

RADOS - Bug #40451

osd/PG.cc: 2410: FAILED ceph_assert(scrub_queued)

06/19/2019 07:13 PM - Sage Weil

Status:	Resolved	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	nautilus	Component(RADOS):	
Regression:	No	Pull request ID:	28660
Severity:	3 - minor	Crash signature (v1):	
Reviewed:		Crash signature (v2):	

Description

```
...
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 494'361 scrubbing mbc={} clear_primary_state
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} release_backoffs [2:20000000::
::0,2:28000000:::head
)
2019-06-19T18:19:46.445+0000 7f7ceaed7700 20 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} agent_stop
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} on_change
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} cancel_copy_ops
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} cancel_copy 2:2a602f7f::smith
i08114715-44:head from
 2:8e5alael::smithi08114715-28:head @2 v284
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} kick_object_context_blocked 2:
2a602f7f::smithi08114
715-44:head requeuing 1 requests
2019-06-19T18:19:46.445+0000 7f7ceaed7700 20 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} requeue_op 0x5564d99238c0
2019-06-19T18:19:46.445+0000 7f7ceaed7700 20 osd.2 op_wq(4) _enqueue_front OpQueueItem(2.4 PGOpte
m(op=osd_op(client.4292.0:5853 2.4 2:2a602f7f::smithi08114715-44:head [stat] snapc 0=[] ondisk+re
ad+rworderred+known_if_redirected e516) v8) prio 63 cost 0 e532)
2019-06-19T18:19:46.445+0000 7f7ceaed7700 10 osd.2 pg_epoch: 532 pg[2.4( v 514'363 (0'0,514'363] 1
ocal-lis/les=490/491 n=85 ec=19/19 lis/c 490/490 les/c/f 491/491/0 486/532/490) [2,1] r=0 lpr=532
pi=[490,532)/1 crt=514'363 lcod 494'361 mlcod 0'0 scrubbing mbc={} requeue_scrub: queueing
2019-06-19T18:19:46.445+0000 7f7ceaed7700 20 osd.2 op_wq(4) _enqueue OpQueueItem(2.4 PGScrub(pgid=
2.4epoch_queued=532) prio 5 cost 52428800 e532)
...
-203> 2019-06-19T18:19:46.487+0000 7f7ce6ecf700 -1 /home/jenkins-build/build/workspace/ceph-dev-
```

```
new-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZE/huge/
release/15.0.0-1795-g2851aac/rpm/el7/BUILD/ceph-15.0.0-1795-g2851aac/src/osd/PG.cc: In function 'v
oid PG::scrub(epoch_t,
  ThreadPool::TPHandle&)' thread 7f7ce6ecf700 time 2019-06-19T18:19:46.483066+0000
/home/jenkins-build/build/workspace/ceph-dev-new-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE
_DIST/centos7/DIST/centos7/MACHINE_SIZE/huge/release/15.0.0-1795-g2851aac/rpm/el7/BUILD/ceph-15.0.
0-1795-g2851aac/src/osd/PG.cc: 2410: FAILED ceph_assert (scrub_queued)
```

```
ceph version 15.0.0-1795-g2851aac (2851aac1b33ac18f9f5295c6e80bb395d621676e) octopus (dev)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x14a) [0x5564bc54e78f]
2: (()+0x4c2957) [0x5564bc54e957]
3: (PG::scrub(unsigned int, ThreadPool::TPHandle&)+0x4ef) [0x5564bc792b9f]
4: (PGScrub::run(OSD*, OSDShard*, boost::intrusive_ptr<PG>&, ThreadPool::TPHandle&)+0x12) [0x5564
bc913c22]
5: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x1508) [0x5564bc702b08]
6: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x5b6) [0x5564bcc9546]
7: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x5564bccab6a0]
8: (()+0x7dd5) [0x7f7d0f2badd5]
```

/a/sage-2019-06-19_13:01:10-rados:thrash-wip-sloppy-snaps-distro-basic-smithi/4048437

Related issues:

Copied to RADOS - Backport #40537: nautilus: osd/PG.cc: 2410: FAILED ceph_ass...

Resolved

History

#1 - 06/19/2019 07:14 PM - Sage Weil

It looks to me like this happened as a side-effect of unblocking the op.

```
if (obc->requeue_scrub_on_unblock) {
    obc->requeue_scrub_on_unblock = false;
    requeue_scrub();
}
```

A simple fix should be to make requeue_scrub() a no-op if the PG is not active?

#2 - 06/19/2019 07:21 PM - Sage Weil

- Status changed from 12 to Fix Under Review

- Backport set to nautilus

<https://github.com/ceph/ceph/pull/28660>

#3 - 06/25/2019 05:15 AM - Kefu Chai

- Status changed from Fix Under Review to Pending Backport

- Pull request ID set to 28660

#4 - 06/25/2019 10:18 AM - Nathan Cutler

- Copied to Backport #40537: nautilus: osd/PG.cc: 2410: FAILED ceph_assert(scrub_queued) added

#5 - 08/19/2019 02:38 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved