

rbd - Bug #14510

TestLibRBD.ObjectMapConsistentSnap librbd/Journal.cc: 376: FAILED assert(m_state == STATE_READY)

01/26/2016 09:57 AM - Mykola Golub

Status:	Resolved	Start date:	01/26/2016
Priority:	Normal	Due date:	
Assignee:	Jason Dillaman	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	other	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

I saw this in a jenkins test: <https://jenkins.ceph.com/job/ceph-pull-requests/545/console>

And running the test manually in loop for about 10 minutes:

```
% while RBD_FEATURES=109 ./unittest_librbd --gtest_filter=TestLibRBD.ObjectMapConsistentSnap; do sleep 0; done
...
```

```
[ RUN ] TestLibRBD.ObjectMapConsistentSnap
using new format!
```

```
./librbd/Journal.cc: In function 'void librbd::Journal< <template-parameter-1-1> >::append_op_event(uint64_t, librbd::journal::EventEntry&&, Context*) [with ImageCtxT = librbd::ImageCtx; uint64_t = long unsigned int]' thread 7f138b5ab700 time 2016-01-26 11:41:23.472702
```

```
./librbd/Journal.cc: 376: FAILED assert(m_state == STATE_READY)
```

```
ceph version 10.0.2-1074-g831cc7a (831cc7a7a01f6e331f0f7f74e6f0d1e0fe7ff070)
```

```
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x82) [0x7f13a1699ab2]
```

```
2: (librbd::Journal<librbd::ImageCtx>::append_op_event(unsigned long, librbd::journal::EventEntry&&, Context*)+0x33b) [0x7f13a1354ffb]
```

```
3: (librbd::operation::Request<librbd::ImageCtx>::append_op_event(Context*)+0x213) [0x7f13a12ebdf3]
```

```
4: (librbd::operation::SnapshotCreateRequest<librbd::ImageCtx>::send_append_op_event()+0x12b) [0x7f13a1401b5b]
```

```
5: (librbd::operation::SnapshotCreateRequest<librbd::ImageCtx>::handle_suspend_aio(int*)+0x268) [0x7f13a1401f58]
```

```
6: (librbd::util::detail::C_StateCallbackAdapter<librbd::operation::SnapshotCreateRequest<librbd::ImageCtx>, &librbd::operation::SnapshotCreateRequest<librbd::ImageCtx>::handle_suspend_aio, true>::complete(int)+0x1a) [0x7f13a1401fea]
```

```
7: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa9f) [0x7f13a168a7df]
```

```
8: (ThreadPool::WorkThread::entry()+0x10) [0x7f13a168b700]
```

```
9: ((+0x80a4) [0x7f13973fa0a4]
```

```
10: (clone()+0x6d) [0x7f1395b5b04d]
```

```
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

History

#1 - 01/26/2016 10:29 AM - Mykola Golub

Below is a debug log for the crash. librbd::SnapshotCreateRequest may be called when the journal has not finished the initialization yet.

```
2016-01-26 12:07:43.873044 7fe7f1ce3100 0 lockdep start
```

```
2016-01-26 12:07:43.874851 7fe7f1ce3100 -1 WARNING: the following dangerous and experimental features are enabled: *
```

```
2016-01-26 12:07:43.876436 7fe7f1ce3100 -1 WARNING: the following dangerous and experimental features are enabled: *
```

```
2016-01-26 12:07:43.876512 7fe7f1ce3100 20 librbd: create 0x7fff933492b0 name = image1 size = 1048576 old_format = 0 features = 109 order = 12 stripe_unit = 0 stripe_count = 0
```

```
2016-01-26 12:07:43.876770 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_hello.so: undefined symbol: _ZTI10PGLSFilter
```

```
2016-01-26 12:07:43.876955 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_user.so: undefined symbol: _Z23cls_cxx_map_read_headerPvPN4ceph6buffer4listE
```

```
2016-01-26 12:07:43.877136 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_version.so: undefined symbol: _Z19cls_gen_rand_base64Pci
```

