rbd - Bug #15690

librbd: resize hang - use after free

05/02/2016 11:24 PM - Josh Durgin

Status:	Resolved	**Start date:**	05/02/2016
Priority:	Urgent	**Due date:**	
Assignee:	Jason Dillaman	**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:		**Spent time:**	0.00 hour
Source:	Q/A	**Reviewed:**	
Tags:		**Affected Versions:**	
Backport:	jewel	**ceph-qa-suite:**	
Regression:	Yes	**Pull request ID:**	
Severity:	3 - minor		

Description

From running RBD_FEATURES=1 ./unittest_librbd

```
Thread 1 (Thread 0x7f284cd5d6c0 (LWP 8649)):
#0  __lll_lock_wait () at ../nptl/sysdeps/unix/sysv/linux/x86_64/lowlevellock.S:135
#1  0x00007f28416d1970 in pthread_rwlock_unlock () at ../nptl/sysdeps/unix/sysv/linux/x86_64/pthread_rwlock_unlock.S:103

#2  0x00007f284c424e90 in RWLock::unlock (this=0x7f2840a7c950 <main_arena+496>, lockdep=<optimized out>) at ./common/RWLock.h:93
#3  0x00007f284c727545 in put_read (this=<optimized out>) at ./common/RWLock.h:116
#4  librbd::(anonymous namespace)::C_InvokeAsyncRequest<librbd::ImageCtx>::send_acquire_exclusive_lock (this=0x7f2856b94490) at librbd/Operations.cc:179
#5  0x00007f284c727748 in librbd::(anonymous namespace)::C_InvokeAsyncRequest<librbd::ImageCtx>::send_refresh_image (this=0x7f2856b94490) at librbd/Operations.cc:128
#6  0x00007f284c72c733 in send (this=0x7f2856b94490) at librbd/Operations.cc:123
#7  librbd::Operations<librbd::ImageCtx>::invoke_async_request(std::string const&, bool, boost::function<void (Context*)> const&, boost::function<void (Context*)> const&) (this=this@entry=0x7f2856ba38f0,
    request_type="resize", permit_snapshot=permit_snapshot@entry=false, local_request=..., remote_request=...) at librbd/Operations.cc:1052
#8  0x00007f284c72daed in librbd::Operations<librbd::ImageCtx>::resize (this=0x7f2856ba38f0, size=size@entry=8388608, prog_ctx=...) at librbd/Operations.cc:499
#9  0x00007f284c6b13be in librbd::Image::resize (this=this@entry=0x7fff3d498eb0, size=size@entry=8388608) at librbd/librbd.cc:581
#10 0x00007f284c5cf156 in TestLibRBD_ResizeAndStatPP_Test::TestBody (this=0x7f2856b939b0) at test/librbd/test_librbd.cc:460
#11 0x00007f284c81f16c in testing::internal::HandleSehExceptionsInMethodIfSupported<testing::Test, void> (object=0x7f2856b939b0, method=&virtual testing::Test::TestBody(),
    location=0x7f284c9caadb "the test body") at ./src/gtest.cc:2078
---Type <return> to continue, or q <return> to quit---
#12 0x00007f284c81a318 in testing::internal::HandleExceptionsInMethodIfSupported<testing::Test, void> (object=0x7f2856b939b0, method=&virtual testing::Test::TestBody(), location=0x7f284c9caadb "the test body")
    at ./src/gtest.cc:2114
#13 0x00007f284c801ad3 in testing::Test::Run (this=0x7f2856b939b0) at ./src/gtest.cc:2151
#14 0x00007f284c8022cc in testing::TestInfo::Run (this=0x7f2856b671a0) at ./src/gtest.cc:2326
#15 0x00007f284c802990 in testing::TestCase::Run (this=0x7f2856b66880) at ./src/gtest.cc:2444
#16 0x00007f284c809458 in testing::internal::UnitTestImpl::RunAllTests (this=0x7f2856b52fb0) at ./src/gtest.cc:4315
#17 0x00007f284c820552 in testing::internal::HandleSehExceptionsInMethodIfSupported<testing::internal::UnitTestImpl, bool> (object=0x7f2856b52fb0,
    method=(bool (testing::internal::UnitTestImpl::*)(testing::internal::UnitTestImpl * const)) 0x7f284c8091c0 <testing::internal::UnitTestImpl::RunAllTests()>,
```

```
    location=0x7f284c9cb330 "auxiliary test code (environments or event listeners)") at ./src/gtes
t.cc:2078
#18 0x00007f284c81b160 in testing::internal::HandleExceptionsInMethodIfSupported<testing::internal
::UnitTestImpl, bool> (object=0x7f2856b52fb0,
    method=(bool (testing::internal::UnitTestImpl::*)(testing::internal::UnitTestImpl * const)) 0x
7f284c8091c0 <testing::internal::UnitTestImpl::RunAllTests()>,
    location=0x7f284c9cb330 "auxiliary test code (environments or event listeners)") at ./src/gtes
t.cc:2114
#19 0x00007f284c808028 in testing::UnitTest::Run (this=0x7f284cdd0760 <testing::UnitTest::GetInsta
nce()::instance>) at ./src/gtest.cc:3929
#20 0x00007f284c413b88 in RUN_ALL_TESTS () at ../src/gmock/gtest/include/gtest/gtest.h:2288
```

Guesing this is a use-after-free since the lock has bogus nr_readers etc.:

```
2  0x00007f284c424e90 in RWLock::unlock (this=0x7f2840a7c950 <main_arena+496>, lockdep=<optimized
out>) at ./common/RWLock.h:93
93          int r = pthread_rwlock_unlock(&L);
(gdb) p *this
$1 = {L = {__data = {__lock = 2, __nr_readers = 32552, __readers_wakeup = 1454987840, __writer_wak
eup = 32552, __nr_readers_queued = 1455408816, __nr_writers_queued = 32552, __writer = 1455105392,

      __shared = 32552, __pad1 = 139811230689168, __pad2 = 139811230536656, __flags = 1454934976},

    __size = "\002\000\000\000(\177\000\000@Z\271V(\177\000\000\260pV(\177\000\000p%\273V(\177\000
\000\220\247\275V(\177\000\000\320S\273V(\177\000\000\300\213\270V(\177\000", __align = 1398097754
19394},
  name = "\311\000\000\000\000\000\000\000\321\001\000\000\000\000\000\000\220\n\273V(\177\000\000
xђ@(\177\000\000'", id = 1455595696, nrlock = {val = 139811230851328}, nwlock = {val = 13981123042
3231},
  track = 112, lockdep = 212}
```

Related issues:

Copied to rbd - Backport #15713: jewel: librbd: resize hang - use after free	**Resolved**

History

#1 - 05/02/2016 11:40 PM - Josh Durgin

I think this hit https://jenkins.ceph.com/job/ceph-pull-requests/5026/console

#2 - 05/03/2016 11:43 AM - Jason Dillaman

- Status changed from New to Need Review

- Assignee set to Jason Dillaman

- Backport set to jewel

PR: https://github.com/ceph/ceph/pull/8907

#3 - 05/03/2016 05:09 PM - Jason Dillaman

- Status changed from Need Review to Pending Backport

#4 - 05/03/2016 05:35 PM - Nathan Cutler

- Copied to Backport #15713: jewel: librbd: resize hang - use after free added

#5 - 05/10/2016 06:03 PM - Jason Dillaman

- Status changed from Pending Backport to Resolved