
Subject: Re: [patch00/05]: Containers(V2)- Introduction
Posted by [Peter Zijlstra](#) on Wed, 20 Sep 2006 18:37:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2006-09-20 at 10:50 -0700, Rohit Seth wrote:
> On Thu, 2006-09-21 at 03:00 +1000, Nick Piggin wrote:
> > (this time to the lists as well)
> >
> > Peter Zijlstra wrote:
> >
> > > I'd much rather containerize the whole reclaim code, which should not
> > > be too hard since he already adds a container pointer to struct page.
> >
> >
>
> Right now the memory handler in this container subsystem is written in
> such a way that when existing kernel reclaimer kicks in, it will first
> operate on those (container with pages over the limit) pages first. But
> in general I like the notion of containerizing the whole reclaim code.

Patch 5/5 seems to have a horrid deactivation scheme.

> > > I still have to reread what Rohit does for file backed pages, that gave
> > > my head a spin.
>
> Please let me know if there is any specific part that isn't making much
> sense.

Well, the whole over the limit handler is quite painfull, having taken a second reading it isn't all that complex after all, just odd.

You just start invalidating whole files for file backed pages. Granted, this will get you below the threshold. but you might just have destroyed your working set.

Pretty much the same for you anonymous memory handler, you scan through the pages in linear fashion and demote the first that you encounter.

Both things pretty thoroughly destroy the existing kernel reclaim.
