Subject: Re: [RFC/PATCH 1/8]: CGroup Files: Add locking mode to cgroups control files

Posted by Li Zefan on Wed, 14 May 2008 01:30:40 GMT

View Forum Message <> Reply to Message

```
>> > @ @ -1518,16 +1580,21 @ @ static ssize_t cgroup_file_read(struct f
         struct cftype *cft = __d_cft(file->f_dentry);
         struct cgroup *cgrp = __d_cgrp(file->f_dentry->d_parent);
>> >
         if (!cft || cgroup is removed(cgrp))
>> > -
>> > +
         if (cgroup_is_removed(cgrp))
              return -ENODEV;
>> >
>>
>> This check seems redundant now.
>>
> It's not needed for safety, but it doesn't seem to hurt to check
> cgroup is removed() prior to doing any copying, since we'll fail after
> copying anyway if cgroup_is_removed() returns true (once we've taken
> any relevant locks).
It's a bit odd to me that cgroup_is_removed() is removed in cgroup_file_write(),
but it's ressered here.
Containers mailing list
```

Containers@lists.linux-foundation.org

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