

Ceph - Bug #41348

osd: need clear PG_STATE_CLEAN when repair object

08/20/2019 07:19 AM - Zengran Zhang

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category:	
Target version:	
Source:	Reviewed:
Tags:	Affected Versions:
Backport: luminous,mimic,nautilus	ceph-qa-suite:
Regression: No	Pull request ID: 29756
Severity: 3 - minor	Crash signature:
Description	
<pre>2019-08-20 09:32:38.216104 7f6504f12700 10 osd.0 pg_epoch: 133696 pg[3.36(v 133696'863063 (133630'861475,133696'863063] local-lis/les=133695/133696 n=106 ec=69/69 lis/c 133695/133695 les/c/f 133696/133696/0 133695/133695/133421) [0,9,11,4,7] r=0 lpr=133695 crt=133696'863063 mlcod 133696'863061 active+clean+snaptrim snaptrimq=[3bf4~7,3bfc~2d]] do_osd_op read 28672~8192 2019-08-20 09:32:38.216842 7f6504f12700 10 osd.0 pg_epoch: 133696 pg[3.36(v 133696'863063 (133630'861475,133696'863063] local-lis/les=133695/133696 n=106 ec=69/69 lis/c 133695/133695 les/c/f 133696/133696/0 133695/133695/133421) [0,9,11,4,7] r=0 lpr=133695 crt=133696'863063 mlcod 133696'863061 active+clean+snaptrim snaptrimq=[3bf4~7,3bfc~2d]] rep_repair_primary_object 3:6d4d0f80:::1000000feec.00000000:head peers osd.{0,4,7,9,11} 2019-08-20 09:32:38.216925 7f6504f12700 -1 log_channel(cluster) log [ERR] : 3.36 missing primary copy of 3:6d4d0f80:::1000000feec.00000000:head, will try copies on 4,7,9,11 2019-08-20 09:32:38.216940 7f6504f12700 10 osd.0 pg_epoch: 133696 pg[3.36(v 133696'863063 (133630'861475,133696'863063] local-lis/les=133695/133696 n=106 ec=69/69 lis/c 133695/133695 les/c/f 133696/133696/0 133695/133695/133421) [0,9,11,4,7] r=0 lpr=133695 crt=133696'863063 mlcod 133696'863061 active+clean+snaptrim m=1 snaptrimq=[3bf4~7,3bfc~2d]] read got -11 / 8192 bytes from obj 3:6d4d0f80:::1000000feec.00000000:head. try again. 2019-08-20 09:32:38.216996 7f6506715700 10 osd.0 pg_epoch: 133696 pg[3.36(v 133696'863063 (133630'861475,133696'863063] local-lis/les=133695/133696 n=106 ec=69/69 lis/c 133695/133695 les/c/f 133696/133696/0 133695/133695/133421) [0,9,11,4,7] r=0 lpr=133695 crt=133696'863063 mlcod 133696'863061 active+clean+snaptrim m=1 snaptrimq=[3bf4~7,3bfc~2d]] SnapTrimmer state<Trimming/AwaitAsyncWork>: AwaitAsyncWork: trimming snap 3bf4 2019-08-20 09:32:38.218817 7f6506715700 10 osd.0 pg_epoch: 133696 pg[3.36(v 133696'863063 (133630'861475,133696'863063] local-lis/les=133695/133696 n=106 ec=69/69 lis/c 133695/133695 les/c/f 133696/133696/0 133695/133695/133421) [0,9,11,4,7] r=0 lpr=133695 crt=133696'863063 mlcod 133696'863061 active+clean+snaptrim m=1 snaptrimq=[3bf4~7,3bfc~2d]] SnapTrimmer state<Trimming/AwaitAsyncWork>: AwaitAsyncWork react trimming 3:6d4d0f80:::1000000feec.00000000:3bf4 /root/rpmbuild/BUILD/ceph-12.2.7-1326-gdb735a3/src/osd/PrimaryLogPG.cc: 10090: FAILED assert(attrs !pg_log.get_missing().is_missing(soid) (pg_log.get_log().objects.count(soid) && pg_log.get_log().objects.find(soid)->second->op == pg_log_entry_t::LOST_REVERT)) 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x110) [0x56286c50f830] 2: (PrimaryLogPG::get_object_context(hobject_t const&, bool, std::map<std::string, ceph::buffer::list, std::less<std::string>, std::allocator<std::pair<std::string const, ceph::buffer::list> > > const*)+0x9b0) [0x56286c0d9460] 3: (PrimaryLogPG::trim_object(bool, hobject_t const&, std::unique_ptr<PrimaryLogPG::OpContext, std::default_delete<PrimaryLogPG::OpContext> >)+0x19c) [0x56286c0f2efc] 4: (PrimaryLogPG::AwaitAsyncWork::react(PrimaryLogPG::DoSnapWork const&)+0x8ea) [0x56286c0f5a7a] 5: (boost::statechart::simple_state<PrimaryLogPG::AwaitAsyncWork, PrimaryLogPG::Trimming, 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, a, 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)+0xb0) [0x56286c161910] 6: (PrimaryLogPG::snap_trimmer(unsigned int)+0x1fe) [0x56286c0b023e] 7: (OSD::ShardedOpWQ::process(unsigned int, ceph::heartbeat_handle_d*)+0x19b4) [0x56286bf7e864] 8: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x839) [0x56286c515279] 9: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x56286c5171f0] 10: ((+0x7e25) [0x7f6528feae25] 11: (clone()+0x6d) [0x7f65280de34d]</pre>	

Related issues:

Copied to Ceph - Backport #41442: mimic: osd: need clear PG_STATE_CLEAN when ...	Resolved
Copied to Ceph - Backport #41443: nautilus: osd: need clear PG_STATE_CLEAN wh...	Resolved
Copied to Ceph - Backport #41733: luminous: osd: need clear PG_STATE_CLEAN wh...	Resolved

History

#1 - 08/21/2019 03:09 AM - David Zafman

- Status changed from New to In Progress
- Pull request ID set to 29756

#2 - 08/23/2019 05:11 AM - Kefu Chai

- Status changed from In Progress to Pending Backport
- Backport set to mimic, nautilus

#3 - 08/26/2019 02:22 PM - Nathan Cutler

- Copied to Backport #41442: mimic: osd: need clear PG_STATE_CLEAN when repair object added

#4 - 08/26/2019 02:22 PM - Nathan Cutler

- Copied to Backport #41443: nautilus: osd: need clear PG_STATE_CLEAN when repair object added

#5 - 09/06/2019 11:33 PM - David Zafman

This should probably back ported to Luminous

#6 - 09/09/2019 05:00 PM - Neha Ojha

- Backport changed from mimic, nautilus to luminous,mimic,nautilus

#7 - 09/09/2019 07:39 PM - Nathan Cutler

- Copied to Backport #41733: luminous: osd: need clear PG_STATE_CLEAN when repair object added

#8 - 11/05/2019 01:15 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".