Subject: Re: IPv6 and OVZ part deux
Posted by lars.bailey on Thu, 13 Jan 2011 07:30:28 GMT
View Forum Message <> Reply to Message

Before this thread is read any further,I need to clear up some things,in my haste to prove that IPv6 can be set up easily
with OpenVZ.
But,I also wanted to illustrate what not to do.
First off,do not use IFCONFIG,with IPv6 configurations.
IFCONFIG,assumes /64 prefix lengths.for host addresses.
This may pose a dilemma,for some end-users.
IPv6 configurations via scripting,or router software on Node?
We chose to add router software on Nodes,and add *init script to VE's.
The second,was using a /64 prefix length.
There have been inquiries elsewhere on the OVZ forum,on whether a /64 prefix can be subnetted?
Yes,a /64 can be subnetted.(use ipv6gen.pl)
Using the IP command manually;

1. Take a /64 private IPv6 address range from SimplyDNS,to test with.
I'm going to use this /64.

fd22:a075:afd0:e096::/64

Subnet it down to a /96.

FD22:A075:AFD0:E096:0000:0000::/96
FD22:A075:AFD0:E096:0000:0001::/96
FD22:A075:AFD0:E096:0000:0002::/96
FD22:A075:AFD0:E096:0000:0003::/96
FD22:A075:AFD0:E096:0000:0004::/96
FD22:A075:AFD0:E096:0000:0005::/96
FD22:A075:AFD0:E096:0000:0006::/96
FD22:A075:AFD0:E096:0000:0007::/96
FD22:A075:AFD0:E096:0000:0008::/96
FD22:A075:AFD0:E096:0000:0009::/96
FD22:A075:AFD0:E096:0000:000A::/96
=> etc.

I'm going to use this /96.

FD22:A075:AFD0:E096:0000:0001::/96

2. Add IPv6 address,to source-route interface.

[root@stooge network-scripts]# ip address add \
> fd22:a075:afd0:e096:0:1::1a dev eth1

3. Check with IFCONFIG.

```
[root@stooge network-scripts]# ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 00:14:BF:5E:51:3F
          inet addr:192.168.1.64  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::214:bfff:fe5e:513f/64 Scope:Link
          inet6 addr: fd22:a075:afd0:e096:0:1:0:1a/128 Scope:Global
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:477340 errors:0 dropped:0 overruns:0 frame:0
          TX packets:353145 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:245276422 (233.9 MiB)  TX bytes:174651991 (166.5 MiB)
          Interrupt:9 Base address:0x4800
```

Notice,it shows up as a hermit host.
Always add prefix length.
So again,add IPv6 address.

```
[root@stooge network-scripts]# ip address add \
> fd22:a075:afd0:e096:0:1::1a/96 dev eth1
```

4. Check with IFCONFIG.

```
[root@stooge network-scripts]# ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 00:14:BF:5E:51:3F
          inet addr:192.168.1.64  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::214:bfff:fe5e:513f/64 Scope:Link
          inet6 addr: fd22:a075:afd0:e096:0:1:0:1a/96 Scope:Global
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:477443 errors:0 dropped:0 overruns:0 frame:0
          TX packets:353163 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:245282889 (233.9 MiB)  TX bytes:174653483 (166.5 MiB)
          Interrupt:9 Base address:0x4800
```

5. Take the /96,and subnet to a /116.

```
FD22:A075:AFD0:E096:0000:0001:0000:0000/116
FD22:A075:AFD0:E096:0000:0001:0000:1000/116
FD22:A075:AFD0:E096:0000:0001:0000:2000/116
FD22:A075:AFD0:E096:0000:0001:0000:3000/116
FD22:A075:AFD0:E096:0000:0001:0000:4000/116
FD22:A075:AFD0:E096:0000:0001:0000:5000/116
FD22:A075:AFD0:E096:0000:0001:0000:6000/116
FD22:A075:AFD0:E096:0000:0001:0000:7000/116
FD22:A075:AFD0:E096:0000:0001:0000:8000/116
FD22:A075:AFD0:E096:0000:0001:0000:9000/116
FD22:A075:AFD0:E096:0000:0001:0000:A000/116
```

=> etc

As a rule,we do not subnet any given prefix,any further than /116.
Some may feel,this is "address waste",but with IPv6,we have no
reason to get anal about it.
You choose your method.
I'm going to use this /116 prefix.

