

---

Subject: Re: IPv6 auto-configuration issue with virtual Ethernet

Posted by [lars.bailey](#) on Fri, 17 Dec 2010 04:53:18 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

I am going to share my final conclusions, based on late night testing, and close the thread.

This is going to be a specific issue.

To refresh the test setup;

Node server - Fedora 13 Node

Test containers - Fedora 13

I downloaded two additional OS template caches from the OpenVZ wiki, and configured each for IPv4.

debian-5.0-x86

suse-11.1-x86

Both containers were bound to the Ethernet bridge, that sends RA's.

The issue of IPv4 configured containers, pulling auto-configured IPv6 addresses, does not seem to affect containers built with a Debian, or OpenSUSE template cache. (based on their specific type of networking setup)

Here is the Debian container's configuration output.

```
moe:~# ifconfig eth0
```

```
eth0   Link encap:Ethernet  HWaddr 00:18:51:03:78:64
       inet addr:192.168.100.101  Bcast:192.168.100.255  Mask:255.255.255.0
       inet6 addr: fe80::218:XXXX:fe03:7864/64 Scope:Link
       UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
       RX packets:456 errors:0 dropped:0 overruns:0 frame:0
       TX packets:55 errors:0 dropped:0 overruns:0 carrier:0
       collisions:0 txqueuelen:0
       RX bytes:45322 (44.2 KiB)  TX bytes:3704 (3.6 KiB)
```

```
moe:~# ip -4 ro show dev eth0
```

```
192.168.100.0/24 proto kernel scope link src 192.168.100.101
default via 192.168.100.101 scope link
```

```
moe:~# ip -6 ro show dev eth0
```

```
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
moe:~#
```

Here is the OpenSUSE container's configuration output.

```
shemp:~# # ifconfig eth0
```

```
eth0   Link encap:Ethernet  HWaddr 00:18:51:55:4B:36
       inet addr:192.168.100.103  Bcast:192.168.100.255  Mask:255.255.255.0
       inet6 addr: fe80::218:51ff:fe55:4b36/64 Scope:Link
       UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
```

```
RX packets:401 errors:0 dropped:0 overruns:0 frame:0
TX packets:27 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:38955 (38.0 Kb) TX bytes:1405 (1.3 Kb)
```

```
shemp:/ # ip -4 ro show dev eth0
192.168.100.0/24 proto kernel scope link src 192.168.100.103
169.254.0.0/16 scope link
default via 192.168.100.103
shemp:/ # ip -6 ro show dev eth0
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
shemp:/ #
```

And of course, the Fedora 13 container.

```
[root@curly /]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:18:51:78:64:52
          inet addr:192.168.100.102  Bcast:192.168.100.255  Mask:255.255.255.0
          inet6 addr: fd60:1014:9458:4b60:218:51ff:fe78:6452/64 Scope:Global
          inet6 addr: fe80::218:51ff:fe78:6452/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:473 errors:0 dropped:0 overruns:0 frame:0
          TX packets:45 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:46552 (45.4 KiB) TX bytes:2920 (2.8 KiB)
```

```
[root@curly /]# ip -4 ro show dev eth0
192.168.100.0/24 proto kernel scope link src 192.168.100.102
169.254.0.0/16 scope link
default via 192.168.100.102
[root@curly /]# ip -6 ro show dev eth0
fd60:1014:9458:4b60::/64 proto kernel metric 256 expires 2147151sec mtu 1500 advmss 1440
hoplimit 4294967295
fe80::/64 proto kernel metric 256 mtu 1500 advmss 1440 hoplimit 4294967295
default via fe80::214:bfff:fe5e:513f proto kernel metric 1024 expires 24sec mtu 1500 advmss
1440 hoplimit 64
[root@curly /]#
```

Now the question is, "If you use other types operating systems for a Node server, will this same issue arise in using containers, based on the Node operating system type?"

My opinion, is inclined to say "yes".

On a personal note, we have no intentions of using auto-configuration on any Node server, and I have no intentions of performing further auto-configuration testing.

I just wanted to experiment with auto-config IPv6, as we do have a few trivial clients, lying around on the network.

As far as the issue effecting our networking world, it won't.

But I thought I would share the info, as I have not found any docs on the Net/Wiki, pertaining to the

issue at hand.