

History

#1 - 06/08/2017 06:18 PM - Sage Weil

/a/sage-2017-06-08_02:04:29-rados-wip-sage-testing-distro-basic-smithi/1269367 too

#2 - 06/08/2017 06:52 PM - Sage Weil

- Status changed from *Verified* to *Need More Info*

Hmm, I see the `fault_range` call (it's in the new `ec_unclone` code), but it's only dirtying the range including extents touching the unshared blob, so i can't tell how that range could include a shard that is not loaded.

Need logs...

#3 - 06/12/2017 07:26 PM - Sage Weil

- Priority changed from *Immediate* to *Urgent*

#4 - 06/13/2017 02:59 AM - Sage Weil

- Status changed from *Need More Info* to *Verified*

/a/sage-2017-06-12_20:56:37-rados-wip-sage-testing-distro-basic-smithi/1280581
has full log

```
o> 2017-06-12 22:01:13.857186 7f220047c700 -1 *** Caught signal (Aborted) **
in thread 7f220047c700 thread_name:tp_osd_tp

ceph version 12.0.3-1508-g9960ae3 (9960ae3c107a7282e831b8993cf30e3cadf20844) luminous (dev)
1: (()+0x9c752f) [0x7f221e8c852f]
2: (()+0xf370) [0x7f221b53f370]
3: (gsignal()+0x37) [0x7f221a5691d7]
4: (abort()+0x148) [0x7f221a56a8c8]
5: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x284) [0x7f221e906264]
6: (BlueStore::ExtentMap::dirty_range(unsigned int, unsigned int)+0x4be) [0x7f221e77c10e]
7: (BlueStore::_do_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intrusive_ptr<BlueStore::Onode>)+0x7b0) [0x7f221e7cd5e0]
8: (BlueStore::_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intrusive_ptr<BlueStore::Onode>)+0x94) [0x7f221e7ce794]
9: (BlueStore::_txc_add_transaction(BlueStore::TransContext*, ObjectStore::Transaction*)+0x18b8) [0x7f221e7e1b58]
10: (BlueStore::queue_transactions(ObjectStore::Sequencer*, std::vector<ObjectStore::Transaction, std::allocator<ObjectStore::Transaction> >&, boost::intrusive_ptr<TrackedOp>, ThreadPool::TPHandle*)+0x3a0) [0x7f221e7e26f0]
11: (PrimaryLogPG::queue_transactions(std::vector<ObjectStore::Transaction, std::allocator<ObjectStore::Transaction> >&, boost::intrusive_ptr<OpRequest>)+0x65) [0x7f221e573905]
12: (ECBackend::handle_sub_write(pg_shard_t, boost::intrusive_ptr<OpRequest>, ECSubWrite&, ZTracer::Trace const&, Context*)+0x62e) [0x7f221e67ad0e]
13: (ECBackend::try_reads_to_commit()+0x1c97) [0x7f221e689327]
14: (ECBackend::check_ops()+0x1c) [0x7f221e6897fc]
15: (ECBackend::handle_sub_write_reply(pg_shard_t, ECSubWriteReply const&, ZTracer::Trace const&)+0x2ae) [0x7f221e689abe]
16: (ECBackend::handle_message(boost::intrusive_ptr<OpRequest>)+0x2df) [0x7f221e68bbff]
17: (PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x5a2) [0x7f221e5141f2]
]
18: (OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)+0x22f) [0x7f221e3b723f]
19: (PGQueueable::RunVis::operator()(boost::intrusive_ptr<OpRequest> const&)+0x57) [0x7f221e3b7697]
20: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0xfce) [0x7f221e3e1f1e]
21: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x8e9) [0x7f221e90bc29]
```