FD22:A075:AFD0:E096:0000:0001:0000:1000/116

For the VETH and VE interfaces,I will use these host addresses.

VETH - fd22:a075:afd0:e096:0:1:0:1fff/116
VE eth0 - fd22:a075:afd0:e096:0:1:0:1001/116

6. Configure IPv6 address for VETH interface.

```
[root@stooge network-scripts]# ip address add \
> fd22:a075:afd0:e096:0:1:0:1fff/116 dev veth6101.0
```

7. Check with IFCONFIG.

```
[root@stooge network-scripts]# ifconfig veth6101.0
veth6101.0 Link encap:Ethernet  HWaddr 00:18:51:2C:C8:6D
          inet addr:192.168.100.101  Bcast:192.168.100.255  Mask:255.255.255.0
          inet6 addr: fd22:a075:afd0:e096:0:1:0:1fff/116 Scope:Global
          inet6 addr: fe80::218:51ff:fe2c:c86d/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:432 errors:0 dropped:0 overruns:0 frame:0
          TX packets:83 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:16956 (16.5 KiB)  TX bytes:5716 (5.5 KiB)

[root@stooge network-scripts]#
```

8. For the VE.

Configure IPv6 address from this same subnet.

FD22:A075:AFD0:E096:0000:0001:0000:1001/116

```
[root@moe /]# ip address add \
> fd22:a075:afd0:e096:0:1:0:1001/116 dev eth0
```

9. Check with IFCONFIG.

```
[root@moe /]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:18:51:8F:16:2E
```

```
        inet addr:192.168.101.1  Bcast:192.168.101.255  Mask:255.255.255.0
        inet6 addr: fd22:a075:afd0:e096:0:1:0:1001/116 Scope:Global
        inet6 addr: fe80::218:51ff:fe8f:162e/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:83 errors:0 dropped:0 overruns:0 frame:0
        TX packets:446 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:5716 (5.5 KiB)  TX bytes:17784 (17.3 KiB)

[root@moe /]#
```

10. Add default route.

```
[root@moe /]# ip ro add default via \
> fd22:a075:afd0:e096:0:1:0:1fff dev eth0
```

11. Check routing table.

```
[root@moe /]# ip -6 ro show dev eth0
fd22:a075:afd0:e096:0:1:0:1000/116  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:e096:0:1:0:1fff  metric 1024  mtu 1500 advmss 1440 hoplimit
4294967295
```

12. Ping gateway.(VETH)

```
[root@moe /]# ping6 -c 3 fd22:a075:afd0:e096:0:1:0:1fff
PING  fd22:a075:afd0:e096:0:1:0:1fff(fd22:a075:afd0:e096:0:1:0:1ff f) 56 data bytes
64 bytes from fd22:a075:afd0:e096:0:1:0:1fff: icmp_seq=1 ttl=64 time=1.43 ms
64 bytes from fd22:a075:afd0:e096:0:1:0:1fff: icmp_seq=2 ttl=64 time=0.270 ms
64 bytes from fd22:a075:afd0:e096:0:1:0:1fff: icmp_seq=3 ttl=64 time=0.287 ms

--- fd22:a075:afd0:e096:0:1:0:1fff ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2011ms
rtt min/avg/max/mdev = 0.270/0.663/1.433/0.544 ms
```

13. Check your neighbor entry.

```
[root@moe /]# ip -6 neigh show
fd22:a075:afd0:e096:0:1:0:1fff dev eth0 lladdr 00:18:51:2c:c8:6d REACHABLE
fe80::218:51ff:fe2c:c86d dev eth0 lladdr 00:18:51:2c:c8:6d REACHABLE
```

The VETH interface,will always be the container's "neighbor".
14. Now,the GW router.(Node)

```
[root@moe /]# ping6 -c 3 fd22:a075:afd0:e096:0:1:0:1a
PING fd22:a075:afd0:e096:0:1:0:1a(fd22:a075:afd0:e096:0:1:0:1a) 56 data bytes
```

64 bytes from fd22:a075:afd0:e096:0:1:0:1a: icmp_seq=1 ttl=64 time=0.944 ms
64 bytes from fd22:a075:afd0:e096:0:1:0:1a: icmp_seq=2 ttl=64 time=0.282 ms
64 bytes from fd22:a075:afd0:e096:0:1:0:1a: icmp_seq=3 ttl=64 time=0.254 ms

--- fd22:a075:afd0:e096:0:1:0:1a ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 0.254/0.493/0.944/0.319 ms