

---

Subject: Re: [PATCH v3 1/6] measure exec\_clock for rt sched entities

Posted by [Glauber Costa](#) on Wed, 30 May 2012 10:32:07 GMT

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

---

On 05/30/2012 02:29 PM, Peter Zijlstra wrote:

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

>> For symmetry with the cfq tasks, measure exec\_clock for the rt

>> sched entities (rt\_se).

>

> Symmetry methinks..

=p bad me

> anyway, where is the symmetry?, fair.c:update\_curr()

> doesn't do the for\_each\_sched\_entity() thing.

It does implicitly, because fair.c:update\_curr() is called from within enqueue\_task(), that is called for\_each\_sched\_entity in enqueue\_task\_fair().

>

>> 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.

The way I read the rt.c code, it it is called from enqueue\_task\_rt only once.

---