Subject: Re: [PATCH net-2.6.25 1/3]sysctl: make the sys.net.core sysctls per-namespace
Posted by davem on Sat, 08 Dec 2007 08:10:38 GMT
View Forum Message <> Reply to Message

From: Pavel Emelyanov <xemul@openvz.org>
Date: Fri, 07 Dec 2007 16:07:19 +0300

> Making them per-namespace is required for the following
> two reasons:
>
> First, some ctl values have a per-namespace meaning.
> Second, making them writable from the sub-namespace
> is an isolation hole.
>
> So I introduce the pernet operations to create these
> tables. For init_net I use the existing statically
> declared tables, for sub-namespace they are duplicated
> and the write bits are removed from the mode.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Applied to net-2.6.25, but I fear you're going to need to do something similar for ipv4, ipv6, and who knows what other protocols. And as a result we'll end up with all kinds of protocol specific things in the net namespace structure which totally stinks.

Such things will need to be registered dynamically just like the protocols themselves are into the kernel.