
Subject: Re: [PATCH v4 17/25] skip memcg kmem allocations in specified code regions

Posted by [KAMEZAWA Hiroyuki](#) on Mon, 18 Jun 2012 12:19:49 GMT

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

(2012/06/18 19:28), Glauber Costa wrote:

> This patch creates a mechanism that skip memcg allocations during
> certain pieces of our core code. It basically works in the same way
> as preempt_disable()/preempt_enable(): By marking a region under
> which all allocations will be accounted to the root memcg.

>

> We need this to prevent races in early cache creation, when we
> allocate data using caches that are not necessarily created already.

>

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

> CC: Christoph Lameter<cl@linux.com>

> CC: Pekka Enberg<penberg@cs.helsinki.fi>

> CC: Michal Hocko<mhocko@suse.cz>

> CC: Kamezawa Hiroyuki<kamezawa.hiroyu@jp.fujitsu.com>

> CC: Johannes Weiner<hannes@cmpxchg.org>

> CC: Suleiman Souhlal<suleiman@google.com>

I'm ok with this approach.

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