
Subject: Re: [PATCH v4 02/14] memcg: Reclaim when more than one page needed.

Posted by [KAMEZAWA Hiroyuki](#) on Tue, 16 Oct 2012 03:22:34 GMT

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

(2012/10/08 19:06), Glauber Costa wrote:

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

>

> mem_cgroup_do_charge() was written before kmem 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 or 3 pages (and

> both the stack 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.

>

> Fix that by passing down a minsize argument in addition to the csize.

>

> And what to do about that (csize == PAGE_SIZE && ret) retry? If it's

> needed at all (and presumably is since it's there, perhaps to handle

> races), then it should be extended to more than PAGE_SIZE, yet how far?

> And should there be a retry count limit, of what? For now retry up to

> COSTLY_ORDER (as page_alloc.c does) and make sure not to do it if

> __GFP_NORETRY.

>

> [v4: fixed nr pages calculation pointed out by Christoph Lameter]

>

> Signed-off-by: Suleiman Souhlal <suleiman@google.com>

> Signed-off-by: Glauber Costa <glommer@parallels.com>

> Reviewed-by: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>

> Aacked-by: Michal Hocko <mhocko@suse.cz>

> Aacked-by: Johannes Weiner <hannes@cmpxchg.org>

Aacked-by: KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
