
Subject: Re: [PATCH 2/2] CFS CGroup: Report usage
Posted by [Srivatsa Vaddagiri](#) on Tue, 23 Oct 2007 02:30:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, Oct 22, 2007 at 05:49:39PM -0700, Paul Menage wrote:
> +static u64 cpu_usage_read(struct cgroup *cgrp, struct cftype *cft)
> +{
> + struct task_group *tg = cgroup_tg(cgrp);
> + int i;
> + u64 res = 0;
> + for_each_possible_cpu(i) {
> + unsigned long flags;
> + spin_lock_irqsave(&tg->cfs_rq[i]->rq->lock, flags);

Is the lock absolutely required here?

Hmm .. I hope the cgroup code prevents a task group from being destroyed while we are still reading a task group's cpu usage. Is that so?

```
> + res += tg->se[i]->sum_exec_runtime;  
> + spin_unlock_irqrestore(&tg->cfs_rq[i]->rq->lock, flags);  
> + }  
> + /* Convert from ns to ms */  
> + do_div(res, 1000000);  
> + return res;  
> +}
```

--
Regards,
vatsa

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