
Subject: Re: [PATCH 0/7] memcg kernel memory tracking
Posted by [Glauber Costa](#) on Tue, 28 Feb 2012 19:02:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 02/23/2012 04:18 PM, Ying Han wrote:

> On Tue, Feb 21, 2012 at 3:34 AM, Glauber Costa<glommer@parallels.com> wrote:

>> This is a first structured approach to tracking general kernel
>> memory within the memory controller. Please tell me what you think.

>>

>> As previously proposed, one has the option of keeping kernel memory
>> accounted separately, or together with the normal userspace memory.
>> However, this time I made the option to, in this later case, bill
>> the memory directly to memcg->res. It has the disadvantage that it becomes
>> complicated to know which memory came from user or kernel, but OTOH,
>> it does not create any overhead of drawing from multiple res_counters
>> at read time. (and if you want them to be joined, you probably don't care)

>

> Keeping one counter for user and kernel pages makes it easier for
> admins to configure the system. About reporting, we should still
> report the user and kernel memory separately. It will be extremely
> useful when diagnosing the system like heavily memory pressure or OOM.

It will also make us charge two different res_counters, which is not a
cheap operation.

I was wondering if we can do something smarter within the res_counter
itself to avoid taking locks for two different res_counters in the
charge path?
