
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 04:57:52 GMT

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

On Sat, Mar 24, 2007 at 09:45:50PM -0700, Paul Jackson wrote:

> Nice work - thanks. Yes, both an extra cpuset count and a negative
> cpuset count are bad news, opening the door to the usual catastrophes.
>
> Would you like the honor of submitting the patch to add a task_lock
> to cpuset_exit()? If you do, be sure to fix, or at least remove,
> the cpuset_exit comment lines:

I will try to send out a patch later today to fix this bug in mainline cpuset code. I happened to notice this race with my rcfs patch and observed same is true with cpuset/container code also.

> * We don't need to task_lock() this reference to tsk->cpuset,
> * because tsk is already marked PF_EXITING, so attach_task() won't
> * mess with it, or task is a failed fork, never visible to attach_task.

Sure, I had seen that.

> So, in real life, this would be a difficult race to trigger.

Agreed, but good to keep code clean isn't it? :)

> Thanks for finding this.

Wellcome!

--

Regards,
vatsa
