

bluestore - Bug #36099

ObjectStore/StoreTest.BluestoreRepairTest/2 fails with os/bluestore/BlueStore.cc: 5894: FAILED ceph_assert(_kv_only || mounted)

09/20/2018 01:27 PM - Sage Weil

Status:	Resolved	Start date:	09/20/2018
Priority:	Urgent	Due date:	
Assignee:	Kefu Chai	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Reviewed:	
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	mimic	Pull request ID:	
Regression:	No		
Severity:	3 - minor		

Description

```
-4> 2018-09-20 08:25:17.578 7efc2299cc80 2 bluestore(bluestore.test_temp_dir) _fsck 0 objects
, 0 of them sharded.
-3> 2018-09-20 08:25:17.578 7efc2299cc80 2 bluestore(bluestore.test_temp_dir) _fsck 0 extents
to 0 blobs, 0 spanning, 0 shared.
-2> 2018-09-20 08:25:17.578 7efc2299cc80 1 bluestore(bluestore.test_temp_dir) _fsck <<<FINISH
>>> with 22 errors, 0 repaired, 22 remaining in 0.514968 seconds
-1> 2018-09-20 08:25:17.583 7efc2299cc80 -1 /home/sage/src/ceph4/src/os/bluestore/BlueStore.cc
: In function 'virtual int BlueStore::umount()' thread 7efc2299cc80 time 2018-09-20 08:25:17.57905
9
/home/sage/src/ceph4/src/os/bluestore/BlueStore.cc: 5894: FAILED ceph_assert(_kv_only || mounted)

ceph version 14.0.0-3458-g20cfc0212a (20cfc0212a2df0ac787b639d27d141b6036432eb) nautilus (dev)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x158) [0x7efc17f38b66]
2: ((()+0x24ad40) [0x7efc17f38d40]
3: (BlueStore::umount()+0x575) [0x55e64ef4d2f5]
4: (StoreTestFixture::TearDown()+0x5b) [0x55e64ee0576b]
5: (void testing::internal::HandleExceptionsInMethodIfSupported<testing::Test, void>(testing::Tes
t*, void (testing::Test::*)(), char const*)+0x4a) [0x55e64f06d97a]
6: (testing::TestInfo::Run()+0x118) [0x55e64f0649b8]
7: (testing::TestCase::Run()+0xb5) [0x55e64f064a95]
```

bin/ceph_test_objectstore --gtest_filter=ObjectStore/StoreTest.BluestoreRepairTest/2 to reproduce

Related issues:

Copied to bluestore - Backport #36551: mimic: ObjectStore/StoreTest.Bluestore...

Resolved

History

#1 - 09/20/2018 01:28 PM - Sage Weil

/a/sage-2018-09-19_18:44:57-rados-wip-sage4-testing-2018-09-19-1054-distro-basic-smithi/3043638

#2 - 09/21/2018 04:29 AM - Kefu Chai

could be a regression introduced by <https://github.com/ceph/ceph/pull/22739>

#3 - 09/21/2018 10:13 AM - jianpeng ma

using thi command `./bin/ceph_test_objectstore --gtest_catch_exceptions=0 --debug-bluestore=20 --log-to-stderr=true --gtest_filter=*BluestoreRepairTest*` can 100% reprodece this bug.

Add message:

```
-64> 2018-09-22 01:14:42.595 7f2da41eeb40 0 bluestore(bluestore.test_temp_dir) _fsck key
0x7f77ffffffffff0aeb83'(!Object#201(dup))='0xffffffffffffffffffffff'o'
-63> 2018-09-22 01:14:42.595 7f2da41eeb40 -1 bluestore(bluestore.test_temp_dir) _fsck oid.shard_id:255 pgid.shard:255 pool:-1!=0
-62> 2018-09-22 01:14:42.595 7f2da41eeb40 -1 bluestore(bluestore.test_temp_dir) _fsck collection 555.0_head cnode(bits 0)
-61> 2018-09-22 01:14:42.595 7f2da41eeb40 -1 bluestore(bluestore.test_temp_dir) fsck error: stray object #-1:0aeb8328::Object 1(dup):head# not
owned by any collection
```

I think because we don't set pool.

And by git bisect. i found commit 0bd2546eaca72ed0122a9c2648df4bef05b0d5d2 cause this bug.

#4 - 09/21/2018 10:16 AM - jianpeng ma

Add in hobjec_t.h

```
/* Do not use when a particular hash function is needed */
explicit hobject_t(const object_t &o) :
    oid(o.oid), snap(o.snap), max(false), pool(PPOOL_META) {
    set_hash(std::hash<object_t>()(o));
}
```

POOL_META=-1.

In store_test.c

```
TEST_P(StoreTest, BluestoreRepairTest)
ghobject_t hoid(hobject_t(object_t("Object 1", CEPH_NOSNAP)));
ghobject_t hoid_dup(hobject_t(object_t("Object 1(dup)", CEPH_NOSNAP)));
ghobject_t hoid2(hobject_t(object_t("Object 2", CEPH_NOSNAP)));
ghobject_t hoid_cloned = hoid2;
hoid_cloned.hobj.snap = 1;
ghobject_t hoid3(hobject_t(object_t("Object 3", CEPH_NOSNAP)));
```

This make pool=POOL_META.

But i'm not sure why this bug now occur? Or am i missing something?

#5 - 09/21/2018 12:00 PM - Kefu Chai

- Assignee set to Kefu Chai

Jianpeng, thanks for the analysis. i don't think you missed anything. it's just good timing =)

#6 - 09/21/2018 12:01 PM - Kefu Chai

- Status changed from Verified to Need Review

<https://github.com/ceph/ceph/pull/24220>

#7 - 09/22/2018 04:25 PM - Kefu Chai

- Status changed from Need Review to Resolved

#8 - 10/23/2018 05:32 AM - Nathan Cutler

- Backport set to mimic

#9 - 10/23/2018 05:32 AM - Nathan Cutler

- Status changed from Resolved to Pending Backport

#10 - 10/23/2018 05:32 AM - Nathan Cutler

- Copied to Backport #36551: mimic: ObjectStore/StoreTest.BluestoreRepairTest/2 fails with os/bluestore/BlueStore.cc: 5894: FAILED
ceph_assert(_kv_only || mounted) added

#11 - 10/23/2018 05:42 AM - Nathan Cutler

- Status changed from Pending Backport to Resolved