Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by dev on Mon, 21 Aug 2006 09:41:56 GMT
View Forum Message <> Reply to Message

Alan Cox wrote:
> Ar Gwe, 2006-08-18 am 12:32 -0700, ysgrifennodd Dave Hansen:
>
>>>It ought to be cheap. Given each set of page structs is an array its a
>>>simple subtract and divide (or with care and people try to pack them
>>>nicely for cache lines - shift) to get to the parallel accounting array.
>>
>>I wish page structs were just a simple array. ;)
>
>
> Note I very carefully said "each set of"
>
>
>>It will just be a bit more code, but we'll need this for the two other
>>memory models: sparsemem and discontigmem. For discontig, we'll just
>>need pointers in the pg_data_ts and, for sparsemem, we'll likely need
>>another pointer in the 'struct mem_section'.
>
>
> Actually I don't believe this is true in either case. Change the code
> which allocates the page arrays to allocate (+ sizeof(void *) *
> pages_in_array on the end of each array when using UBC. The rest then
> seems to come out naturally.
I only doubt what gain we will have in this situation.
boot-time selectable vs. CONFIG-selectable?

Kirill