
Subject: Re: [PATCH] ext3: wrong error behavior
Posted by [vaverin](#) on Wed, 06 Sep 2006 06:53:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric Sandeen wrote:

> Vasily Averin wrote:

>> 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.

>

> I've also noticed this differing behavior,

>

> http://marc.theaimsgroup.com/?l=linux-ext4&m=11537696682_1953&w=2

>

> It passed w/o comment. :)

I would note that one of our developers has noticed it 3 year ago.

http://marc.theaimsgroup.com/?l=linux-kernel&m=104824948_712104&w=2
it was w/o comments too.

> Unless Ted had a specific reason for changing the behavior, 2.4 and 2.6
> should probably be brought into line.

I did not found any arguments explaining Ted's reasons, I believe he did not noticed that he had inversed this condition. If otherwise he had some important reasons, there is another bug: he must change the documentation at least.

I would like to add that our customers are really discontented by this change. Now they pays a high price for any minor ext3 errors: they should stop any processes that uses this partition, check the partition, restart the processes again... Downtime is too long and there is not any workarounds.

> Calling ext3_commit_super()

> before the panic may defeat (some of) the purpose of the panic option,

> though, which is to preserve as much state as possible at the time of

> the error for later analysis...

On the other hand if error is not saved on the disc, fsck will not check this filesystem automatically on the next reboot and therefore node will mount corrupted filesystem, it's dangerous too. In the worst scenario the node find the same error on the disc and rebooted again and it is real nightmare for remote admins.

I would note that Eric Lammerts has tried to fix this issue, however I'm not

sure that he has achieved a success:

<http://marc.theaimsgroup.com/?l=ext3-users&m=110651773907882&w=2>

Also I would like to add that Andrew Morton has signed off the patch and added it to the -mm tree.

Thank you,
Vasily Averin

SWsoft Virtuozzo/OpenVZ Linux kernel team
