
Subject: Re: [RFC][PATCH 3/5] Container Freezer: Implement freezer cgroup subsystem

Posted by [Matt Helsley](#) on Wed, 30 Apr 2008 22:30:49 GMT

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

On Wed, 2008-04-30 at 14:28 -0700, Matt Helsley wrote:

> On Wed, 2008-04-23 at 23:47 -0700, Matt Helsley wrote:

>

> <snip>

> > + cgroup_lock();

> > +

> > + if (cgroup_is_removed(cgroup)) {

> > + retval = -ENODEV;

> > + goto unlock;

> > + }

>

> I think this was copy/paste'd from cgroup_common_file_write() which
> modifies the cgroup hierarchy. However this function does not modify the
> cgroup hierarchy and we're not getting the cgroup from the task. So I
> don't think cgroup_lock()/unlock() are needed here. Paul, do you agree?

Oops! I didn't track the call chain correctly. Sorry for the noise...

Cheers,

-Matt Helsley

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
