
Subject: Re: [PATCH 4/5] blk-throttle: track buffered and anonymous pages

Posted by [Chad Talbott](#) on Tue, 22 Feb 2011 18:42:41 GMT

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

On Tue, Feb 22, 2011 at 9:12 AM, Andrea Righi <arighi@develer.com> wrote:

> Add the tracking of buffered (writeback) and anonymous pages.

```
...
> ---
> block/blk-throttle.c | 87 ++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
> include/linux/blkdev.h | 26 ++++++
> 2 files changed, 111 insertions(+), 2 deletions(-)
>
> diff --git a/block/blk-throttle.c b/block/blk-throttle.c
> index 9ad3d1e..a50ee04 100644
> --- a/block/blk-throttle.c
> +++ b/block/blk-throttle.c
...
> +int blk_throtl_set_anonpage_owner(struct page *page, struct mm_struct *mm)
> +int blk_throtl_set_filepage_owner(struct page *page, struct mm_struct *mm)
> +int blk_throtl_copy_page_owner(struct page *npage, struct page *opage)
```

It would be nice if these were named `blk_cgroup_*`. This is arguably more correct as the id comes from the blkio subsystem, and isn't specific to blk-throttle. This will be more important very shortly, as CFQ will be using this same cgroup id for async IO tracking soon.

`is_kernel_io()` is a good idea, it avoids a bug that we've run into with CFQ async IO tracking. Why isn't `PF_KTHREAD` sufficient to cover all kernel threads, including `kswapd` and those marked `PF_MEMALLOC`?

Thanks,
Chad

Containers mailing list

Containers@lists.linux-foundation.org

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