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

Posted by [Peter Zijlstra](#) on Wed, 30 May 2012 10:29:19 GMT

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

On Wed, 2012-05-30 at 13:48 +0400, Glauber Costa wrote:

> For simetry with the cfq tasks, measure exec_clock for the rt
> sched entities (rt_se).

Symmetry methinks.. anyway, where is the symmetry?, fair.c:update_curr()
doesn't do the for_each_sched_entity() thing.

> This can be used in a number of fashions. For instance, to
> compute total cpu usage in a cgroup that is generated by
> rt tasks.

>

> Signed-off-by: Glauber Costa <glommer@parallels.com>

> CC: Peter Zijlstra <a.p.zijlstra@chello.nl>

> CC: Paul Turner <pjt@google.com>

> ---

> kernel/sched/rt.c | 5 +++++

> kernel/sched/sched.h | 1 +

> 2 files changed, 6 insertions(+)

>

> diff --git a/kernel/sched/rt.c b/kernel/sched/rt.c

> index c5565c3..30ee4e2 100644

> --- a/kernel/sched/rt.c

> +++ b/kernel/sched/rt.c

> @@ -919,6 +919,11 @@ static void update_curr_rt(struct rq *rq)

>

> sched_rt_avg_update(rq, delta_exec);

>

> + for_each_sched_rt_entity(rt_se) {

> + rt_rq = rt_rq_of_se(rt_se);

> + schedstat_add(rt_rq, exec_clock, delta_exec);

> + }

> +

> if (!rt_bandwidth_enabled())

> return;

See, this just makes me sad.. you now have a double
for_each_sched_rt_entity() loop.
