## Subject: Re: [PATCH v3 15/16] memcg/sl[au]b: shrink dead caches Posted by Glauber Costa on Fri, 21 Sep 2012 08:40:04 GMT

View Forum Message <> Reply to Message

```
On 09/21/2012 08:48 AM, JoonSoo Kim wrote:
> Hi Glauber.
>
Hi
> 2012/9/18 Glauber Costa <glommer@parallels.com>:
>> diff --git a/mm/slub.c b/mm/slub.c
>> index 0b68d15..9d79216 100644
>> --- a/mm/slub.c
>> +++ b/mm/slub.c
>> @ @ -2602,6 +2602,7 @ @ redo:
>>
       } else
               _slab__free(s, page, x, addr);
>>
>>
>> +
        kmem cache verify dead(s);
>> }
> As far as u know, I am not a expert and don't know anything about memcg.
> IMHO, this implementation may hurt system performance in some case.
> In case of memcg is destoried, remained kmem_cache is marked "dead".
> After it is marked.
> every free operation to this "dead" kmem_cache call
> kmem cache verify dead() and finally call kmem cache shrink().
```

As long as it is restricted to that cache, this is a non issue. dead caches are exactly what they name imply: dead.

Means that we actively want them to go away, and just don't kill them right away because they have some inflight objects - which we expect not to be too much.

- > kmem\_cache\_shrink() do invoking kmalloc and flush\_all() and taking a
- > lock for online node and invoking kfree.
- > Especially, flush\_all() may hurt performance largely, because it call
- > has\_cpu\_slab() against all the cpus.

Again, this is all right, but being a dead cache, it shouldn't be on any hot path.

>

- > And, I found one case that destroying memcg's kmem\_cache don't works properly.
- > If we destroy memcg after all object is freed, current implementation
- > doesn't destroy kmem cache.

- > kmem\_cache\_destroy\_work\_func() check "cachep->memcg\_params.nr\_pages == 0",
- > but in this case, it return false, because kmem\_cache may have
- > cpu\_slab, and cpu\_partials\_slabs.
- > As we already free all objects, kmem\_cache\_verify\_dead() is not invoked forever.
- > I think that we need another kmem\_cache\_shrink() in
- > kmem\_cache\_destroy\_work\_func().

I'll take a look here. What you describe makes sense, and can potentially happen. I tried to handle this case with care in destroy all caches, but I may have always made a mistake...

Did you see this actively happening, or are you just assuming this can happen from your read of the code?