Subject: IPv6 with venet : is it possible ?
Posted by Romain Riviere on Tue, 24 Jun 2008 19:32:11 GMT
View Forum Message <> Reply to Message

Hello crowd,

The title pretty much sums it up.
I've read multiple posts all over the place. Some of them suggest that IPv6 and venet don't mix, others seem to indicate that it works well.

So far, my own experience has been rather miserable : my VE can ping the host but not the outside world (the host itself is fine).

If you have IPv6 running along with venet, I would love to hear about it. If you know for a fact that it is completely and utterly impossible, I'd be glad if you could share it too

Many thanks and best regards,

Romain

Subject: Re: IPv6 with venet : is it possible ?
Posted by swindmill on Wed, 25 Jun 2008 19:34:11 GMT
View Forum Message <> Reply to Message

It would seem as though there is just enough information on the forums and elsewhere to suggest that this should work, but there are also many people experiencing the same issue as the post above.

Can we confirm whether ipv6 should work inside of a CT using venet? If so, is there something we're missing in regards to getting ipv6 to properly route from a CT to the outside?

Subject: Re: IPv6 with venet : is it possible ?
Posted by swindmill on Thu, 17 Jul 2008 17:26:40 GMT
View Forum Message <> Reply to Message

On my CentOS 5.2 host nodes I do not see that ipv6 forwarding is enabled by default.

According to this page:

http://pages.citebite.com/y6x6c2h3misp

Linux does not allow you to turn on ipv6 forwarding for a single interface. It must be enabled or disabled globally.

Can you check `cat /proc/sys/net/ipv6/conf/all/forwarding` on your host node and report back?

Subject: Re: IPv6 with venet : is it possible ?
Posted by Romain Riviere on Thu, 17 Jul 2008 17:48:06 GMT
View Forum Message <> Reply to Message

I have tested virtually everything, including echo 1 > /proc/sys/net/ipv6/conf/all/forwarding

So far my conclusion is that it is impossible to run IPv6 with venet, or that those who managed to do it won't let others hear about it ...

Subject: Re: IPv6 with venet : is it possible ?
Posted by qermit on Tue, 22 Jul 2008 09:34:04 GMT
View Forum Message <> Reply to Message

i have similar problem.

IPv6 works good on Hardware Node when i add address to my eth0 device

```
vpsbox ~ # ping6 ipv6.google.com
PING ipv6.google.com(2001:4860:0:1001::68) 56 data bytes
64 bytes from 2001:4860:0:1001::68: icmp_seq=1 ttl=53 time=85.7 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=2 ttl=53 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=3 ttl=53 time=21.0 ms
```

when i move ip address to dummy0 it works for some time

```
vpsbox ~ # ip a d  2001:41d0:1:f11c::1/64 dev eth0
vpsbox ~ # ip a a  2001:41d0:1:f11c::1/64 dev dummy0
vpsbox ~ # ping6 ipv6.google.com
PING ipv6.google.com(2001:4860:0:1001::68) 56 data bytes
64 bytes from 2001:4860:0:1001::68: icmp_seq=1 ttl=53 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=2 ttl=53 time=21.4 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=3 ttl=53 time=21.4 ms
....
....
....
64 bytes from 2001:4860:0:1001::68: icmp_seq=21 ttl=53 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=22 ttl=53 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=23 ttl=53 time=21.0 ms

--- ipv6.google.com ping statistics ---
42 packets transmitted, 23 received, 45% packet loss, time 40996ms
rtt min/avg/max/mdev = 20.925/21.220/21.938/0.239 ms
```

