
Subject: Re: Revert for cgroups CPU accounting subsystem patch

Posted by [Balbir Singh](#) on Tue, 13 Nov 2007 07:29:59 GMT

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

Paul Menage wrote:

> On Nov 12, 2007 11:00 PM, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>> 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.

>

> Could we get around this with, say, a flag that always treats a CFS
> schedulable entity as having a weight equal to the number of runnable
> tasks in it? So CPU bandwidth would be shared between groups in
> proportion to the number of runnable tasks, which would distribute the
> cycles approximately equivalently to them all being separate
> schedulable entities.

>

I think it's a good hack, but not sure about the complexity to implement the code. I worry that if the number of tasks increase (say run into thousands for one or more groups and a few groups have just a few tasks), we'll lose out on accuracy.

>> cpu_acct fills this gap.

>

> Agreed, but not in the right way IMO.

>

I think we already have the code, we need to make it more useful and reusable.

--

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>
