
Subject: Re: [PATCH][RFC] incorrect direct io error handling (v3)
Posted by [Dmitriy Monakhov](#) on Fri, 26 Jan 2007 07:59:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton <akpm@osdl.org> writes:

```
> On Wed, 24 Jan 2007 22:05:06 +0300
> Dmitriy Monakhov <dmonakhov@sw.ru> wrote:
>
>> incorrect direct io error handling (v3)
>> Changes from v2:
>> - Remove BUG_ON(!mutex_is_locked(..)) for non blkdev.
>> - vmtruncate() called from generic_file_aio_write().
>> - depends on patch titled:
>>   [PATH][RFC] mm: Move common segments checks to separate function
>
> drat, I skipped that patch due to rejects, and because Nick is working on
> things in the same area.
>
[skip]
>> if ((written >= 0 || written == -EIOCBQUEUED) &&
>>     ((file->f_flags & O_SYNC) || IS_SYNC(inode))) {
>> @@ -2365,6 +2366,17 @@ ssize_t generic_file_aio_write(struct ki
>>   &iocb->ki_pos);
>>   mutex_unlock(&inode->i_mutex);
>>
>> + if (unlikely(ret < 0 && (file->f_flags & O_DIRECT))) {
>> +   ssize_t cnt = generic_segment_checks(nr_segs, iov, VERIFY_READ);
>> +   loff_t isize = i_size_read(inode);
>> +   /*
>> +    * generic_file_direct_write() may have instantiated a few
>> +    * blocks outside i_size. Trim these off again.
>> +    */
>> +   if (cnt > 0 && (pos + cnt > isize))
>> +     vmtruncate(inode, isize);
>> + }
>> +
>
> vmtruncate() really wants i_mutex to be held. Can't we do that here?
> Yepp 110% true, baaahh ..... it looks like my brain wasn't clear at the time i
> wrote this. We have to do vmtruncate() before dropping i_mutex, right after
> __generic_file_aio_write_nolock() call, i'm sorry to waste your time.
```
