

Ceph - Bug #8011

osd/ReplicatedPG.cc: 5244: FAILED assert(soid < scrubber.start || soid >= scrubber.end)

04/07/2014 10:38 AM - Samuel Just

Status:	Resolved	Start date:	04/07/2014
Priority:	Urgent	Due date:	
Assignee:	Samuel Just	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Q/A	Reviewed:	
Tags:		Affected Versions:	
Backport:	giant,firefly	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	2 - major		

Description

osd/ReplicatedPG.cc: 5244: FAILED assert(soid < scrubber.start || soid >= scrubber.end)

ceph version 0.78-600-g19f50b9 (19f50b9d7bbbb2cce3b599f3ed8a9fa32c3d4e53)

```
1: (ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool)+0x1c86) [0x7f8976]
2: (ReplicatedPG::try_flush_mark_clean(boost::shared_ptr<ReplicatedPG::FlushOp>)+0x72f) [0x7fa5ff]
3: (ReplicatedPG::finish_flush(hobject_t, unsigned long, int)+0x2da) [0x7fb0ea]
4: (C_Flush::finish(int)+0xa7) [0x856e77]
5: (Context::complete(int)+0x9) [0x66ed59]
6: (Finisher::finisher_thread_entry()+0x1c0) [0x9a7050]
7: (()+0x7e9a) [0x7f14f9c60e9a]
8: (clone()+0x6d) [0x7f14f82213fd]
```

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

Related issues:

Related to Ceph - Bug #10689: osd/ReplicatedPG.cc: FAILED assert(soid < scrub...	Duplicate	01/29/2015
Related to Ceph - Bug #10693: FAILED assert(soid < scrubber.start soid >= ...	Resolved	01/29/2015
Duplicated by Ceph - Bug #8747: OSD crash on scrub:osd/ReplicatedPG.cc: 5297:...	Duplicate	07/04/2014
Duplicated by Ceph - Bug #10433: OSD osd/ReplicatedPG.cc: 5540: FAILED assert...	Duplicate	12/26/2014

Associated revisions

Revision e66f2e36 - 04/30/2014 10:11 PM - Samuel Just

ReplicatedPG: block scrub on blocked object contexts

Fixes: #8011

Signed-off-by: Samuel Just <sam.just@inktank.com>

Reviewed-by: Sage Weil <sage@inktank.com>

Revision 0f3235d4 - 04/30/2014 10:13 PM - Samuel Just

ReplicatedPG: block scrub on blocked object contexts

Fixes: #8011

Signed-off-by: Samuel Just <sam.just@inktank.com>

Reviewed-by: Sage Weil <sage@inktank.com>

(cherry picked from commit e66f2e36c06ca00c1147f922d3513f56b122a5c0)

Revision db4ccb04 - 05/08/2014 07:41 PM - Samuel Just

ReplicatedPG: block scrub on blocked object contexts

Fixes: #8011

Backport: firefly

Signed-off-by: Samuel Just <sam.just@inktank.com>

Revision 74114771 - 05/12/2014 10:31 PM - Samuel Just

ReplicatedPG: block scrub on blocked object contexts

Fixes: #8011

Backport: firefly

Signed-off-by: Samuel Just <sam.just@inktank.com>

Revision 29ee6fae - 05/19/2014 08:55 PM - Samuel Just

ReplicatedPG: block scrub on blocked object contexts

Fixes: #8011

Backport: firefly

Signed-off-by: Samuel Just <sam.just@inktank.com>

(cherry picked from commit 7411477153219d66625a74c5886530029c516036)

Revision 9b26de3f - 12/11/2014 09:30 PM - Samuel Just

ReplicatedPG: fail a non-blocking flush if the object is being scrubbed

Fixes: #8011

Backport: firefly, giant

Signed-off-by: Samuel Just <sjust@redhat.com>

Revision f8567398 - 03/11/2015 12:45 AM - Samuel Just

ReplicatedPG: fail a non-blocking flush if the object is being scrubbed

Fixes: #8011

Backport: firefly, giant

Signed-off-by: Samuel Just <sjust@redhat.com>

(cherry picked from commit 9b26de3f3653d38dcdfc5b97874089f19d2a59d7)

Revision 681c99fe - 03/17/2015 06:03 PM - Samuel Just

ReplicatedPG: fail a non-blocking flush if the object is being scrubbed

Fixes: #8011

Backport: firefly, giant

Signed-off-by: Samuel Just <sjust@redhat.com>

(cherry picked from commit 9b26de3f3653d38dcdfc5b97874089f19d2a59d7)

History

#1 - 04/08/2014 10:12 PM - Sage Weil

ubuntu@teuthology:/var/lib/teuthworker/archive/sage-2014-04-08_14:01:14-rados:thrash-wip-7891-testing-basic-plana/178972

#2 - 04/13/2014 09:42 PM - Sage Weil

ubuntu@teuthology:/var/lib/teuthworker/archive/sage-2014-04-13_09:43:35-rados:thrash-testing-testing-basic-plana/189166

#3 - 04/13/2014 09:47 PM - Sage Weil

ubuntu@teuthology:/var/lib/teuthworker/archive/sage-2014-04-13_09:43:35-rados:thrash-testing-testing-basic-plana/189348

#4 - 04/17/2014 01:56 PM - Samuel Just

- Status changed from New to In Progress

- Assignee set to Samuel Just

#5 - 04/17/2014 02:00 PM - Samuel Just

ReplicatedPG::do_op already does the right thing as far as blocking ops which may flush. What remains is to avoid flushing objects with blocked obcs.

