
Subject: Re: [PATCH] Show slab memory usage on OOM and SysRq-M
Posted by [Eric Dumazet](#) on Wed, 18 Apr 2007 12:07:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 18 Apr 2007 09:17:19 +0300 (EEST)
Pekka J Enberg <penberg@cs.helsinki.fi> wrote:

> On Tue, 17 Apr 2007, Eric Dumazet wrote:
> > This nr_pages should be in struct kmem_list3, not in struct kmem_cache,
> > or else you defeat NUMA optimizations if touching a field in kmem_cache
> > at kmem_getpages()/kmem_freepages() time.
>
> We already touch ->flags, ->gfpflags, and ->gfporder in kmem_getpages().
> Sorry for my ignorance, but how is this different?
>

Those fields are read. Thats OK, because several CPUS might share all those without problem.

But modifying one field in kmem_cache would invalidate one cache line for all cpus that would have to reload it later.

This is what we call "false sharing" or cache line ping pongs
