

---

Subject: Re: [PATCH v3 13/28] slub: create duplicate cache  
Posted by [Glauber Costa](#) on Tue, 29 May 2012 17:05:45 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 05/29/2012 08:05 PM, Christoph Lameter wrote:

> 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.

Again: for the page allocation itself, we could do it. (maybe we still  
can, keeping the rest of the approach, so as to simplify that particular  
piece of code, and reduce the churn inside the cache - I am all for it)

But that only solves the page allocation part. I would still need to  
make sure the caller is directed to the right page.

>>> 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?

Well, consistency being one, because it sounded weird to have the parent  
cache being deleted while the kids are around. You wouldn't be able to  
reach them, for once.

But one can argue that if you deleted the cache, why would you want to  
ever reach it?

I can try to remove the refcount.

>

>> 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.

Well, I still think that duplication of the structure is better - a lot less intrusive - than messing with the cache internals.

I will try to at least have the page accounting done in a consistent way. How about that?

---