Subject: [PATCH v3 0/6] per cgroup /proc/stat statistics
Posted by Glauber Costa on Wed, 30 May 2012 09:48:31 GMT
View Forum Message <> Reply to Message

Hi

This is a rework of my v2 series taking into account Paul
Turner's comments. All the series is the same except for
the patch that collets nr_switches that is now completely reworked.

It is out of schedule(), relegated to a scheduler hook. Only difference
from Paul's suggestion is that I added a new one, instead of messing with
put_prev_task + pick_next_task pairs. In the patchset that adds it,
I'll try to argue for my solution.

Hope this is acceptable, and please let me know of any other concerns.

v2:
* completely reworked nr_switches gathering
* separated per-se sleep_start to be more clear about it

Glauber Costa (6):
  measure exec_clock for rt sched entities
  account guest time per-cgroup as well.
  expose fine-grained per-cpu data for cpuacct stats
  add a new scheduler hook for context switch
  Also record sleep start for a task group
  expose per-taskgroup schedstats in cgroup

```
 include/linux/sched.h |    1 +
 kernel/sched/core.c   |  166 ++++++++++++++++++++++++++++++++++++++++++++++++--
 kernel/sched/fair.c   |   42 ++++++++++++-
 kernel/sched/rt.c     |   20 ++++++
 kernel/sched/sched.h  |    6 ++
 5 files changed, 228 insertions(+), 7 deletions(-)
```

--
1.7.10.2