
Subject: Re: [RFC][mm] [1/2] Simple stats for cpu resource controller
Posted by [Dhaval Giani](#) on Sat, 05 Apr 2008 21:01:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, Apr 06, 2008 at 02:01:52AM +0530, Balaji Rao wrote:

> On Sunday 06 April 2008 01:10:41 am Dhaval Giani wrote:
> > > +};
> > > +
> > > +struct cpu_cgroup_stat_cpu {
> > > + s64 count[CPU_CGROUP_STAT_NSTATS];
> >
> > u64? time does not go negative :)
> Right. But these stats are not only going to measure time. We need the same
> variables for measuring other stats as well. I'm not sure if we would
> encounter scheduler stats that would count negative.
>
> Balbir, what do you say ?

I would prefer to keep the stats logically separate. So something like
struct cpu_cgroup_stat_cpu {
 u64 time[];
 s64 some_other_stat;
}
and so on. (I am not sure, is there some advantage gained by using
structs?) Makes the code more maintainable imho.

>
> > count also is not very clear? Can you give a more descriptive name?
> >
> ok. How does 'value' look ?
>
> <snip>
>
> > > +static s64 cpu_cgroup_read_stat(struct cpu_cgroup_stat *stat,
> > > + enum cpu_cgroup_stat_index idx)
> > > +{
> > > + int cpu;
> > > + s64 ret = 0;
> > > + unsigned long flags;
> >
> > > +
> > > + local_irq_save(flags);
> >
> > I am just wondering. Is local_irq_save() enough?
> >
> Hmm.. You are right. This does not prevent concurrent updates on other CPUs
> from crossing a 32bit boundary. Am not sure how to do this in a safe way. I
> can only think of using atomic64_t now..

>

I am going to answer that one when I am awake :-)

--

regards,
Dhaval

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