

---

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

---