
Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by [Alan Cox](#) on Fri, 18 Aug 2006 14:39:01 GMT

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

Ar Gwe, 2006-08-18 am 12:52 +0400, ysgrifennodd Kirill Korotaev:

> > hmm, not sure why it is simpler.

> because introducing additonal lookups/hashes etc. is harder and

> adds another source for possible mistakes.

> we can always optimize it out if people insist (by cost of slower accounting).

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.

Alan
