

---

Subject: Re: [PATCH v2 06/11] memcg: kmem controller infrastructure  
Posted by [Glauber Costa](#) on Wed, 15 Aug 2012 10:44:15 GMT

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

---

On 08/15/2012 01:42 PM, Glauber Costa wrote:

>> Also, as I

>> > have mentioned in the other email in this thread. Why should we reclaim

>> > just because of kernel allocation when we are not reclaiming any of it

>> > because shrink\_slab is ignored in the memcg reclaim.

>

> Don't get too distracted by the fact that shrink\_slab is ignored. It is

> temporary, and while this being ignored now leads to suboptimal

> behavior, it will 1st, only affect its users, and 2nd, not be disastrous.

>

> I see it this as more or less on par with the soft limit reclaim

> problem we had. It is not ideal, but it already provided functionality

>

Okay, I sent the e-mail before finishing it... duh

What I meant in this last sentence, is that the situation while the memcg-aware shrinkers doesn't land in the kernel is more or less the same (obviously not exactly) as with the soft reclaim work. It is an evolutionary approach that provides some functionality that is not yet perfect but already solves lots of problems for people willing to live with its temporary drawbacks.

---