

## rbd - Bug #24141

### [rbd-mirror] potential deadlock when running asok 'flush' command

05/16/2018 06:40 AM - Mykola Golub

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>	Mykola Golub		
<b>Category:</b>			
<b>Target version:</b>			
<b>Source:</b>	Q/A	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>	mimic, luminous	<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature:</b>	

#### Description

A case was observed during testing when rbd-mirror got stuck on admin socket "flush" command. The threads that deadlock were the admin socket tread executing "flush" command and an ImageReplayer thread processing a journal replay entry that triggered reregister admin socket command. The second thread was waiting for the admin socket command (flush) to complete, while the "flush" could not proceed because they share the same thread pool.

Thread 4 (Thread 0x7f4fe1b12700 (LWP 24517)):

```
#0 pthread_cond_wait@@GLIBC_2.3.2 () at ../sysdeps/unix/sysv/linux/x86_64/pthread_cond_wait.S:185
#1 0x0000556f5e1bea0b in Cond::Wait (mutex=..., this=0x7f4fe1b0c988) at /home/mgolub/ceph/ceph.ci/src/common/Cond.h:48
#2 C_SaferCond::wait (this=0x7f4fe1b0c910) at /home/mgolub/ceph/ceph.ci/src/common/Cond.h:195
#3 rbd::mirror::(anonymous namespace)::FlushCommand<librbd::ImageCtx>::call (this=<optimized out>, f=<optimized out>, ss=0x7f4fe1b0ca50) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:155
#4 0x0000556f5e1bde96 in rbd::mirror::(anonymous namespace)::ImageReplayerAdminSocketHook<librbd::ImageCtx>::call (this=<optimized out>, command=..., cmdmap=..., format=..., out=...) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:210
#5 0x00007f4fe645913f in AdminSocket::do_accept (this=this@entry=0x556f6071e7f0) at /home/mgolub/ceph/ceph.ci/src/common/admin_socket.cc:396
#6 0x00007f4fe645a2a8 in AdminSocket::entry (this=0x556f6071e7f0) at /home/mgolub/ceph/ceph.ci/src/common/admin_socket.cc:247
#7 0x00007f4fe58e16f in ?? () from /usr/lib/x86_64-linux-gnu/libstdc++.so.6
#8 0x00007f4fe5fcd6ba in start_thread (arg=0x7f4fe1b12700) at pthread_create.c:333
#9 0x00007f4fe535841d in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:109
```

Thread 81 (Thread 0x7f4ebc9f9700 (LWP 25731)):

```
#0 pthread_cond_wait@@GLIBC_2.3.2 () at ../sysdeps/unix/sysv/linux/x86_64/pthread_cond_wait.S:185
#1 0x00007f4fe58e7e2c in std::condition_variable::wait(std::unique_lock<std::mutex>&) () from /usr/lib/x86_64-linux-gnu/libstdc++.so.6
#2 0x00007f4fe645472b in std::condition_variable::wait<AdminSocket::unregister_command(std::string_view)::<lambda()>> (&__p=..., &__lock=..., this=0x556f6071e830) at /usr/include/c++/7/condition_variable:99
#3 AdminSocket::unregister_command (this=0x556f6071e7f0, command=...) at /home/mgolub/ceph/ceph.ci/src/common/admin_socket.cc:474
#4 0x0000556f5e1c1e81 in rbd::mirror::(anonymous namespace)::ImageReplayerAdminSocketHook<librbd::ImageCtx>::~~ImageReplayerAdminSocketHook (this=0x7f4fcc0633b0, __in_chrg=<optimized out>) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:197
#5 rbd::mirror::(anonymous namespace)::ImageReplayerAdminSocketHook<librbd::ImageCtx>::~~ImageReplayerAdminSocketHook (this=0x7f4fcc0633b0, __in_chrg=<optimized out>) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:202
#6 rbd::mirror::ImageReplayer<librbd::ImageCtx>::unregister_admin_socket_hook (this=this@entry=0x7f4fcc065ab0) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:1800
#7 0x0000556f5e1c61bd in rbd::mirror::ImageReplayer<librbd::ImageCtx>::reregister_admin_socket_hook (this=this@entry=0x7f4fcc065ab0) at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:1813
```

```
#8 0x0000556f5e1c6387 in rbd::mirror::ImageReplayer<librbd::ImageCtx>::handle_process_entry_ready
(this=0x7f4fcc065ab0, r=<optimized out>)
  at /home/mgolub/ceph/ceph.ci/src/tools/rbd_mirror/ImageReplayer.cc:1224
#9 0x0000556f5e157939 in Context::complete (this=0x7f4f680563e0, r=<optimized out>) at /home/mgol
ub/ceph/ceph.ci/src/include/Context.h:77
#10 0x00007f4fe6482d37 in ThreadPool::worker (this=0x7f4ed800b3d0, wt=<optimized out>) at /home/mg
olub/ceph/ceph.ci/src/common/WorkQueue.cc:120
#11 0x00007f4fe6483c60 in ThreadPool::WorkThread::entry (this=<optimized out>) at /home/mgolub/cep
h/ceph.ci/src/common/WorkQueue.h:448
#12 0x00007f4fe5fcd6ba in start_thread (arg=0x7f4ebcff9700) at pthread_create.c:333
#13 0x00007f4fe535841d in clone () at ../sysdeps/unix/sysv/linux/x86_64/clone.S:109
```

#### Related issues:

Copied to rbd - Backport #24155: mimic: [rbd-mirror] potential deadlock when ...	<b>Resolved</b>
Copied to rbd - Backport #24156: luminous: [rbd-mirror] potential deadlock wh...	<b>Resolved</b>

#### History

##### #1 - 05/16/2018 11:35 AM - Mykola Golub

- Status changed from *In Progress* to *Fix Under Review*

PR: <https://github.com/ceph/ceph/pull/22027>

##### #2 - 05/16/2018 06:09 PM - Jason Dillaman

- Status changed from *Fix Under Review* to *Pending Backport*

##### #3 - 05/17/2018 08:28 AM - Nathan Cutler

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

##### #4 - 05/17/2018 08:40 AM - Nathan Cutler

- Copied to Backport #24155: *mimic*: [rbd-mirror] potential deadlock when running *asok 'flush'* command added

##### #5 - 05/17/2018 08:41 AM - Nathan Cutler

- Copied to Backport #24156: *luminous*: [rbd-mirror] potential deadlock when running *asok 'flush'* command added

##### #6 - 07/06/2018 04:34 AM - Nathan Cutler

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