
Subject: Re: [PATCH v5 07/14] mm: Allocate kernel pages to the right memcg
Posted by [akpm](#) on Wed, 17 Oct 2012 22:12:21 GMT

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

On Tue, 16 Oct 2012 14:16:44 +0400

Glauber Costa <glommer@parallels.com> wrote:

> When a process tries to allocate a page with the `__GFP_KMEMCG` flag, the
> page allocator will call the corresponding memcg functions to validate
> the allocation. Tasks in the root memcg can always proceed.
>
> To avoid adding markers to the page - and a kmem flag that would
> necessarily follow, as much as doing `page_cgroup` lookups for no reason,
> whoever is marking its allocations with `__GFP_KMEMCG` flag is responsible
> for telling the page allocator that this is such an allocation at
> `free_pages()` time.

Well, why? Was that the correct decision?

> This is done by the invocation of
> `__free_accounted_pages()` and `free_accounted_pages()`.

These are very general-sounding names. I'd expect the identifiers to contain "memcg" and/or "kmem", to identify what's going on.
