
Subject: Re: [PATCH 2/2] CFS CGroup: Report usage
Posted by [Srivatsa Vaddagiri](#) on Tue, 23 Oct 2007 17:30:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Oct 23, 2007 at 09:41:49AM -0700, Paul Menage wrote:

> > > Adds a cpu.usage file to the CFS cgroup that reports CPU usage in
> > > milliseconds for that cgroup's tasks
> >
> > It would be nice to split this into user and sys time at some point.
>
> Sounds reasonable - but does CFS track this?

No, not for a group. We could extend `account_user_time()` and `account_systime_time()` in this regard.

> > We have also received request to provide idle time for a
> > container/cgroup.
>
> The semantics of "idle time" for a cgroup on a shared system seem a
> bit fuzzy. How would you define it?

I think the percentage of time when it didn't have any runnable task in its runqueues.

> Suppose you have two cgroups that would each want to use, say, 55% of
> a CPU - technically they should each be regarded as having 45% idle
> time, but if they run on a the same CPU the chances are that they will
> both always have some processes on their runqueue due to contention
> with the other group. So how would you measure the difference between
> this and a cgroup that really is trying to use 100%?

Good point. I think we need to subtract out the time it was waiting on runqueue when calculating idle time.

|----- -----ZZZZZZZZZZZZ.....-----|
t0 t1 t2 t3 t4 t5 t6

---- -> Running time

.... -> Waiting time (to get on the cpu)

zzzz -> Sleeping time (when it didnt want to run because of
lack of tasks)

So, in this case,

idle time = $(t4 - t3) / [(t6 - t1) - (t2 - t1) - (t5 - t4)]$

?

This idle time will be a per-cpu stat for every cgroup and needs to be consolidated across cpus into a single idle-stat number, just like how top does it.

--

Regards,
vatsa

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