but, after some time it fails. I did the same thing with ipv6 address assigned to VE

```
vpsbox ~ # ip a a  2001:41d0:1:f11c::1:1/64 dev eth0
vpsbox ~ # ping6 ipv6.google.com
PING ipv6.google.com(2001:4860:0:1001::68) 56 data bytes
64 bytes from 2001:4860:0:1001::68: icmp_seq=1 ttl=53 time=23.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=2 ttl=53 time=21.6 ms

--- ipv6.google.com ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 21.627/22.368/23.110/0.756 ms
vpsbox ~ # ip a d 2001:41d0:1:f11c::1:1/64 dev eth0



vpsbox ~ # vzctl enter 777
entered into VE 777
vpsbox:/# ping6 ipv6.google.com
PING ipv6.google.com(2001:4860:0:1001::68) 56 data bytes
64 bytes from 2001:4860:0:1001::68: icmp_seq=1 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=2 ttl=52 time=21.5 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=3 ttl=52 time=21.7 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=4 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=5 ttl=52 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=6 ttl=52 time=21.4 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=7 ttl=52 time=20.8 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=8 ttl=52 time=21.4 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=9 ttl=52 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=10 ttl=52 time=21.9 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=11 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=12 ttl=52 time=21.5 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=13 ttl=52 time=20.9 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=14 ttl=52 time=21.7 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=15 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=16 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=17 ttl=52 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=18 ttl=52 time=21.2 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=19 ttl=52 time=21.3 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=20 ttl=52 time=21.0 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=21 ttl=52 time=21.1 ms
64 bytes from 2001:4860:0:1001::68: icmp_seq=22 ttl=52 time=21.4 ms

--- ipv6.google.com ping statistics ---
30 packets transmitted, 22 received, 26% packet loss, time 29000ms
rtt min/avg/max/mdev = 20.852/21.289/21.986/0.323 ms
vpsbox:/#
```

Subject: Re: IPv6 with venet : is it possible ?
Posted by broquea on Thu, 24 Jul 2008 18:21:18 GMT
View Forum Message <> Reply to Message

I've got OpenVZ working with native IPv4/IPv6 dual-stack connectivity. I'm certain it can work with a tunnel, but for my experience with this I have native available.

HN runs CentOS 5.1 32bit

ISP provides me with a single /64 allocation and provide gateway on their router. So for this example, I have A:B:C:D::/64, and A:B:C:D::1 is configured on their router, and A:B:C:D::2 is assigned to eth0 on the HN, and A:B:C:D::3+ I can assign to my VEs.

Inside the /etc/sysconfig/network-scripts/ifcfg-eth0:

IPV6INIT=yes
IPV6ADDR=A:B:C:D::2/64
IPV6FORWARDING=yes
IPV6_AUTOCONF=no

Inside /etc/sysconfig/networking:

NETWORKING_IPV6=yes
IPV6_DEFAULTGW=A:B:C:D::1
IPV6_DEFAULTDEV=eth0

Inside /etc/sysctl.conf:

net.ipv6.conf.all.forwarding = 1

Now I've got 2 machines running OpenVZ, one running CentOS 5.0 and the other 5.1. The 5.1 requires after a reboot, to run:
ifconfig venet0 0
Otherwise the VEs with IPv6 configured in them do not start up on that machine. There is probably some fix for this already that I haven't paid attention to, but for now I have modified my system to run that before starting OpenVZ.

I've also seen in the past that under CentOS as the HN, you had to run the following to set the default IPv6 route:
route -A inet6 2000::/3 dev eth0

I found that using the "vzctl set VEID --ipadd" command supports IPv6, and that is how I assign IPv6 addresses to my VEs.

A trick for those who find that their IPv6 times out after, well.... no use, is to configure NTP on the system to use an IPv6 time server. NTP uses very little bandwidth, and checks frequently enough

to act as a keepalive.

If this helps, awesome! If you still need help, let me know, and I'll see what I can cook up.

Subject: Re: IPv6 with venet : is it possible ?
Posted by qermit on Fri, 25 Jul 2008 06:07:46 GMT
View Forum Message <> Reply to Message

i was writing some article on wiki but my browser hangs :/

fist thing is to find default router. I used tcpdump icmp6 and waited for router advertisement packets. (i set ipv6 forwarding first)

Next thing was adding routing entry ip r a ::/0 via <router address> dev <interface>

broquea wrote on Thu, 24 July 2008 14:21A trick for those who find that their IPv6 times out after, well.... no use, is to configure NTP on the system to use an IPv6 time server. NTP uses very little bandwidth, and checks frequently enough to act as a keepalive.

