
Subject: Re: IPv6 and OVZ part deux
Posted by [Jean-Marc Pigeon](#) **on Tue, 18 Jan 2011 20:45:36 GMT**
[View Forum Message](#) <> [Reply to Message](#)

Bonjour,

Found difficult to set up IPV6 on a container, the key part to catch is the bridge position within the stack. I was already using veth with IPV4 config but it was not working properly with IPV6.

Found lack of detailed example setup making "life difficult", so I am providing my test setup, such it can be used as starting point to some of you and a base for advise or suggestion for those with an IPV6/IPV4 setup already working.

On an standard IPV4 with VETH here is my production config (on the host side)

```
ifcfg-eth0=
  DEVICE=eth0
  ONBOOT=yes
  BOOTPROTO=static
  HWADDR=00:26:B9:67:A7:E1
  IPADDR=192.2Y.X.248
  NETMASK=255.255.255.224
```

```
ifconfig eth0=
eth0    Link encap:Ethernet HWaddr 00:26:B9:67:A7:E1
        inet addr:192.2Y.X.Y Bcast:192.2Y.X.255 Mask:255.255.X.224
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:19359552 errors:0 dropped:0 overruns:0 frame:0
        TX packets:38144778 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:5174070630 (4.8 GiB) TX bytes:35554807003 (33.1 GiB)
        Interrupt:177 Memory:dxXf0000-dff00000
```

```
ifcfg-br0=
  DEVICE=br0
  ONBOOT=yes
  TYPE=Bridge
  BOOTPROTO=static
  IPADDR=192.0.2.1
  NETMASK=255.255.255.255
```

```
ifconfig br0=
br0    Link encap:Ethernet HWaddr 00:18:51:08:8D:E1
        inet addr:192.0.2.1 Bcast:192.0.2.1 Mask:255.255.255.255
```

```
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:135837164 errors:0 dropped:0 overruns:0 frame:0
TX packets:144774563 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:98367107659 (91.6 GiB) TX bytes:160370520042 (149.3 GiB)
```

with ifconfig VE as setup by vectl start

```
testve Link encap:Ethernet HWaddr 00:18:51:EE:2A:1C
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:48523 errors:0 dropped:0 overruns:0 frame:0
TX packets:279080 errors:0 dropped:336 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:2799240 (2.6 MiB) TX bytes:24558278 (23.4 MiB)
```

This configuration is workin alright on IPV4, but adding IPV6 make it
not working alright, the host was able to ping6 the VE but not the other way.

Found IPV6 Working config as such:

```
ifcfg-eth0=
DEVICE=eth0
BOOTPROTO=none
ONBOOT=yes
BRIDGE=br0
HWADDR=00:19:D1:12:41:87
```

```
ifconfig eth0=
eth0 Link encap:Ethernet HWaddr 00:19:D1:12:41:87
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:2262 errors:0 dropped:0 overruns:0 frame:0
TX packets:976 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:512724 (500.7 KiB) TX bytes:123021 (120.1 KiB)
Memory:e8000000-e8020000
```

```
ifcfg-br0=
DEVICE=br0
ONBOOT=yes
TYPE=Bridge
BOOTPROTO=dhcp
```

```
IPV6INIT=yes
IPV6_AUTOCONF=yes
DHCpv6C=yes

ifconfig br0=
br0    Link encap:Ethernet HWaddr 00:19:D1:12:41:87
        inet addr:192.2Y.X.242 Bcast:192.2Y.X.255 Mask:255.255.255.224
        inet6 addr: 2002::c02xY:Xf2/64 Scope:Global
        inet6 addr: xX80::2Y:d1ff:xX12:4187/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:2514 errors:0 dropped:0 overruns:0 frame:0
        TX packets:1140 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:514621 (502.5 KiB) TX bytes:142714 (139.3 KiB)
```

you have both IPV4 and IPV6 assigned via DHCP.

I added an alias to br0 to keep IPV4 routing capability.

```
ifcfg-br0:vegw=
DEVICE=br0:vegw
ONBOOT=no
IPADDR=192.0.2.1
NETMASK=255.255.255.255
```

```
ifconfig br0:vegw=
br0:vegw Link encap:Ethernet HWaddr 00:19:D1:12:41:87
        inet addr:192.0.2.1 Bcast:192.0.2.1 Mask:255.255.255.255
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

This completes the HOST side setup.

On the VE side

```
ifcfg-eth0=
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
IPADDR=192.2Y.X.77
NETMASK=255.255.255.255
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6ADDR=2001:2xY8::192.2Y.X.77/64
IPV6ADDR_SECONDARIES=2002::192.2Y.X.77/64
```

```
ifconfig eth0=
eth0    Link encap:Ethernet HWaddr 00:18:51:1E:6D:F4
        inet addr:192.2Y.X.77 Bcast:192.2Y.X.77 Mask:255.255.255.255
        inet6 addr: xX80::218:51ff:xX1e:6df4/64 Scope:Link
```

```
inet6 addr: 2002::c02xY:xX4d/64 Scope:Global  
inet6 addr: 2001:2xY8::c02xY:xX4d/64 Scope:Global  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:2671 errors:0 dropped:0 overruns:0 frame:0  
TX packets:257 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:585756 (572.0 KiB) TX bytes:27772 (27.1 KiB)
```

As usual you have file route-eth0=
192.0.2.0/24 dev eth0 scope link
default via 192.0.2.1 dev eth0

on the VE

```
[root@testvz network-scripts]# ping6 -n -c 3 2002::c02xY:xXf2  
PING 2002::c02xY:xXf2(2002::c02xY:xXf2) 56 data bytes  
64 bytes from 2002::c02xY:xXf2: icmp_seq=0 ttl=64 time=0.445 ms  
64 bytes from 2002::c02xY:xXf2: icmp_seq=1 ttl=64 time=0.035 ms  
64 bytes from 2002::c02xY:xXf2: icmp_seq=2 ttl=64 time=0.034 ms
```

on the HOST

```
ping6 -c 3 -n 2002::c02xY:xX4d  
PING 2002::c02xY:xX4d(2002::c02xY:xX4d) 56 data bytes  
64 bytes from 2002::c02xY:xX4d: icmp_seq=0 ttl=64 time=1.77 ms  
64 bytes from 2002::c02xY:xX4d: icmp_seq=1 ttl=64 time=0.037 ms  
64 bytes from 2002::c02xY:xX4d: icmp_seq=2 ttl=64 time=0.041 ms
```

Sure enough IPV4 stack is still fully fonctionnal

my 2 cents, have fun and keep me posted.