

crimson - Bug #47212

out-of-order "Error: finished tid 3 when last_acked_tid was 5"

08/31/2020 11:16 AM - Kefu Chai

Status:	Pending Backport	% Done:	0%
Priority:	Normal		
Assignee:	Kefu Chai		
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	36919
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

```
2020-08-31T08:46:39.826 INFO:tasks.rados.rados.0.smithi116.stdout:191: writing smithi11620526-191
from 793330 to 1483112 tid 1
2020-08-31T08:46:39.834 INFO:tasks.rados.rados.0.smithi116.stdout:191: writing smithi11620526-191
from 2035461 to 2653192 tid 2
2020-08-31T08:46:39.835 INFO:tasks.rados.rados.0.smithi116.stdout:191: writing smithi11620526-191
from 3101545 to 3193330 tid 3
...
2020-08-31T08:46:40.860 INFO:tasks.rados.rados.0.smithi116.stdout:191: finishing write tid 1 to s
mithi11620526-191
2020-08-31T08:46:40.920 INFO:tasks.rados.rados.0.smithi116.stdout:191: finishing write tid 2 to s
mithi11620526-191
2020-08-31T08:46:41.480 INFO:tasks.rados.rados.0.smithi116.stdout:191: finishing write tid 3 to s
mithi11620526-191
2020-08-31T08:46:41.480 INFO:tasks.rados.rados.0.smithi116.stderr:Error: finished tid 3 when last_
acked_tid was 5
2020-08-31T08:46:41.481 INFO:tasks.rados.rados.0.smithi116.stderr:/home/jenkins-build/build/worksp
ace/ceph-dev-new-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos8/DIST/centos8/MACHI
NE_SIZE/
gigantic/release/16.0.0-4903-g268a7471a80/rpm/el8/BUILD/ceph-16.0.0-4903-g268a7471a80/src/test/osd
/RadosModel.h: In function 'virtual void WriteOp::_finish(TestOp::CallbackInfo*)' thread 7f4cd37fe
700 time
2020-08-31T08:46:41.474872+0000
2020-08-31T08:46:41.481 INFO:tasks.rados.rados.0.smithi116.stderr:/home/jenkins-build/build/worksp
ace/ceph-dev-new-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos8/DIST/centos8/MACHI
NE_SIZE/
gigantic/release/16.0.0-4903-g268a7471a80/rpm/el8/BUILD/ceph-16.0.0-4903-g268a7471a80/src/test/osd
/RadosModel.h: 902: ceph_abort_msg("abort() called")
2020-08-31T08:46:41.481 INFO:tasks.rados.rados.0.smithi116.stderr: ceph version 16.0.0-4903-g268a7
471a80 (268a7471a8046235899274beb3fc7ebd8e2e731) pacific (dev)
2020-08-31T08:46:41.481 INFO:tasks.rados.rados.0.smithi116.stderr: 1: (ceph::__ceph_abort(char con
st*, int, char const*, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<cha
r> > con
st&)+0xe0) [0x7f4ce700873c]
2020-08-31T08:46:41.482 INFO:tasks.rados.rados.0.smithi116.stderr: 2: (WriteOp::_finish(TestOp::Ca
llbackInfo*)+0x6e2) [0x55be6ef68ef2]
2020-08-31T08:46:41.482 INFO:tasks.rados.rados.0.smithi116.stderr: 3: (write_callback(void*, void*
)+0x1e) [0x55be6ef88b8e]
```

History

#1 - 08/31/2020 11:35 AM - Kefu Chai

- Description updated

#2 - 08/31/2020 11:37 AM - Kefu Chai

```
$ zgrep smithi1620526-191 remote/*/log/ceph-osd* | grep do_osd_ops|grep reply
remote/smithi186/log/ceph-osd.2.log.gz:DEBUG 2020-08-31 08:46:39,896 [shard 0] osd - do_osd_ops: osd_op(client
.4228.0:835 3.e 3:7d6f3947::smithi1620526-191:head {write 793330~689782 in=689782b, stat out=16b} snapc 0={}
ondisk+write+known_if_redirected e16) v8 - object 3:7d6f3947::smithi1620526-191:head sending reply
remote/smithi186/log/ceph-osd.2.log.gz:DEBUG 2020-08-31 08:46:39,938 [shard 0] osd - do_osd_ops: osd_op(client
.4228.0:836 3.e 3:7d6f3947::smithi1620526-191:head {write 2035461~617731 in=617731b} snapc 0={} ondisk+write
+known_if_redirected e16) v8 - object 3:7d6f3947::smithi1620526-191:head sending reply
remote/smithi186/log/ceph-osd.2.log.gz:DEBUG 2020-08-31 08:46:39,978 [shard 0] osd - do_osd_ops: osd_op(client
.4228.0:839 3.e 3:7d6f3947::smithi1620526-191:head {read 0~1 out=1b} snapc 0={} ondisk+read+rworderred+known_
if_redirected e16) v8 - object 3:7d6f3947::smithi1620526-191:head sending reply
remote/smithi186/log/ceph-osd.3.log.gz:DEBUG 2020-08-31 08:46:41,033 [shard 0] osd - do_osd_ops: osd_op(client
.4228.0:837 3.e 3:7d6f3947::smithi1620526-191:head {write 3101545~91785 in=91785b} snapc 0={} RETRY=1 ondisk
+retry+write+known_if_redirected e18) v8 - object 3:7d6f3947::smithi1620526-191:head sending reply
remote/smithi186/log/ceph-osd.3.log.gz:DEBUG 2020-08-31 08:46:41,054 [shard 0] osd - do_osd_ops: osd_op(client
.4228.0:838 3.e 3:7d6f3947::smithi1620526-191:head {setattr (57), truncate 3193330} snapc 0={} RETRY=1 ondi
sk+retry+write+known_if_redirected e18) v8 - object 3:7d6f3947::smithi1620526-191:head sending reply
```

#3 - 08/31/2020 12:58 PM - Kefu Chai

i think we should drop the read op after acquiring the obc lock if osd is not able to serve this request either because

- it's not the primary or
- it's not a replica and request is not a balanced read.
- or recovery_state indicates that the obc of the object is stale according to the pglog.

#4 - 09/01/2020 08:55 AM - Kefu Chai

- Status changed from New to Fix Under Review

- Assignee set to Kefu Chai

- Pull request ID set to 36919

#5 - 09/05/2020 07:03 AM - Kefu Chai

- Status changed from Fix Under Review to Resolved

#6 - 04/07/2021 06:12 AM - Deepika Upadhyay

- Status changed from Resolved to Pending Backport

- Backport set to *pacific, octopus*

```
2021-04-03T19:25:59.853 INFO:tasks.rados.rados.0.smithi046.stdout:603: finishing write tid 1 to smithi04631543-27
```

```
2021-04-03T19:25:59.854 INFO:tasks.rados.rados.0.smithi046.stderr:Error: finished tid 1 when last_acked_tid was 5
```

/ceph/teuthology-archive/yuriw-2021-04-03_17:21:06-rados-wip-yuri2-testing-2021-04-01-1440-octopus-distro-basic-smithi/6018206/teuthology.log

#7 - 05/11/2021 09:03 AM - Loic Dachary

- Backport deleted (*pacific, octopus*)

Removing the *pacific/octopus* backport request because there is no backport tracker for *crimson* at the moment. Please revert my change if I'm mistaken and there should be backports of this particular issue.