```
2016-01-26 12:07:43.877178 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_log.so: undefined symbol: _Z23cls_cxx_map_read_headerPvPN4ceph6buffer4listE
```

2016-01-26 12:07:43.878402 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_rgw.so: undefined symbol: _Z23cls_cxx_map_read_headerPvPN4ceph6buffer4listE
2016-01-26 12:07:43.878542 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_cephfs.so: undefined symbol: _ZTI10PGLSFilter
2016-01-26 12:07:43.879710 7fe7f1ce3100 -1 Failed to load class: .libs/libcls_refcount.so: undefined symbol: _Z14cls_cxx_removePv
2016-01-26 12:07:43.879775 7fe7f1ce3100 2 librbd: adding rbd image to directory...
2016-01-26 12:07:43.879902 7fe7f1ce3100 5 librbd::Journal: create: image=06a785fd
2016-01-26 12:07:43.879919 7fe7f1ce3100 5 Journaler: creating new journal: journal.06a785fd
2016-01-26 12:07:43.879955 7fe7f1ce3100 10 JournalMetadata: register_client:
2016-01-26 12:07:43.879985 7fe7f1ce3100 10 JournalMetadata: notifying journal header update
2016-01-26 12:07:43.888854 7fe7f1ce3100 2 librbd: done.
2016-01-26 12:07:43.888999 7fe7f1ce3100 5 librbd::AioImageRequestWQ: 0x7fe7fcc19010 : ictx=0x7fe7fcc1f310
2016-01-26 12:07:43.889021 7fe7f1ce3100 20 librbd::ImageState: open
2016-01-26 12:07:43.889023 7fe7f1ce3100 10 librbd::ImageState: 0x7fe7fcc1fd00 send_open
2016-01-26 12:07:43.889024 7fe7f1ce3100 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_v2_detect_header
2016-01-26 12:07:43.889096 7fe7e10fa700 10 librbd::image::OpenRequest: handle_v2_detect_header: r=0
2016-01-26 12:07:43.889101 7fe7e10fa700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_v2_get_id
2016-01-26 12:07:43.889147 7fe7e10fa700 10 librbd::image::OpenRequest: handle_v2_get_id: r=0
2016-01-26 12:07:43.889150 7fe7e10fa700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_v2_get_immutable_metadata
2016-01-26 12:07:43.889192 7fe7e10fa700 10 librbd::image::OpenRequest: handle_v2_get_immutable_metadata: r=0
2016-01-26 12:07:43.889194 7fe7e10fa700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_v2_get_stripe_unit_count
2016-01-26 12:07:43.889225 7fe7e10fa700 10 librbd::image::OpenRequest: handle_v2_get_stripe_unit_count: r=-8
2016-01-26 12:07:43.889231 7fe7e10fa700 10 librbd::ImageCtx: init_layout stripe_unit 4096 stripe_count 1 object_size 4096 prefix rbd_data.06a785fd format rbd_data.06a785fd.%016llx
2016-01-26 12:07:43.889233 7fe7e10fa700 20 librbd::ImageCtx: apply_metadata_confs
2016-01-26 12:07:43.889307 7fe7e10fa700 20 librbd::ImageCtx: enabling caching...
2016-01-26 12:07:43.889310 7fe7e10fa700 20 librbd::ImageCtx: Initial cache settings: size=33554432 num_objects=10 max_dirty=0 target_dirty=16777216 max_dirty_age=1
2016-01-26 12:07:43.889358 7fe7e10fa700 10 librbd::ImageCtx: cache bytes 33554432 -> about 1048 objects
2016-01-26 12:07:43.889388 7fe7e10fa700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_register_watch
2016-01-26 12:07:43.889447 7fe7db7d6700 10 librbd::ImageWatcher: 0x7fe7c4000d90 registering image watcher
2016-01-26 12:07:43.889471 7fe7db7d6700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 handle_register_watch: r=0
