
Subject: Re: [RFC][PATCH 4/5] utsname namespaces: sysctl hack

Posted by [serue](#) on Wed, 19 Apr 2006 20:24:14 GMT

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

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

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

>

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

> >> Kirill Korotaev <dev@sw.ru> writes:

> >>

> >> > Serge,

> >> >

> >> > can we do nothing with sysctls at this moment, instead of committing hacks?

> >>

> >> Except that we modify a static table changing the uts behaviour in

> >> proc_doutsstring isn't all that bad.

> >>

> >> I'm just about to start on something more comprehensive, in

> >> the sysctl case.

> >

> > So assuming that I take out the switch(), leaving that for a better

> > solution by Eric (or Dave, or whoever),

> >

> > Is it time to ask for the utsname namespace patch to be tried out

> > in -mm?

>

> Can we please suggest a syscall interface?

We can, but I was hoping that would be a separate patch, separate discussion.

Are you asking for a new syscall, specifically to unshare utsname()? Or for discussion over whether we want to do

one syscall per namespace

extend CLONE_NEWns flags

use unshare

use namespacefs

-serge
