
Subject: Re: [RFC PATCH ext3/ext4] orphan list corruption due bad inode
Posted by [vaverin](#) on Tue, 05 Jun 2007 06:11:12 GMT

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

Andrew Morton wrote:

> On Mon, 04 Jun 2007 09:19:10 +0400 Vasily Averin <vvs@sw.ru> wrote:

>> diff --git a/fs/ext3/namei.c b/fs/ext3/namei.c

>> index 9bb046d..e3ac8c3 100644

>> --- a/fs/ext3/namei.c

>> +++ b/fs/ext3/namei.c

>> @@ -1019,6 +1019,11 @@ static struct dentry *ext3_lookup(struct inode * dir, struct dentry *dentry, str

>>

>> if (!inode)

>> return ERR_PTR(-EACCES);

>> +

>> + if (is_bad_inode(inode)) {

>> + iput(inode);

>> + return ERR_PTR(-ENOENT);

>> + }

>> }

>> return d_splice_alias(inode, dentry);

>> }

> Seems reasonable. So this prevents the bad inodes from getting onto the

> orphan list in the first place?

make_bad_inode() is called from ext3_read_inode() that is called from iget() only.

When we found that inode is bad we will call iput ASAP. I expect it should be enough to exclude any chances to call any functions that can include this inode into orphan list.

ext3_lookup

iget

ext3_read_inode

make_bad_inode

is_bad_inode? --> iput

However am I probably err?

Probably is better to add is_bad_inode check to ext3_orphan_add()?

thank you,

Vasily Averin