2016-01-26 12:07:43.889472 7fe7db7d6700 10 librbd::image::OpenRequest: 0x7fe7fcc19a10 send_refresh
2016-01-26 12:07:43.889473 7fe7db7d6700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 send_v2_get_mutable_metadata
2016-01-26 12:07:43.889532 7fe7e10fa700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 handle_v2_get_mutable_metadata: r=0
2016-01-26 12:07:43.889535 7fe7e10fa700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 send_v2_get_flags
2016-01-26 12:07:43.889555 7fe7e10fa700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 handle_v2_get_flags: r=0
2016-01-26 12:07:43.889562 7fe7e10fa700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 send_v2_init_exclusive_lock
2016-01-26 12:07:43.889567 7fe7e10fa700 10 librbd::ExclusiveLock: 0x7fe7bc003250 init
2016-01-26 12:07:43.889575 7fe7e10fa700 5 librbd::AioImageRequestWQ: block_writes: 0x7fe7fcc1f310, num=1
2016-01-26 12:07:43.889595 7fe7db7d6700 10 librbd::ExclusiveLock: 0x7fe7bc003250 handle_init_complete
2016-01-26 12:07:43.889598 7fe7db7d6700 10 librbd::image::RefreshRequest: 0x7fe7c4002410 handle_v2_init_exclusive_lock: r=0
2016-01-26 12:07:43.889599 7fe7db7d6700 20 librbd::image::RefreshRequest: 0x7fe7c4002410 apply
2016-01-26 12:07:43.889640 7fe7db7d6700 10 librbd::image::OpenRequest: handle_refresh: r=0
2016-01-26 12:07:43.889641 7fe7db7d6700 10 librbd::ImageState: 0x7fe7fcc1fd00 handle_open: r=0
2016-01-26 12:07:43.889815 7fe7d8fd1700 20 librbd: info 0x7fe7fcc1f310
2016-01-26 12:07:43.889818 7fe7f1ce3100 5 librbd::Operations: 0x7fe7fcb15d0 snap_create: snap_name=snap0
2016-01-26 12:07:43.889970 7fe7d8fd1700 20 librbd::AioImageRequestWQ: write: ictx=0x7fe7fcc1f310, off=0, len = 3
2016-01-26 12:07:43.889976 7fe7d8fd1700 20 librbd::AioImageRequestWQ: aio_write: ictx=0x7fe7fcc1f310, completion=0x7fe7b0000d40, off=0, len=3, flags=0
2016-01-26 12:07:43.889978 7fe7d8fd1700 20 librbd::AioImageRequestWQ: queue: ictx=0x7fe7fcc1f310, req=0x7fe7b0000b60
2016-01-26 12:07:43.890028 7fe7d8fd1700 20 librbd::ExclusiveLock: 0x7fe7bc003250 is_lock_owner=0
2016-01-26 12:07:43.890031 7fe7d8fd1700 10 librbd::ExclusiveLock: 0x7fe7bc003250 request_lock
2016-01-26 12:07:43.890032 7fe7d8fd1700 10 librbd::ExclusiveLock: 0x7fe7bc003250 send_acquire_lock
2016-01-26 12:07:43.890073 7fe7d8fd1700 10 librbd::exclusive_lock::AcquireRequest: send_lock
2016-01-26 12:07:43.890118 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: handle_lock: r=0
2016-01-26 12:07:43.890120 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: send_open_object_map
2016-01-26 12:07:43.890123 7fe7e10fa700 10 librbd::object_map::RefreshRequest: 0x7fe7bc0039b0 send_load: oid=rbd_object_map.06a785fd
2016-01-26 12:07:43.890193 7fe7e10fa700 10 librbd::object_map::RefreshRequest: 0x7fe7bc0039b0 handle_load: r=0
2016-01-26 12:07:43.890201 7fe7e10fa700 20 librbd::object_map::RefreshRequest: refreshed object map: num_objs=256
2016-01-26 12:07:43.890204 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: handle_open_object_map: r=0
2016-01-26 12:07:43.890205 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: send_lock_object_map
