
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)"
