
Subject: Re: Re: Hang with fair cgroup scheduler (reproducer is attached.)
Posted by [Dmitry Adamushko](#) on Sun, 16 Dec 2007 10:06:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 16/12/2007, Ingo Molnar <mingo@elte.hu> wrote:

```
>
> * Dmitry Adamushko <dmitry.adamushko@gmail.com> wrote:
>
> > --- a/kernel/sched.c
> > +++ b/kernel/sched.c
> > @@ -7360,7 +7360,7 @@ void sched_move_task(struct task_struct *tsk)
> >
> >     update_rq_clock(rq);
> >
> > -    running = task_running(rq, tsk);
> > +    running = (rq->curr == tsk);
> >     on_rq = tsk->se.on_rq;
> >
>
> thanks, i've queued this up (pending more testing).
```

btw., sched_setscheduler() and rt_mutex_setprio() are also affected
(in general, anything that may call put_prev_task/set_curr_task()
relying task_running()).

Will see, maybe we may come up with smth better than just replacing
task_running() with (rq->curr == tsk) there.

```
> Btw., you should be able to force the ia64 scheduling by adding this to
> the very top of include/linux/sched.h:
>
> #define __ARCH_WANT_UNLOCKED_CTXSW
> #define __ARCH_WANT_INTERRUPTS_ON_CTXSW
```

Yeah, with both we even get ARM behavior. Can be a good test indeed.

```
>
>     Ingo
>
```

--
Best regards,
Dmitry Adamushko

Containers mailing list
Containers@lists.linux-foundation.org

