

---

Subject: [PATCH v4 1/4] account guest time per-cgroup as well.

Posted by [Glauber Costa](#) on Tue, 05 Jun 2012 14:49:38 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](mailto:glommer@parallels.com)>

CC: Peter Zijlstra <[a.p.zijlstra@chello.nl](mailto:a.p.zijlstra@chello.nl)>

CC: Paul Turner <[pjt@google.com](mailto: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

---