
Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by [Andi Kleen](#) on Thu, 17 Aug 2006 17:35:36 GMT

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

On Thursday 17 August 2006 19:28, Rohit Seth wrote:

> 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".

You would need to be careful. Both the rewritten slab and the new
tree network allocator use struct page fields already. There might
be conflicts already.

-Andi (who still doesn't see what's so bad about a separate table)
