Subject: Re: [PATCH 0/5] Kernel memory accounting container (v5) Posted by Paul Menage on Fri, 05 Oct 2007 07:11:30 GMT

View Forum Message <> Reply to Message

On 10/2/07, Pavel Emelyanov < xemul@openvz.org> wrote:

>

- > Such a lookup would require a hastable or something similar. We already
- > have such a bad experience (with OpenVZ RSS fractions accounting for
- > example). Hash lookups imply the CPU caches screwup and hurt the performance.
- > See also the comment below.

I think you could do it with an array lookup if you assigned an index to each cache as it was created, and used that as an offset into a per-cgroup array.

>

- > I thought the same some time ago and tried to make a per-beancounter kmem
- > caches. The result was awful the memory waste was much larger than in the
- > case of pointer-per-object approach.

OK, fair enough.

Was this with a large number of bean counters? I imagine that with a small number, the waste might be rather more reasonable.

Paul