
Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [Paul Menage](#) on Thu, 08 Mar 2007 09:10:27 GMT

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

On 3/7/07, Sam Vilain <sam@vilain.net> wrote:

>
> Ok, they share this characteristic with namespaces: that they group
> processes. So, they conceptually hang off task_struct. But we put them
> on ns_proxy because we've got this vague notion that things might be
> better that way.

Remember that I'm not the one pushing to move them into ns_proxy.
These patches are all Srivatsa's work. Despite that fact that they say
"Signed-off-by: Paul Menage", I'd never seen them before they were
posted to LKML, and I'm not sure that they're the right approach.
(Although some form of unification might be good).

>
> >> about this you still insist on calling this sub-system specific stuff
> >> the "container",
> >>
> > Uh, no. I'm trying to call a *grouping* of processes a container.
> >
>
> Ok, so is this going to supplant the namespaces too?

I don't know. It would be nice to have a single object hanging off the
task struct that contains all the various grouping pointers. Having
something that was flexible enough to handle all the required
behaviours, or else allowing completely different behaviours for
different subsets of that structure, could be the fiddly bit.

See my expanded reply to Eric' earlier post for a possible way of
unifying them, and simplifying the nsproxy and container.c code in the
process.

>

Resource groups isn't a terrible name for them (although I'd be
wondering whether the BeanCounters folks would object :-)) but the
intention is that they're more generic than purely for resource
accounting. (E.g. see my other email where I suggested that things
like task->mempolicy and task->user could potentially be treated in
the same way)

Task Group is a good name, except for the fact that it's too easily

confused with process group.

>

> And do we bother changing IPC namespaces or let that one slide?

>

I think that "namespace" is a fine term for the IPC id virtualization/restriction that ipc_ns provides. (Unless I'm totally misunderstanding the concept).

Paul

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
