

Ceph - Bug #45537

librados calls after fork() result in lockdep aborts

05/13/2020 06:35 PM - David Disseldorp

Status:	New	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	David Disseldorp		
Category:	librados		
Target version:			
Source:	Community (dev)	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			
While working on Samba / Ceph integration, I noticed that librados doesn't play nicely with fork(). I wrote a simple reproducer (to be attached) which performs the following: - rados_create(... &ctx1)->rados_connect(ctx1 ...) - fork() - child process: rados_create(... &ctx2)->rados_connect(ctx2 ...) This results in an abort in the lockdep code path.			

History

#1 - 05/13/2020 07:26 PM - David Disseldorp

- File 0001-test-librados-fork-test.patch added

The attached unit test atop current master (3252b19084bf7b6864733b4183eb8ef4d937d691) results in the following abort from the ms_local thread:

```
#2 0x00007fffee266739 in ceph::__ceph_abort (file=0x7fffeebea1398 "/home/david/zram/work/ceph/src/common/lockdep.cc", line=336,
func=0x7fffeebea3080 <lockdep_will_lock(char const*, int, bool, bool)::__PRETTY_FUNCTION__> "int lockdep_will_lock(const char*, int, bool, bool)",
msg=...)
at /home/david/zram/work/ceph/src/common/assert.cc:198
#3 0x00007fffee35e348 in lockdep_will_lock (name=0x555555ce1960 "Messenger::DispatchQueue::local_delivery_lockradosclient", id=20,
force_backtrace=false, recursive=false)
at /home/david/zram/work/ceph/src/common/lockdep.cc:336
#4 0x00007fffee369d6a in ceph::mutex_debug_detail::mutex_debugging_base::_will_lock (this=0x555555cddc00, recursive=false)
at /home/david/zram/work/ceph/src/common/mutex_debug.cc:49
#5 0x00007fffee105882 in ceph::mutex_debug_detail::mutex_debug_impl<false>::lock (this=0x555555cddc00, no_lockdep=false)
at /home/david/zram/work/ceph/src/common/mutex_debug.h:180
#6 0x00007fffee105980 in std::unique_lock<ceph::mutex_debug_detail::mutex_debug_impl<false>>>::lock (this=0x7fffe94930b0) at
/usr/include/c++/7/bits/std_mutex.h:267
#7 0x00007fffee103856 in std::unique_lock<ceph::mutex_debug_detail::mutex_debug_impl<false>>>::unique_lock (this=0x7fffe94930b0, __m=...)
at /usr/include/c++/7/bits/std_mutex.h:197
#8 0x00007fffee43cbba in DispatchQueue::run_local_delivery (this=0x555555cdda48) at
/home/david/zram/work/ceph/src/msg/DispatchQueue.cc:114
#9 0x00007fffee59b436 in DispatchQueue::LocalDeliveryThread::entry (this=0x555555cddcf0) at
/home/david/zram/work/ceph/src/msg/DispatchQueue.h:115
#10 0x00007fffee201ed6 in Thread::entry_wrapper (this=0x555555cddcf0) at /home/david/zram/work/ceph/src/common/Thread.cc:84
#11 0x00007fffee201e54 in Thread::_entry_func (arg=0x555555cddcf0) at /home/david/zram/work/ceph/src/common/Thread.cc:71
#12 0x00007ffff776c4f9 in start_thread () from /lib64/libpthread.so.0
#13 0x00007fffeb0b6f2f in clone () from /lib64/libc.so.6
```

While the main thread is calling _rados_connect():

```
(gdb) thread 2.1
[Switching to thread 2.1 (Thread 0x7ffff7fd84c0 (LWP 86009))]
#0 0x00007ffff7775c3d in __lll_lock_wait () from /lib64/libpthread.so.0
(gdb) bt
#0 0x00007ffff7775c3d in __lll_lock_wait () from /lib64/libpthread.so.0
#1 0x00007ffff776ee95 in pthread_mutex_lock () from /lib64/libpthread.so.0
```

#2 0x00007fffee35e9ed in lockdep_will_unlock (name=0x555555bbfaf0 "AsyncMessenger::lock", id=22) at /home/david/zram/work/ceph/src/common/lockdep.cc:383

