
Subject: Re: [PATCH v3 1/6] measure exec_clock for rt sched entities

Posted by [Glauber Costa](#) on Wed, 30 May 2012 12:09:54 GMT

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

On 05/30/2012 03:00 PM, Paul Turner wrote:

> Does this even need to be hierarchical? While it's natural for it to
> be in the CFS case, it feels forced here.

>

> You could instead make this rt_rq->local_exec_clock charging only to
> the parenting rt_rq and post-aggregate when you want to report. The
> only thing you'd need to be careful of is also accounting children
> somewhere on the parent on destruction (reaped_exec_clock?).

>

> Harking back to symmetry, local_exec_clock is also a potentially
> useful stat on the CFS side of things since it allows you to usefully
> disambiguate versus your children (common case where this is useful is
> calculating usage of threads in the root cgroup); so it wouldn't need
> to be unique to rt_rq.

I can try this approach.
