
Subject: Re: How to draw values for /proc/stat
Posted by [Glauber Costa](#) on Mon, 12 Dec 2011 09:35:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 12/12/2011 01:33 PM, Peter Zijlstra wrote:

> On Fri, 2011-12-09 at 12:55 -0200, Glauber Costa wrote:

>> On 12/09/2011 12:03 PM, Peter Zijlstra wrote:

>

>>> Namespaces seem to be about limiting visibility, cgroups about

>>> controlling resources.

>>>

>>> The two things are hopelessly disjoint atm, but I believe someone was

>>> looking at this mess.

>>

>> I did take a look at this (if anyone else was, I'd like to know so we
>> can share some ideas), but I am not convinced we should do anything to
>> join them anymore. We virtualization people are to the best of my
>> knowledge the only ones doing namespaces. Cgroups, OTOH, got a lot bigger.

>>

>> What I am mostly concerned about now, is how consistent they will be.

>> /proc always being always global indeed does make sense, but my question

>> still stands: if you live in a resource-controlled world, why should you

>> even see resources you will never own ?

>

> Since without namespaces you can still see the rest of the world. So it
> makes sense to me to still see all resources too.

>

> Also, proportional controllers might not see a consistent slice of the
> resource, making the stats rather awkward to interpret.

>

> Furthermore, not everybody might care about these statistics at all and
> I know pj1 objected to being subjected to the extra accounting (pj1 do
> speak up etc.).

>

>> If it is not co-mounted, we draw the global value. If you don't mount
>> it, I someone does not mount it, I can assure you he doesn't care about
>> it. We for sure will.

>

> Anyway, looking at the rest of the emails in this thread the current
> proposal is a cgroup mount option that indicates if you want these
> per-cgroup stats or not, right?

Well, it is something in this direction. I don't think it's entirely
clear what exactly it will look like, but it seems we're making progress.
