Subject: Re: [PATCH v9 1/9] Basic kernel memory functionality for the Memory Controller

Posted by Michal Hocko on Fri, 16 Dec 2011 13:30:49 GMT

View Forum Message <> Reply to Message On Fri 16-12-11 17:02:51, Glauber Costa wrote: > On 12/16/2011 04:32 PM. Michal Hocko wrote: [...] >>So why do we need kmem accounting when tcp (the only user at the moment) > >doesn't use it? > Well, a bit historical. I needed a basic placeholder for it, since > it top is officially kmem. As the time passed, I took most of the > stuff out of this patch to leave just the basics I would need for > Turns out I ended up focusing on the rest, and some of the stuff was > left here. > At one point I merged top data into kmem, but then reverted this > behavior. the kmem counter stayed. > I agree deferring the whole behavior would be better. >>>In summary, we still never do non-independent accounting. When we >>>start doing it for the other caches, We will have to add a test at >>>charge time as well. >>So we shouldn't do it as a part of this patchset because the further > >usage is not clear and I think there will be some real issues with

> >user+kmem accounting (e.g. a proper memcg-oom implementation). > > Can you just drop this patch?

> > Yes, but the whole set is in the net tree already.

Isn't it only in some for-next branch? Can that one be updated?

- > (All other patches are tcp-related but this) Would you mind if I'd
- > send a follow up patch removing the kmem files, and leaving just the
- > registration functions and basic documentation? (And sorry for that as
- > well in advance)

Yes a followup patch would work as well.

Michal Hocko SUSE Labs SUSE LINUX s.r.o. Lihovarska 1060/12 Page 2 of 2 ---- Generated from OpenVZ Forum