
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [ebiederm](#) on Wed, 28 Jun 2006 14:51:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dave Hansen <haveblue@us.ibm.com> writes:

> On Wed, 2006-06-28 at 00:52 +0200, Herbert Poetzl wrote:
>> seriously, what I think Eric meant was that it
>> might be nice (especially for migration purposes)
>> to keep the device namespace completely virtualized
>> and not just isolated ...
>
> It might be nice, but it is probably unneeded for an initial
> implementation. In practice, a cluster doing
> checkpoint/restart/migration will already have a system in place for
> assigning unique IPs or other identifiers to each container. It could
> just as easily make sure to assign unique network device names to
> containers.
>
> The issues really only come into play when you have an unstructured set
> of machines and you want to migrate between them without having prepared
> them with any kind of unique net device names beforehand.
>
> It may look weird, but do application really **need** to see eth0 rather
> than eth858354?

Actually there is a very practical reason we don't need to preserve device names across a migration event between machines, is the only sane thing to do is to generate a hotplug event that says you have removed the old interface and added a new interface.

My expectation is that during migration you will wind up with add and remove events for all of your hardware devices. But most applications because they do not access hardware devices directly will not care.

I haven't looked closely but I suspect this is one area where a container style approach will be noticeably different from a Xen or Vmware style approach.

Eric