#5 - 06/16/2017 01:40 PM - Sage Weil

- Priority changed from Urgent to Immediate

```
2017-06-16T02:00:53.940 INFO:tasks.ceph.osd.4.smithi113.stderr:/build/ceph-12.0.3-1748-g62eaa13/src/os/bluestore/BlueStore.cc: 2633: FAILED assert(0 == "can't mark unloaded shard dirty")
2017-06-16T02:00:53.944 INFO:tasks.ceph.osd.0.smithi154.stderr:2017-06-16 02:00:53.947564 7fa473a93700 -1 received signal: Hangup from PID: 15152 task name: /usr/bin/python /bin/daemon-helper kill ceph-osd -f --cluster ceph -i 0 UID: 0
2017-06-16T02:00:53.945 INFO:tasks.ceph.osd.4.smithi113.stderr: ceph version 12.0.3-1748-g62eaa13 (62eaa139b3add8abb500dc17b522ce9bbfb9fac3) luminous (dev)
2017-06-16T02:00:53.945 INFO:tasks.ceph.osd.4.smithi113.stderr: 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x102) [0x55baea8f9712]
2017-06-16T02:00:53.945 INFO:tasks.ceph.osd.4.smithi113.stderr: 2: (BlueStore::ExtentMap::dirty_range(unsigned int, unsigned int)+0x54a) [0x55baea74f3fa]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 3: (BlueStore::_do_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intrusive_ptr<BlueStore::Onode>)+0xd07) [0x55baea7a2a37]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 4: (BlueStore::_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intrusive_ptr<BlueStore::Onode>&)+0x7b) [0x55baea7a36eb]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 5: (BlueStore::_txc_add_transaction(BlueStore::TransContext*, ObjectStore::Transaction*)+0x1be7) [0x55baea7ba337]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 6: (BlueStore::queue_transactions(ObjectStore::Sequencer*, std::vector<ObjectStore::Transaction, std::allocator<ObjectStore::Transaction> >&, boost::intrusive_ptr<TrackedOp>, ThreadPool::TPHandle*)+0x52e) [0x55baea7bb33e]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 7: (PrimaryLogPG::queue_transactions(std::vector<ObjectStore::Transaction, std::allocator<ObjectStore::Transaction> >&, boost::intrusive_ptr<OpRequest>)+0x66) [0x55baea4ff926]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 8: (ECBackend::handle_sub_write(pg_shard_t, boost::intrusive_ptr<OpRequest>, ECSubWrite&, ZTracer::Trace const&, Context*)+0x886) [0x55baea62a796]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 9: (ECBackend::handle_message(boost::intrusive_ptr<OpRequest>)+0x331) [0x55baea644961]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 10: (PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x541) [0x55baea4a1751]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 11: (OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)+0x1c9) [0x55baea335a89]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 12: (PGQueueable::RunVis::operator() (boost::intrusive_ptr<OpRequest> const&)+0x57) [0x55baea335f07]
2017-06-16T02:00:53.946 INFO:tasks.ceph.osd.4.smithi113.stderr: 13: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x130e) [0x55baea35c06e]
2017-06-16T02:00:53.947 INFO:tasks.ceph.osd.4.smithi113.stderr: 14: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x933) [0x55baea8fe543]
2017-06-16T02:00:53.947 INFO:tasks.ceph.osd.4.smithi113.stderr: 15: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x55baea901780]
2017-06-16T02:00:53.947 INFO:tasks.ceph.osd.4.smithi113.stderr: 16: (()+0x770a) [0x7f9f5612d70a]
2017-06-16T02:00:53.947 INFO:tasks.ceph.osd.4.smithi113.stderr: 17: (clone()+0x6d) [0x7f9f551a482d]
```

/a/sage-2017-06-16_00:46:50-rados-wip-sage-testing-distro-basic-smithi/1292301

#6 - 06/16/2017 09:10 PM - Sage Weil

/a/sage-2017-06-16_18:45:23-rados-wip-sage-testing-distro-basic-smithi/1293630

#7 - 06/17/2017 01:36 PM - Sage Weil

http://pulcito.ceph.com/sage-2017-06-16_19:23:03-rbd:qemu-wip-19882---basic-smithi/

