
Subject: Re: [PATCH 0/5] per-cpu/cpuacct cgroup scheduler statistics
Posted by [Serge E. Hallyn](#) on Tue, 14 Feb 2012 22:31:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

Quoting Glauber Costa (glommer@parallels.com):

> On 02/02/2012 06:19 PM, Glauber Costa wrote:
> >Hi,
> >
> >Here is my new attempt to get a per-container version of some
> >/proc data such as /proc/stat and /proc/uptime.
> >
> >In this series I solved the visibility problem, which is,
> >the problem of how and when to show /proc/stat data per-cgroup,
> >by declaring it not a problem.
> >
> >This can probably be done in userspace with other aids, like mounting
> >a fuse overlay that simulates /proc from outside a container, to a
> >container location.
> >
> >Here, we should have most of the data needed to do that. They are drawn
> >from both the cpu cgroup, and cpuacct. Each cgroup exports the data it
> >knows better, and I am not really worried here about bindings between them.
> >
> >In this first version, I am using clock_t units, being quite proc-centric.
> >It made my testing easier, but I am happy to show any units you guys would
> >prefer.
> >
> >Besides that, it still has some other minor issues to be sorted out.
> >But I verified the general direction to be working, and would like to know
> >what you think.
> >
> >
> >
> Hi,
> >
> Did someone had any chance to take a look at this already?
> >
> Thanks

Hi,

By declaring proc visibility not a problem and sticking to io stats,
you sort of left me where I don't know what I'm talking about :) So
let me just say, on patch 2, "store number of iowait events in a task_group",
my initial reaction is "boy that's a lot more work. What is the performance
impact?"

It'd be possible to move the extra processing out of the hot-path by
only changing the # for the deepest cgroup, and pulling it into

ancestor cgroups only when someone is viewing the stats or the child cgroup goes away. But if you have #s showing statistically negligible performance impact anyway then that wouldn't be worth it.

-serge
