
Subject: Re: [PATCH 8/8] Per-container pages reclamation
Posted by [Andrew Morton](#) on Fri, 01 Jun 2007 07:49:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 01 Jun 2007 11:44:25 +0400 Pavel Emelianov <xemul@openvz.org> wrote:

> Andrew Morton wrote:
> > On Thu, 31 May 2007 14:35:43 +0400
> > Pavel Emelianov <xemul@openvz.org> wrote:
> >
> >>> Would I be correct in guessing that pages which are on the
> >>> per-rss-container lists are also eligible for reclaim off the traditional
> >>> page LRUs? If so, how does that work? When a page gets freed off the
> >> Yes. All the pages are accessible from booth - global and per-container
> >> LRU lists and reclamation can be performed from booth.
> >>
> >>> per-zone LRUs does it also get removed from the per-rss_container LRU? But
> >>> how can this be right?
> >> I don't get your idea here.
> >
> > If we have a page which is on the zone LRU with refcount=1 and someone does
> > put_page() on it, we will take that page off the zone LRU and then actually
> > free the page.
> >
> > I am assuming that your patches change that logic so that we will also
> > remove that page from the per-container LRU at the same time?
>
> Oh, I see. No that will work another way. Page stays in per-container
> LRU lists as long as it is mapped. When the last process is unmapping
> the page, it is removed, but stays in global LRU till its refcount
> drops to 0.

ho hum, OK. This wasn't obvious from the code and it isn't something
I should have learned by emailing you guys! Please have a think about
preparing an overall decription of the design and implementation of this
rather important kernel change?
