
Subject: Re: [RFC][PATCH 5/7] UBC: kernel memory accounting (core)

Posted by [Dave Hansen](#) on Wed, 16 Aug 2006 18:47:09 GMT

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

On Wed, 2006-08-16 at 19:40 +0400, Kirill Korotaev wrote:

```
> --- ./include/linux/mm.h.kmemcore      2006-08-16 19:10:38.000000000
> +0400
> +++ ./include/linux/mm.h      2006-08-16 19:10:51.000000000 +0400
> @@ -274,8 +274,14 @@ struct page {
>     unsigned int gfp_mask;
>     unsigned long trace[8];
> #endif
> +#ifdef CONFIG_USER_RESOURCE
> +     union {
> +         struct user_beancounter *page_ub;
> +     } bc;
> +#endif
> };
```

Is everybody OK with adding this accounting to the 'struct page'? Is there any kind of noticeable performance penalty for this? I thought that we had this aligned pretty well on cacheline boundaries.

How many things actually use this? Can we have the slab ubcs without the struct page pointer?

-- Dave
