
Subject: Re: [RFC] [PATCH] Re: Prefixing cgroup generic control filenames with "cgroup."

Posted by [Paul Menage](#) on Fri, 29 Feb 2008 09:34:37 GMT

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

On Thu, Feb 28, 2008 at 10:20 PM, Paul Jackson <pj@sgi.com> wrote:

>
> For example, what would have been cgroup:
>
> /mnt/cgroup/user_created_groupname1/user_created_groupname2
>
> now becomes:
>
> /mnt/cgroup/cgroups/user_created_groupname1/cgroups/user_created_groupname2
>
> Right?

Well, the additional components are called "groups" not "cgroups", but basically yes.

>
> You're doing this just to "protect the user from themselves", to make
> it more difficult for them to rely on some name that in a future
> version is no longer available.

Correct. "Future-proofing" and "Forward planning" are two alternatively-nuanced ways of describing this ...

>
> There is really a trivial resolution to this ... stake out what
> additional kernel generated names might ever be added ... some
> pattern(s) of characters which all future names will match, which
> leave wide swaths of names safely available, in perpetuity, for
> user created names, with no risk of future collision.

Yes, we could just say "the kernel reserves the right to use any names that begin with a lower-case letter, and no others", and be done with it. That leaves a bit of an ugly taste in my mouth, but if people seem to prefer that approach we can go for it.

>
> And did I say incompatible with released versions?

Not at all incompatible if it requires a mount option to enable it ...

Paul

Containers mailing list

