
Subject: Re: [PATCH v2 04/11] kmem accounting basic infrastructure
Posted by [James Bottomley](#) on Wed, 15 Aug 2012 11:12:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2012-08-15 at 13:33 +0400, Glauber Costa wrote:

> > This can
> > be quite confusing. I am still not sure whether we should mix the two
> > things together. If somebody wants to limit the kernel memory he has to
> > touch the other limit anyway. Do you have a strong reason to mix the
> > user and kernel counters?
>
> This is funny, because the first opposition I found to this work was
> "Why would anyone want to limit it separately?" =p
>
> It seems that a quite common use case is to have a container with a
> unified view of "memory" that it can use the way he likes, be it with
> kernel memory, or user memory. I believe those people would be happy to
> just silently account kernel memory to user memory, or at the most have
> a switch to enable it.
>
> What gets clear from this back and forth, is that there are people
> interested in both use cases.

Haven't we already had this discussion during the Prague get together?
We discussed the use cases and finally agreed to separate accounting for
k and then k+u mem because that satisfies both the Google and Parallels
cases. No-one was overjoyed by k and k+u but no-one had a better
suggestion ... is there a better way of doing this that everyone can
agree to?

We do need to get this nailed down because it's the foundation of the
patch series.

James
