

RADOS - Bug #38483

FAILED ceph_assert(p != pg_slots.end()) in OSDShard::register_and_wake_split_child(PG*)

02/26/2019 12:40 PM - Sage Weil

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Sage Weil		
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	nautilus	Component(RADOS):	
Regression:	No	Pull request ID:	30018
Severity:	3 - minor	Crash signature (v1):	
Reviewed:		Crash signature (v2):	
Description			
<pre>2019-02-26 03:04:15.748 7fefccfc0700 10 osd.4:7._attach_pg 1.7 0x55aec8689000 2019-02-26 03:04:15.748 7fefccfc0700 20 osd.4:7._wake_pg_slot _wake_pg_slot 1.7 to_process <> wait ing <> waiting_peering {} ... 2019-02-26 03:04:15.749 7fefbf7a5700 20 osd.4 572 advance_pg 1.7 is merge target, sources are 1.f 2019-02-26 03:04:15.749 7fefbf7a5700 1 osd.4 572 advance_pg 1.f is merge source, target is 1.7 2019-02-26 03:04:15.750 7fefbf7a5700 10 osd.4 572 add_merge_waiter added merge_waiter 1.f for 1.7, have 1/1 ... 2019-02-26 03:04:15.750 7fefbf7a5700 10 osd.4 pg_epoch: 547 pg[1.7(DNE empty local-lis/les=0/0 n= 0 ec=0/0 lis/c 0/0 les/c/f 0/0/0 0/0/0) [6,3,7] r=-1 lpr=547 crt=0'0 unknown mbc={}] merge_from fr om {1.f=0x55aec9916000} split_bits 3 2019-02-26 03:04:15.750 7fefbf7a5700 10 osd.4 pg_epoch: 547 pg[1.7(DNE empty local-lis/les=0/0 n= 0 ec=0/0 lis/c 0/0 les/c/f 0/0/0 0/0/0) [6,3,7] r=-1 lpr=547 crt=0'0 unknown mbc={}] merge_from ta rget incomplete 2019-02-26 03:04:15.750 7fefbf7a5700 10 osd.4 pg_epoch: 547 pg[1.7(DNE empty local-lis/les=0/0 n= 0 ec=0/0 lis/c 0/0 les/c/f 0/0/0 0/0/0) [6,3,7] r=-1 lpr=547 crt=0'0 unknown mbc={}] merge_from ta king source's past_intervals ... 2019-02-26 03:04:15.751 7fefbf7a5700 10 osd.4 572 split_pgs splitting pg[1.7(empty lb MIN (NIBBLE WISE) local-lis/les=0/547 n=0 ec=33/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 l pr=560 pi=[374,560)/2 crt=0'0 unknown NOTIFY mbc={}] into 1.f 2019-02-26 03:04:15.751 7fefbf7a5700 10 osd.4 pg_epoch: 560 pg[1.7(empty lb MIN (bitwise) local-l is/les=0/547 n=0 ec=33/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=560 pi=[37 4,560)/2 crt=0'0 unknown NOTIFY mbc={}] release_backoffs [MIN,MAX) 2019-02-26 03:04:15.751 7fefbf7a5700 10 osd.4 572 split_pgs splitting pg[1.7(empty lb MIN (bitwis e) local-lis/les=0/547 n=0 ec=33/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr= 560 pi=[374,560)/2 crt=0'0 unknown NOTIFY mbc={}] into 1.17 ... 2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4 572 _finish_splits pg[1.17(empty lb MIN (bitwise) l ocal-lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=0 p i=[374,560)/2 crt=0'0 unknown NOTIFY mbc={}] 2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4 pg_epoch: 560 pg[1.17(empty lb MIN (bitwise) local- lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=0 pi=[37 4,560)/2 crt=0'0 unknown NOTIFY mbc={}] handle_initialize 2019-02-26 03:04:15.752 7fefbf7a5700 5 osd.4 pg_epoch: 560 pg[1.17(empty lb MIN (bitwise) local- lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=0 pi=[37 4,560)/2 crt=0'0 unknown NOTIFY mbc={}] exit Initial 0.001060 0 0.000000 2019-02-26 03:04:15.752 7fefbf7a5700 5 osd.4 pg_epoch: 560 pg[1.17(empty lb MIN (bitwise) local- lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=0 pi=[37 4,560)/2 crt=0'0 unknown NOTIFY mbc={}] enter Reset</pre>			

```
2019-02-26 03:04:15.752 7fefbf7a5700 20 osd.4 pg_epoch: 560 pg[1.17( empty lb MIN (bitwise) local-
lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=0 pi=[37
4,560)/2 crt=0'0 unknown NOTIFY mbc={}] set_last_peering_reset 560
2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4 pg_epoch: 560 pg[1.17( empty lb MIN (bitwise) local-
lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=560 pi=[
374,560)/2 crt=0'0 unknown NOTIFY mbc={}] Clearing blocked outgoing recovery messages
2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4 pg_epoch: 560 pg[1.17( empty lb MIN (bitwise) local-
lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=560 pi=[
374,560)/2 crt=0'0 unknown NOTIFY mbc={}] Not blocking outgoing recovery messages
2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4 pg_epoch: 560 pg[1.17( empty lb MIN (bitwise) local-
lis/les=0/547 n=0 ec=560/14 lis/c 374/374 les/c/f 547/547/0 536/560/374) [6,3,7] r=-1 lpr=560 pi=[
374,560)/2 crt=0'0 unknown NOTIFY mbc={}] null
2019-02-26 03:04:15.752 7fefbf7a5700 15 osd.4 572 enqueue_peering_evt 1.17 epoch_sent: 560 epoch_r
equested: 560 NullEvt
2019-02-26 03:04:15.752 7fefbf7a5700 20 osd.4 op_wq(7) _enqueue OpQueueItem(1.17 PGPeeringEvent(ep
och_sent: 560 epoch_requested: 560 NullEvt) prio 255 cost 10 e560)
2019-02-26 03:04:15.752 7fefbf7a5700 10 osd.4:7.register_and_wake_split_child 1.17 0x55aec8234000
0> 2019-02-26 03:04:15.756 7fefbf7a5700 -1 *** Caught signal (Aborted) **
in thread 7fefbf7a5700 thread_name:tp_osd_tp
```

```
ceph version 14.1.0-125-g8b98d22 (8b98d22533def4c768359c2efe9496780b036d22) nautilus (dev)
1: (()+0xf5d0) [0x7fefe558e5d0]
2: (gsignal()+0x37) [0x7fefe4385207]
3: (abort()+0x148) [0x7fefe43868f8]
4: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x199) [0x55aebaf52a1b]
5: (ceph::__ceph_assertf_fail(char const*, char const*, int, char const*, char const*, ...)+0) [0
x55aebaf52b9a]
6: (OSDShard::register_and_wake_split_child(PG*)+0x7e3) [0x55aebb0bf503]
7: (OSD::_finish_splits(std::set<boost::intrusive_ptr<PG>, std::less<boost::intrusive_ptr<PG> >,
std::allocator<boost::intrusive_ptr<PG> > >&)+0x121) [0x55aebb0bf671]
8: (Context::complete(int)+0x9) [0x55aebb0c6349]
9: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x67c) [0x55aebb0abd7c]
10: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x433) [0x55aebb69f003]
11: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x55aebb6a20a0]
```

