
Subject: Re: [PATCH v5 07/14] mm: Allocate kernel pages to the right memcg
Posted by [David Rientjes](#) on Wed, 17 Oct 2012 22:43:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 16 Oct 2012, 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()`.
>
> [v2: inverted test order to avoid a `memcg_get` leak,
> `free_accounted_pages` simplification]
> [v4: test for `TIF_MEMDIE` at `newpage_charge`]
>
> Signed-off-by: Glauber Costa <glommer@parallels.com>
> Acked-by: Michal Hocko <mhocko@suse.cz>
> Acked-by: Mel Gorman <mgorman@suse.de>
> Acked-by: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
> CC: Christoph Lameter <cl@linux.com>
> CC: Pekka Enberg <penberg@cs.helsinki.fi>
> CC: Johannes Weiner <hannes@cmpxchg.org>
> CC: Suleiman Souhlal <suleiman@google.com>
> CC: Tejun Heo <tj@kernel.org>

Acked-by: David Rientjes <rientjes@google.com>
