

---

Subject: Re: [PATCH v3 6/6] expose per-taskgroup schedstats in cgroup  
Posted by [Peter Zijlstra](#) on Wed, 30 May 2012 11:22:47 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2012-05-30 at 13:48 +0400, Glauber Costa wrote:

```
> +static u64 tg_idle(struct task_group *tg, int cpu)
> +{
> +    u64 val;
> +
> +    if (tg != &root_task_group) {
> +        val = cfs_read_sleep(tg->se[cpu]);
> +        /* If we have rt tasks running, we're not really idle */
> +        val -= rt_rq(exec_clock, tg, cpu);
> +    } else
> +        /*
> +         * There are many errors here that we are accumulating.
> +         * However, we only provide this in the interest of having
> +         * a consistent interface for all cgroups. Everybody
> +         * probing the root cgroup should be getting its figures
> +         * from system-wide files as /proc/stat. That would be faster
> +         * to begin with...
> +         *
> +         * Ditto for steal.
> +         */
> +        val = kcpustat_cpu(cpu).cpustat[CPUTIME_IDLE] * TICK_NSEC;
```

You just violated 2 coding style rules in one go :-)

If one side of the if-else has braces, the other side should have too.  
If a block is multi-line (regardless of multi-stmt) it should have braces.

/me hands you a bucket full of {}.

```
> +    return val;
> +}
```

---