


```
444790476-2017-02-14T01:31:35.400 INFO:tasks.ceph.osd.1.smithi171.stderr: 8: (OSD::process_peering_events(std::list<PG*, std::allocator<PG*> > const&, ThreadPool::TPHandle&)+0x1a5) [0x7f4ede798665]
444790664-2017-02-14T01:31:35.400 INFO:tasks.ceph.osd.1.smithi171.stderr: 9: (OSD::PeeringWQ::_process(std::list<PG*, std::allocator<PG*> > const&, ThreadPool::TPHandle&)+0x17) [0x7f4ede7e67c7]
444790848-2017-02-14T01:31:35.400 INFO:tasks.ceph.osd.1.smithi171.stderr: 10: (ThreadPool::worker(ThreadPool::WorkThread*)+0xb65) [0x7f4edec8ee15]
444790985-2017-02-14T01:31:35.400 INFO:tasks.ceph.osd.1.smithi171.stderr: 11: (ThreadPool::WorkThread::entry()+0x10) [0x7f4edec8fde0]
444791109-2017-02-14T01:31:35.401 INFO:tasks.ceph.osd.1.smithi171.stderr: 12: (()+0x8184) [0x7f4edca25184]
444791206-2017-02-14T01:31:35.402 INFO:tasks.ceph.osd.1.smithi171.stderr: 13: (clone()+0x6d) [0x7f4edbb1537d]
444791306-2017-02-14T01:31:35.402 INFO:tasks.ceph.osd.1.smithi171.stderr: NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

Related issues:

Copied to Ceph - Backport #18999: kraken: "osd/PG.cc: 6896: FAILED assert(pg->is_acting(osd_with_shard) pg->is_up(osd_with_shard))" thread 7f4ec3904700 time 2017-02-14 01:31:35.350665	Resolved
Copied to Ceph - Backport #19000: jewel: "osd/PG.cc: 6896: FAILED assert(pg->is_acting(osd_with_shard) pg->is_up(osd_with_shard))" thread 7f4ec3904700 time 2017-02-14 01:31:35.350665	Resolved

History

#1 - 02/14/2017 04:46 PM - Sage Weil

```
2017-02-14 01:31:34.313088 7f4ec3904700 10 osd.1 pg_epoch: 2605 pg[2.5d( empty local-les=2605 n=0 ec=1620 les/c/f 2605/2600/0 2604/2604/2604) []/[1] r=0 lpr=2604 pi=2596-2603/3 crt=0'0 mlcod 0'0 active+undersized+degraded+remapped] choose_acting want [1,5] != acting [1], requesting pg_temp change
...
2017-02-14 01:31:35.350650 7f4ec3904700 10 osd.1 pg_epoch: 2605 pg[2.5d( empty local-les=2605 n=0 ec=1620 les/c/f 2605/2605/0 2604/2604/2604) []/[1] r=0 lpr=2604 crt=0'0 mlcod 0'0 active+undersized+degraded+remapped] handle_advance_map [3,4]/[1] -- 3/1
2017-02-14 01:31:35.350661 7f4ec3904700 10 osd.1 pg_epoch: 2606 pg[2.5d( empty local-les=2605 n=0 ec=1620 les/c/f 2605/2605/0 2604/2604/2604) []/[1] r=0 lpr=2604 crt=0'0 mlcod 0'0 active+undersized+degraded+remapped] state<Started/Primary/Active>: Active advmap
2017-02-14 01:31:35.353391 7f4ec3904700 -1 /build/ceph-12.0.0-279-gd8a7d0d/src/osd/PG.cc: In function 'boost::statechart::result PG::RecoveryState::Active::react(const PG::AdvMap&)' thread 7f4ec3904700 time 2017-02-14 01:31:35.350665
/build/ceph-12.0.0-279-gd8a7d0d/src/osd/PG.cc: 6896: FAILED assert(pg->is_acting(osd_with_shard) || pg->is_up(osd_with_shard))
```

The code is

```
for (size_t i = 0; i < pg->want_acting.size(); i++) {
    int osd = pg->want_acting[i];
    if (!advmap.osdmap->is_up(osd)) {
        pg_shard_t osd_with_shard(osd, shard_id_t(i));
        assert(pg->is_acting(osd_with_shard) || pg->is_up(osd_with_shard));
    }
}
```

I can't tell what that loop and assert is intended to check. The acting set can change arbitrarily, with no regard to what want_acting in the prior interval was!

#2 - 02/14/2017 07:13 PM - Sage Weil

- Status changed from *New* to *Testing*

<https://github.com/ceph/ceph/pull/13420>

#3 - 02/14/2017 08:30 PM - Sage Weil

- Project changed from *ceph-qa-suite* to *Ceph*

#4 - 02/15/2017 09:54 PM - Samuel Just

samuelj@teuthology:/a/samuelj-2017-02-15_01:03:44-rados-wip-sam-testing---basic-smithi/816292 also

#5 - 02/15/2017 09:55 PM - Samuel Just

I don't understand why this is not popping up. Sage's patch is correct, but something else is going on. Why is the up set empty here?

#6 - 02/16/2017 06:58 PM - Sage Weil

- Priority changed from *Normal* to *Immediate*

#7 - 02/17/2017 06:13 PM - Samuel Just

- Assignee set to *Greg Farnum*

#8 - 02/18/2017 06:57 PM - Sage Weil

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

- Backport set to *kraken,jewel*

#9 - 02/20/2017 10:51 AM - Loic Dachary

- Copied to Backport #18999: *kraken: "osd/PG.cc: 6896: FAILED assert(pg->is_acting(osd_with_shard) || pg->is_up(osd_with_shard))" in rados/upgrade added*

#10 - 02/20/2017 10:51 AM - Loic Dachary

- Copied to Backport #19000: *jewel: "osd/PG.cc: 6896: FAILED assert(pg->is_acting(osd_with_shard) || pg->is_up(osd_with_shard))" in rados/upgrade added*

#11 - 02/23/2017 03:23 PM - Sage Weil

- Priority changed from *Immediate* to *Urgent*

#12 - 07/06/2017 07:18 AM - Nathan Cutler

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