
Subject: Routing Stops to work(VE's are no longer reachable)

Posted by [alarmar](#) on Thu, 19 Feb 2009 17:06:41 GMT

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

Hi everybody,

I'm relatively new to OpenVZ but I'm trying my best to describe my problem.

I lately installed an openvz-patched kernel on a debian HN.(linux-image-2.6.26-1-openvz-amd64)
Everything seemed to work fine except for IPv6,

I wasn't able to configure a SIT tunnel. (ioctl error - no such device)

I then tried 2.6.24 from the repository linked from the openvz wiki. This time the server didn't finish booting (though I couldn't find an error explaining this in /var/log/kernel.log - are there other places to look?)

Well I then downloaded the kernel sources for 2.6.24 and configured&compiled the kernel myself.
This time the kernel booted fine, ipv6 SIT tunnel worked fine
and even OpenVZ seemed to work fine (at first).

I created about ten containers. After a while they lost connectivity to the internet. The routing table didn't change, I didn't configure netfilter yet, no cronjobs that would do anything like that are running (at least I can't find any).

For better understanding an example:

on the HN:

vzlist

gw:/home/julian# vzctl start 103

Starting container ...

Container is mounted

Adding IP address(es): 91.143.93.205 2a01:30:100d:cafe::1 2001:1638:18ff:2:aaaa::1

Setting CPU units: 1000

Configure meminfo: 227892

Set hostname: kugel.kontextfrei.de

Setting quota ugidlimit: 100

Container start in progress...

gw:/home/julian# vzlist -H 103

103 56 running 91.143.93.205 kugel.kontextfrei.de

On my homepc:

Quote:

alarmar@stronghold ~ \$ ping 91.143.93.205 -c 10

PING 91.143.93.205 (91.143.93.205) 56(84) bytes of data.

--- 91.143.93.205 ping statistics ---

10 packets transmitted, 0 received, 100% packet loss, time 9014ms

Traceroute showed as last working hop the router my server is connected to.

So I thought maybe routing rules were missing/deleted

(EDIT:

I think I (maybe?) just solved this by adding the following route)

Quote:

```
91.143.93.0/24 dev eth0 proto kernel scope link src 85.31.187.154
```

Can somebody explain me why this route is necessary and how it comes that it wasn't set but the containers did work for a while?)

Quote:

```
# ip ro show
```

```
91.143.93.205 dev venet0 scope link
```

```
85.31.186.0/23 dev eth0 proto kernel scope link src 85.31.187.154
```

```
default via 85.31.186.1 dev eth0
```

and:

```
net.ipv4.conf.venet0.forwarding = 1
```

```
net.ipv4.conf.eth0.forwarding = 1
```

```
net.ipv4.conf.all.forwarding = 1
```

```
net.ipv4.conf.lo.forwarding = 1
```

```
net.ipv4.conf.default.forwarding = 1
```

I then started tcpdump to see if the packets `_do_` arrive at the HN or if the router is the problem.

Quote:

```
gw:/home/julian# tcpdump -i eth0 -n icmp
```

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

```
listening on eth0, link-type EN10MB (Ethernet), capture size 96 bytes
```

Didn't show anything when I started pinging.

I rechecked that no netfilter rules are set and just in case did a flush - nothing changed.

Then I tried to add the VE's Ip address to the eth0 interface of the HN and look if it works and indeed it did:

Quote:

HN:

```
gw:/home/julian# ip addr add 91.143.93.205 dev eth0
```

Home:

```
alar@stronghold ~ $ ping 91.143.93.205
```

```
PING 91.143.93.205 (91.143.93.205) 56(84) bytes of data.
```

```
64 bytes from 91.143.93.205: icmp_seq=2 ttl=60 time=44.1 ms
```

64 bytes from 91.143.93.205: icmp_seq=3 ttl=60 time=44.3 ms
^C

After removing the ip address from the interface again suddenly packets arrived at the VE.

Quote:

HN:

gw:/home/julian# ip addr del 91.143.93.205/32 dev eth0

Home:

alarmar@stronghold ~ \$ ping 91.143.93.205

PING 91.143.93.205 (91.143.93.205) 56(84) bytes of data.

64 bytes from 91.143.93.205: icmp_seq=1 ttl=60 time=44.9 ms

^C

Routing tables still looked the same.

From the VE I then can connect to the internet.

But after a while - and I can't see why/when or what triggers it, the connectivity disappears.

Adding the IP to eth0, pinging it, and removing it again works everytime. (Without a ping/or any other arriving packet for the address it doesn't work)

If any relevant information is missing I'm sorry and will add it ASAP.

Kernel Options related to openvz:

Quote:

gw:/usr/src/linux# grep -ie '_v\(z|e\)_' .config

CONFIG_VZ_FAIRSCHED=y

CONFIG_VE_CALLS=m

CONFIG_VZ_GENCALLS=y

CONFIG_VE_NETDEV=m

CONFIG_VE_ETHDEV=m

CONFIG_VZ_DEV=m

CONFIG_VE_IPTABLES=y

CONFIG_VZ_WDOG=m

CONFIG_VZ_CHECKPOINT=m

CONFIG_VZ_QUOTA=m

CONFIG_VZ_QUOTA_UNLOAD is not set

CONFIG_VZ_QUOTA_UGID=y

Sorry if the topic title is misleading or unprecise, I didn't know how to better name it(As I'm unsure what the problem is), if there is a more fitting title I'll change it.