Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by dev on Mon, 21 Aug 2006 08:56:08 GMT
View Forum Message <> Reply to Message

Dave Hansen wrote:
> On Fri, 2006-08-18 at 12:12 +0400, Kirill Korotaev wrote:
>
>>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?
>
>
> Nope. ;)
so what is your proposal then? Sorry, probably missed it due to lots of emails :)

>>the same applies to fdset, fdarray, ipc ids and iptables entries.
>
>
> The vmalloc area, along with all of those other structures _have_ other
> data structures. Now, it will take a wee bit more patching to directly
> tag those thing with explicit container pointers (or accounting
> references), but I would much prefer that, especially for the things
> that are larger than a page.
do you mean that you prefer adding a explicit pointer to the structures
itself?

> I worry that this approach was used instead of patching all of the
> individual subsystems because this was easier to maintain as an
> out-of-tree patch, and it isn't necessarily the best approach.
:) if we were to optimize for patch size then we would select vserver
approach and be happy...

Dave, we used to add UBC pointers on each data structure and then do
a separate accounting in the places where objects are allocated.
We spent a lot of time and investigation on how to make it better,
because it was leading to often accounting errors, wrong error paths etc.
The approach provided in this patchset proved to be much more efficient
and more error prone. And it is much much more elegant!

Thanks,
Kirill