
Subject: Re: [RFC][-mm] Simple stats for cpu resource controller v4
Posted by [Balaji Rao](#) on Tue, 13 May 2008 03:30:27 GMT

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

On Tuesday 13 May 2008 05:18:33 am Andrew Morton wrote:

> 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.

>

Here it is. 'stat' can tell us if percpu_counters are ready or not.

```
+static void __cpu_cgroup_stat_add(struct cpu_cgroup_stat *stat,
```

```
+{
```

```
    }
```

```
+#endif
```

It's here as well.

```
+static s64 cpu_cgroup_read_stat(struct cpu_cgroup_stat *stat,
```

```
+{
```

```
    
```

```
    }
```

```
    
```

> > @@ -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.

```

OK. Will do.

--
Warm Regards,

Balaji Rao
Dept. of Mechanical Engineering
NITK

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