

---

Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)  
Posted by [Rohit Seth](#) on Fri, 18 Aug 2006 17:06:16 GMT

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

---

On Fri, 2006-08-18 at 12:29 +0400, Kirill Korotaev wrote:

> 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?

> in this case no memory fractions accounting is possible :/

> please, note, this field added by this patchset is in union

> and used by user pages accounting as well.

>

> > Kernel pages are not using all of their struct page entries,

> > so we could overload a field.

> yeah, we can. probably mapping.

> but as I said we use the same pointer for user pages accounting as well.

>

> > It all depends on how much we really care about not growing

> > struct page :)

> so what is your opinion?

> Kernel compiled w/o UBC do not introduce additional pointer.

As Andi pointed out earlier that slab and network codes are going to use the mapping field (and you pointed out that some of this is allocated out of context), so seems like for kernel accounting we will need another field in page structure.

-rohit

---