

---

Subject: Re: [PATCH v2 09/11] memcg: propagate kmem limiting information to children

Posted by [Glauber Costa](#) on Tue, 21 Aug 2012 10:01:24 GMT

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

---

On 08/21/2012 02:00 PM, Michal Hocko wrote:

> On Tue 21-08-12 13:22:09, Glauber Costa wrote:

>> On 08/21/2012 11:54 AM, Michal Hocko wrote:

> [...]

>>> But maybe you have a good use case for that?

>>>

>> Honestly, I don't. For my particular use case, this would be always on,  
>> and end of story. I was operating under the belief that being able to  
>> say "Oh, I regret", and then turning it off would be beneficial, even at  
>> the expense of the - self contained - complication.

>>

>> For the general sanity of the interface, it is also a bit simpler to say  
>> "if kmem is unlimited, x happens", which is a verifiable statement, than  
>> to have a statement that is dependent on past history.

>

> OK, fair point. We shouldn't rely on the history. Maybe

> memory.kmem.limit\_in\_bytes could return some special value like -1 in

> such a case?

>

Way I see it, this is simplifying the code at the expense of complicating the interface.

>> But all of those need of course, as you pointed out, to be traded off  
>> by the code complexity.

>>

>> I am fine with either, I just need a clear sign from you guys so I don't  
>> keep deimplementing and reimplementing this forever.

>

> I would be for make it simple now and go with additional features later  
> when there is a demand for them. Maybe we will have runtime switch for  
> user memory accounting as well one day.

>

Since this would change a then established behavior, the same discussions about compatibility we ever get to will rise. It is a pain we'd better avoid if we can.

> But let's see what others think?

>

Absolutely. Hello others, what do you think ?

---