Subject: Re: [PATCH v2 2/5] account guest time per-cgroup as well.
Posted by Glauber Costa on Tue, 29 May 2012 10:34:55 GMT
View Forum Message <> Reply to Message

On 05/28/2012 05:26 PM, Glauber Costa wrote:
>
> I plan to measure this today, but an extra branch cost for the common
> case of a task in the root cgroup + O(depth) for tasks inside cgroups
> may be acceptable, given the simplification it brings.
>
> Let me know what you think.

Numbers:

benchmark is hackbench -pipe 1 thread 4000

task sitting in the root cgroup
==========
Without this patch:
4.857700 (0.69 %)
With this patch:
4.828733 (0.55 %)

Difference between them: 0.59 %, very close to the standard deviation,
no real difference.

task sitting in a 3-level cgroup
==========
Without this patch
5.120867 (1.60 %)
With this patch
5.126267 (1.30 %)

Difference between them: 0.10 %, way within the standard deviation

Task sitting in a level-30 cgroup: (total crazy)
==========
Without this patch:
8.829385 (2.63 %)
With this patch:
9.347846 (2.25 %)

Difference is about 5.8 %, way out of the standard deviation, so it is
really worse. But who uses 30-level hierarchy?

I believe depth-3 is close to a practical worst case, for the very
majority of the workloads out there. Therefore I don't see the loop here

as a big problem. It does degrade, but not in any use case that matters.