#6 - 04/17/2014 02:00 PM - Samuel Just

and to check that agent_work also does the right thing

#7 - 04/17/2014 05:35 PM - Samuel Just

- Status changed from In Progress to Testing

#8 - 04/19/2014 08:15 AM - Sage Weil

- Source changed from other to Q/A

ubuntu@teuthology:/var/lib/teuthworker/archive/sage-2014-04-18_21:29:10-rados:thrash-testing-testing-basic-plana/202157

#9 - 05/01/2014 10:14 AM - Samuel Just

- Status changed from Testing to Resolved

#10 - 05/02/2014 04:06 PM - Sage Weil

- Status changed from Resolved to Verified

this triggered again on [c6ada53a146f3196e11f545cfc968fc21657aec6](#)

```
O> 2014-05-02 11:26:12.053553 7fd1ee1b2700 -1 osd/ReplicatedPG.cc: In function 'void ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool)' thread 7fd1ee1b2700 time 2014-05-02 11:26:12.039772
osd/ReplicatedPG.cc: 5282: FAILED assert(soid < scrubber.start || soid >= scrubber.end)
```

ubuntu@teuthology:/var/lib/teuthworker/archive/teuthology-2014-05-02_02:30:10-rados-master-testing-basic-plana/229437

#11 - 05/08/2014 01:36 PM - Samuel Just

- Status changed from Verified to Need Review

#12 - 05/08/2014 01:49 PM - Samuel Just

- Status changed from *Need Review* to *Pending Backport*

#13 - 05/20/2014 01:12 PM - Samuel Just

- Status changed from *Pending Backport* to *Resolved*

#14 - 07/15/2014 01:36 PM - Sage Weil

- Status changed from *Resolved* to *Verified*

see [#8747](#) for a log of this happening on 0.80.3

#15 - 07/15/2014 01:37 PM - Sage Weil

- Assignee deleted (*Samuel Just*)

#16 - 09/19/2014 04:43 PM - Sage Weil

- Status changed from *Verified* to *Can't reproduce*

Pinged Dmitry to see if he is still seeing this or has a log

#17 - 10/02/2014 05:34 AM - Dmitry Smirnov

- Status changed from *Can't reproduce* to *Resolved*

I'm unable to reproduce it any more, assuming fixed.

#18 - 11/20/2014 08:38 AM - Sage Weil

- Status changed from *Resolved* to *Verified*

this popped up again: `ubuntu@teuthology:/var/lib/teuthworker/archive/teuthology-2014-11-17_02:32:01-rados-giant-distro-basic-multi/604740`

#19 - 11/20/2014 01:56 PM - Samuel Just

- Assignee set to *Samuel Just*

#20 - 11/20/2014 02:20 PM - Samuel Just

Urgh, non-blocking flushes do not cause scrub to pause. I think the simplest solution is to fail a non-blocking scrub in `try_flush_mark_clean` if the object is being scrubbed.

#21 - 11/20/2014 02:34 PM - Samuel Just

- Status changed from *Verified* to *Testing*

#22 - 12/12/2014 02:24 PM - Samuel Just

- Status changed from *Testing* to *Pending Backport*

- Backport set to *giant,firefly*

#23 - 01/27/2015 04:35 PM - Sage Weil

- Status changed from *Pending Backport* to *Verified*

happened again, i believe with the latest fix applied.

`ubuntu@teuthology:/a/teuthology-2015-01-25_23:10:02-knfs-next-testing-basic-multi/722944`

#24 - 01/28/2015 04:43 PM - Yuri Weinstein

Also see in run: http://pulpito.ceph.com/teuthology-2015-01-27_17:13:01-upgrade:firefly-x-next-distro-basic-multi/

Job: [726068]