If this helps, awesome! If you still need help, let me know, and I'll see what I can cook up.

my trick is to send ICMP6 answers for "who has" requests. I use modified parasite6 (thc-ipv6).
vpsbox thc-ipv6-0.7 # ./parasite6 eth0 00:1c:c0:5b:8d:e3
Remember to enable routing (ip_forwarding), you will denial service otherwise!
Started ICMP6 Neighbor Solitication Interceptor (Press Control-C to end) ...
Spoofed packet to fe80:0000:0000:0000:021c:c0ff:fe5b:8de3 as
2001:41d0:0001:f11c:0000:0000:0000:0003
Spoofed packet to fe80:0000:0000:0000:021c:c0ff:fe5b:8de3 as
2001:41d0:0001:f11c:0000:0000:0000:0003
Spoofed packet to fe80:0000:0000:0000:021c:c0ff:fe5b:8de3 as
2001:41d0:0001:f11c:0000:0000:0000:0003
Spoofed packet to fe80:0000:0000:0000:021c:c0ff:fe5b:8de3 as
2001:41d0:0001:f11c:0000:0000:0000:0003

last thing was adding ip address to dummy0 interface or assign to container via --ipadd command.

Subject: Re: IPv6 with venet : is it possible ?
Posted by Romain Riviere on Sun, 27 Jul 2008 16:54:36 GMT
View Forum Message <> Reply to Message

Thanks to those who posted.
It seems that IPv6 support is not exactly stable nor well-documented, from what I read here. I'll steer well clear for the time being.

Cheers !

Subject: Re: IPv6 with venet : is it possible ?
Posted by broquea on Tue, 29 Jul 2008 00:06:21 GMT
View Forum Message <> Reply to Message

FYI, I checked and using a tunnelbroker.net tunnel (and perhaps another brokers might work) I terminate the tunnel on the HN, and then allocate and assign one of the IPv6 addresses out of the routed allocation to the VE.

The VE can ping6/etc out and can be reached. The HN shows up as a hop in the traceroute6 as expected.

So the steps are (and I'll use my tunnels' information)

1) create tunnel interface with commands on the HN:
ifconfig sit0 up
ifconfig sit0 inet6 tunnel ::66.220.18.42
ifconfig sit1 up
ifconfig sit1 inet6 add 2001:470:c:29::2/64
route -A inet6 add ::/0 dev sit1

2) add ipv6 packet forwarding (edit /etc/sysctl.conf for permanent setting):
echo 1 > /proc/sys/net/ipv6/conf/all/forwarding

3) using vzctl, add an IPv6 address out of your routed allocation (lets just say the /64) to the running configuration of a VE:
vzctl set 1001 --ipadd 2001:470:d:29::2 --save

4) test connectivity from inside the VE:

[root@vps0010 ~]# vzctl enter 1001
root@testve:/# traceroute6 -n ipv6.google.com
traceroute to ipv6.google.com (2001:4860:0:2001::68), 30 hops max, 40 byte packets
 1 2001:470:c:29::2 0.091 ms 0.032 ms 0.019 ms
 2 2001:470:0:9d::1 10.832 ms 10.909 ms 10.986 ms
 3 2001:470:0:3a::2 42.412 ms 42.374 ms 42.418 ms
 4 2001:470:0:3b::2 77.478 ms 77.435 ms 77.520 ms
 5 2001:504:0:2:0:1:5169:1 78.224 ms !X 78.172 ms !X 78.382 ms !X

You'll see that the HN will appear as the first hop in the trace, since the /64 is routed behind it. The above commands are issued one after the other, no other magic needed that I can tell. And the reason the trace didn't reach the destination: "!X (communication administratively prohibited)"

Subject: Re: IPv6 with venet : is it possible ?
Posted by broquea on Mon, 04 Aug 2008 07:56:13 GMT
View Forum Message <> Reply to Message

Working with another OpenVZ user, we noted that he had to use 2000::/3 as their default gateway due to the kernel version.

Subject: Re: IPv6 with venet : is it possible ?
Posted by Romain Riviere on Mon, 04 Aug 2008 08:40:21 GMT
View Forum Message <> Reply to Message

I have switched to Linux-VServer in the meantime.

I will try that "ifconfig venet0 0" for completeness's sake, but apart from that, my attempts at using IPv6 were conducted exactly as described elsewhere (IPv6 forwarding, default route via HN etc).

Thanks for your feedback.