
Subject: Re: IPv6 and OVZ part deux
Posted by [lars.bailey](#) on Fri, 21 Jan 2011 04:24:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

@ Jean-Marc

This is what I now use on a test Node,for IPv6/OVZ/Ethernet bridging.

The Node server;

```
* /etc/sysconfig/network
```

```
NETWORKING=yes  
HOSTNAME=stooge  
GATEWAYDEV=virtbr0  
NETWORKING_IPV6=yes  
IPV6INIT=yes  
IPV6FORWARDING=yes  
IPV6_AUTOCONF=no  
IPV6_AUTOTUNNEL=no
```

This is the source-route bridge network configuration.

```
* ifcfg-virtbr0
```

```
DEVICE=virtbr0  
TYPE=Bridge  
ONBOOT=yes  
STP=off  
DELAY=0  
BOOTPROTO=static  
IPADDR=192.168.1.72  
NETMASK=255.255.255.0  
IPV6ADDR=fd22:a075:afd0:e096::101/64  
IPV6_DEFAULTGW=fd22:a075:afd0:e096::1FF
```

I used a private IPv6 address range,for internal testing.
Node IPv6 routing table.

```
# ip -6 ro show dev virtbr0  
fd22:a075:afd0:e096::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295  
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295  
default via fd22:a075:afd0:e096::1ff metric 1 mtu 1500 advmss 1440 hoplimit 4294967295  
#
```

A test VE container was created,and bound to the Ethernet bridge.

```
# brctl show
```

```
bridge name bridge id STP enabled interfaces
virtbr0 8000.001851a86b76 no eth0
      veth6101.0
#
```

This is the test VE network configuration.

```
DEVICE=eth0
TYPE=Ethernet
ONBOOT=yes
IPV6INIT=yes
IPV6ADDR=fd22:a075:afd0:e096:65::65/64
```

From Node,the VE is reachable.

```
# ping6 -c 3 fd22:a075:afd0:e096:65::65
PING fd22:a075:afd0:e096:65::65(fd22:a075:afd0:e096:65::65) 56 data bytes
64 bytes from fd22:a075:afd0:e096:65::65: icmp_seq=1 ttl=64 time=1.34 ms
64 bytes from fd22:a075:afd0:e096:65::65: icmp_seq=2 ttl=64 time=0.375 ms
64 bytes from fd22:a075:afd0:e096:65::65: icmp_seq=3 ttl=64 time=0.372 ms
```

```
--- fd22:a075:afd0:e096:65::65 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 0.372/0.696/1.342/0.457 ms
#
```

From VE,Node is reachable.

```
# vzctl enter 6101
entered into VE 6101
[root@moe /]# ping6 -c 3 fd22:a075:afd0:e096::101
PING fd22:a075:afd0:e096::101(fd22:a075:afd0:e096::101) 56 data bytes
64 bytes from fd22:a075:afd0:e096::101: icmp_seq=1 ttl=64 time=1.50 ms
64 bytes from fd22:a075:afd0:e096::101: icmp_seq=2 ttl=64 time=0.403 ms
64 bytes from fd22:a075:afd0:e096::101: icmp_seq=3 ttl=64 time=0.401 ms
```

```
--- fd22:a075:afd0:e096::101 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 0.401/0.769/1.503/0.519 ms
[root@moe /]#
```

It's pretty straight-forward,and no more worries on link-local breakage.
For IPv4,it was pretty straight-forward too.
This is the IPv4 container's routing.

```
[root@curly /]# ip ro show
192.168.1.0/24 dev eth0 proto kernel scope link src 192.168.1.64
169.254.0.0/16 dev eth0 scope link
```

default via 192.168.1.254 dev eth0

This should be self-explanatory.
Ping OpenVZ website.

```
[root@curly /]# ping -c 3 www.openvz.org
PING www.openvz.org (64.131.90.7) 56(84) bytes of data.
64 bytes from openvz.org (64.131.90.7): icmp_seq=1 ttl=48 time=38.3 ms
64 bytes from openvz.org (64.131.90.7): icmp_seq=2 ttl=48 time=40.1 ms
64 bytes from openvz.org (64.131.90.7): icmp_seq=3 ttl=48 time=38.3 ms
```

```
--- www.openvz.org ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2042ms
rtt min/avg/max/mdev = 38.381/38.982/40.184/0.879 ms
[root@curly /]#
```

What differs here now is, the IPv4 containers are directly routed, and I'm a little concerned about DHCP time-outs, due to X bridged interfaces on source-route.

In using DHCP on the Node for network assignment, required the use of NAT rules, as the IPv4 containers resided on their own subnet.

This simplified route management for a VE.

Since, I have never used a source-route Ethernet bridge setup with DHCP, I think NAT rules still applies, but I will give your aliased interface a shot, and see what happens.

In truth, I'm not a big fan of source-route bridging.

But with IPv6/Ethernet bridging, this is going to be a common practice.

One arena I want to play around with, is using VDE.

I downloaded "openVswitch", compiled, and installed on the test Node.

Time is not a premium for me right now, and my technology mistress is going to accompany me, in finding a good divorce lawyer. (LOL)