
Subject: Re: Revert for cgroups CPU accounting subsystem patch

Posted by [Balbir Singh](#) on Tue, 13 Nov 2007 07:00:20 GMT

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

Paul Menage wrote:

> On Nov 12, 2007 10:00 PM, Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com> wrote:

>> On second thoughts, this may be a usefull controller of its own.

>> Say I just want to "monitor" usage (for accounting purpose) of a group of

>> tasks, but don't want to control their cpu consumption, then cpuacct

>> controller would come in handy.

>>

>

> That's plausible, but having two separate ways of tracking and

> reporting the CPU usage of a cgroup seems wrong.

>

> How bad would it be in your suggested case if you just give each

> cgroup the same weight? So there would be fair scheduling between

> cgroups, which seems as reasonable as any other choice in the event

> that the CPU is contended.

>

Right now, one of the limitations of the CPU controller is that the moment you create another control group, the bandwidth gets divided by the default number of shares. We can't create groups just for monitoring. cpu_acct fills this gap. I think in the long run, we should move the helper functions into cpu_acct.c and the interface logic into kernel/sched.c (cpu controller).

--

Warm Regards,

Balbir Singh

Linux Technology Center

IBM, ISTL

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
