
Subject: Re: [PATCH v2 04/29] slub: always get the cache from its page in kfree
Posted by [Christoph Lameter](#) on Fri, 11 May 2012 18:17:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 11 May 2012, Glauber Costa wrote:

> On 05/11/2012 03:06 PM, Christoph Lameter wrote:
> > On Fri, 11 May 2012, Glauber Costa wrote:
> >
> > > Adding a VM_BUG_ON may be useful to make sure that kmem_cache_free is
> > > always passed the correct slab cache.
> > >
> > > Well, problem is , it isn't always passed the "correct" slab cache.
> > > At least not after this series, since we'll have child caches associated
> > > with
> > > the main cache.
> > >
> > > So we'll pass, for instance, kmem_cache_free(dentry_cache...), but will in
> > > fact free from the memcg copy of the dentry cache.
> >
> > Urg. But then please only do this for the MEMCG case and add a fat big
> > warning in kmem_cache_free.
>
> I can do that, of course.
> Another option if you don't oppose, is to add another field in the kmem_cache
> structure (I tried to keep them at a minimum),
> to record the parent cache we got created from.
>
> Then, it gets trivial to do the following:
>
> VM_BUG_ON(page->slab != s && page->slab != s->parent_cache);

Sounds ok but I need to catch up on what this whole memcg thing in slab
allocators should accomplish in order to say something definite.
