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

Posted by [Andrea Righi](#) on Tue, 22 Feb 2011 19:12:47 GMT

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

On Tue, Feb 22, 2011 at 10:42:41AM -0800, Chad Talbott wrote:

> 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.

Sounds reasonable. Will do in the next version.

>

> 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?

With PF_MEMALLOC we're sure we don't add the page tracking overhead also to non-kernel threads when memory gets low.

PF_KSWAPD is not probably needed, AFAICS it is only used by kswapd, that is created by kthread_create() and so it has the PF_KTHREAD flag set.

Let's see if someone can give more details about that. In the while I'll investigate and try to do some tests only with PF_KTHREAD.

Thanks,
-Andrea

Containers mailing list
Containers@lists.linux-foundation.org

