Subject: Re: [ckrm-tech] [PATCH 4/6] containers: Simple CPU accounting container subsystem
Posted by Paul Menage on Fri, 12 Jan 2007 00:33:48 GMT
View Forum Message <> Reply to Message

On 1/10/07, Balbir Singh <balbir@in.ibm.com> wrote:
>
> I have run into a problem running this patch on a powerpc box. Basically,
> the machine panics as soon as I mount the container filesystem with

This is a multi-processor system?

My guess is that it's a race in the subsystem API that I've been
meaning to deal with for some time - basically I've been using
(<foo>_subsys.subsys_id != -1) to indicate that <foo> is ready for
use, but there's a brief window during subsystem registration where
that's not actually true.

I'll add an "active" field in the container_subsys structure, which
isn't set until registration is completed, and subsystems should use
that instead. container_register_subsys() will set it just prior to
releasing callback_mutex, and cpu_acct.c (and other subsystems) will
check <foo>_subsys.active rather than (<foo>_subsys.subsys_id != -1)

> I am trying to figure out the reason for the panic and trying to find
> a fix. Since the introduction of whole hierarchy system, the debugging
> has gotten a bit harder and taking longer, hence I was wondering if you
> had any clues about the problem
>

Yes, the multi-hierarchy support does make the whole code a little
more complex - but people presented reasonable scenarios where a
single container tree for all resource controllers just wasn't
flexible enough.

Paul