Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by davem on Sat, 23 Jun 2007 20:57:37 GMT
View Forum Message <> Reply to Message

From: ebiederm@xmission.com (Eric W. Biederman)
Date: Sat, 23 Jun 2007 11:19:34 -0600

> Further and fundamentally all a global achieves is removing the need
> for the noise patches where you pass the pointer into the various
> functions. For long term maintenance it doesn't help anything.

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.

At least the hypervisor approach you have some chance to fully harden in some verifyable and truly protected way, with namespaces it's just a pipe dream and everyone who works on these namespace approaches knows that very well.

The only positive thing that came out of this work is the great auditing that the openvz folks have done and the bugs they have found, but it basically ends right there.