Subject: [PATCH 0/5] per-cpu/cpuacct cgroup scheduler statistics
Posted by Glauber Costa on Thu, 02 Feb 2012 14:19:27 GMT
View Forum Message <> Reply to Message

Hi,

Here is my new attempt to get a per-container version of some
/proc data such as /proc/stat and /proc/uptime.

In this series I solved the visibility problem, which is,
the problem of how and when to show /proc/stat data per-cgroup,
by declaring it not a problem.

This can probably be done in userspace with other aids, like mounting
a fuse overlay that simulates /proc from outside a container, to a
container location.

Here, we should have most of the data needed to do that. They are drawn
from both the cpu cgroup, and cpuacct. Each cgroup exports the data it
knows better, and I am not really worried here about bindings between them.

In this first version, I am using clock_t units, being quite proc-centric.
It made my testing easier, but I am happy to show any units you guys would
prefer.

Besides that, it still has some other minor issues to be sorted out.
But I verified the general direction to be working, and would like to know
what you think.

Thanks

Glauber Costa (5):
  make steal time's to-tick routine generic
  store number of iowait events in a task_group
  account guest time per-cgroup as well.
  expose fine-grained per-cpu data for cpuacct stats
  expose per-taskgroup schedstats in cgroup

 include/linux/sched.h | 1 +
 kernel/sched/core.c | 207 ++----
 kernel/sched/fair.c | 45 ++++++++++
 kernel/sched/sched.h | 3 +
 4 files changed, 242 insertions(+), 14 deletions(-)

--
1.7.7.4