Logs: http://qa-proxy.ceph.com/teuthology/teuthology-2015-01-27_17:13:01-upgrade:firefly-x-next-distro-basic-multi/726068/teuthology.log

```
2015-01-28T07:11:26.339 INFO:tasks.ceph.osd.0.burnupil6.stderr:osd/ReplicatedPG.cc: In function 'void ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool, bool)' thread 7fd534fd4700 time 2015-01-28 07:11:26.502343
2015-01-28T07:11:26.339 INFO:tasks.ceph.osd.0.burnupil6.stderr:osd/ReplicatedPG.cc: 5943: FAILED assert(soid < scrubber.start || soid >= scrubber.end)
2015-01-28T07:11:26.339 INFO:tasks.ceph.osd.0.burnupil6.stderr: ceph version 0.91-388-g5064787 (50647876971a2fe65a02e4de3c0bc62fec4887c4)
2015-01-28T07:11:26.339 INFO:tasks.ceph.osd.0.burnupil6.stderr: 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x7f) [0xadb3df]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 2: (ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool, bool)+0x1bd3) [0x826f93]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 3: (ReplicatedPG::prepare_transaction(ReplicatedPG::OpContext*)+0x16d) [0x845d6d]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 4: (ReplicatedPG::execute_ctx(ReplicatedPG::OpContext*)+0xa13) [0x846963]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 5: (ReplicatedPG::do_op(std::tr1::shared_ptr<OpRequest>&)+0x2987) [0x8519a7]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 6: (ReplicatedPG::do_request(std::tr1::shared_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x63f) [0x7ea05f]
2015-01-28T07:11:26.340 INFO:tasks.ceph.osd.0.burnupil6.stderr: 7: (OSD::dequeue_op(boost::intrusive_ptr<PG>, std::tr1::shared_ptr<OpRequest>, ThreadPool::TPHandle&)+0x17f) [0x6647ef]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: 8: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x65f) [0x66524f]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: 9: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x65c) [0xacab4c]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: 10: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0xacd720]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: 11: (()+0x7e9a) [0x7fd54fe2de9a]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: 12: (clone()+0x6d) [0x7fd54e5d8ccd]
2015-01-28T07:11:26.341 INFO:tasks.ceph.osd.0.burnupil6.stderr: NOTE: a copy of the executable, or `objdump -r ds <executable>` is needed to interpret this.
```

#25 - 01/29/2015 12:49 PM - Irek Fasikhov

Yes, I have the same error.

```
[root@ceph04 ceph]# ceph -v
ceph version 0.80.8 (69eaad7f8308f21573c604f121956e64679a52a7)
```

Fix will be in version 0.80.9? Thanks

```
2015-01-28 09:50:11.082466 7fbc1437b700 -1 osd/ReplicatedPG.cc: In function 'void ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool)' thread 7fbc1437b700 time 2015-01-28 09:50:10.852829
osd/ReplicatedPG.cc: 5318: FAILED assert(soid < scrubber.start || soid >= scrubber.end)
```

```
ceph version 0.80.8 (69eaad7f8308f21573c604f121956e64679a52a7)
1: (ReplicatedPG::finish_ctx(ReplicatedPG::OpContext*, int, bool)+0x32f5) [0x889275]
2: (ReplicatedPG::finish_promote(int, std::tr1::shared_ptr<OpRequest>, ReplicatedPG::CopyResults*, std::tr1::shared_ptr<ObjectContext>)+0x110f) [0x8903ef]
3: (PromoteCallback::finish(boost::tuples::tuple<int, ReplicatedPG::CopyResults*, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type>)+0x78) [0x8e29b8]
4: (GenContext<boost::tuples::tuple<int, ReplicatedPG::CopyResults*, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type> >::complete(boost::tuples::tuple<int, ReplicatedPG::CopyResults*, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type, boost::tuples::null_type>)+0x15) [0x8b34f5]
5: (ReplicatedPG::process_copy_chunk(hobject_t, unsigned long, int)+0x747) [0x885407]
6: (C_Copyfrom::finish(int)+0xb7) [0x8e2777]
7: (Context::complete(int)+0x9) [0x667209]
8: (Finisher::finisher_thread_entry()+0x1d8) [0x9ed148]
9: (()+0x79d1) [0x7fbc36e5f9d1]
10: (clone()+0x6d) [0x7fbc35dd88fd]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

```
--- begin dump of recent events ---
```

#26 - 01/29/2015 09:48 PM - Samuel Just

Irek, you are seeing the previous incarnation of this bug. The relevant fix has not yet been backported to firefly.

#27 - 01/29/2015 09:57 PM - Samuel Just

This most recent incarnation is due to oi digests: we block_writes until the end of COMPARE_MAPS. The assert should not fire if !block_writes. Previously, we would always change scrubber.start to be the same as scrubber.end as we changed block_writes.

#28 - 01/29/2015 09:57 PM - Samuel Just

Making patch.

#29 - 01/29/2015 09:58 PM - Samuel Just

Also, this is an entirely distinct bug, so I'm making this Pending Backport again and opening a new one.

#30 - 01/29/2015 10:06 PM - Samuel Just

- Status changed from Verified to Pending Backport

#31 - 01/30/2015 07:38 AM - Irek Fasikhov

Samuel.

yes, this is another mistake. <http://tracker.ceph.com/issues/10433#note-3>

#32 - 03/10/2015 07:57 AM - Loic Dachary

- Severity changed from 3 - minor to 2 - major

#33 - 03/11/2015 12:47 AM - Loic Dachary

- firefly backport <https://github.com/ceph/ceph/pull/3943>

#34 - 03/17/2015 06:04 PM - Loic Dachary

- giant backport <https://github.com/ceph/ceph/pull/4053>

#35 - 03/26/2015 12:11 PM - Loic Dachary

- Status changed from Pending Backport to Resolved