

---

Subject: Re: [PATCH v5 13/18] memcg/sl[au]b Track all the memcg children of a kmem\_cache.

Posted by [Glauber Costa](#) on Tue, 30 Oct 2012 11:31:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

On 10/29/2012 07:26 PM, JoonSoo Kim wrote:

> 2012/10/19 Glauber Costa <glommer@parallels.com>:

>> +void kmem\_cache\_destroy\_memcg\_children(struct kmem\_cache \*s)

>> +{

>> + struct kmem\_cache \*c;

>> + int i;

>> +

>> + if (!s->memcg\_params)

>> + return;

>> + if (!s->memcg\_params->is\_root\_cache)

>> + return;

>> +

>> + /\*

>> + \* If the cache is being destroyed, we trust that there is no one else

>> + \* requesting objects from it. Even if there are, the sanity checks in

>> + \* kmem\_cache\_destroy should caught this ill-case.

>> + \*

>> + \* Still, we don't want anyone else freeing memcg\_caches under our

>> + \* noses, which can happen if a new memcg comes to life. As usual,

>> + \* we'll take the set\_limit\_mutex to protect ourselves against this.

>> + \*/

>> + mutex\_lock(&set\_limit\_mutex);

>> + for (i = 0; i < memcg\_limited\_groups\_array\_size; i++) {

>> + c = s->memcg\_params->memcg\_caches[i];

>> + if (c)

>> + kmem\_cache\_destroy(c);

>> + }

>> + mutex\_unlock(&set\_limit\_mutex);

>> +}

>

> It may cause NULL deref.

> Look at the following scenario.

>

> 1. some memcg slab caches has remained object.

> 2. start to destroy memcg.

> 3. schedule\_delayed\_work(kmem\_cache\_destroy\_work\_func, @delay 60hz)

> 4. all remained object is freed.

> 5. start to destroy root cache.

> 6. kmem\_cache\_destroy makes 's->memcg\_params->memcg\_caches[i]' NULL!!

> 7. Start delayed work function.

> 8. cachep in kmem\_cache\_destroy\_work\_func() may be NULL

>

> Thanks.

>

Thanks for spotting. This is the same problem we have in `memcg_cache_destroy()`, which I solved by not respawning the worker.

In here, I believe it should be possible to just cancel all remaining pending work, since we are now in the process of deleting the caches ourselves.

---