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

Posted by [Srivatsa Vaddagiri](#) on Thu, 05 Apr 2007 14:13:21 GMT

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

On Thu, Apr 05, 2007 at 06:13:25PM +0530, Srivatsa Vaddagiri wrote:

> Lets go back to the f_bc example here for a moment. Lets say T1 was in C1 and
> opened file f1. f1->f_bc points to C1->beancounter.

>

> T1 moves from C1 -> C2, but f1 is not migrated.

> C1->beancounter.count stays at 1 (to account for f1->f_bc).

Actually C1->beancounter.count should be at 2 (C1->beancounter and f1->f_bc are pointing to it).

> File f1 is closed. C1->beancounter.count becomes zero.

C1->beancounter.count should go to 1 ..

> Now user issues rmdir C1. If rmdir finds (after taking manage_mutex that
> is)

>

> - zero tasks in C1

> - zero refcount in C1->beancounter

s/zero refcount in C1->beancounter/exactly 1 refcount in C1->beancounter

> why is it not safe to assume that C1->beancounter.count will continue to
> stay zero?

s/zero/at one

> Basically I am struggling to answer "How can a zero refcount (beancounter)

> object go non-zero when zero tasks are attached to it" ..

s/zero/one and s/non-zero/>1

Essentially bc_subsys->can_attach(struct bean_counter *b) can return -EBUSY if (atomic_read(&b->count) > 1) ..

--

Regards,
vatsa

Containers mailing list

Containers@lists.linux-foundation.org

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