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

Posted by [Andrea Righi](#) on Tue, 22 Feb 2011 23:03:54 GMT

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

On Tue, Feb 22, 2011 at 03:49:28PM -0500, Vivek Goyal wrote:

> 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.

>

> Should this really be all part of blk-cgroup.c and not blk-throttle.c

> so that it can be used by CFQ code also down the line? Anyway all this

> is not throttle specific as such but blkio controller specific.

Agreed.

>

> Though function naming convention is not great in blk-cgroup.c But

> functions either have blkio_ prefix or blkicg_ prefix.

ok.

>

> Functions which are not directly dealing with cgroups or in general

> are called by blk-throttle.c and/or cfq-iosched.c I have marked as

> prefixed with "blkio_". Functions which directly deal with cgroup stuff

> and register with cgroup subsystem for this controller are generally

> having "blkicg_" prefix.

>

> In this case probably we can use probably blkio_ prefix.

ok.

Thanks,
-Andrea

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
