
Subject: Re: [PATCH v4 14/14] Add documentation about the kmem controller
Posted by [Michal Hocko](#) on Fri, 12 Oct 2012 08:44:32 GMT

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

On Fri 12-10-12 11:53:23, Glauber Costa wrote:

> On 10/11/2012 06:35 PM, Michal Hocko wrote:

> > On Mon 08-10-12 14:06:20, Glauber Costa wrote:

[...]

> >> Kernel memory limits are not imposed for the root cgroup. Usage for the root

> >> -cgroup may or may not be accounted.

> >> +cgroup may or may not be accounted. The memory used is accumulated into

> >> +memory.kmem.usage_in_bytes, or in a separate counter when it makes sense.

> >

> > Which separate counter? Is this about tcp kmem?

> >

>

> So far, yes, this is the only case that makes sense, and the fewer the

> better. In any case it exists, and I wanted to be generic.

Add (currently tcp) or something similar

[...]

> >> + Kernel memory is effectively set as a percentage of the user memory. This

> >

> > not a percentage it is subset of the user memory

> >

> Well, this is semantics. I can change, but for me it makes a lot of

> sense to think of it in terms of a percentage, because it is easy to

> administer. You don't actually write a percentage, which I tried to

> clarify by using the term "effective set as a percentage".

I can still see somebody reading this and wondering why echo 50 > ...limit
didn't set a percentage...

--

Michal Hocko
SUSE Labs
