

---

Subject: Re: [PATCH v2 04/11] kmem accounting basic infrastructure  
Posted by [Glauber Costa](#) on Wed, 15 Aug 2012 15:35:13 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 08/15/2012 07:34 PM, Christoph Lameter 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.

>

Remember we copy over the metadata and create copies of the caches  
per-memcg. Therefore, a dentry belongs to a memcg if it was allocated  
from the slab pertaining to that memcg.

It is not 100 % accurate, but it is good enough.

---