
Subject: Re: [RFC][mm] Simple stats for cpu resource controller v4

Posted by [akpm](#) on Mon, 12 May 2008 23:48:33 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Mon, 12 May 2008 01:48:37 +0530

Balaji Rao <balajirao@gmail.com> wrote:

> Hi Andrew,

>

> Here's a version that uses percpu_counters and which actually works. The

> only evil it contains is the check,

>

> if (percpu_counter_ready) {

> ..

> }

There's no instance of "percpu_counter_ready" in the patch so I'm a bit stumped.

> @@ -3837,6 +3858,16 @@ void account_user_time(struct task_struct *p, cputime_t cputime)

> cpustat->nice = cputime64_add(cpustat->nice, tmp);

> else

> cpustat->user = cputime64_add(cpustat->user, tmp);

> +

> + /* Charge the task's group */

> + #ifdef CONFIG_CGROUP_SCHED

> + {

> + struct task_group *tg;

> + tg = task_group(p);

> + __cpu_cgroup_stat_add(tg->stat, CPU_CGROUP_STAT_UTIME,

> + cputime_to_msecs(cputime));

> + }

> + #endif

> }

>

> /*

> @@ -3892,8 +3923,17 @@ void account_system_time(struct task_struct *p, int hardirq_offset,

> cpustat->irq = cputime64_add(cpustat->irq, tmp);

> else if (softirq_count())

> cpustat->softirq = cputime64_add(cpustat->softirq, tmp);

> - else if (p != rq->idle)

> + else if (p != rq->idle) {

> cpustat->system = cputime64_add(cpustat->system, tmp);

> + #ifdef CONFIG_CGROUP_SCHED

> + {

> + struct task_group *tg;

> + tg = task_group(p);

> + __cpu_cgroup_stat_add(tg->stat, CPU_CGROUP_STAT_STIME,

```
> +    cputime_to_msecs(cputime));  
> + }  
> + }  
> + #endif
```

I'd suggest that the above be turned into calls to a helper function which is a no-op if !CONFIG_CGROUP_SCHED.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
