
Subject: Re: [ckrm-tech] [patch02/05]: Containers(V2)- Generic Linux kernel changes

Posted by [Rohit Seth](#) on Thu, 21 Sep 2006 22:29:45 GMT

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

On Wed, 2006-09-20 at 18:33 -0700, Chandra Seetharaman wrote:

> On Wed, 2006-09-20 at 17:23 -0700, Christoph Lameter wrote:

> > On Wed, 20 Sep 2006, Andi Kleen wrote:

> >

> > > There are lots of different cases. At least for anonymous memory

> > > ->mapping should be free. Perhaps that could be used for anonymous

> > > memory and a separate data structure for the important others.

> >

> > mapping is used for swap and to point to the anon vma.

> >

> > > slab should have at least one field free too, although it might be a different

> > > one (iirc Christoph's rewrite uses more than the current slab, but it would

> > > surprise me if he needed all)

> >

> > slab currently has lots of fields free but my rewrite uses all of them.

> > And AFAICT this patchset does not track slab pages.

> >

> > Hmm.... Build a radix tree with pointers to the pages?

>

> Yes, that would be a way to isolate the overhead.

>

As said earlier, the additional cost is not really that much. And if we do get additional benefit of proper cache alignment for page struct because of this additional pointer then it could be a almost zero cost.

And in mm kernel, I also see PAGE_OWNER defined which puts the whole page struct out of sync with caches.

Christoph, Please let me know if there is still a placeholder for container pointer ONLY for kernel pages. I think it is easy to do away with per page pointer for user land as it can be stored in anon_vma or mappings.

-rohit
