
Subject: Re: [PATCH v3 04/13] kmem accounting basic infrastructure
Posted by [Glauber Costa](#) on Wed, 26 Sep 2012 18:56:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 09/26/2012 10:01 PM, Tejun Heo wrote:

> Hello,

>

> On Wed, Sep 26, 2012 at 09:53:09PM +0400, Glauber Costa wrote:

>> 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.

>

> So, I'm not gonna go as far as pushing for enabling it unconditionally
> but would really like to hear why it's necessary to make it per node
> instead of one global switch. Maybe it has already been discussed to
> hell and back. Care to summarize / point me to it?
>

For me, it is the other way around: it makes perfect sense to have a per-subtree selection of features where it doesn't hurt us, provided it is hierarchical. For the mere fact that not every application is interested in this (Michal is the one that is being so far more vocal about this not being needed in some use cases), and it is perfectly valid to imagine such applications would coexist.

So given the flexibility it brings, the real question is, as I said, backwards: what is it necessary to make it a global switch ?
