
Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by [dev](#) on Fri, 18 Aug 2006 09:29:26 GMT

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

Dave Hansen wrote:

> On Thu, 2006-08-17 at 17:27 +0400, Kirill Korotaev wrote:

>

>> charged kernel objects can't be _reclaimed_. how do you propose

>> to reclaim tasks page tables or files or task struct or vma or etc.?

>

>

> Do you have any statistics on which of these objects are the most

> troublesome? If it _is_ pagetables, for instance, it is quite

> conceivable that we could reclaim them.

they all are troublesome :/

user can create lots of vmas, w/o page tables.

lots of fdsets, ipcids.

These are not reclaimable.

Also consider the following scenario with reclaimable page tables.

e.g. user hit kmemsize limit due to fat page tables.

kernel reclaims some of the page tables and frees user kernel memory.

after that user creates some unclaimable objects like fdsets or ipcids

and then accesses memory with reclaimed page tables.

Sooner or later we kill user with SIGSEGV from page fault due to no memory. This is worse than returning ENOMEM from poll() or mmap() where user allocates kernel objects.

> This one probably deserves a big, fat comment, though. ;)

tell me where to write it and what? :)

Thanks,

Kirill
