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

Posted by [Peter Zijlstra](#) on Wed, 30 May 2012 10:42:35 GMT

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

On Wed, 2012-05-30 at 14:32 +0400, Glauber Costa wrote:

```
>>> + 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.

>

> The way I read the rt.c code, it is called from enqueue_task_rt only

> once.

Ah, what I meant was, right after that !rt_bandwidth_enabled() muck we do another for_each_sched_rt_entity() walk.
