
Subject: Re: [PATCH v2 00/29] kmem limitation for memcg
Posted by [Glauber Costa](#) on Fri, 11 May 2012 18:05:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 05/11/2012 02:44 PM, Glauber Costa wrote:

> Hello All,
>
> This is my new take for the memcg kmem accounting.
> At this point, I consider the series pretty mature - although of course,
> bugs are always there...
>
> As a disclaimer, however, I must say that the slub code is much more stressed
> by me, since I know it better. If you have no more objections to the concepts
> presented, the remaining edges can probably be polished in a rc cycle,
> at the maintainers discretion, of course.
>
> Otherwise, I'll be happy to address any concerns of yours.
>
> Since last submission:
>
> * memcgs can be properly removed.
> * We are not charging based on current->mm->owner instead of current
> * kmem_large allocations for slub got some fixes, specially for the free case
> * A cache that is registered can be properly removed (common module case)
> even if it spans memcg children. Slab had some code for that, now it works
> well with both
> * A new mechanism for skipping allocations is proposed (patch posted
> separately already). Now instead of having kmalloc_no_account, we mark
> a region as non-accountable for memcg.
>

Forgot to mention the ida-based index allocation, instead of keeping our
own bitmap.
