
Subject: Re: [PATCH 07/11] mm: Allocate kernel pages to the right memcg
Posted by [Glauber Costa](#) on Mon, 25 Jun 2012 22:27:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 06/25/2012 10:07 PM, Tejun Heo wrote:

> On Mon, Jun 25, 2012 at 06:15:24PM +0400, Glauber Costa 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. This is done by the invocation of
>> `__free_accounted_pages()` and `free_accounted_pages()`.

>

> Shouldn't we be documenting that in the code somewhere, preferably in
> the function comments?

>

I can certainly do that.
