
Subject: Re: [PATCH][RFC] dirty balancing for cgroups
Posted by [KAMEZAWA Hiroyuki](#) on Fri, 11 Jul 2008 07:09:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 11 Jul 2008 14:59:26 +0900 (JST)
yamamoto@valinux.co.jp (YAMAMOTO Takashi) wrote:

> > > - This looks simple but, could you merge this into memory resource controller ?
> > >
> > > why?
> > >
> > 3 points.
> > 1. Is this useful if used alone ?
>
> it can be. why not?
>
> > 2. memcg requires this kind of feature, basically.
> >
> > 3. I wonder I need more work to make this work well under memcg.
>
> i'm not sure if i understand these points. can you explain a bit?
>
In my understanding, dirty_ratio is for helping memory (reclaim) subsystem.

See comments in fs/page-writeback.c:: determin_dirtyable_memory()

```
==  
/*  
 * Work out the current dirty-memory clamping and background writeout  
 * thresholds.  
 *  
 * The main aim here is to lower them aggressively if there is a lot of mapped  
 * memory around. To avoid stressing page reclaim with lots of unreclaimable  
 * pages. It is better to clamp down on writers than to start swapping, and  
 * performing lots of scanning.  
 *  
 * We only allow 1/2 of the currently-unmapped memory to be dirtied.  
 *  
 * We don't permit the clamping level to fall below 5% - that is getting rather  
 * excessive.  
 *  
 * We make sure that the background writeout level is below the adjusted  
 * clamping level.  
==
```

"To avoid stressing page reclaim with lots of unreclaimable pages"

Then, I think memcg should support this for helping reclaim under memcg.

> my patch penalizes heavy-writer cgroups as task_dirty_limit does
> for heavy-writer tasks. i don't think that it's necessary to be
> tied to the memory subsystem because i merely want to group writers.

>

Hmm, maybe what I need is different from this ;)

Does not seem to be a help for memory reclaim under memcg.

Thanks,

-Kame

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
