

CephFS - Bug #14716

"Thread.cc: 143: FAILED assert(status == 0)" in fs-hammer---basic-smithi

02/09/2016 11:58 PM - Yuri Weinstein

Status: Won't Fix	% Done: 0%
Priority: Normal	
Assignee:	
Category:	
Target version:	
Source: Q/A	Affected Versions:
Tags:	ceph-qa-suite: fs
Backport:	Component(FS):
Regression: No	Labels (FS):
Severity: 3 - minor	Pull request ID:
Reviewed:	Crash signature:

Description

Jobs:

http://qa-proxy.ceph.com/teuthology/teuthology-2016-02-09_08:45:22-fs-hammer---basic-smithi/1624/teuthology.log

http://qa-proxy.ceph.com/teuthology/teuthology-2016-02-09_14:10:44-fs-hammer---basic-mira/1873/teuthology.log

```
2016-02-09T14:17:25.520 INFO:tasks.ceph.mds.a.mira035.stderr:common/Thread.cc: In function 'int Th
read::join(void**)' thread 7f1f3372e700 time 2016-02-09 22:17:25.504583
2016-02-09T14:17:25.520 INFO:tasks.ceph.mds.a.mira035.stderr:common/Thread.cc: 143: FAILED assert (
status == 0)
2016-02-09T14:17:25.520 INFO:tasks.ceph.mds.a.mira035.stderr: ceph version 0.94.5-305-g7abb6ae (7a
bb6ae8f3cba67009bd022aaeee0a87cdfc6477)
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 1: (ceph::__ceph_assert_fail(char co
nst*, char const*, int, char const*)+0x8b) [0x94c72b]
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 2: (Thread::detach()+0) [0x938730]
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 3: (Finisher::stop()+0x9d) [0x87e6cd
]
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 4: (MDS::suicide()+0x85) [0x59fa35]
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 5: (C_MDL_WriteError::finish(int)+0x
65) [0x7da085]
2016-02-09T14:17:25.521 INFO:tasks.ceph.mds.a.mira035.stderr: 6: (MDSIOContextBase::complete(int)+
0x81) [0x7c7971]
2016-02-09T14:17:25.522 INFO:tasks.ceph.mds.a.mira035.stderr: 7: (Finisher::finisher_thread_entry(
)+0x1a0) [0x87f0e0]
2016-02-09T14:17:25.522 INFO:tasks.ceph.mds.a.mira035.stderr: 8: (()+0x8182) [0x7f1f3b4b2182]
2016-02-09T14:17:25.522 INFO:tasks.ceph.mds.a.mira035.stderr: 9: (clone()+0x6d) [0x7f1f39c2147d]
2016-02-09T14:17:25.522 INFO:tasks.ceph.mds.a.mira035.stderr: NOTE: a copy of the executable, or `
objdump -rdS <executable>>` is needed to interpret this.
```

History

#1 - 02/09/2016 11:59 PM - Yuri Weinstein

- Related to Bug #14697: mds: assert in SafeTimer while suiciding added

#2 - 02/11/2016 11:39 PM - Greg Farnum

- Related to deleted (Bug #14697: mds: assert in SafeTimer while suiciding)

#3 - 02/11/2016 11:48 PM - Greg Farnum

This one's odd. The problem in [#14697](#) is different; it's actually calling timer.shutdown() twice there. Here, that isn't happening. Moreover, the assert

it's hitting on line 143 in Thread::join() doesn't match the backtrace, which claims to be in Thread::detach().

But both instances are happening when the Journaler gets ENOSPC. I wonder if MDS::suicide() is getting invoked twice? The finisher is the first one that gets join()ed in that process.

#4 - 02/12/2016 12:17 AM - Greg Farnum

Nope, that's not it directly, we only propagate one at a time thanks to Journaler::handle_write_error().

This is apparently pretty easy to have recurring now? But the earliest failure I can find is http://pulpito.ovh.sepia.ceph.com:8081/teuthology-2016-02-07_17:00:02-fs-hammer---basic-openstack/19350/. The only commit between a passing run¹ and that one that's obviously part of the FS stuff is the fsx qa script, which I don't think caused it. ;) There is also 2817ffcf4e57f92551b86388681fc0fe70c386ec in ReplicatedPG which changes the full behavior some; I wonder if that's broken the semantics or activated a new path in the MDS which is causing issues?

[1]:http://pulpito.ovh.sepia.ceph.com:8081/teuthology-2016-01-31_17:00:01-fs-hammer---basic-openstack/14718/

#5 - 02/16/2016 05:31 PM - Yuri Weinstein

Tried on diff machines all failed in similar fashion:

http://pulpito.ceph.com/teuthology-2016-02-15_16:52:37-fs-hammer---basic-smithi/
http://pulpito.ceph.com/teuthology-2016-02-15_14:34:12-fs-hammer---basic-vps/
http://pulpito.ceph.com/teuthology-2016-02-15_13:41:15-fs-hammer---basic-mira/

#6 - 02/18/2016 01:11 AM - Greg Farnum

Well at least it's consistent. Can you also try commit:2c8e57934284dae0ae92d1aa0839a87092ec7c51 against smithi/mira? If that passes, a commit bisect should tell us which patch broke stuff.

#7 - 02/18/2016 10:31 PM - Yuri Weinstein

Greg the test passed on that commit.

http://pulpito.ceph.com/teuthology-2016-02-18_13:42:45-fs-wip-test-14716-2---basic-smithi/

Here is what I did for memory and make sure that it was good.

