Ceph - Bug #3613

Objecter::scan_requests crash

12/12/2012 05:35 PM - Greg Farnum

Status:	Resolved	**% Done:**	0%
Priority:	Urgent	**Spent time:**	0.00 hour
Assignee:	Sage Weil		
Category:	Objecter		
Target version:	v0.56		
Source:	Development	**Affected Versions:**	
Tags:		**ceph-qa-suite:**	
Backport:		**Pull request ID:**	
Regression:	No	**Crash signature (v1):**	
Severity:	3 - minor	**Crash signature (v2):**	
Reviewed:			

Description

One of the failures in #3459 was due to testrados_watch_notify crashing:

```
#0  0x00007fa9efc199ec in Objecter::scan_requests (this=0xdaef80, skipped_map=false, need_resend=..., need_resend_linger=...) at osdc/Objecter.cc:445
#1  0x00007fa9efc1a9e8 in Objecter::handle_osd_map (this=0xdaef80, m=0x7fa9e0000bf0) at osdc/Objecter.cc:528
#2  0x00007fa9efbf8a0e in librados::RadosClient::_dispatch (this=0xda9b30, m=0x7fa9e0000bf0) at librados/RadosClient.cc:313
#3  0x00007fa9efbf87df in librados::RadosClient::ms_dispatch (this=0xda9b30, m=0x7fa9e0000bf0) at librados/RadosClient.cc:280
#4  0x00007fa9efd25fe7 in Messenger::ms_deliver_dispatch (this=0xdae1b0, m=0x7fa9e0000bf0) at msg/Messenger.h:549
#5  0x00007fa9efd256f5 in DispatchQueue::entry (this=0xdae298) at msg/DispatchQueue.cc:107
#6  0x00007fa9efdc074c in DispatchQueue::DispatchThread::entry (this=0xdae3b0) at msg/DispatchQueue.h:85
#7  0x00007fa9efe2b119 in Thread::_entry_func (arg=0xdae3b0) at common/Thread.cc:41
#8  0x00007fa9ee9dbe9a in start_thread () from /lib/x86_64-linux-gnu/libpthread.so.0
#9  0x00007fa9eece34bd in klogctl () from /lib/x86_64-linux-gnu/libc.so.6
#10 0x0000000000000000 in ?? ()
```

See ubuntu@teuthology:/a/sage-2012-12-11_19:47:13-rados-wip-3459-testing-basic/11922. I cannot for my life make sense out of the claimed segfault, although the op does look a little suspect:

```
(gdb) p op->pgid
$1 = {m_pool = 33, m_seed = 0, m_preferred = 0}
```

I do notice that the watch-notify handling in the Objecter just got changed though, so maybe that's involved.

Associated revisions

Revision 4bf90782 - 12/16/2012 01:45 AM - Sage Weil

osdc/Objecter: prevent pool dne check from invalidating scan_requests iterator

We iterate over ops and, if the pool dne and other conditions are true,
we will immediately return ENOENT and cancel an op. Increment the
iterator at the top of the loop to avoid invalidating it.

We also need to switch to a map<>, because hash_map<> mutations may
invalidate any/all iterators.

Fixes: #3613
Signed-off-by: Sage Weil <sage@inktank.com>

History

#1 - 12/13/2012 10:33 AM - Sage Weil

- Priority changed from High to Urgent

#2 - 12/13/2012 12:43 PM - Josh Durgin

I don't see any evidence pointing to the recent objecter changes yet. The op->ops vector seems to be invalid though.

#3 - 12/15/2012 05:44 PM - Sage Weil

- Status changed from New to 7

- Assignee changed from Josh Durgin to Sage Weil

the pool dne check invalidated the iterator. switching to map<> and incrementing hte iterator at hte top of the loop

#4 - 12/15/2012 08:23 PM - Sage Weil

- Status changed from 7 to Resolved

4bf9078286d58c2cd4e85cb8b31411220a377092

passed 100 iterations of the test (previously failed after ~15).