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

Posted by [Srivatsa Vaddagiri](#) on Thu, 08 Mar 2007 10:31:05 GMT

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

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)?

> Apart from that, yes, it may well be possible to move callback lock
> entirely inside cpusets.

Yes, that way only the hierarchy hosting cpusets takes the hit of double-locking. cpuset_subsys->create/destroy can take this additional lock inside cpuset.c.

--

Regards,
vatsa
