
Subject: Re: [PATCH v3 05/28] memcg: Reclaim when more than one page needed.

Posted by [Christoph Lameter](#) on Tue, 29 May 2012 14:19:02 GMT

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

On Fri, 25 May 2012, Glauber Costa wrote:

> From: Suleiman Souhlal <ssouhlal@FreeBSD.org>

>

> mem_cgroup_do_charge() was written before slab accounting, and expects
> three cases: being called for 1 page, being called for a stock of 32 pages,
> or being called for a hugepage. If we call for 2 pages (and several slabs
> used in process creation are such, at least with the debug options I had),
> it assumed it's being called for stock and just retried without reclaiming.

Slab pages are allocated up to order 3 (PAGE_ALLOC_COSTLY_ORDER). That is 8 pages.

> * unlikely to succeed so close to the limit, and we fall back

> * to regular pages anyway in case of failure.

> */

> - if (nr_pages == 1 && ret)

> + if (nr_pages <= (PAGE_SIZE << PAGE_ALLOC_COSTLY_ORDER) && ret) {

> + cond_resched();

> return CHARGE_RETRY;

> + }

>

> /*

> * At task move, charge accounts can be doubly counted. So, it's

Ok. That looks correct.
