
Subject: Re: [PATCH v2 4/5] expose fine-grained per-cpu data for cpuacct stats
Posted by [Glauber Costa](#) on Wed, 18 Apr 2012 16:14:48 GMT

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

On 04/18/2012 09:30 AM, Sha Zhengju wrote:

> On Mon, Apr 9, 2012 at 6:25 PM, Glauber Costa<glommer@parallels.com> wrote:

>> > The cpuacct cgroup already exposes user and system numbers in a per-cgroup
>> > fashion. But they are a summation along the whole group, not a per-cpu figure.

>> > Also, they are coarse-grained version of the stats usually shown at places
>> > like /proc/stat.

>> >

>> > I want to have enough cgroup data to emulate the /proc/stat interface. To
>> > achieve that, I am creating a new file "stat_percpu" that displays the
>> > fine-grained per-cpu data. The original data is left alone.

>> >

>> > The format of this file resembles the one found in the usual cgroup's stat
>> > files. But of course, the fields will be repeated, one per cpu, and prefixed
>> > with the cpu number.

>> >

>> > Therefore, we'll have something like:

>> >

>> > cpu0.user X

>> > cpu0.system Y

>> > ...

>> > cpu1.user X1

>> > cpu1.system Y1

>> > ...

>> >

> Why not show the all-cpu data together with the per-cpu one? I think
> the total one

> is an usual concern in most cases.

>

Because that is a trivial operation that can be done in userspace.

In general, I see no value in formatting this file any further if we'll
have to get to userspace for the final solution anyway.
