
Subject: Re: How to draw values for /proc/stat

Posted by [KAMEZAWA Hiroyuki](#) on Tue, 06 Dec 2011 00:05:43 GMT

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

On Mon, 5 Dec 2011 07:32:33 -0200

Glauber Costa <glommer@parallels.com> wrote:

> Hi,

>

> Specially Peter and Paul, but all the others:

>

> As you can see in <https://lkml.org/lkml/2011/12/4/178>, and in my answer
> to that, there is a question - one I've asked before but without that
> much of an audience - of whether /proc files read from process living on
> cgroups should display global or per-cgroup resources.

>

> In the past, I was arguing for a knob to control that, but I recently
> started to believe that a knob here will only overcomplicate matters:
> if you live in a cgroup, you should display only the resources you can
> possibly use. Global is for whoever is in the main cgroup.

>

Hm. I have a suggestion and a concern.

(A suggestion)

How about having a mount option for procfs ?

For example,

`mount -t proc -o cgroup_virtualized`

Then, /proc/stat etc shows per-cgroup information.

(A concern)

/proc/stat will be a mixture of virtualized values and not-virtualized values.

1. Don't users need to know whether each value is virtualized or not ?

2. Can we have a way to show "this value is virtualized!" annotation ?

> Now, it comes two questions:

> 1) Do you agree with that, for files like /proc/stat ? I think the most
> important part is to be consistent inside the system, regardless of what
> is done

>

I think some kind of care for users are required as I wrote above.

> 2) Will cpuacct stay? I think if it does, that becomes almost mandatory
> (at least the bind mount idea is pretty much over here), because drawing
> value for /proc/stat becomes quite complex.
> The cpuacct cgroup can provide user, sys, etc values. But we also have:

>

If virtualized /proc/stat works, I don't think 'account only' cgroup is necessary. It can be obsolete.

Thanks,
-Kame
