

---

Subject: Re: [PATCH v2 09/11] memcg: propagate kmem limiting information to children

Posted by [Michal Hocko](#) on Tue, 21 Aug 2012 07:54:30 GMT

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

---

On Fri 17-08-12 14:36:00, Glauber Costa wrote:

> On 08/17/2012 02:35 PM, Michal Hocko wrote:

> >> > But I never said that can't happen. I said (ok, I meant) the static

> >> > branches can't be disabled.

> > Ok, then I misunderstood that because the comment was there even before

> > static branches were introduced and it made sense to me. This is

> > inconsistent with what we do for user accounting because even if we set

> > limit to unlimited we still account. Why should we differ here?

>

> Well, we account even without a limit for user accounting. This is a

> fundamental difference, no ?

Yes, user memory accounting is either on or off all the time (switchable at boot time).

My understanding of kmem is that the feature is off by default because it brings an overhead that is worth only special use cases. And that sounds good to me. I do not see a good reason to have runtime switch off. It makes the code more complicated for no good reason. E.g. how do you handle charges you left behind? Say you charged some pages for stack?

But maybe you have a good use case for that?

--

Michal Hocko

SUSE Labs

---