
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [Paul Menage](#) on Tue, 03 Apr 2007 16:52:35 GMT

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

On 4/3/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

> On Tue, Apr 03, 2007 at 08:45:37AM -0700, Paul Menage wrote:
> > Whilst I've got no objection in general to using nsproxy rather than
> > the container_group object that I introduced in my latest patches,
>
> So are you saying lets (re-)use task->nsproxy but also retain 'struct
> container' to store general per-group state? If so, I think that would
> address my main concern of redundant/avoidable new pointers in
> task_struct introduced in the container patches ..

I'm not saying "let's use nsproxy" - I'm not yet convinced that the lifetime/mutation/correlation rate of a pointer in an nsproxy is likely to be the same as for a container subsystem; if not, then reusing nsproxy could actually increase space overheads (since you'd end up with more, larger nsproxy objects, compared to smaller numbers of smaller nsproxy objects and smaller numbers of smaller container_group objects), even though it saved (just) one pointer per task_struct.

Reusing nsproxy versus using a separate container_group object as the aggregator is mostly an implementation detail, i think, and it would be pretty easy to switch from one to the other without any user-visible changes. So I'm inclined to keep them separate at this point.

Paul
