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

Posted by [Glauber Costa](#) on Wed, 26 Sep 2012 17:53:09 GMT

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

On 09/26/2012 09:44 PM, Tejun Heo wrote:

> Hello, Glauber.

>

> On Wed, Sep 26, 2012 at 10:36 AM, Glauber Costa <glommer@parallels.com> wrote:

>> This was discussed multiple times. Our interest is to preserve existing

>> deployed setup, that were tuned in a world where kmem didn't exist.

>> Because we also feed kmem to the user counter, this may very well

>> disrupt their setup.

>

> So, that can be served by .kmem_accounted at root, no?

>

>> User memory, unlike kernel memory, may very well be totally in control

>> of the userspace application, so it is not unreasonable to believe that

>> extra pages appearing in a new kernel version may break them.

>>

>> It is actually a much worse compatibility problem than flipping

>> hierarchy, in comparison

>

> Again, what's wrong with one switch at the root?

>

I understand your trauma about over flexibility, and you know I share of it. But I don't think there is any need to cap it here. Given kmem accounted is perfectly hierarchical, and there seem to be plenty of people who only care about user memory, I see no reason to disallow a mixed use case here.

I must say that for my particular use case, enabling it unconditionally would just work, so it is not that what I have in mind.