2016-01-26 12:07:43.890206 7fe7e10fa700 10 librbd::object_map::LockRequest: 0x7fe7bc004550 send_lock: oid=rbd_object_map.06a785fd
2016-01-26 12:07:43.890331 7fe7e10fa700 10 librbd::object_map::LockRequest: 0x7fe7bc004550 handle_lock: r=0
2016-01-26 12:07:43.890334 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: handle_lock_object_map: r=0
2016-01-26 12:07:43.890335 7fe7e10fa700 10 librbd::ExclusiveLock: 0x7fe7bc003250 handle_acquiring_lock
2016-01-26 12:07:43.890337 7fe7e10fa700 10 librbd::exclusive_lock::AcquireRequest: send_open_journal
2016-01-26 12:07:43.890345 7fe7e10fa700 5 librbd::Journal: 0x7fe7bc005050: ictx=0x7fe7fcc1f310
2016-01-26 12:07:43.890346 7fe7e10fa700 20 librbd::Journal: 0x7fe7bc005050 open
2016-01-26 12:07:43.890348 7fe7e10fa700 20 librbd::Journal: 0x7fe7bc005050 wait_for_steady_state: on_state=0x7fe7bc0042c0
2016-01-26 12:07:43.890348 7fe7e10fa700 20 librbd::Journal: 0x7fe7bc005050 create_journaler
2016-01-26 12:07:43.890349 7fe7e10fa700 20 librbd::Journal: 0x7fe7bc005050 transition_state: new state=1
2016-01-26 12:07:43.890458 7fe7f1ce3100 20 librbd::ExclusiveLock: 0x7fe7bc003250 is_lock_owner=1
2016-01-26 12:07:43.890462 7fe7f1ce3100 20 librbd::ExclusiveLock: 0x7fe7bc003250 is_lock_owner=1
2016-01-26 12:07:43.890464 7fe7f1ce3100 20 librbd::ExclusiveLock: 0x7fe7bc003250 is_lock_owner=1
2016-01-26 12:07:43.890465 7fe7f1ce3100 5 librbd::Operations: 0x7fe7fcb15d0 snap_create: snap_name=snap0

2016-01-26 12:07:43.890471 7fe7e10fa700 10 JournalMetadata: initialized immutable metadata
2016-01-26 12:07:43.890471 7fe7f1ce3100 5 librbd::SnapshotCreateRequest: 0x7fe7fcc1b640 send_suspend_requests
2016-01-26 12:07:43.890472 7fe7e10fa700 10 JournalMetadata: refreshing mutable metadata
2016-01-26 12:07:43.890472 7fe7f1ce3100 5 librbd::SnapshotCreateRequest: 0x7fe7fcc1b640 send_suspend_aio
2016-01-26 12:07:43.890474 7fe7f1ce3100 5 librbd::AioImageRequestWQ: block_writes: 0x7fe7fcc1f310, num=2
2016-01-26 12:07:43.890513 7fe7e10fa700 10 JournalMetadata: refreshed mutable metadata: r=0
2016-01-26 12:07:43.890515 7fe7e10fa700 20 Journaler: using image pool for journal data
2016-01-26 12:07:43.890617 7fe7db7d6700 5 librbd::SnapshotCreateRequest: 0x7fe7fcc1b640 handle_suspend_aio: r=0
2016-01-26 12:07:43.890678 7fe7db7d6700 10 librbd::Request: 0x7fe7fcc1b640 append_op_event

#2 - 01/26/2016 02:56 PM - Jason Dillaman

I think you have an old version of code -- "assert(m_state == STATE_READY)" doesn't exist in the codebase.

#3 - 01/26/2016 02:57 PM - Jason Dillaman

- Status changed from New to Need More Info

#4 - 01/26/2016 03:18 PM - Mykola Golub

I see it in today master: <https://github.com/ceph/ceph/blob/master/src/librbd/Journal.cc#L376>

#5 - 01/27/2016 05:58 PM - Jason Dillaman

Whoops -- I was thinking this was a different issue I just fixed with Journal::handle_op_event_safe.

#6 - 01/27/2016 05:58 PM - Jason Dillaman

- Status changed from Need More Info to In Progress

- Assignee set to Jason Dillaman

#7 - 01/27/2016 07:17 PM - Jason Dillaman

- Status changed from In Progress to Need Review

master PR: <https://github.com/ceph/ceph/pull/7382>

#8 - 01/28/2016 02:11 AM - Jason Dillaman

- Status changed from Need Review to Resolved