

Ceph - Bug #1609

osd: failed assert(info.last_complete == info.last_update)

10/07/2011 06:05 PM - Josh Durgin

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Sage Weil	Reviewed:	
Category:	OSD	Affected Versions:	
Target version:	v0.38	ceph-qa-suite:	
Source:		Pull request ID:	
Tags:		Crash signature:	
Backport:			
Regression:	No		
Severity:	3 - minor		

Description

From teuthology:~teuthworker/archive/nightly_coverage_2011-10-07/246/remote/ubuntu@sepia67.ceph.dreamhost.com/log/osd.0.log.gz:

```
osd/PG.cc: In function 'void PG::finish_recovery(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)', in thread '0x7f4573df7700'
osd/PG.cc: 1800: FAILED assert(info.last_complete == info.last_update)
ceph version 0.36-251-g6e29c28 (commit:6e29c2826066a7723ed05b60b8ac0433a04c3c13)
1: (PG::finish_recovery(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)+0x8d) [0x6ff0ed]
2: (PG::RecoveryState::Active::react(PG::RecoveryState::ActMap const&)+0x316) [0x729196]
3: (boost::statechart::simple_state<PG::RecoveryState::Active, PG::RecoveryState::Primary, boost::mpl::list<mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na, mpl::_na>, (boost::statechart::history_mode)0>::react_impl(boost::statechart::event_base const&, void const*+0x21b) [0x759c0b]
4: (boost::statechart::state_machine<PG::RecoveryState::RecoveryMachine, PG::RecoveryState::Initial, std::allocator<void>, boost::statechart::null_exception_translator>::process_event(boost::statechart::event_base const&)+0x8d) [0x7423dd]
5: (PG::RecoveryState::handle_activate_map(PG::RecoveryCtx*+0x183) [0x711f43]
6: (OSD::activate_map(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)+0x674) [0x579884]
7: (OSD::handle_osd_map(MOSDMap*+0x2270) [0x57bd50]
8: (OSD::_dispatch(Message*+0x4d0) [0x596bb0]
9: (OSD::ms_dispatch(Message*+0x17b) [0x59803b]
10: (SimpleMessenger::dispatch_entry()+0x9c2) [0x617562]
11: (SimpleMessenger::DispatchThread::entry()+0x2c) [0x4a3dec]
12: (Thread::_entry_func(void*+0x12) [0x611a92]
13: (()+0x7971) [0x7f457f87b971]
14: (clone()+0x6d) [0x7f457e10b92d]
```

Associated revisions

Revision 03ad5a28 - 10/25/2011 05:50 AM - Sage Weil

osd: fix last_complete adjustment after recovering an object

After we recover each object, we try to raise the last_complete value (and matching complete_to iterator). If our log was purely a backlog, this won't necessarily end up bringing last_complete all the way up to the last_update value, and we'll fail an assert later.

If complete_to does reach the end of the log, then we fast-forward last_complete to last_update.

The crash we were hitting was in finish_recovery(), and looked something

like

osd/PG.cc: In function 'void PG::finish_recovery(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)', in thread '0x7f4573df7700'

osd/PG.cc: 1800: FAILED assert(info.last_complete == info.last_update)

ceph version 0.36-251-g6e29c28 (commit:6e29c2826066a7723ed05b60b8ac0433a04c3c13)

1: (PG::finish_recovery(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)+0x8d) [0x6ff0ed]

2: (PG::RecoveryState::Active::react(PG::RecoveryState::ActMap const&)+0x316) [0x729196]

3: (boost::statechart::simple_state<PG::RecoveryState::Active, PG::RecoveryState::Primary, boost::mpl::list<mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na, mpl::na>, (boost::statechart::history_mode)0>::react_impl(boost::statechart::event_base const&, void const*)+0x21b) [0x759c0b]

4: (boost::statechart::state_machine<PG::RecoveryState::RecoveryMachine, PG::RecoveryState::Initial, std::allocator<void>, boost::statechart::null_exception_translator>::process_event(boost::statechart::event_base const&)+0x8d) [0x7423dd]

5: (PG::RecoveryState::handle_activate_map(PG::RecoveryCtx*)+0x183) [0x711f43]

6: (OSD::activate_map(ObjectStore::Transaction&, std::list<Context*, std::allocator<Context*> >&)+0x674) [0x579884]

7: (OSD::handle_osd_map(MOSDMap*)+0x2270) [0x57bd50]

8: (OSD::_dispatch(Message*)+0x4d0) [0x596bb0]

9: (OSD::ms_dispatch(Message*)+0x17b) [0x59803b]

10: (SimpleMessenger::dispatch_entry()+0x9c2) [0x617562]

11: (SimpleMessenger::DispatchThread::entry()+0x2c) [0x4a3dec]

12: (Thread::_entry_func(void*)+0x12) [0x611a92]

13: (()+0x7971) [0x7f457f87b971]

14: (clone()+0x6d) [0x7f457e10b92d]

Fixes: #1609

Signed-off-by: Sage Weil <sage@newdream.net>

History

#1 - 10/09/2011 08:39 PM - Sage Weil

- Target version changed from v0.37 to v0.38

#2 - 10/09/2011 08:41 PM - Sage Weil

- translation missing: en.field_position set to 58

#3 - 10/21/2011 10:55 AM - Sage Weil

- Status changed from New to Need More Info

#4 - 10/25/2011 03:29 PM - Sage Weil

- Status changed from Need More Info to Resolved

- Assignee set to Sage Weil

lots of stuff, mainly [03ad5a28eee2328eb2419c48a14df1a3624fc4c7](#)