Subject: Re: [PATCH v3 09/16] sl[au]b: always get the cache from its page in kfree
Posted by Christoph Lameter on Wed, 19 Sep 2012 14:14:23 GMT

On Wed, 19 Sep 2012, Glauber Costa wrote:

> > This is an extremely hot path of the kernel and you are adding significant
> > processing. Check how the benchmarks are influenced by this change.
> > virt_to_cache can be a bit expensive.
> Would it be enough for you to have a separate code path for
> !CONFIG_MEMCG_KMEM?

Yes, at least add an #ifdef around this.

> I don't really see another way to do it, aside from deriving the cache
> from the object in our case. I am open to suggestions if you do.

Rethink the whole memcg approach and find some other way to do it? This
whole scheme is very intrusive and is likely to increase instability in
the VM given the explosion of lru lists, duplication of slab caches and
significantly more complex processing throughout the VM.