
Subject: Re: [ckrm-tech] [PATCH 1/7] containers (V7): Generic container system abstracted from cpusets code

Posted by [Paul Menage](#) on Thu, 08 Mar 2007 10:40:30 GMT

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

On 3/8/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

> On Wed, Mar 07, 2007 at 12:50:03PM -0800, Paul Menage wrote:

> > The callback mutex (which is what container_lock() actually locks) is

> > also used to synchronize fork/exit against subsystem additions, in the

> > event that some subsystem has registered fork or exit callbacks. We

> > could probably have a separate subsystem_mutex for that instead.

>

> Why can't manage_mutex itself be used there (to serialize fork/exit callbacks

> against modification to hierarchy)?

Because manage_mutex can be held for very long periods of time. I

think that a combination of a new lock that's only taken by fork/exit

and register_subsys, plus task_lock (which prevents the current task

from being moved) would be more lightweight.

Paul
