
Subject: Re: [PATCH 4/4] net: Implement the per network namespace sysctl infrastructure

Posted by [serue](#) on Sat, 01 Dec 2007 00:01:50 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Eric W. Biederman (ebiederm@xmission.com):

> "Serge E. Hallyn" <serue@us.ibm.com> writes:

>

> >

> > Hey Eric,

> >

> > the patches look nice.

> >

> > The hand-forcing of the passed-in net_ns into a copy of current->nsproxy

> > does make it seem like nsproxy may not be the best choice of what to

> > pass in. Doesn't only net_sysctl_root->lookup() look at the argument?

>

> Yes. Although I call it from __register_sysctl_paths.

>

> > But I assume you don't want to be more general than sending in a

> > nsproxy so as to dissuade abuse of this interface for needlessly complex

> > sysctl interfaces?

>

> A bit of that. I would love to pass in a task_struct so you can use

> anything from a task. The trouble is I don't have any task_structs or

> nsproxys with the proper value at the point where I am first setting

> this up. Further I have to have the full sysctl lookup working or I

> could not call sysctl_check.

>

> > (Well I expect that'll become clear once the the patches using this

> > come out.)

> >

> > Are you planning to use this infrastructure for the uts and ipc

> > sysctls as well?

>

> Yes. Where it comes in especially useful, is I can move /proc/sys

> to /proc/sys/<tgid>/task/<pid>/sys. And get a particular processes

> view of sysctl.

>

> We also get a little more reuse of common functions.

>

> Otherwise Pavel does have a point that using this for uts and ipc

> is not a savings lines of code wise.

>

> After having seen Pavel changes I am asking myself if there is a sane

> way to remove the ctl_name argument from the ctl_path.

>

> Anyway where I am with the nsproxy question was that I don't

> see anything easily better. What I have works and gets the job
> done, and doesn't have any module unload races or holes where a sloppy
> programmer can mess up the sysctl tree. We needed a solution.
> Trying any harder to find something better would take ages. So
> I figured this implementation was good enough.

I agree. So it's already in -mm but still

Acked-by: Serge Hallyn <serue@us.ibm.com>

thanks,
-serge

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