reliably reproduced by rbd/qemu

#8 - 06/17/2017 09:27 PM - Greg Farnum

- Project changed from Ceph to RADOS

- Component(RADOS) BlueStore added

#9 - 06/19/2017 06:45 PM - Sage Weil

reliably triggered, it seems, by rbd/qemu xfstests workload

#10 - 06/20/2017 02:50 AM - Sage Weil

/a/sage-2017-06-19_18:44:38-rbd:qemu-master---basic-smithi/1301319

#11 - 06/20/2017 02:54 AM - Sage Weil

```
-941> 2017-06-19 22:10:27.505224 7f630d537700 30 bluestore.OnodeSpace(0x555e9e282168 in 0x555e9b638460) look
up 0#1:3f5d397a:::rbd_data.0.104669e373.0000000000000283:head#28205 hit 0x555ea10ff8c0
-940> 2017-06-19 22:10:27.505226 7f630d537700 15 bluestore(/var/lib/ceph/osd/ceph-1) _remove 1.0s0_head 0#1:
3f5d397a:::rbd_data.0.104669e373.0000000000000283:head#28205
-939> 2017-06-19 22:10:27.505229 7f630d537700 15 bluestore(/var/lib/ceph/osd/ceph-1) _do_truncate 1.0s0_head
0#1:3f5d397a:::rbd_data.0.104669e373.0000000000000283:head#28205 0x0
-938> 2017-06-19 22:10:27.505231 7f630d537700 30 bluestore(/var/lib/ceph/osd/ceph-1) _dump_onode 0x555ea10ff
8c0 0#1:3f5d397a:::rbd_data.0.104669e373.0000000000000283:head#28205 nid 212716 size 0x74000 (475136) expected
_object_size 0 expected_write_size 0 in 0 shards, 0 spanning blobs
-937> 2017-06-19 22:10:27.505234 7f630d537700 30 bluestore(/var/lib/ceph/osd/ceph-1) _dump_extent_map 0x730
00~1000: 0x3000~1000 Blob(0x555ecc18cbd0 blob([0x3e098000~4000,0x3dff0000~4000] csum+shared crc32c/0x1000) use
_tracker(0x2*0x4000 0x[1000,0]) SharedBlob(0x555ea03ccae0 loaded (sbid 0x3d542 ref_map(0x3dff000
0~4000=2,0x3e098000~4000=1))))
-936> 2017-06-19 22:10:27.505240 7f630d537700 30 bluestore(/var/lib/ceph/osd/ceph-1) _dump_extent_map c
sum: [0,0,0,6706be76,5bf25118,b2f40631,3df868d,c999f301]
-935> 2017-06-19 22:10:27.505241 7f630d537700 30 bluestore(/var/lib/ceph/osd/ceph-1) _dump_extent_map
0x3000~1000 buffer(0x555e9f0a3980 space 0x555ea03ccaf8 0x3000~1000 clean)
-934> 2017-06-19 22:10:27.505243 7f630d537700 30 bluestore.extentmap(0x555ea10ffa10) fault_range 0x0~74000
-933> 2017-06-19 22:10:27.505244 7f630d537700 20 bluestore.blob(0x555ecc18cbd0) put_ref 0x3000~1000 Blob(0x5
55ecc18cbd0 blob([0x3e098000~4000,0x3dff0000~4000] csum+shared crc32c/0x1000) use_tracker(0x2*0x4000 0x[1000,0
]) SharedBlob(0x555ea03ccae0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=2,0x3e098000~4000=1)))
)
-932> 2017-06-19 22:10:27.505247 7f630d537700 30 bluestore.extentmap(0x555ea10ffa10) dirty_range 0x0~74000
-931> 2017-06-19 22:10:27.505248 7f630d537700 20 bluestore.extentmap(0x555ea10ffa10) dirty_range mark inline
shard dirty
-930> 2017-06-19 22:10:27.505249 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _wctx_finish lex_old 0x
73000~1000: 0x3000~1000 Blob(0x555ecc18cbd0 blob([!~8000] csum+shared crc32c/0x1000) use_tracker(0x2*0x4000 0x
[0,0]) SharedBlob(0x555ea03ccae0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=2,0x3e098000~4000=
1))))
-929> 2017-06-19 22:10:27.505251 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _wctx_finish blob rele
ase [0x3e098000~4000,0x3dff0000~4000]
-928> 2017-06-19 22:10:27.505253 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _wctx_finish shared_bl
ob release [0x3e098000~4000] from SharedBlob(0x555ea03ccae0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=1)))
-927> 2017-06-19 22:10:27.505254 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _wctx_finish release 0
x3e098000~4000
-926> 2017-06-19 22:10:27.505257 7f630d537700 10 bluestore(/var/lib/ceph/osd/ceph-1) _do_remove gen and mayb
e_unshared_blobs 0x555ea03ccae0
-925> 2017-06-19 22:10:27.505258 7f630d537700 30 bluestore.OnodeSpace(0x555e9e282168 in 0x555e9b638460) look
up
-924> 2017-06-19 22:10:27.505259 7f630d537700 30 bluestore.OnodeSpace(0x555e9e282168 in 0x555e9b638460) look
up 0#1:3f5d397a:::rbd_data.0.104669e373.0000000000000283:head# hit 0x555ec9319c80
-923> 2017-06-19 22:10:27.505261 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _do_remove checking for
unshareable blobs on 0x555ec9319c80 0#1:3f5d397a:::rbd_data.0.104669e373.0000000000000283:head#
-922> 2017-06-19 22:10:27.505265 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) ? SharedBlob(0x555ea03
```

```

ccea0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=1)) vs ref_map(0x3dff0000~4000=1)
-921> 2017-06-19 22:10:27.505266 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _do_remove unsharing S
haredBlob(0x555ea03ccea0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=1)))
-920> 2017-06-19 22:10:27.505267 7f630d537700 10 bluestore(/var/lib/ceph/osd/ceph-1).collection(1.0s0_head 0
x555e9e282000) make_blob_unshared SharedBlob(0x555ea03ccea0 loaded (sbid 0x3d542 ref_map(0x3dff0000~4000=1)))
-919> 2017-06-19 22:10:27.505269 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1).collection(1.0s0_head 0
x555e9e282000) make_blob_unshared now SharedBlob(0x555ea03ccea0 sbid 0x0)
-918> 2017-06-19 22:10:27.505270 7f630d537700 20 bluestore(/var/lib/ceph/osd/ceph-1) _do_remove 0x74000~400
0: 0x4000~4000 Blob(0x555ea03caf50 blob([!~4000,0x3dff0000~4000] csum+shared crc32c/0x1000) use_tracker(0x2*0x
4000 0x[0,4000]) SharedBlob(0x555ea03ccea0 sbid 0x0))
-917> 2017-06-19 22:10:27.505272 7f630d537700 30 bluestore.extentmap(0x555ec9319dd0) dirty_range 0x74000~780
00
-916> 2017-06-19 22:10:27.505273 7f630d537700 20 bluestore.extentmap(0x555ec9319dd0) dirty_range mark shard
0x60000 dirty
-915> 2017-06-19 22:10:27.505274 7f630d537700 20 bluestore.extentmap(0x555ec9319dd0) dirty_range shard 0xa00
00 is not loaded, can't mark dirty

