
Subject: [PATCH v3 2/6] account guest time per-cgroup as well.
Posted by [Glauber Costa](#) on Wed, 30 May 2012 09:48:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

We already track multiple tick statistics per-cgroup, using the task_group_account_field facility. This patch accounts guest_time in that manner as well.

Signed-off-by: Glauber Costa <glommer@parallels.com>
CC: Peter Zijlstra <a.p.zijlstra@chello.nl>
CC: Paul Turner <pjt@google.com>

kernel/sched/core.c | 10 +++++-----
1 file changed, 4 insertions(+), 6 deletions(-)

diff --git a/kernel/sched/core.c b/kernel/sched/core.c
index 39eb601..220d416 100644

--- a/kernel/sched/core.c

+++ b/kernel/sched/core.c

@@ -2717,8 +2717,6 @@ void account_user_time(struct task_struct *p, cputime_t cputime,
static void account_guest_time(struct task_struct *p, cputime_t cputime,
cputime_t cputime_scaled)

{
- u64 *cpustat = kcpustat_this_cpu->cpustat;
-

/* Add guest time to process. */

p->utime += cputime;
p->utimescaled += cputime_scaled;

@@ -2727,11 +2725,11 @@ static void account_guest_time(struct task_struct *p, cputime_t
cputime,

/* Add guest time to cpustat. */

if (TASK_NICE(p) > 0) {
- cpustat[CPUTIME_NICE] += (__force u64) cputime;
- cpustat[CPUTIME_GUEST_NICE] += (__force u64) cputime;
+ task_group_account_field(p, CPUTIME_NICE, (__force u64) cputime);
+ task_group_account_field(p, CPUTIME_GUEST, (__force u64) cputime);
} else {
- cpustat[CPUTIME_USER] += (__force u64) cputime;
- cpustat[CPUTIME_GUEST] += (__force u64) cputime;
+ task_group_account_field(p, CPUTIME_USER, (__force u64) cputime);
+ task_group_account_field(p, CPUTIME_GUEST, (__force u64) cputime);
}
}

--

1.7.10.2
