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

Posted by [ebiederm](#) on Wed, 19 Apr 2006 16:52:54 GMT

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

Dave Hansen <haveblue@us.ibm.com> writes:

> Besides ipc and utsnames, can anybody think of some other things in
> sysctl that we really need to virtualize?

All of the networking entries.

> It seems to me that most of the other stuff is kernel-global and we
> simply won't allow anything in a container to touch it.
>
> That said, there may be things in the future that need to get added as
> we separate out different subsystems. Things like min_free_kbytes could
> have a container-centric meaning (although I think that is probably a
> really bad one to mess with).
>
> I have a slightly revamped way of doing the sysv namespace sysctl code.
> I've attached a couple of (still pretty raw) patches. Do these still
> fall in the "hacks" category?

Only in that you attacked the wrong piece of the puzzle.
The strategy table entries simply need to die, or be rewritten
to use the appropriate proc entries.

The proc entries are the real interface, and the two pieces
don't share an implementation unfortunately.

Eric
