
Subject: Re: [ckrm-tech] [PATCH 1/1] Fix a panic while mouting containers on powerpc and some other small cle

Posted by [Balbir Singh](#) on Mon, 15 Jan 2007 10:10:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On 1/15/07, Balbir Singh <balbir@in.ibm.com> wrote:

>> While writing/extending the cpuacct container, I found it useful to

>> know if the container resource group we are controlling is really mounted.

>> Controllers can try and avoid doing work when not mounted and start

>> when the subsystem is mounted. Also, without these callbacks, one has no

>> definite way of checking if the top_container is dummy or for real.

>>

>

> That's somewhat intentional - my aim was that the controllers

> shouldn't really care whether they're connected to the default

> hierarchy or have been bound to some mounted hierarchy. Having said

> that, they can determine it by checking <foo>_subsys.hierarchy if they

> really want to. If that's 0 then they're in the default hierarchy (and

> can assume that all tasks are in one top-level container).

>

That makes sense, the only additional thing required is to know when the subsystem really got mounted (we cannot keep polling hierarchy for it:-))

> Paul

--

Balbir Singh,
Linux Technology Center,
IBM Software Labs
