Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by Andrey Savochkin on Wed, 28 Jun 2006 16:58:48 GMT
View Forum Message <> Reply to Message

On Wed, Jun 28, 2006 at 12:17:35PM -0400, jamal wrote:
>
> On Wed, 2006-28-06 at 18:19 +0400, Andrey Savochkin wrote:
> >
> > Seeing guestXX-eth0 interfaces by standard tools has certain attractive
> > sides. But it creates a lot of undesired side effects.
> >
>
> I apologize because i butted into the discussion without perhaps reading
> the full thread.

Your comments are quite welcome

>
> > For example, ntpd queries all network devices by the same ioctls as ifconfig,
> > and creates separate sockets bound to IP addresses of each device, which is
> > certainly not desired with namespaces.
> >
>
> Ok, so the problem is that ntp in this case runs on the host side as

yes

> opposed to the guest? This would explain why Eric is reacting vehemently
> to the suggestion.

:)

And I actually do not want to distinguish host and guest sides much.
They are namespaces in the first place.
Parent namespace may have some capabilities to manipulate its child
namespaces, like donate its own device to one of its children.

But it comes secondary to having namespace isolation borders.
In particular, because most cases of cross-namespace interaction lead to
failures of formal security models and inability to migrate
namespaces between computers.

>
> > Or more subtle question: do you want hotplug events to be generated when
> > guest0-eth0 interface comes up in the root namespace, and standard scripts
> > to try to set some IP address on this interface?..
> >
>

> yes, thats what i was thinking. Even go further and actually create
> guestxx-eth0 on the host (which results in creating eth0 on the guest)
> and other things.

This actually goes in the opposite direction to what I keep in mind.
I want to offload as much as possible of network administration work to
guests. Delegation of management is one of the motivating factors
behind covering not only sockets but devices, routes, and so on by the
namespace patches.

>
> > In my opinion, the downside of this scheme overweights possible advantages,
> > and I'm personally quite happy with running commands with switched namespace,
> > like
> > vzctl exec guest0 ip addr list
> > vzctl exec guest0 ip link set eth0 up
> > and so on.
>
> Ok, above may be good enough and doesnt require any state it seems on
> the host side.
> I got motivated when the word "migration" was mentioned. I understood it
> to be meaning that a guest may become inoperative for some reason and
> that its info will be transfered to another guest which may be local or
> even remote. In such a case, clearly one would need a protocol and the
> state of all guests sitting at the host. Maybe i am over-reaching.

Migration will work inside the kernel, so it has full access
to whatever state information it needs.

Best regards

Andrey