Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by Andrey Savochkin on Fri, 30 Jun 2006 07:45:51 GMT
View Forum Message <> Reply to Message

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]
> > We plan to have them separate - so for
> > that to work, each network namespace could have an arbitrary "prefix"
> > that determines what the interface name will look like from the outside
> > when combined. We'd have to be careful about length limits.
> >
> > And guest0-eth0 doesn't necessarily make sense; it's not really an
> > ethernet interface, more like a tun or something.
> >
>
> it wouldnt quiet fit as a tun device. More like a mirror side of the
> guest eth0 created on the host side
> i.e a sort of passthrough device with one side visible on the host (send
> from guest0-eth0 is received on eth0 in the guest and vice-versa).
>
> Note this is radically different from what i have heard Andrey and co
> talk about and i dont wanna disturb any shit because there seems to be
> some agreement. But if you address me i respond because it is very
> interesting a topic;->

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.

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.

>
> > So, an equally good convention might be to use sequential prefixes on
> > the host, like "tun", "dummy", or a new prefix - then a property of that
> > is what the name of the interface is perceived to be to those who are in
> > the corresponding network namespace.
> >
> > 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?

Regards,

Andrey