#3 0x00007fffee369df8 in ceph::mutex_debug_detail::mutex_debugging_base::_will_unlock (this=0x555555cddde0) at /home/david/zram/work/ceph/src/common/mutex_debug.cc:55

#4 0x00007fffee105917 in ceph::mutex_debug_detail::mutex_debug_impl<false>::unlock (this=0x555555cddde0, no_lockdep=false) at /home/david/zram/work/ceph/src/common/mutex_debug.h:194

#5 0x00007fffee1246e8 in std::lock_guard<ceph::mutex_debug_detail::mutex_debug_impl<false> >::~lock_guard (this=0x7fffffd270, __in_chrg=<optimized out>) at /usr/include/c++/7/bits/std_mutex.h:168

#6 0x00007fffee592e28 in AsyncMessenger::ready (this=0x555555cdd6b0) at /home/david/zram/work/ceph/src/msg/async/AsyncMessenger.cc:334

#7 0x00007ffff7a4a461 in Messenger::add_dispatcher_head (this=0x555555cdd6b0, d=0x555555cd6080) at /home/david/zram/work/ceph/src/msg/Messenger.h:390

#8 0x00007ffff7a41778 in librados::v14_2_0::RadosClient::connect (this=0x555555cd5840) at /home/david/zram/work/ceph/src/librados/RadosClient.cc:275

#9 0x00007ffff79b871e in _rados_connect (cluster=0x555555cd5840) at /home/david/zram/work/ceph/src/librados/librados_c.cc:178

#10 0x000055555566463c in forker_start_proc (cluster_parent=0x555555bb43a0) at /home/david/zram/work/ceph/src/test/librados/misc.cc:383

#11 0x0000555555664c8e in LibRadosMisc_Forker_Test::TestBody (this=0x555555aa09b0) at /home/david/zram/work/ceph/src/test/librados/misc.cc:421

#12 0x00005555556e7f01 in testing::internal::HandleSehExceptionsInMethodIfSupported<testing::Test, void> (object=0x555555aa09b0, method=&virtual testing::Test::TestBody(), location=0x555555710ffb "the test body") at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2593

#13 0x00005555556e0c59 in testing::internal::HandleExceptionsInMethodIfSupported<testing::Test, void> (object=0x555555aa09b0, method=&virtual testing::Test::TestBody(), location=0x555555710ffb "the test body") at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2629

#14 0x00005555556bcb10 in testing::Test::Run (this=0x555555aa09b0) at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2668

#15 0x00005555556bd4a7 in testing::TestInfo::Run (this=0x5555559a55e0) at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2845

#16 0x00005555556bdbba in testing::TestSuite::Run (this=0x5555559e4940) at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2977

#17 0x00005555556c97ea in testing::internal::UnitTestImpl::RunAllTests (this=0x5555559e7c30) at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:5518

#18 0x00005555556e9121 in testing::internal::HandleSehExceptionsInMethodIfSupported<testing::internal::UnitTestImpl, bool> (object=0x5555559e7c30, method=(bool (testing::internal::UnitTestImpl::*)(testing::internal::UnitTestImpl * const)) 0x5555556c93d0 <testing::internal::UnitTestImpl::RunAllTests()>, location=0x5555557119e8 "auxiliary test code (environments or event listeners)") at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2593

#19 0x00005555556e1cfc in testing::internal::HandleExceptionsInMethodIfSupported<testing::internal::UnitTestImpl, bool> (object=0x5555559e7c30, method=(bool (testing::internal::UnitTestImpl::*)(testing::internal::UnitTestImpl * const)) 0x5555556c93d0 <testing::internal::UnitTestImpl::RunAllTests()>, location=0x5555557119e8 "auxiliary test code (environments or event listeners)") at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:2629

#20 0x00005555556c8049 in testing::UnitTest::Run (this=0x555555962a80 <testing::UnitTest::GetInstance()::instance>) at /home/david/zram/work/ceph/src/googletest/googletest/src/gtest.cc:5105

#21 0x00005555556986c2 in RUN_ALL_TESTS () at /home/david/zram/work/ceph/src/googletest/googletest/include/gtest/gtest.h:2472

#22 0x0000555555698577 in main (argc=1, argv=0x7fffffe408) at /home/david/zram/work/ceph/src/test/unit.cc:45

Files

0001-test-librados-fork-test.patch

2.93 KB

05/13/2020

David Disseldorp