Subject: Re: [PATCH 3/7] Containers (V8): Add generic multi-subsystem API to containers
Posted by Srivatsa Vaddagiri on Sat, 07 Apr 2007 04:11:36 GMT
View Forum Message <> Reply to Message

On Fri, Apr 06, 2007 at 04:32:24PM -0700, menage@google.com wrote:
> +static int attach_task(struct container *cont, struct task_struct *tsk)
> {

[snip]

> + task_lock(tsk);

You need to check here if task state is PF_EXITING and fail with
-ESRCH if so? Otherwise we risk breaking refcount on
init_container_group.

> + rcu_assign_pointer(tsk->containers, newcg);
> + task_unlock(tsk);

--
Regards,
vatsa

Subject: Re: [PATCH 3/7] Containers (V8): Add generic multi-subsystem API to containers
Posted by Paul Menage on Sat, 07 Apr 2007 17:30:27 GMT
View Forum Message <> Reply to Message

On 4/6/07, Srivatsa Vaddagiri <vatsa@in.ibm.com> wrote:
> On Fri, Apr 06, 2007 at 04:32:24PM -0700, menage@google.com wrote:
> > +static int attach_task(struct container *cont, struct task_struct *tsk)
> > {
>
> [snip]
>
> > + task_lock(tsk);
>
> You need to check here if task state is PF_EXITING and fail with
> -ESRCH if so? Otherwise we risk breaking refcount on
> init_container_group.
>

Yes, I think you're right; I've now changed it to this in my tree:

```
task_lock(tsk);
if (tsk->flags & PF_EXITING) {
```

```
                task_unlock(tsk);
                put_container_group(newcg);
                return -ESRCH;
        }
        rcu_assign_pointer(tsk->containers, newcg);
        task_unlock(tsk);
```

Thanks,

Paul
