2.1       0.0.0.0         UG    0      0        0 venet0


Node:

asgard:~# route -n
Kernel-IP-Routentabelle
Ziel            Router          Genmask         Flags Metric Ref    Use Iface
172.16.6.4      0.0.0.0         255.255.255.255 UH    0      0        0 venet0
172.16.6.7      0.0.0.0         255.255.255.255 UH    0      0        0 venet0
217.0.118.101   0.0.0.0         255.255.255.255 UH    0      0        0 ppp0
172.16.6.1      0.0.0.0         255.255.255.255 UH    0      0        0 venet0
172.16.6.2      0.0.0.0         255.255.255.255 UH    0      0        0 venet0
172.16.6.3      0.0.0.0         255.255.255.255 UH    0      0        0 venet0
172.16.2.0      172.16.6.7      255.255.255.0   UG    0      0        0 venet0
172.16.1.0      0.0.0.0         255.255.255.0   U     0      0        0 eth0
0.0.0.0         0.0.0.0         0.0.0.0         U     0      0        0 ppp0
```

Subject: Re: Routing blues
Posted by maratrus on Mon, 15 Dec 2008 10:12:31 GMT
View Forum Message <> Reply to Message

Hello,

What does it mean:
Quote:
The problem is to ping eth3 attached to the container from the node.

Have you used --netdev_add or --netif_add option of vzctl?
Have you read http://wiki.openvz.org/Veth ?

Subject: Re: Routing blues
Posted by laotse on Sun, 21 Dec 2008 12:42:30 GMT
View Forum Message <> Reply to Message

The setup is as follows:

Node FQDN mapped to eth0: 172.16.1.66
Container FQDN mapped to venet0:0: 172.16.6.7
eth3 mapped to container using: vzctl --netdev_add
used network.template to configure eth3 static as: 172.16.2.1
external device hooked to eth3: 172.16.2.11
other system on the LAN: 172.16.1.3

On the container I can so successfully:
ping 172.16.1.66
ping 172.16.1.3
ping 172.16.6.7
ping 172.16.2.1
ping 172.16.2.11

On the node I can do successfully:
ping 172.16.1.66
ping 172.16.1.3
ping 172.16.6.7

but the following (ping to eth3) fail:
ping 172.16.2.1
ping 172.16.2.11

I don't even see ICMP incoming to the container in tcpdump, despite the routing table shown at the beginning of this thread.

Regards,
- lars.

Subject: Re: Routing blues
Posted by maratrus on Mon, 22 Dec 2008 10:29:38 GMT
View Forum Message <> Reply to Message

Hello,

Quote:
eth3 mapped to container using: vzctl --netdev_add

so, that's mean you have moved eth3 interface from the HN to your VE i.e. you are not allowed to access eth3 from HN directly.
Thus we have situation when two nodes (HN and VE) has two cards (eth0 and eth3) from different subnets. That's mean that you have to care about possibility to communicate from one subnet to another.

venet driver drops all packets with destination addresses which are not assigned to venet0 interface inside VE.
In your case (172.16.2.1 is belong to eth3 not venet0).

Why don't you use veth interface
http://wiki.openvz.org/Veth

Subject: Re: Routing blues
Posted by laotse on Mon, 22 Dec 2008 11:43:26 GMT
View Forum Message <> Reply to Message

Quote:venet driver drops all packets with destination addresses which are not assigned to venet0 interface inside VE.
In your case (172.16.2.1 is belong to eth3 not venet0).

Ah, this explains it. I guess that this is a feature and probably has no configuration option.

Quote:Why don't you use veth interface

If I understand correctly this produces an additional internal interface.

The container is expected to be the endpoint of subnets of medium and no trust. So, if someone should manage breaking e.g. OpenVPN and inject code, he'll end up in a container, which does not even have a root user. Dead End!

The rest of the machine - both node and containers - do not have to be aware of these physical interfaces. In security language: eth3 shall not be accessible except for this one container.

So the second best solution will probably be to add another venet and NAT eth3 to it. Or would it drop the post routed packets?

Setting up a veth bridged to eth3 adds another level of configuration and therefore errors. Is there anything, which I am missing in this picture?

Regards,
- lars.

Subject: Re: Routing blues
Posted by maratrus on Mon, 22 Dec 2008 12:34:24 GMT
View Forum Message <> Reply to Message

Quote:
The rest of the machine - both node and containers - do not have to be aware of these physical interfaces. In security language: eth3 shall not be accessible except for this one container.

If so why do you expect HN to communicate to your VE via eth3 ip address? Why do you need this possibility?

Subject: Re: Routing blues
Posted by laotse on Tue, 23 Dec 2008 11:58:49 GMT
View Forum Message <> Reply to Message

Quote:If so why do you expect HN to communicate to your VE via eth3 ip address? Why do you need this possibility?

I do not expect the node to directly communicate with eth3, I'd like to configure the VE as a gateway. Some information will run from other VE to the subnet on eth3, e.g. DHCP through a relay (bad example, since that actually works). Furthermore, OpenVPN will establish a virtual subnet, say 172.16.128.0, through that IF. So the node shall route that traffic to the VE running OpenVPN. As I understand using venet for that communication will not work.

Subject: Re: Routing blues <SOLVED>
Posted by laotse on Tue, 23 Dec 2008 22:12:26 GMT
View Forum Message <> Reply to Message

Okay the trouble was that venet does not transport packets, which are not destinned for the container.

The idea to put eth3 on the container using --netdev_add is fine, but veth has to be used instead of venet to connect to the node. Setting up veth for the container allows it to operate as a router.

Thanks maratrus.

Further helpers were:

http://wiki.openvz.org/Using_private_IPs_for_Hardware_Nodes
http://affix.sourceforge.net/affix-newdoc/Affix-enduser/x199 .html
http://vireso.blogspot.com/2008/02/2-veth-with-2-brindges-on -openvz-at.html