

Signed-off-by: Monakhov Dmitriy <dmonakhov@openvz.org>

```
---
mm/filemap.c | 28 ++++++-----
1 files changed, 24 insertions(+), 4 deletions(-)

diff --git a/mm/filemap.c b/mm/filemap.c
index 8bd1ea4..95d49fe 100644
--- a/mm/filemap.c
+++ b/mm/filemap.c
@@ -1932,8 +1932,10 @@ generic_file_direct_write(struct kiocb *iocb, const struct iovec *iov,
/*
 * Sync the fs metadata but not the minor inode changes and
 * of course not the data as we did direct DMA for the IO.
- * i_mutex is held, which protects generic_osync_inode() from
- * livelocking. AIO O_DIRECT ops attempt to sync metadata here.
+ * i_mutex is held in case of DIO_LOCKING, which protects
+ * generic_osync_inode() from livelocking. If it is not held, then
+ * the filesystem must prevent this livelock. AIO O_DIRECT ops
+ * attempt to sync metadata here.
 */
if ((written >= 0 || written == -EIOCBQUEUED) &&
    ((file->f_flags & O_SYNC) || IS_SYNC(inode))) {
@@ -2155,8 +2157,26 @@ __generic_file_aio_write_nolock(struct kiocb *iocb, const struct iovec
*iov,
    loff_t endbyte;
    ssize_t written_buffered;

+ /*
+ * In case of non blockdev we may fail to buffered I/O.
+ * So i_mutex must be held.
+ */
+ if (!S_ISBLK(inode->i_mode))
+ BUG_ON(!mutex_is_locked(&inode->i_mutex));
+
    written = generic_file_direct_write(iocb, iov, &nr_segs, pos,
        ppos, count, ocount);
+ /*
+ * If host is not S_ISBLK generic_file_direct_write() may
+ * have instantiated a few blocks outside i_size files
+ * Trim these off again.
+ */
+ if (unlikely(written < 0) && !S_ISBLK(inode->i_mode)) {
+ loff_t isize = i_size_read(inode);
+ if (pos + count > isize)
+ vmtruncate(inode, isize);
+ }
}
```

```
+
  if (written < 0 || written == count)
    goto out;
  /*
@@ -2261,8 +2281,8 @@ ssize_t generic_file_aio_write(struct kiocb *iocb, const struct iovec
*iov,
EXPORT_SYMBOL(generic_file_aio_write);

/*
- * Called under i_mutex for writes to S_ISREG files. Returns -EIO if something
- * went wrong during pagecache shutdown.
+ * Called under i_mutex for writes to S_ISREG files in case of DIO_LOCKING.
+ * Returns -EIO if something went wrong during pagecache shutdown.
*/
static ssize_t
generic_file_direct_IO(int rw, struct kiocb *iocb, const struct iovec *iov,
--
1.5.0.1
```
