
Subject: Re: [PATCH 0/3] Sysctl shadow management
Posted by [Pavel Emelianov](#) on Wed, 21 Nov 2007 09:52:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:
> Pavel Emelyanov <xemul@openvz.org> writes:
>
>> Eric W. Biederman wrote:
>>> To be very very very clear.
>>>
>>> This is the way I think we should do the core sysctl infrastructure.
>>>
>>> On top of the register_sysctl_table patch, getting all of the
>>> infrastructure in at once.
>>>
>>> With a list of lists so we don't kill ourselves when we try to
>>> implement sysctls that are per network devices.
>> Hm... My patch looks to do very very same thing, but in
>> a bit simpler manner. Except for the absence of the
>> sysctl paths, but they are just cleanups. I think I can
>> port them on top of my shadows :) Thanks
>
> Then slow down and listen.

Press my ear close to the monitor.

> Your code simply does not address module unload races. BAD BAD BAD.

Can you unload IPC or UTS module? =)

> Your code skips sysctl_check allowing buggy sysctl tables to continue
> to exist BAD BAD BAD.

The tables are cloned from those who passed these checks. ;)

> Your code does not implement a list of lists needed to implement
> a dynamic list of network devices, instead leaving it to the user of
> sysctl registration code to do. (Which is where the bulk of your
> simplicity seems to come from). BAD BAD BAD.

Do we need this stuff for IPC? UTS? Stick to the patches please.

Network sysctls are more complex, I know it. Multiple shadows
per-head are to be done. I agree with that. I'm not trying
to create the most generic code that can solve all our needs
at once.

These patches do not provide any facility to display the

information via sysfs. So? Does anybody care about it?

I already have a bad experience with sending
doing-it-all-in-one-go patches to the community.

> So I'm sorry. Your code is simpler because it is WORSE.
>
> Your code is race prone, encourages buggy users, and doesn't even
> solve the same problem. So no I don't think it makes a bit
> of sense.

I decidedly object against such an outrageous statements.

> Eric
>

Pavel

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