
Subject: Re: Revert for cgroups CPU accounting subsystem patch
Posted by [Srivatsa Vaddagiri](#) on Tue, 13 Nov 2007 07:49:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Nov 13, 2007 at 12:59:59PM +0530, Balbir Singh wrote:

> 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 agree that it would be adding unnecessary complexity, just to meet the
accounting needs.

Thinking of it more, this requirement to "group tasks for only accounting
purpose" may be required for other resources (mem, io, network etc) as well?
Should we have a generic accounting controller which can provide these
various resource usage stats for a group (cpu, mem etc) by iterating thr' the
task list for the group and summing up the corresponding stats already present
in task structure?

--

Regards,
vatsa

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