
Subject: Re: [PATCH 08/10] memcg: Add
CONFIG_CGROUP_MEM_RES_CTLR_KMEM_ACCT_ROOT.
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 28 Feb 2012 23:54:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 28 Feb 2012 15:36:27 -0800
Suleiman Souhlal <suleiman@google.com> wrote:

> On Tue, Feb 28, 2012 at 5:34 AM, Glauber Costa <glommer@parallels.com> wrote:
> > On 02/27/2012 07:58 PM, Suleiman Souhlal wrote:
> >>
> >> This config option dictates whether or not kernel memory in the
> >> root cgroup should be accounted.
> >>
> >> This may be useful in an environment where everything is supposed to be
> >> in a cgroup and accounted for. Large amounts of kernel memory in the
> >> root cgroup would indicate problems with memory isolation or accounting.
> >
> >
> > I don't like accounting this stuff to the root memory cgroup. This causes
> > overhead for everybody, including people who couldn't care less about memcg.
> >
> > If it were up to me, we would simply not account it, and end of story.
> >
> > However, if this is terribly important for you, I think you need to at
> > least make it possible to enable it at runtime, and default it to disabled.
>
> Yes, that is why I made it a config option. If the config option is
> disabled, that memory does not get accounted at all.
>
> Making it configurable at runtime is not ideal, because we would
> prefer slab memory that was allocated before cgroups are created to
> still be counted toward root.
>

I never like to do accounting in root cgroup.
If you want to do this, add config option to account `_all_` memory in the root.
Accounting only slab seems very dirty hack.

Thanks,
-Kame
