
Subject: Re: [PATCH v3 01/13] memcg: Make it possible to use the stock for more than one page.

Posted by [Johannes Weiner](#) on Mon, 01 Oct 2012 18:48:12 GMT

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

On Tue, Sep 18, 2012 at 06:03:58PM +0400, Glauber Costa wrote:

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

>

> We currently have a percpu stock cache scheme that charges one page at a

> time from memcg->res, the user counter. When the kernel memory

> controller comes into play, we'll need to charge more than that.

>

> This is because kernel memory allocations will also draw from the user

> counter, and can be bigger than a single page, as it is the case with

> the stack (usually 2 pages) or some higher order slabs.

>

> [glommer@parallels.com: added a changelog]

>

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

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

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

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

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

Independent of how the per-subtree enable-through-setting-limit discussion pans out, we're going to need the charge cache, so:

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