

---

Subject: Re: [PATCH v3 04/13] kmem accounting basic infrastructure

Posted by [Glauber Costa](#) on Wed, 26 Sep 2012 17:34:00 GMT

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

---

On 09/26/2012 08:01 PM, Michal Hocko wrote:

> On Wed 26-09-12 18:33:10, Glauber Costa wrote:

>> On 09/26/2012 06:03 PM, Michal Hocko wrote:

>>> On Tue 18-09-12 18:04:01, Glauber Costa wrote:

> [...]

>>>> @@ -4961,6 +5015,12 @@ mem\_cgroup\_create(struct cgroup \*cont)

>>>> int cpu;

>>>> enable\_swap\_cgroup();

>>>> parent = NULL;

>>>> +

>>>> + #ifdef CONFIG\_MEMCG\_KMEM

>>>> + WARN\_ON(cgroup\_add\_cftypes(&mem\_cgroup\_subsys,

>>>> + kmem\_cgroup\_files));

>>>> + #endif

>>>> +

>>>> if (mem\_cgroup\_soft\_limit\_tree\_init())

>>>> goto free\_out;

>>>> root\_mem\_cgroup = memcg;

>>>> @@ -4979,6 +5039,7 @@ mem\_cgroup\_create(struct cgroup \*cont)

>>>> if (parent && parent->use\_hierarchy) {

>>>> res\_counter\_init(&memcg->res, &parent->res);

>>>> res\_counter\_init(&memcg->memsw, &parent->memsw);

>>>> + res\_counter\_init(&memcg->kmem, &parent->kmem);

>>>

>>> Haven't we already discussed that a new memcg should inherit kmem\_accounted

>>> from its parent for use\_hierarchy?

>>> Say we have

>>> root

>>> |

>>> A (kmem\_accounted = 1, use\_hierarchy = 1)

>>> \

>>> B (kmem\_accounted = 0)

>>> \

>>> C (kmem\_accounted = 1)

>>>

>>> B find's itself in an awkward situation becuae it doesn't want to

>>> account u+k but it ends up doing so becuae C.

>>>

>>

>> Ok, I haven't updated it here. But that should be taken care of in the

>> lifecycle patch.

>

> I am not sure which patch you are thinking about but I would prefer to

> have it here because it is safe wrt. races and it is more obvious as

> well.

>

The patch where I make kmem\_accounted into a bitfield. So any code here will eventually disappear.

But BTW, I am not saying I won't update the patch - I like that all patches work and make sense in their own, I am just saying that I forgot to update this patch, because I added the code in its final version to the end and then squashed it.

---