
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [jamal](#) on Fri, 30 Jun 2006 13:50:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Andrey,

BTW - I was just looking at openvz, very impressive. To the other folks,
I am not putting down any of your approaches - just haven't
had time to study them. Andrey, this is the same thing you guys have
been working on for a few years now, you just changed the name, correct?

Ok, since you guys are encouraging me to speak, here goes ;->
Hopefully this addresses the other email from Herbert et al.

On Fri, 2006-30-06 at 11:45 +0400, Andrey Savochkin wrote:

> Hi Jamal,

>

> On Thu, Jun 29, 2006 at 08:15:52PM -0400, jamal wrote:

> > On Fri, 2006-30-06 at 09:07 +1200, Sam Vilain wrote:

> [snip]

>

> I do not have anything against guest-eth0 - eth0 pairs _if_ they are set up

> by the host administrators explicitly for some purpose.

> For example, if these guest-eth0 and eth0 devices stay as pure virtual ones,

> i.e. they don't have any physical NIC, host administrator may route traffic

> to guestXX-eth0 interfaces to pass it to the guests.

>

Well there will be purely virtual of course. Something along the lines
for openvz:

```
// create the guest
```

```
[host-node]# vzctl create 101 --ostemplate fedora-core-5-minimal
```

```
// create guest101::eth0, seems to only create config to boot up with
```

```
[host-node]# vzctl create 101 --netdev eth0
```

```
// bootup guest101
```

```
[host-node]# vzctl start 101
```

As soon as bootup of guest101 happens, creating guest101::eth0 should activate
creation of the host side netdevice. This could be triggered for example by
the netlink event message seen on host which is a result of creating guest101::eth0
Which means control sits purely in user space.

at that point if i do ifconfig on host i see g101-eth0
on guest101 i see just name eth0.

My earlier suggestion was that instead of:

```
host-node]# vzctl set 101 --ipadd 10.1.2.3
```

you do:

```
host-node]# ip addr add g101-eth0 10.1.2.3/32
```

you should still use vzctl to save config for next bootup

> However, I oppose the idea of automatic mirroring of _all_ devices appearing
> inside some namespaces ("guests") to another namespace (the "host").
> This clearly goes against the concept of namespaces as independent realms,
> and creates a lot of problems with applications running in the host, hotplug
> scripts and so on.
>

I was thinking that the host side is the master i.e you can peek at namespaces in the guest from the host.

Also note that having the pass through device allows for guests to be connected via standard linux schemes in the host side (bridge, point routes, tc redirect etc); so you dont need a speacial device to hook them together.

> > > Then the pragmatic question becomes how to correlate what you see from
> > > `ip addr list` to guests.
> >
> > on the host ip addr and the one seen on the guest side are the same.
> > Except one is seen (on the host) on guest0-eth0 and another is seen
> > on eth0 (on guest).
>
> Then what to do if the host system has 10.0.0.1 as a private address on eth3,
> and then interfaces guest1-tun0 and guest2-tun0 both get address 10.0.0.1
> when each guest has added 10.0.0.1 to their tun0 device?
>

Yes, that would be a conflict that needs to be resolved. If you look at ip addresses as also belonging to namespaces, then it should work, no? i am assuming a tag at the ifa table level.

cheers,
jamal
