

---

Subject: IPv6 and OVZ part deux  
Posted by [lars.bailey](#) on Mon, 03 Jan 2011 10:53:38 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Getting IPv6 working with OpenVZ, is either going to be heaven, or hell, depending on several factors.  
For IPv6 in general, it mostly involves;

- \* unfamiliar with IPv6 numeration(s)
- \* improper subnetting schemas
- \* improper routing configuration(s)
- \* OS specific configurations for IPv6

These are easily remedied.

For OpenVZ, it is going to depend on the Node server operating system, and/or OS template cache.

If you think I'm on crack in my remark above, think again.

My testing has been done on F13 and Oracle Enterprise(EL5)

On F13, IPv6 and OpenVZ is a snap, and was tested with RedHat, Debian, and OpenSUSE containers.

On Oracle, or any EL5 flavor, it gets really interesting.

Especially using EL5 template caches(CentOS)

This also goes for CentOS templates on F13.

You have to enable IPv6 in a EL5 container, just like you would on the Node server.

To prove IPv6 can easily be setup, this is a live IPv6 F13 container, running on F13 Node server.

=> Taken from my testing notes

The prefix below, is a private global range, used in testing IPv6

The IPv6 router's source route interface(Node), is configured with a IPv6 address from the "/64" prefix.

```
fd60:1014:9458:4b60::9
```

A test VE container was created, with virtual Ethernet, and its VETH interface, was also configured with an IPv6 address, from the available "/64" prefix.

```
fd60:1014:9458:4b60::a
```

Test for reachability.

```
# ping6 -c 3 fd60:1014:9458:4b60::9
```

```
PING fd60:1014:9458:4b60::9(fd60:1014:9458:4b60::9) 56 data bytes
```

```
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=1 ttl=64 time=0.369 ms
```

```
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=2 ttl=64 time=0.315 ms
```

```
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=3 ttl=64 time=0.242 ms
```

```
--- fd60:1014:9458:4b60::9 ping statistics ---
```

```
3 packets transmitted, 3 received, 0% packet loss, time 2001ms  
rtt min/avg/max/mdev = 0.242/0.308/0.369/0.055 ms
```

```
# ping6 -c 3 fd60:1014:9458:4b60::a
```

```
PING fd60:1014:9458:4b60::a(fd60:1014:9458:4b60::a) 56 data bytes
```

```
64 bytes from fd60:1014:9458:4b60::a: icmp_seq=1 ttl=64 time=0.329 ms
```

```
64 bytes from fd60:1014:9458:4b60::a: icmp_seq=2 ttl=64 time=0.342 ms
```

```
64 bytes from fd60:1014:9458:4b60::a: icmp_seq=3 ttl=64 time=0.346 ms
```

```
--- fd60:1014:9458:4b60::a ping statistics ---
```

```
3 packets transmitted, 3 received, 0% packet loss, time 2001ms  
rtt min/avg/max/mdev = 0.329/0.339/0.346/0.007 ms
```

```
#
```

"ping6" results show IPv6 networking is working on both interfaces.  
/68 from /64 prefix, cut to a /96, and then down to "/120" prefix length.  
May not be standard, but works in a pinch.  
Secondary IPv6 address, from the the "/120" subnet.

```
FD60:1014:9458:4B60:E003:5000:0010:0101/128 VETH interface  
FD60:1014:9458:4B60:E003:5000:0010:01FF/128 VE
```

```
# ifconfig veth6101.0
```

```
veth6101.0 Link encap:Ethernet HWaddr 00:18:51:40:DF:09  
  inet6 addr: fe80::218:51ff:fe40:df09/64 Scope:Link  
  inet6 addr: fd60:1014:9458:4b60:e003:5000:10:101/64 Scope:Global  
  inet6 addr: fd60:1014:9458:4b60::a/64 Scope:Global  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:48 errors:0 dropped:0 overruns:0 frame:0  
TX packets:99 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0  
RX bytes:2274 (2.2 KiB) TX bytes:7664 (7.4 KiB)
```

