Subject: Re: [ckrm-tech] [PATCH 1/7] containers (V7): Generic container system abstracted from cpusets code
Posted by Srivatsa Vaddagiri on Sun, 25 Mar 2007 04:16:02 GMT
View Forum Message <> Reply to Message

On Sun, Mar 25, 2007 at 07:58:16AM +0530, Srivatsa Vaddagiri wrote:
> Not just this, continuing further we have more trouble:
>
> ---------- --------
> CPU0 (attach_task T1 to CS2) CPU1 (T1 is exiting)
> ---------- --------
>
> synchronize_rcu()
> atomic_dec(&CS1->count);
> [CS1->count = 0]
>
> if atomic_dec_and_test(&oldcs->count))
> [CS1->count = -1]
>
>
>
> We now have CS1->count negative. Is that good? I am uncomfortable ..
>
> We need a task_lock() in cpuset_exit to avoid this race.

2nd race is tricky. We probably need to do this to avoid it:

```
 task_lock(tsk);

 /* Check if tsk->cpuset is still same. We may have raced with
  * cpuset_exit changing tsk->cpuset again under our feet.
  */
 if (tsk->cpuset == cs && atomic_dec_and_test(&oldcs->count)) {
  task_unlock(tsk);
                check_for_release(oldcs, ppathbuf);
  goto done;
 }

 task_unlock(tsk);

done:
 return 0;
```

--
Regards,
vatsa