
Subject: Re: [ckrm-tech] [RFC][PATCH 5/7] UBC: kernel memory accounting (core)
Posted by [Dave Hansen](#) on Thu, 17 Aug 2006 16:37:39 GMT

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

On Thu, 2006-08-17 at 16:01 +0100, Alan Cox wrote:

> Ar lau, 2006-08-17 am 07:26 -0700, ysgrifennodd Dave Hansen:
> > My main thought is that `_everybody_` is going to have to live with the
> > entry in the 'struct page'. Distros ship one kernel for everybody, and
> > the cost will be paid by those not even using any kind of resource
> > control or containers.
> >
> > That said, it sure is simpler to implement, so I'm all for it!
>
> I don't see any good way around that. For the page struct it is a
> material issue, for the others its not a big deal providing we avoid
> accounting dumb stuff like dentries.

The only way I see around it is using other mechanisms to more loosely attribute ownership of a page to particular containers. I know that my suggestion of traversing the rmap chains at page reclaim time is going appears to be slow compared to the regular reclaim path, but I'm not sure it really matters.

Let's say you have 20 containers sharing 20 pages which are on the LRU, and those pages are evenly distributed so that each container owns one page. Only one of those 20 containers is over its limit. With something that strictly assigns container ownership to one and only one container, you're going to have to walk half of the LRU to find the page.

With a "loose ownership" scheme, you'd hit on the first page, and you'd pay the cost by having to walk halfway through the 20 rmap entries. Admittedly, the rmap walk is much slower than an LRU walk.

-- Dave
