
Subject: VETH ipv6 subnet wont work inside container
Posted by [bacarrdy](#) on Tue, 17 Oct 2017 09:48:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

I'm trying to route ipv6 subnet in to container with VETH interface, but i can't understand why it wont work.
IPV4 works fine, IPV6 with VENET works fine to, but i need veth because with venet interface you cant route whole subnet in to container.

How i'm routing ipv4 ip and ipv6 subnet:

INSIDE NODE WHERE CONTAINER IS CREATED

```
vzctl create 123 --ostemplate debian-7.0-x86_64 --config vps4
vzctl set 123 --hostname vps123.hostname.com --save
vzctl set 123 --nameserver 8.8.8.8 --nameserver 4.4.4.4--nameserver 3.3.3.3 --save
vzctl set 123 --features bridge:on --save
vzctl set 123 --netif_add eth0 --save
vzctl set 123 --netfilter full --save
vzctl mount 123
echo "\"JoinControllers=cpu,cpuacct,cpuset freezer,devices\" >>
/vz/root/123/etc/systemd/system.conf
vzctl start 123
echo 1 > /proc/sys/net/ipv4/conf/veth123.0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/veth123.0/proxy_arp
echo 1 > /proc/sys/net/ipv4/conf/eth0/forwarding
echo 1 > /proc/sys/net/ipv4/conf/eth0/proxy_arp
echo 1 > /proc/sys/net/ipv6/conf/veth123.0/forwarding
echo 1 > /proc/sys/net/ipv6/conf/eth0/forwarding
echo 1 > /proc/sys/net/ipv6/conf/veth123.0/proxy_ndp
echo 1 > /proc/sys/net/ipv6/conf/eth0/proxy_ndp
ip route add xx.xx.xx.xx dev veth123.0
ip -6 route add 2a05:yy:yy:xx:xx:xx:xx:1/112 dev veth123.0
```

INSIDE CT:

```
ip link set dev eth0 up
ip addr add xx.xx.xx.xx dev eth0
ip route add default dev eth0 table 1
ip route add default dev eth0
ip rule add from xx.xx.xx.xx table 1

/sbin/ip -6 addr add 2a05:yy:yy:xx:xx:xx:xx:1/112 dev eth0
/sbin/ip -6 route add default dev eth0 table 1
```

```
/sbin/ip -6 route add default dev eth0
/sbin/ip -6 rule add from 2a05:yy:yy:xx:xx:xx:xx:1/112 table 1
```

ifconfig inside CT:

```
root@vps123:/# ifconfig
eth0  Link encap:Ethernet HWaddr 00:18:51:3e:ac:17
      inet addr:xx.xx.xx.xx Bcast:0.0.0.0 Mask:255.255.255.255
      inet6 addr: 2a05:yy:yy:xx:xx:xx:xx:1/112 Scope:Global
      inet6 addr: fe80::218:51ff:fe3e:ac17/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:0 errors:0 dropped:0 overruns:0 frame:0
      TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:0 (0.0 B) TX bytes:564 (564.0 B)

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:65536 Metric:1
      RX packets:0 errors:0 dropped:0 overruns:0 frame:0
      TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
```

```
root@vps1239:/#
```

route answer inside container:

```
root@vps123:/# route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
default * 0.0.0.0 U 0 0 0 eth0
root@vps123:/#
```

route -6 answer inside container:

```
root@vps123:/# route -6
Kernel IPv6 routing table
Destination Next Hop Flag Met Ref Use If
::/0 :: U 1024 0 0 eth0
```

```

::/0                ::                !n -1 1 1 lo
2a05:yy:yy:xx:xx:xx:0/112 ::          U 256 0 0 eth0
fe80::/64           ::                U 256 0 0 eth0
::/0                ::                U 1024 0 0 eth0
::/0                ::                !n -1 1 1 lo
::1/128             ::                Un 0 1 5 lo
2a05:yy:yy:xx:xx:xx:1/128 ::          Un 0 1 0 lo
fe80::218:51ff:fe3e:ac17/128 ::        Un 0 1 0 lo
ff00::/8           ::                U 256 0 0 eth0
::/0                ::                !n -1 1 1 lo
root@vps123:/#

```

pings from container:

```

root@vps123:/# ping google.com
PING google.com (216.58.211.142) 56(84) bytes of data.
64 bytes from arn09s10-in-f14.1e100.net (216.58.211.142): icmp_req=1 ttl=57 time=11.9 ms
64 bytes from arn09s10-in-f14.1e100.net (216.58.211.142): icmp_req=2 ttl=57 time=12.1 ms
^C
--- google.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 11.947/12.042/12.137/0.095 ms
root@vps123:/#

```

```

root@vps123:/# ping6 ipv6.google.com
PING ipv6.google.com(waw02s07-in-x0e.1e100.net) 56 data bytes
From 2a05:yy:yy:xx:xx:xx:1 icmp_seq=1 Destination unreachable: Address unreachable
From 2a05:yy:yy:xx:xx:xx:1 icmp_seq=2 Destination unreachable: Address unreachable
From 2a05:yy:yy:xx:xx:xx:1 icmp_seq=3 Destination unreachable: Address unreachable
^C
--- ipv6.google.com ping statistics ---
5 packets transmitted, 0 received, +3 errors, 100% packet loss, time 4002ms

```

root@vps123:/#

```

root@vps123:/# ping6 2a05:yy:yy:xx:xx:xx:1
PING 2a05:yy:yy:xx:xx:xx:1(2a05:yy:yy:xx:xx:xx:1) 56 data bytes
64 bytes from 2a05:yy:yy:xx:xx:xx:1: icmp_seq=1 ttl=64 time=0.015 ms
^C
--- 2a05:yy:yy:xx:xx:xx:1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.015/0.015/0.015/0.000 ms
root@vps123:/#

```

PING NODE WHERE CONTAINER IS CREATED

```
root@vps123:/# ping6 2a05:yy:yy:bb:bb:bb:221
PING 2a05:yy:yy:bb:bb:bb:221(2a05:yy:yy:bb:bb:bb:221) 56 data bytes
From 2a05:yy:yy:xx:xx:xx:xx:1 icmp_seq=2 Destination unreachable: Address unreachable
From 2a05:yy:yy:xx:xx:xx:xx:1 icmp_seq=3 Destination unreachable: Address unreachable
From 2a05:yy:yy:xx:xx:xx:xx:1 icmp_seq=4 Destination unreachable: Address unreachable
^C
--- 2a05:yy:yy:bb:bb:bb:221 ping statistics ---
4 packets transmitted, 0 received, +3 errors, 100% packet loss, time 2999ms
```

```
root@vps123:/#
```

What i can see from pings that node where CT is created can't be reached from container.

Routes in node where container is created:

```
node.someserver.com:~$ route | grep veth123.0
xx.xx.xx.xx *          255.255.255.255 UH  0  0  0 veth123.0
node.someserver.com:~$
```

```
node.someserver.com:~$ route -6 | grep veth123.0
2a05:yy:yy:xx:xx:xx:xx:0/112 ::          U  1024 0  0 veth123.0
fe80::/64                ::          U  256 0  0 veth123.0
ff00::/8                 ::          U  256 0  0 veth123.0
node.someserver.com:~$
```

i dont know or it is required, but i will provide:

```
node.someserver.com:~$ vzctl --version
vzctl version 4.9.4
node.someserver.com:~$ uname -a
Linux node 2.6.32-042stab116.2 #1 SMP Fri Jun 24 15:33:57 MSK 2016 x86_64 GNU/Linux
node.someserver.com:~$
```

May be some one could point me where i`m doing mistake?

Thank you for your help :)
