Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by dev on Fri, 18 Aug 2006 08:12:57 GMT
View Forum Message <> Reply to Message

Dave Hansen wrote:
> On Thu, 2006-08-17 at 17:31 +0400, Kirill Korotaev wrote:
>
>>>How many things actually use this? Can we have the slab ubcs
>>
>>without
>>
>>>the struct page pointer?
>>
>>slab doesn't use this pointer on the page.
>>It is used for pages allocated by buddy
>>alocator implicitly (e.g. LDT pages, page tables, ...).
>
>
> Hmmm. There aren't _that_ many of those cases, right? Are there any
> that absolutely need raw access to the buddy allocator? I'm pretty sure
> that pagetables can be moved over to a slab, as long as we bump up the
> alignment.
LDT takes from 1 to 16 pages. and is allocated by vmalloc.
do you propose to replace it with slab which can fail due to memory
fragmentation?

the same applies to fdset, fdarray, ipc ids and iptables entries.

> It does seem a wee bit silly to have the pointer in _all_ of the struct
> pages, even the ones for which we will never do any accounting (and even
> on kernels that never need it). But, a hashing scheme sounds like a
> fine idea.
It seems a silly for you since 2nd patchset accounting user pages
is not here yet. As you can see we added a union into page,
which is shared between kernel memory and user memory accounting.

THERE IS NOT USER ACCOUNTING HERE YET GUYS! :) THIS FIELD WILL BE USED!!!

Thanks,
Kirill