
Subject: Re: [RFC][-mm] Simple stats for cpu resource controller v3
Posted by [akpm](#) on Fri, 02 May 2008 19:53:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sat, 3 May 2008 01:10:28 +0530
Balaji Rao <balajirao@gmail.com> wrote:

> On Friday 02 May 2008 02:30:26 am Andrew Morton wrote:
> <snip>
>
> Hi Andrew,
>
> Thank you for the review.
> >
> > Did you consider using include/linux/percpu_counter.h?
> >
> > If so, what was wrong with it?
> >
> > Because it would be much better to fix per-cpu counters than to invent new
> > stuff.
> No, I hadn't consider using the percpu_counters infrastructure. But today when
> I tried using it, I got an early exception. I guess its because I tried
> calling percpu_counter_init from within sched_init, which I perhaps shouldn't
> do, because percpu_counter_init expects cpu hotplug code to be initialized by
> then. Right ? Correct me if I'm wrong.

I don't see any reason why we cannot run percpu_counter_init() prior to
running percpu_counter_startup(). And it is desirable that we be able to
start using the percpu-counters quite early.

Can you debug it a bit please? It's probably some silly little thing,
perhaps fixable by calling percpu_counter_startup() earlier.

> How about we start collecting statistics at a later stage i.e, after
> percpu_counter becomes usable ?

It would be better to make the core infrastructure more robust, rather
than working around problems it might have.

It's rather nice that percpu_counters internally take care of
cpu-hotplugging, and use cpu_online_map. I was amazed at how easily that
was added. I still expect it to break somehow..

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