
Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by [Jeff Garzik](#) on Sat, 23 Jun 2007 22:15:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Miller wrote:

> I don't accept that we have to add another function argument
> to a bunch of core routines just to support this crap,
> especially since you give no way to turn it off and get
> that function argument slot back.
>
> To be honest I think this form of virtualization is a complete
> waste of time, even the openvz approach.
>
> We're protecting the kernel from itself, and that's an endless
> uphill battle that you will never win. Let's do this kind of
> stuff properly with a real minimal hypervisor, hopefully with
> appropriate hardware level support and good virtualized device
> interfaces, instead of this namespace stuff.

Strongly seconded. This containerized virtualization approach just
bloats up the kernel for something that is inherently fragile and IMO
less secure -- protecting the kernel from itself.

Plenty of other virt approaches don't stir the code like this, while
simultaneously providing fewer, more-clean entry points for the
virtualization to occur.

And that's speaking WITHOUT my vendor hat on...

Jeff

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
