Subject: Re: [PATCH] incorrect error handling inside generic\_file\_direct\_write Posted by Christoph Hellwig on Tue, 02 Jan 2007 11:17:46 GMT View Forum Message <> Reply to Message

On Fri, Dec 15, 2006 at 10:53:18AM -0800, Chen, Kenneth W wrote: > Christoph Hellwig wrote on Friday, December 15, 2006 2:44 AM > > So we're doing the sync\_page\_range once in \_\_generic\_file\_aio\_write > > with i mutex held. > > > > >>> mutex\_lock(&inode->i\_mutex); >> - ret = \_\_generic\_file\_aio\_write\_nolock(iocb, iov, nr\_segs, >>> - &iocb->ki\_pos); >> + ret = \_\_generic\_file\_aio\_write(iocb, iov, nr\_segs, pos); >>> mutex\_unlock(&inode->i\_mutex); >>> >>> if (ret > 0 && ((file->f flags & O SYNC) || IS SYNC(inode))) { > > And then another time after it's unlocked, this seems wrong. > > I didn't invent that mess though. > I should've ask the question first: in 2.6.20-rc1, generic\_file\_aio\_write > will call sync\_page\_range twice, once from \_\_generic\_file\_aio\_write\_nolock > and once within the function itself. Is it redundant? Can we delete the > one in the top level function? Like the following?

Really? I'm looking at -rc3 now as -rc1 is rather old and it's definitly not the case there. I also can't remember ever doing this - when I started the generic read/write path untangling I had exactly the same situation that's now in -rc3:

- generic\_file\_aio\_write\_nolock calls sync\_page\_range\_nolock
- generic\_file\_aio\_write calls sync\_page\_range
- \_\_generic\_file\_aio\_write\_nolock doesn't call any sync\_page\_range variant