
Subject: Re: [PATCH 2/2] CFS CGroup: Report usage
Posted by [Paul Menage](#) on Tue, 23 Oct 2007 07:21:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 10/22/07, Paul Menage <menage@google.com> wrote:

>
> Using cgroup_mutex is certainly possible for now, although more
> heavy-weight than I'd like long term. Using css_get isn't the right
> approach, I think - we shouldn't be able to cause an rmdir to fail due
> to a concurrent read.
>

OK, the obvious solution is to use the same approach for subsystem state objects as we do for the struct cgroup itself - move the calls to the subsystem destroy methods to cgroup_diput. A control file dentry will keep alive the parent dir's dentry, which will keep alive the cgroup and (with this change) the subsystem state objects too.

The only potential drawback that I can see is that an open fd on a cgroup directory or a control file will keep more memory alive than it would have done previously.

Paul

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