Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by Rohit Seth on Thu, 17 Aug 2006 17:28:20 GMT
View Forum Message <> Reply to Message

On Thu, 2006-08-17 at 11:19 -0400, Rik van Riel wrote:
> Dave Hansen wrote:
>
> > My main thought is that _everybody_ is going to have to live with the
> > entry in the 'struct page'. Distros ship one kernel for everybody, and
> > the cost will be paid by those not even using any kind of resource
> > control or containers.
>
> Every userspace or page cache page will be in an object
> though. Could we do the pointer on a per object (mapping,
> anon vma, ...) basis?
>
> Kernel pages are not using all of their struct page entries,
> so we could overload a field.
>

Bingo. Even though it has the word "overload".

> It all depends on how much we really care about not growing
> struct page :)
>

Besides, if we have the container pointer based on address_space (for example) then it will also allow file based tracking...

I think page based container pointer makes more sense when you have container as the central part of page lists (in place of nodes) deciding which list the free page is going to come from, and when freed which list it is going to go back to.

-rohit