Subject: Re: [PATCH] task containersv11 add tasks file interface fix for cpusets Posted by Paul Menage on Sat, 06 Oct 2007 20:53:53 GMT

View Forum Message <> Reply to Message

On 10/6/07, Paul Jackson <pj@sgi.com> wrote:

>

> This isn't working for me.

>

- > The key kernel routine for updating a tasks cpus_allowed
- > cannot be called while holding a spinlock.

>

- > But the above loop holds a spinlock, css set lock, between
- > the cgroup_iter_start and the cgroup_iter_end.

>

- > I end up generating complaints of:
- > BUG: scheduling while atomic
- > when I invoke the set_cpus_allowed() above.

>

- > Should css_set_lock be a mutex? Locking changes like that
- > can be risky.

css_set_lock is an rwlock currently; I'd rather not turn it into an rw mutex since there are places that it gets taken where we can't afford to sleep.

>

- > Or perhaps there should be another callback, called only by
- > attach() requests back to the same group. Likely cpusets would
- > be the only subsystem interested in plugging that callback.

>

- > That, or my original patch, which calls the attach routine
- > even if re-attaching to the current cgroup ...

I'd prefer David's solution of grabbing references to tasks during the iteration and then doing set_cpus_allowed outside the tasklist_lock.

Paul

Contain are realling list

Containers mailing list

Containers@lists.linux-foundation.org

https://lists.linux-foundation.org/mailman/listinfo/containers