Subject: [PATCH v3 1/6] measure exec_clock for rt sched entities
Posted by Glauber Costa on Wed, 30 May 2012 09:48:32 GMT
View Forum Message <> Reply to Message

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

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;

diff --git a/kernel/sched/sched.h b/kernel/sched/sched.h
index ba9dccf..cd2f1e1 100644
--- a/kernel/sched/sched.h
+++ b/kernel/sched/sched.h
@@ -295,6 +295,7 @@ struct rt_rq {
  struct plist_head pushable_tasks;
 #endif
  int rt_throttled;
+ u64 exec_clock;
  u64 rt_time;
  u64 rt_runtime;
  /* Nests inside the rq lock: */
--

1.7.10.2