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

Posted by [Paul Menage](#) on Wed, 07 Mar 2007 20:50:03 GMT

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

On 3/7/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

> If that is the case, I think we can push container_lock entirely inside
> cpuset.c and not have others exposed to this double-lock complexity.
> This is possible because cpuset.c (build on top of containers) still has
> cpuset->parent and walking cpuset->parent list safely can be made
> possible with a second lock which is local to only cpuset.c.
>

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.

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

Paul
