Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by Andrey Savochkin on Wed, 28 Jun 2006 14:19:00 GMT
View Forum Message <> Reply to Message

Hi Jamal,

On Wed, Jun 28, 2006 at 09:53:23AM -0400, jamal wrote:
>
> On Wed, 2006-28-06 at 15:36 +0200, Herbert Poetzl wrote:
>
> > note: personally I'm absolutely not against virtualizing
> > the device names so that each guest can have a separate
> > name space for devices, but there should be a way to
> > 'see' _and_ 'identify' the interfaces from outside
> > (i.e. host or spectator context)
> >
>
> Makes sense for the host side to have naming convention tied
> to the guest. Example as a prefix: guest0-eth0. Would it not
> be interesting to have the host also manage these interfaces
> via standard tools like ip or ifconfig etc? i.e if i admin up
> guest0-eth0, then the user in guest0 will see its eth0 going
> up.

Seeing guestXX-eth0 interfaces by standard tools has certain attractive
sides. But it creates a lot of undesired side effects.

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.

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?..

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.

Best regards

Andrey

Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by jamal on Wed, 28 Jun 2006 16:17:35 GMT
View Forum Message <> Reply to Message

Andrey,

On Wed, 2006-28-06 at 18:19 +0400, Andrey Savochkin wrote:
> Hi Jamal,
>
> On Wed, Jun 28, 2006 at 09:53:23AM -0400, jamal 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.

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

> 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.

> 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.

cheers,
jamal

Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by Herbert Poetzl on Wed, 28 Jun 2006 17:04:15 GMT
View Forum Message <> Reply to Message

On Wed, Jun 28, 2006 at 06:19:00PM +0400, Andrey Savochkin wrote:
> Hi Jamal,
>
> On Wed, Jun 28, 2006 at 09:53:23AM -0400, jamal wrote:
> >
> > On Wed, 2006-28-06 at 15:36 +0200, Herbert Poetzl wrote:
> >
> > > note: personally I'm absolutely not against virtualizing
> > > the device names so that each guest can have a separate
> > > name space for devices, but there should be a way to
> > > 'see' _and_ 'identify' the interfaces from outside
> > > (i.e. host or spectator context)
> >
> > Makes sense for the host side to have naming convention tied
> > to the guest. Example as a prefix: guest0-eth0. Would it not
> > be interesting to have the host also manage these interfaces
> > via standard tools like ip or ifconfig etc? i.e if i admin up
> > guest0-eth0, then the user in guest0 will see its eth0 going
> > up.
>
> Seeing guestXX-eth0 interfaces by standard tools has certain
> attractive sides. But it creates a lot of undesired side effects.

which all can be avoided by not using the host
context for that, but a special 'all seeing'
context (as we have in Linux-VServer) which
can see (and probably manipulate) those interfaces
from the 'admin' PoV without entering the guest
context

> 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.

applications scanning the interfaces at startup
are broken by design and should probably be

fixed instead of worked around ...

> 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?..

why not, would it do any harm when the hotplug
scripts on the host would take the appropriate
actions (i.e. do the required config for the guest)
for special guest interfaces?

but now that you mention it, what about hotplug
events inside the guest?

> 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

I do not consider this the best solution, especially
from the security PoV. don't forget you basically
enter the guest and execute arbitrary programs
(which might have been compromised) to do a setup
task you actually want to happen on the host

best,
Herbert

> and so on.
>
> Best regards
>
> Andrey