
Subject: Re: [PATCH v2 04/11] kmem accounting basic infrastructure
Posted by [Ying Han](#) on Wed, 15 Aug 2012 18:11:49 GMT

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

On Wed, Aug 15, 2012 at 8:34 AM, Christoph Lameter <cl@linux.com> wrote:

> On Wed, 15 Aug 2012, Glauber Costa wrote:

>

>> On 08/15/2012 06:47 PM, Christoph Lameter wrote:

>> > On Wed, 15 Aug 2012, Michal Hocko wrote:

>> >

>> >>> That is not what the kernel does, in general. We assume that if he wants

>> >>> that memory and we can serve it, we should. Also, not all kernel memory

>> >>> is unreclaimable. We can shrink the slabs, for instance. Ying Han

>> >>> claims she has patches for that already...

>> >>

>> >> Are those patches somewhere around?

>> >

>> > You can already shrink the reclaimable slabs (dentries / inodes) via

>> > calls to the subsystem specific shrinkers. Did Ying Han do anything to

>> > go beyond that?

>> >

>> That is not enough for us.

>> We would like to make sure that the objects being discarded belong to

>> the memcg which is under pressure. We don't need to be perfect here, and

>> an occasional slip is totally fine. But if in general, shrinking from

>> memcg A will mostly wipe out objects from memcg B, we harmed the system

>> in return for nothing good.

>

> How can you figure out which objects belong to which memcg? The ownerships

> of dentries and inodes is a dubious concept already.

I figured it out based on the kernel slab accounting.

obj->page->kmem_cache->memcg

--Ying

>

> --

> To unsubscribe, send a message with 'unsubscribe linux-mm' in

> the body to majordomo@kvack.org. For more info on Linux MM,

> see: <http://www.linux-mm.org/> .

> Don't email: <[a href="mailto:dont@kvack.org">href=mailto:"dont@kvack.org"](mailto:dont@kvack.org)> email@kvack.org