
Subject: [PATCH] ext3: wrong error behavior
Posted by [vaverin](#) on Sat, 02 Sep 2006 11:41:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

SWsoft Virtuozzo/OpenVZ Linux kernel team has discovered that ext3 error behavior was broken in linux kernels since 2.5.x versions by the following patch:

2002/10/31 02:15:26-05:00 tytso@snap.thunk.org
Default mount options from superblock for ext2/3 filesystems
<http://linux.bkbits.net:8080/linux-2.6/gnupatch@3dc0d88eKbV9 ivV4ptRNM8fBuA3JBQ>

In case ext3 file system is mounted with errors=continue (EXT3_ERRORS_CONTINUE) errors should be ignored when possible. However at present in case of any error kernel aborts journal and remounts filesystem to read-only. Such behavior was hit number of times and noted to differ from that of 2.4.x kernels.

This patch fixes this:

- do nothing in case of EXT3_ERRORS_CONTINUE,
- set EXT3_MOUNT_ABORT and call journal_abort() in all other cases
- panic() should be called after ext3_commit_super() to save sb marked as EXT3_ERROR_FS

Signed-off-by: Vasily Averin <vvs@sw.ru>
Ack-by: Kirill Korotaev <dev@sw.ru>

Thank you,
Vasily Averin

SWsoft Virtuozzo/OpenVZ Linux kernel team

```
--- linux-2.6.18-rc5/fs/ext3/super.c.orig 2006-09-02 12:54:01.000000000 +0400
+++ linux-2.6.18-rc5/fs/ext3/super.c 2006-09-02 13:10:02.000000000 +0400
@@ -159,20 +159,21 @@ static void ext3_handle_error(struct sup
    if (sb->s_flags & MS_RDONLY)
        return;

- if (test_opt (sb, ERRORS_RO)) {
-     printk (KERN_CRIT "Remounting filesystem read-only\n");
-     sb->s_flags |= MS_RDONLY;
- } else {
+ if (!test_opt (sb, ERRORS_CONT)) {
    journal_t *journal = EXT3_SB(sb)->s_journal;

    EXT3_SB(sb)->s_mount_opt |= EXT3_MOUNT_ABORT;
    if (journal)
        journal_abort(journal, -EIO);
    }
+ if (test_opt (sb, ERRORS_RO)) {
```

```
+ printk (KERN_CRIT "Remounting filesystem read-only\n");
+ sb->s_flags |= MS_RDONLY;
+ }
+ ext3_commit_super(sb, es, 1);
  if (test_opt(sb, ERRORS_PANIC))
    panic("EXT3-fs (device %s): panic forced after error\n",
        sb->s_id);
- ext3_commit_super(sb, es, 1);
}
```

```
void ext3_error (struct super_block * sb, const char * function,
```
