Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem
Posted by Paul Menage on Tue, 03 Apr 2007 17:10:35 GMT
View Forum Message <> Reply to Message

On 4/3/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:
> On Tue, Apr 03, 2007 at 09:52:35AM -0700, Paul Menage wrote:
> > 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.
>
> Even if nsproxy objects are made larger a bit, the number of such object will

You're not making them "a bit" larger, you're adding N+M pointers
where N is the number of container hierarchies and M is the number of
subsystem slots.

Basically, it means that anyone that uses containers without
namespaces or vice versa ends up paying the space overheads for both.

> be -much- lesser compared to number of task_structs I would think, so
> the win/lose in space savings would need to take that into account.

Agreed. So I'm not saying it's fundamentally a bad idea - just that
merging container_group and nsproxy is a fairly simple space
optimization that could easily be done later.

Paul