The container.

```
[root@moe /]# ifconfig eth0
```

```
eth0 Link encap:Ethernet HWaddr 00:18:51:0C:72:E2  
  inet6 addr: fd60:1014:9458:4b60:e003:5000:10:1ff/64 Scope:Global  
  inet6 addr: fe80::218:51ff:fe0c:72e2/64 Scope:Link  
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
RX packets:99 errors:0 dropped:0 overruns:0 frame:0  
TX packets:61 errors:0 dropped:0 overruns:0 carrier:0  
collisions:0 txqueuelen:0
```

RX bytes:7664 (7.4 KiB) TX bytes:3166 (3.0 KiB)

A default route to the VETH interface was added

```
[root@moe /]# ip -6 ro show dev eth0
fd60:1014:9458:4b60::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 0
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 0
default via fd60:1014:9458:4b60:e003:5000:10:101 metric 1 mtu 1500 advmss 1440 hoplimit 0
```

Ping the default gateway.

```
# ping6 -c 3 fd60:1014:9458:4b60:e003:5000:10:101
PING fd60:1014:9458:4b60:e003:5000:10:101(fd60:1014:9458:4b60:e003:5000:10:101) 56 data
bytes
64 bytes from fd60:1014:9458:4b60:e003:5000:10:101: icmp_seq=1 ttl=64 time=0.989 ms
64 bytes from fd60:1014:9458:4b60:e003:5000:10:101: icmp_seq=2 ttl=64 time=0.375 ms
64 bytes from fd60:1014:9458:4b60:e003:5000:10:101: icmp_seq=3 ttl=64 time=0.381 ms

--- fd60:1014:9458:4b60:e003:5000:10:101 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 0.375/0.581/0.989/0.289 ms
```

From the IPv6 router,to the container.

```
# ping6 -c 3 fd60:1014:9458:4b60:e003:5000:10:1ff
PING fd60:1014:9458:4b60:e003:5000:10:1ff(fd60:1014:9458:4b60:e003:5000:10:1ff) 56 data
bytes
64 bytes from fd60:1014:9458:4b60:e003:5000:10:1ff: icmp_seq=1 ttl=64 time=1.04 ms
64 bytes from fd60:1014:9458:4b60:e003:5000:10:1ff: icmp_seq=2 ttl=64 time=0.364 ms
64 bytes from fd60:1014:9458:4b60:e003:5000:10:1ff: icmp_seq=3 ttl=64 time=0.382 ms

--- fd60:1014:9458:4b60:e003:5000:10:1ff ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 0.364/0.597/1.046/0.317 ms
#
```

From the IPv6 router,VE6101 is reachable.

Router not reachable from VE,\*may need route on Node - add /120 => VETH

Added direct route to the IPv6 router.

"route6-eth0" file,in container's /etc/sysconfig/network-scripts directory.

```
fd60:1014:9458:4b60::9 via fd60:1014:9458:4b60:e003:5000:10:101
```

```
# ip -6 ro show dev eth0
fd60:1014:9458:4b60::9 via fd60:1014:9458:4b60:e003:5000:10:101 metric 1024 mtu 1500
advmss 1440 hoplimit 0
fd60:1014:9458:4b60::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 0
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 0
```

#

Ping IPv6 router.

```
# ping6 -c 3 fd60:1014:9458:4b60::9
PING fd60:1014:9458:4b60::9(fd60:1014:9458:4b60::9) 56 data bytes
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=1 ttl=64 time=1.48 ms
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=2 ttl=64 time=0.377 ms
64 bytes from fd60:1014:9458:4b60::9: icmp_seq=3 ttl=64 time=0.375 ms
```

```
--- fd60:1014:9458:4b60::9 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 0.375/0.746/1.488/0.525 ms
```

#

Test container shows as reachable neighbor,without Proxy\_NDP.  
IPv6 transit over GRE tunnel6 successful.

""

Like I said,it couldn't have gotten any easier.  
On Oracle,this same setup,breaks the IPv6 stack.  
Loss of link-local,is common on source-route interface,and is causing me grief,although we do not use Oracle as a rule for OVZ.  
If you are going to transition,better think ahead.  
I'm glad I put a little effort first,before any assumptions.

---