
Subject: Re: [PATCH 09/23] kmem slab accounting basic infrastructure
Posted by [Glauber Costa](#) on Wed, 25 Apr 2012 14:38:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 04/24/2012 10:32 PM, KAMEZAWA Hiroyuki wrote:

> (2012/04/21 6:57), Glauber Costa wrote:

>

>> This patch adds the basic infrastructure for the accounting of the slab
>> caches. To control that, the following files are created:

>>

>> * memory.kmem.usage_in_bytes

>> * memory.kmem.limit_in_bytes

>> * memory.kmem.failcnt

>> * memory.kmem.max_usage_in_bytes

>>

>> They have the same meaning of their user memory counterparts. They reflect
>> the state of the "kmem" res_counter.

>>

>> The code is not enabled until a limit is set. This can be tested by the flag
>> "kmem_accounted". This means that after the patch is applied, no behavioral
>> changes exists for whoever is still using memcg to control their memory usage.

>>

>

> Hmm, res_counter never goes naeative ?

Why would it?

This one has more or less the same logic as the sock buffers.

If we are not accounted, the caches don't get created. If the caches
don't get created, we don't release them. (this is modulo bugs, of course)
