
Subject: Re: [PATCH 1/5] Revert recent removal of set_curr_task()

Posted by [Ingo Molnar](#) on Mon, 24 Sep 2007 16:35:25 GMT

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

* Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:

> > (3) rework enqueue/dequeue_entity() to get rid of
> > sched_class::set_curr_task()
>
> Dmitry/Ingo,
> I am sorry for not having reviewed this change properly, but I
> think we need to revert this.

ah, i was wondering about that already. We can certainly skip that optimization.

> In theory its possible to solve these problems w/o reintroducing
> set_curr_task(). I tried doing so, but found it clutters
> dequeue_entity and enqueue_entity a lot and makes it less readable. It
> will duplicate what put_prev_entity() and set_next_entity() are
> supposed to do. Moreover it is slightly inefficient to do all these
> in dequeue_entity() if we consider that dequeue_entity can be called
> on current task for other reasons as well (like when it is abt to
> sleep or change its nice value).

yeah, it's not worth it. I'd go for keeping the code unified even if
adds a few instructions runtime overhead, as i'd expect most distros to
enable fair-group-scheduling by default in the future. (once all the
containers infrastructure and tools has trickled down to them)

Ingo

Containers mailing list

Containers@lists.linux-foundation.org

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