
Subject: Re: [ckrm-tech] [PATCH 2/7] UBC: core (structures, API)

Posted by [Matt Helsley](#) on Fri, 18 Aug 2006 08:26:03 GMT

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

On Fri, 2006-08-18 at 11:35 +0400, Andrey Savochkin wrote:

> On Thu, Aug 17, 2006 at 10:31:37PM -0700, Andrew Morton wrote:

> > On Thu, 17 Aug 2006 15:53:40 +0400

> > Kirill Korotaev <dev@sw.ru> wrote:

> >

> > > >+struct user_beancounter

> > > >+{

> > > >+ atomic_t ub_refcount;

> > > >+ spinlock_t ub_lock;

> > > >+ uid_t ub_uid;

> > >

> > >

> > > > Why uid? Will it be possible to club processes belonging to different
> > > > users to same bean counter.

> > > oh, its a misname. Should be ub_id. it is ID of user_beancounter

> > > and has nothing to do with user id.

> >

> > But it uses a uid_t. That's more than a misnaming?

>

> It used to be uid-related in ancient times when the notion of container
> hadn't formed up.

> "user" part of user_beancounter name has the same origin :)

Is it similarly irrelevant now? If so perhaps a big rename could be used
to make the names clearer (s/user_/, s/ub_/bc_, ...).

<snip>

Cheers,

-Matt Helsley
