
Subject: Re: [PATCH v3 13/28] slub: create duplicate cache
Posted by [Christoph Lameter](#) on Tue, 29 May 2012 16:05:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 29 May 2012, Glauber Costa wrote:

> Accounting pages seems just crazy to me. If new allocators come in the future,
> organizing the pages in a different way, instead of patching it here and
> there, we need to totally rewrite this.

Quite to the contrary. We could either pass a `THIS_IS_A_SLAB` page flag to the page allocator call or have a special call that does the accounting and then calls the page allocator. The code could be completely in cgroups. There would be no changes to the allocators aside from setting the flag or calling the alternate page allocator functions.

> > Why do you need to increase the refcount? You made a full copy right?
>
> Yes, but I don't want this copy to go away while we have other caches around.

You copied all metadata so what is there that you would still need should the other copy go away?

> So, in the memcg internals, I used a different reference counter, to avoid
> messing with this one. I could use that, and leave the original refcnt alone.
> Would you prefer this?

The refcounter is really not the issue.

I am a bit worried about the various duplicate features here and there.
The approach is not tightened down yet.