Git:

```
git clone https://github.com/ceph/ceph/
git checkout 2c8e57934284dae0ae92d1aa0839a87092ec7c51
git checkout -b wip-test-14716-2
git push origin wip-test-14716-2
```

Gitbuilders:

<http://gitbuilder.sepia.ceph.com/gitbuilder-ceph-rpm-centos7-amd64-basic/rebuild.cgi?log=2c8e57934284dae0ae92d1aa0839a87092ec7c51>

<http://gitbuilder.sepia.ceph.com/gitbuilder-ceph-deb-trusty-amd64-basic/rebuild.cgi?log=2c8e57934284dae0ae92d1aa0839a87092ec7c51>

teuthology:

```
filter="fs/recovery/{clusters/2-remote-clients.yaml debug/mds_client.yaml mounts/ceph-fuse.yaml tasks/mds-full.yaml}"
```

#8 - 02/19/2016 04:01 PM - Yuri Weinstein

Tests on 2817ffc4e57f92551b86388681fc0fe70c386ec in ReplicatedPG commit failed all in similar way => .

```
2016-02-18T17:23:08.607 INFO:tasks.ceph.mds.a.smithi051.stderr:2016-02-19 01:23:08.603875 7f10be93c700 -1 mds.0.journaler(rw) _finish_flush got (28) No space left on device
2016-02-18T17:23:08.607 INFO:tasks.ceph.mds.a.smithi051.stderr:2016-02-19 01:23:08.603894 7f10be93c700 -1 mds.0.journaler(rw) handle_write_error (28) No space left on device
2016-02-18T17:23:08.608 INFO:tasks.ceph.mds.a.smithi051.stderr:2016-02-19 01:23:08.603918 7f10be93c700 -1 mds.0.log unhandled error (28) No space left on device, shutting down...
2016-02-18T17:23:08.609 INFO:tasks.ceph.mds.a.smithi051.stderr:common/Thread.cc: In function 'int Thread::join(void**)' thread 7f10be93c700 time 2016-02-19 01:23:08.603946
2016-02-18T17:23:08.609 INFO:tasks.ceph.mds.a.smithi051.stderr:common/Thread.cc: 143: FAILED assert(status == 0)
2016-02-18T17:23:08.610 INFO:tasks.ceph.mds.a.smithi051.stderr: ceph version 0.94.5-243-g2817ffc (2817ffc4e57f92551b86388681fc0fe70c386ec)
2016-02-18T17:23:08.610 INFO:tasks.ceph.mds.a.smithi051.stderr: 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0x94c72b]
2016-02-18T17:23:08.610 INFO:tasks.ceph.mds.a.smithi051.stderr: 2: (Thread::detach()+0) [0x938730]
2016-02-18T17:23:08.611 INFO:tasks.ceph.mds.a.smithi051.stderr: 3: (Finisher::stop()+0x9d) [0x87e6cd]
2016-02-18T17:23:08.611 INFO:tasks.ceph.mds.a.smithi051.stderr: 4: (MDS::suicide()+0x85) [0x59fa35]
2016-02-18T17:23:08.611 INFO:tasks.ceph.mds.a.smithi051.stderr: 5: (C_MDL_WriteError::finish(int)+0x65) [0x7da085]
2016-02-18T17:23:08.612 INFO:tasks.ceph.mds.a.smithi051.stderr: 6: (MDSIOContextBase::complete(int)+0x81) [0x7c7971]
2016-02-18T17:23:08.612 INFO:tasks.ceph.mds.a.smithi051.stderr: 7: (Finisher::finisher_thread_entry()+0x1a0) [0x87f0e0]
2016-02-18T17:23:08.612 INFO:tasks.ceph.mds.a.smithi051.stderr: 8: (()+0x8182) [0x7f10c66c0182]
2016-02-18T17:23:08.612 INFO:tasks.ceph.mds.a.smithi051.stderr: 9: (clone()+0x6d) [0x7f10c4e2f47d]
2016-02-18T17:23:08.612 INFO:tasks.ceph.mds.a.smithi051.stderr: NOTE: a copy of the executable, or `objdump -r dS <executable>` is needed to interpret this.
```

http://pulpito.ceph.com/teuthology-2016-02-18_17:17:37-fs-wip-test-14716-3---basic-smithi/

http://pulpito.ceph.com/teuthology-2016-02-18_16:58:42-fs-wip-test-14716-3---basic-vps/

http://pulpito.ceph.com/teuthology-2016-02-18_16:55:48-fs-wip-test-14716-3---basic-vps/

#9 - 02/19/2016 04:49 PM - Yuri Weinstein

Next try commit 951339103d35bc8ee2de880f77aada40d15b592a

passed

http://pulpito.ceph.com/teuthology-2016-02-19_10:07:15-fs-wip-test-14716-4---basic-smithi/
http://pulpito.ceph.com/teuthology-2016-02-19_10:06:33-fs-wip-test-14716-4---basic-vps/
http://pulpito.ceph.com/teuthology-2016-02-19_10:41:01-fs-wip-test-14716-4---basic-mira/

#10 - 02/19/2016 07:39 PM - Yuri Weinstein

- Project changed from CephFS to Ceph

#11 - 02/19/2016 07:40 PM - Yuri Weinstein

- Project changed from Ceph to CephFS

corresponding issue [#14824](#)

#12 - 04/14/2016 10:28 PM - Greg Farnum

- Status changed from New to Won't Fix

This was a result of the OSD full handling changes, which got partly backported. I think the resolution we ended up at was "too bad"?

#13 - 08/19/2016 04:18 PM - Yuri Weinstein

Same in hammer 0.94.8

http://qa-proxy.ceph.com/teuthology/yuriw-2016-08-18_20:11:00-fs-master---basic-smithi/373246/teuthology.log