```

#12 - 06/20/2017 02:56 AM - Sage Weil

- Status changed from Verified to Need Review

<https://github.com/ceph/ceph/pull/15766>

#13 - 06/20/2017 08:24 PM - Sage Weil

- Status changed from Need Review to Resolved

#14 - 08/04/2017 08:41 AM - Josh Durgin

Hit the same assert in

http://qa-proxy.ceph.com/teuthology/joshd-2017-08-04_06:16:52-rados-wip-20904-distro-basic-smithi/1482581/remote/smithi060/log/ceph-osd.2.log.gz

```

2017-08-04 06:34:49.217813 7fe8fac07700 -1 /build/ceph-12.1.2-87-gb9439b5/src/os/bluestore/BlueStore.cc: In fu
nction 'void BlueStore::ExtentMap::dirty_range(uint32_t, uint32_t)' thread 7fe8fac07700 time 2017-08-04 06:34:
49.209121
/build/ceph-12.1.2-87-gb9439b5/src/os/bluestore/BlueStore.cc: 2690: FAILED assert(0 == "can't mark unloaded sh
ard dirty")

```

```

ceph version 12.1.2-87-gb9439b5 (b9439b59b42f1d32573da461a557ad41a37ce799) luminous (rc)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x10e) [0x7fe9153c8a7e]
2: (BlueStore::ExtentMap::dirty_range(unsigned int, unsigned int)+0x48e) [0x7fe91524427e]
3: (BlueStore::_do_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intr
usive_ptr<BlueStore::Onode>)+0x7bb) [0x7fe91529531b]
4: (BlueStore::_remove(BlueStore::TransContext*, boost::intrusive_ptr<BlueStore::Collection>&, boost::intrusi
ve_ptr<BlueStore::Onode>&)+0x83) [0x7fe915296603]
5: (BlueStore::_txc_add_transaction(BlueStore::TransContext*, ObjectStore::Transaction*)+0x17ef) [0x7fe9152a8
ccf]
6: (BlueStore::queue_transactions(ObjectStore::Sequencer*, std::vector<ObjectStore::Transaction, std::allocat
or<ObjectStore::Transaction> >&, boost::intrusive_ptr<TrackedOp>, ThreadPool::TPHandle*)+0x388) [0x7fe9152a99a
8]
7: (PrimaryLogPG::queue_transactions(std::vector<ObjectStore::Transaction, std::allocator<ObjectStore::Transa
ction> >&, boost::intrusive_ptr<OpRequest>)+0x55) [0x7fe91502dff5]
8: (ECBackend::handle_sub_write(pg_shard_t, boost::intrusive_ptr<OpRequest>, ECSubWrite&, ZTracer::Trace cons
t&, Context*)+0x617) [0x7fe9151439b7]
9: (ECBackend::_handle_message(boost::intrusive_ptr<OpRequest>)+0x2f7) [0x7fe915154047]

```

```
10: (PGBackend::handle_message(boost::intrusive_ptr<OpRequest>)+0x78) [0x7fe915061ab8]
11: (PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x55e) [0x7fe914fd0c7e
]
12: (OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHandle&)+0x3e6)
[0x7fe914e6e8c6]
13: (PGQueueable::RunVis::operator()(boost::intrusive_ptr<OpRequest> const&)+0x47) [0x7fe9150cb9f7]
14: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0xff5) [0x7fe914e99225]
15: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x83f) [0x7fe9153ce23f]
16: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x7fe9153d0190]
17: (()+0x8184) [0x7fe912e8c184]
18: (clone()+0x6d) [0x7fe911f7cbed]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```