

---

Subject: Re: [PATCH][RFC] dirty balancing for cgroups  
Posted by Peter Zijlstra on Mon, 14 Jul 2008 13:37:17 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2008-07-09 at 15:00 +0900, YAMAMOTO Takashi wrote:

> hi,  
>  
> the following patch is a simple implementation of  
> dirty balancing for cgroups. any comments?  
>  
> it depends on the following fix:  
> <http://lkml.org/lkml/2008/7/8/428>  
>  
> YAMAMOTO Takashi  
>  
>  
> Signed-off-by: YAMAMOTO Takashi <yamamoto@valinux.co.jp>  
> ---

Yamamoto-san,

```
> @@ -408,7 +412,11 @@ get_dirty_limits(long *pbbackground, long *pdirty, long *pbdi_dirty,
>
>     *pbdi_dirty = bdi_dirty;
>     clip_bdi_dirty_limit(bdi, dirty, pbdi_dirty);
> - task_dirty_limit(current, pbdi_dirty);
> + task_dirty = *pbdi_dirty;
> + task_dirty_limit(current, &task_dirty);
> + cgroup_dirty = *pbdi_dirty;
> + memdirtylimitcgroup_dirty_limit(current, &cgroup_dirty);
> + *pbdi_dirty = min(task_dirty, cgroup_dirty);
> }
> }
```

I think this is wrong - is basically breaks task dirty throttling within groups. You'd need a multiplicative operation, something like:

$$\text{bdi\_dirty} = \text{dirty} * p(\text{bdi}) * p(\text{cgrou}) * (1 - p(\text{task}))$$

However then we still have problems... see the next email further down the thread.

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---