

RADOS - Bug #38377

OpTracker destruct assert when OSD destruct

02/19/2019 09:21 AM - bing lin

Status: Resolved	Start date: 02/19/2019
Priority: High	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	Spent time: 0.00 hour
Source:	Reviewed:
Tags:	Affected Versions:
Backport: mimic	ceph-qa-suite:
Regression: No	Component(RADOS):
Severity: 3 - minor	Pull request ID: 26504
Description coredump (gdb) bt #0 0x00007f506d3b623b in raise (sig=6) at ../nptl/sysdeps/unix/sysv/linux/pt-raise.c:37 #1 0x0000563eafebba6 in reraise_fatal (signum=6) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/global/signal_handler.cc:74 #2 handle_fatal_signal (signum=6) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/global/signal_handler.cc:138 #3 <signal handler called> #4 0x00007f506c3e01d7 in __GI_raise (sig=sig@entry=6) at ../nptl/sysdeps/unix/sysv/linux/raise.c:56 #5 0x00007f506c3e18c8 in __GI_abort () at abort.c:90 #6 0x0000563eaff2ab14 in ceph::__ceph_assert_fail (assertion=assertion@entry=0x563eb04695b8 " (sharded_in_flight_list.back())->ops_in_flight_sharded.empty()", file=file@entry=0x563eb04694e0 " /root/rpmbuild/BUILD/ceph-12.2.10-469-g57b4c2d/src/common/TrackedOp.cc", line=line@entry=153, func=func@entry=0x563eb0469ee0 <OpTracker::~OpTracker()::__PRETTY_FUNCTION__> " OpTracker::~OpTracker()" at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/common/assert.cc:66 #7 0x0000563eafb2ee8 in OpTracker::~OpTracker (this=0x563eba8f3150, __in_chrg=<optimized out>) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/common/TrackedOp.cc:153 #8 0x0000563eaf981332 in OSD::~OSD (this=0x563eba8f2000, __in_chrg=<optimized out>) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/osd/OSD.cc:2042 #9 0x0000563eaf9815c9 in OSD::~OSD (this=0x563eba8f2000, __in_chrg=<optimized out>) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/osd/OSD.cc:2052 #10 0x0000563eaf870b68 in main (argc=<optimized out>, argv=<optimized out>) at /usr/src/debug/ceph-12.2.10-469-g57b4c2d/src/ceph_osd.cc:675 when Optracker destruct, ops_in_flight_sharded is not empty	
Related issues:	
Duplicated by RADOS - Bug #38592: mon,osd: src/common/TrackedOp.cc: 163: FAIL...	Duplicate 03/05/2019
Duplicated by RADOS - Bug #36546: common/TrackedOp.cc: 163: FAILED ceph_asser...	Duplicate 10/22/2018
Copied to RADOS - Backport #38646: mimic: OpTracker destruct assert when OSD ...	Resolved

History

#1 - 02/20/2019 10:37 PM - Greg Farnum

Is this a custom build? Where did it come from?

#2 - 02/21/2019 06:33 AM - bing lin

Greg Farnum wrote:

Is this a custom build? Where did it come from?

aha ceph version is Luminous 12.2.10,
see [<https://github.com/ceph/ceph/pull/26504>]

1. waiting_for_osdmap is list of OpRequestRef, so when push OpRequest to waiting_for_osdmap, OpRequest ++nref and when take_waiter, Op will be push to finished from waiting_for_osdmap, until Op take from finished and nref put to zero OpRequest will unregister from optracker.

2. when osd shutdown, finished will be clear, that will put Op nref(--nref), but when op has not been take to finished, Op will still in waiting_for_osdmap, however OpTracker destruct before waiting_for_osdmap, so that will cause OpTracker::~OpTracker
assert((sharded_in_flight_list.back()->ops_in_flight_sharded.empty());

timeline below:

t1 push Op to waiting_for_osdmap(++nref)

t2 osd shutdown

t3 delete osd

t4 delete OpTracker

t5 got assert

#3 - 02/25/2019 06:53 PM - Greg Farnum

- Status changed from New to Need Review

- Pull request ID set to 26504

#4 - 03/06/2019 02:50 AM - Kefu Chai

- Backport set to nautilus,mimic

#5 - 03/07/2019 03:02 PM - Sage Weil

- Duplicated by Bug #38592: mon,osd: src/common/TrackedOp.cc: 163: FAILED
ceph_assert((sharded_in_flight_list.back()->ops_in_flight_sharded.empty()) on shutdown added

#6 - 03/07/2019 03:03 PM - Sage Weil

- Priority changed from Normal to Urgent

#7 - 03/07/2019 03:03 PM - Sage Weil

- Duplicated by Bug #36546: common/TrackedOp.cc: 163: FAILED ceph_assert((sharded_in_flight_list.back()->ops_in_flight_sharded.empty()) added

#8 - 03/08/2019 04:12 AM - Sage Weil

- Status changed from *Need Review* to *Pending Backport*

- Priority changed from *Urgent* to *High*

#9 - 03/08/2019 02:46 PM - Nathan Cutler

- Copied to Backport #38646: *mimic: OpTracker destruct assert when OSD destruct added*

#10 - 03/08/2019 02:48 PM - Nathan Cutler

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

master is still being merged into nautilus AFAICT

#11 - 05/03/2019 09:31 AM - Nathan Cutler

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