
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [Herbert Poetzl](#) on Wed, 28 Jun 2006 13:36:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Jun 27, 2006 at 09:38:14PM -0600, Eric W. Biederman wrote:

> Alexey Kuznetsov <kuznet@ms2.inr.ac.ru> writes:

>

> > Hello!

> >

> >> It may look weird, but do application really *need* to see eth0 rather
> >> than eth858354?

> >

> > Applications do not care, humans do. :-)

> >

> > What's about applications they just need to see exactly the same
> > device after migration. Not only name, but f.e. also its ifindex.

> > If you do not create a separate namespace for netdevices, you will
> > inevitably end up with some strange hack sort of VPIDs to translate
> > (or to partition) ifindices or to tell that "ping -I eth858354 xxx"
> > is too complicated application to survive migration.

>

>

> Actually there are applications with peculiar licensing practices that
> do look at devices like eth0 to verify you have the appropriate mac, and
> do really weird things if you don't have an eth0.

>

> Plus there are other cases where it can be simpler to hard code things
> if it is allowable. (The human factor) Otherwise your configuration
> must be done through hotplug scripts.

>

> But yes there are misguided applications that care.

last time I pointed to such 'misguided' apps which
made assumptions that are not necessarily true
inside a virtual environment (e.g. pstree, initpid)
the general? position was that those apps should
be fixed instead adding a 'workaround'

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)

best,
Herbert

> Eric
