Subject: Re: [PATCH v5 14/14] Add documentation about the kmem controller
Posted by Christoph Lameter on Tue, 16 Oct 2012 18:25:06 GMT
View Forum Message <> Reply to Message

On Tue, 16 Oct 2012, Glauber Costa wrote:

>
> + memory.kmem.limit_in_bytes      # set/show hard limit for kernel memory
> + memory.kmem.usage_in_bytes      # show current kernel memory allocation
> + memory.kmem.failcnt             # show the number of kernel memory usage hits limits
> + memory.kmem.max_usage_in_bytes  # show max kernel memory usage recorded

Does it actually make sense to limit kernel memory? The user generally has
no idea how much kernel memory a process is using and kernel changes can
change the memory footprint. Given the fuzzy accounting in the kernel a
large cache refill (if someone configures the slab batch count to be
really big f.e.) can account a lot of memory to the wrong cgroup. The
allocation could fail.

Limiting the total memory use of a process (U+K) would make more sense I
guess. Only U is probably sufficient? In what way would a limitation on
kernel memory in use be good?