
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [Srivatsa Vaddagiri](#) on Thu, 05 Apr 2007 06:39:06 GMT

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

On Thu, Apr 05, 2007 at 12:09:50PM +0530, Srivatsa Vaddagiri wrote:

> IF (note that /if/ again) the count returned is zero, then my point was
> we can destroy nsproxy behind foo and also B1, not worrying about a
> 'struct file' still pointing to B1. This stems from the fact that you
> cannot have a task's file->f_bc pointing to B1 w/o the task itself
> pointing to B1 also (task->nsproxy->ctrl_data[BC_ID] == B1). I also
> assume f_bc will get migrated with its owner task across beancounters
> (which seems reasonable to me atleast from 'struct file' context).
>
> If there was indeed a file object still pointing to B1, then that can
> only be true if rcfs_task_count() returns non-zero value. Correct?

arghh .. s/B1/A1 pls

--

Regards,
vatsa
