
Subject: Re: [PATCH 3/3][RFC] Containers: Pagecache controller reclaim
Posted by [Aubrey Li](#) on Tue, 27 Mar 2007 08:41:33 GMT

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

On 3/27/07, Vaidyanathan Srinivasan <svaidy@linux.vnet.ibm.com> wrote:

> Correct, shrink_page_list() is called from shrink_inactive_list() but
> the above code is patched in shrink_active_list(). The
> 'force_reclaim_mapped' label is from function shrink_active_list() and
> not in shrink_page_list() as it may seem in the patch file.
>
> While removing pages from active_list, we want to select only
> pagecache pages and leave the remaining in the active_list.
> page_mapped() pages are _not_ of interest to pagecache controller
> (they will be taken care by rss controller) and hence we put it back.
> Also if the pagecache controller is below limit, no need to reclaim
> so we put back all pages and come out.

Oh, I just read the patch, not apply it to my local tree, I'm working
on 2.6.19 now.

So the question is, when vfs pagecache limit is hit, the current
implementation just reclaim few pages, so it's quite possible the
limit is hit again, and hence the reclaim code will be called again
and again, that will impact application performance.

-Aubrey
