
Subject: Re: [PATCH v5 11/14] memcg: allow a memcg with kmem charges to be destroyed.

Posted by [akpm](#) on Wed, 17 Oct 2012 22:12:35 GMT

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

On Tue, 16 Oct 2012 14:16:48 +0400

Glauber Costa <glommer@parallels.com> wrote:

> Because the ultimate goal of the kmem tracking in memcg is to track slab
> pages as well,

It is? For a major patchset such as this, it's pretty important to discuss such long-term plans in the top-level discussion. Covering things such as expected complexity, expected performance hit, how these plans affected the current implementation, etc.

The main reason for this is that if the future plans appear to be of doubtful feasibility and the current implementation isn't sufficiently useful without the future stuff, we shouldn't merge the current implementation. It's a big issue!

> we can't guarantee that we'll always be able to point a
> page to a particular process, and migrate the charges along with it -
> since in the common case, a page will contain data belonging to multiple
> processes.
>
> Because of that, when we destroy a memcg, we only make sure the
> destruction will succeed by discounting the kmem charges from the user
> charges when we try to empty the cgroup.
>
> ...
>
