

---

Subject: Re: [PATCH 4/4] cpuacct.stat: re-use scheduler statistics for the root cgroup

Posted by [KAMEZAWA Hiroyuki](#) on Fri, 25 Nov 2011 02:37:20 GMT

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

---

On Thu, 24 Nov 2011 23:33:26 -0200

Glauber Costa <[glommer@parallels.com](mailto:glommer@parallels.com)> wrote:

> Right now, after we collect tick statistics for user and system and store them  
> in a well known location, we keep the same statistics again for cpuacct.  
> Since cpuacct is hierarchical, the numbers for the root cgroup should be  
> absolutely equal to the system-wide numbers.

>

> So it would be better to just use it: this patch changes cpuacct accounting  
> in a way that the cpustat statistics are kept in a struct kernel\_cpustat percpu  
> array. In the root cgroup case, we just point it to the main array. The rest of  
> the hierarchy walk can be totally disabled later with a static branch - but I am  
> not doing it here.

>

> Signed-off-by: Glauber Costa <[glommer@parallels.com](mailto:glommer@parallels.com)>

> CC: Paul Tuner <[pjt@google.com](mailto:pjt@google.com)>

> CC: Peter Zijlstra <[a.p.zijlstra@chello.nl](mailto:a.p.zijlstra@chello.nl)>

seems reasonable to me.

Reviewed-by: KAMEZAWA Hiroyuki <[kamezawa.hiroyu@jp.fujitsu.com](mailto:kamezawa.hiroyu@jp.fujitsu.com)>

---