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

Posted by [Paul Menage](#) on Mon, 12 Feb 2007 19:26:34 GMT

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

On 2/12/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:

> On Mon, Feb 12, 2007 at 12:15:22AM -0800, menage@google.com wrote:

> > +void container_fork(struct task_struct *child)

> > +{

> > + task_lock(current);

>

> Can't this be just rcu_read_lock()?

>

In this particular patch (which is an almost verbatim extraction/renaming of the generic bits of the cpuset code into container.c) it probably could - but the main patch that adds the container_group support would lose it since we use kref to refcount container_group objects, and hence they're freed when their refcount reaches zero. RCU is still fine for reading the container_group pointers, but it's no good for updating them, since by the time you update it it may no longer be your container_group structure, and may instead be about to be deleted as soon as the other thread's rcu_synchronize() completes.

Paul
