

---

Subject: Re: [PATCH 0/4] Kernel memory accounting container (v3)

Posted by [Pavel Emelianov](#) on Tue, 18 Sep 2007 06:25:48 GMT

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

---

Christoph Lameter wrote:

> On Mon, 17 Sep 2007, Pavel Emelyanov wrote:

>

>> As I have already told kmalloc caches cannot be accounted easily

>> so turning the accounting on for them will fail with -EINVAL.

>> Turning the accounting off is possible only if the cache has

>> no objects. This is done so because turning accounting off

>> implies unaccounting of all the objects in the cache, but due

>> to full-pages in slub are not stored in any lists (usually)

>> this is impossible to do so, however I'm open for discussion

>> of how to make this work.

>

> Where can I find more information why it would not be possible to

> account kmalloc caches?

It is possible, but the problem is that we want to account only allocations explicitly caused by the user request. E.g. the vfstmount name is kmalloc-ed by explicit user request, but such things as buffer heads are not.

So we have to explicitly specify which calls to kmalloc() do we want to be accounted and which we do not by passing additional flag (GFP\_ACCT?) to kmalloc, but this is another patch.

Yet again - as soon as we agree on the implementation of kmem caches accounting, I will proceed with working on kmalloc, vmalloc and buddy allocator.

>

Thanks,  
Pavel

---