/a/sage-2019-02-26_00:43:29-rados-wip-sage-testing-2019-02-25-1642-distro-basic-smithi/3638207

looks like the merge -> split sequence doesn't prime the merge target

Related issues:

Duplicates RADOS - Bug #36304: FAILED ceph_assert(p != pg_slots.end()) in OSD...	New
Copied to RADOS - Backport #41712: nautilus: FAILED ceph_assert(p != pg_slots...	Resolved

History

#1 - 02/26/2019 01:21 PM - Sage Weil

- Status changed from 12 to Fix Under Review

<https://github.com/ceph/ceph/pull/26651>

#2 - 03/01/2019 09:45 PM - Sage Weil

- Status changed from Fix Under Review to In Progress

- Assignee set to Sage Weil

#3 - 04/25/2019 08:13 PM - Neha Ojha

/a/nojha-2019-04-25_05:43:35-rados-wip-39441-distro-basic-smithi/3892156/

#4 - 08/20/2019 08:29 PM - Greg Farnum

- Duplicates Bug #36304: FAILED ceph_assert(p != pg_slots.end()) in OSDShard::register_and_wake_split_child(PG*) added

#5 - 08/20/2019 09:12 PM - Josh Durgin

- Priority changed from Urgent to Normal

Sage says the PR is buggy, and this case is very hard to hit, so moving to normal priority.

#6 - 09/04/2019 12:32 AM - xie xingguo

- Status changed from In Progress to Pending Backport

- Backport set to nautilus

#7 - 09/04/2019 12:49 PM - xie xingguo

- Pull request ID set to 30018

#8 - 09/07/2019 09:29 AM - Nathan Cutler

- Copied to Backport #41712: nautilus: FAILED `ceph_assert(p != pg_slots.end())` in `OSDShard::register_and_wake_split_child(PG*)` added

#9 - 11/14/2019 03:53 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved

While running with `--resolve-parent`, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".