

fs - Bug #14256

mds: objecter assert on shutdown

01/06/2016 06:18 AM - Greg Farnum

Status:	Resolved	Start date:	01/06/2016
Priority:	Normal	Due date:	
Assignee:	Adam Emerson	% Done:	0%
Category:	Correctness/Safety	Estimated time:	0.00 hour
Target version:		Affected Versions:	
Source:	Development	ceph-qa-suite:	
Tags:		Component(FS):	MDS, osdc
Backport:		Labels (FS):	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Reviewed:			

Description

http://pulpito.ceph.com/gregf-2015-12-21_23:08:59-fs-master---basic-smithi/1782/

Only saw this once so far and it might have a cause elsewhere, but I didn't see any similar reports so logging this for reference.

```
2015-12-22T21:54:33.134 INFO:tasks.ceph.mds.a-s.smithi015.stderr:osdc/Objecter.cc: In function 'void Objecter::shutdown()' thread e46e700 time 2015-12-23 00:54:33.102493
2015-12-22T21:54:33.134 INFO:tasks.ceph.mds.a-s.smithi015.stderr:osdc/Objecter.cc: 477: FAILED assert(tick_event == 0)
2015-12-22T21:54:33.156 INFO:tasks.ceph.mds.a-s.smithi015.stderr: ceph version 10.0.1-721-g09a3f69 (09a3f69e6f6a42d3a517b843d98706e8850edfac)
2015-12-22T21:54:33.156 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x85) [0x6d2775]
2015-12-22T21:54:33.156 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 2: (()+0x481d0a) [0x589d0a]
2015-12-22T21:54:33.157 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 3: (MDSRankDispatcher::shutdown()+0x261) [0x3036e1]
2015-12-22T21:54:33.157 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 4: (MDSDaemon::suicide()+0x22f) [0x2edb6f]
2015-12-22T21:54:33.157 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 5: (MDSDaemon::handle_signal(int)+0x8b) [0x2edd0b]
2015-12-22T21:54:33.157 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 6: (SignalHandler::entry()+0x127) [0x5de4d7]
2015-12-22T21:54:33.158 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 7: (()+0x7df5) [0x539fdf5]
2015-12-22T21:54:33.158 INFO:tasks.ceph.mds.a-s.smithi015.stderr: 8: (clone()+0x6d) [0x66f01ad]
2015-12-22T21:54:33.158 INFO:tasks.ceph.mds.a-s.smithi015.stderr: NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
2015-12-22T21:54:33.190 INFO:tasks.ceph.mds.a-s.smithi015.stderr:2015-12-23 00:54:33.145574 e46e700 -1 osdc/Objecter.cc: In function 'void Objecter::shutdown()' thread e46e700 time 2015-12-23 00:54:33.102493
```

Associated revisions

Revision 9179ce8c - 01/07/2016 09:55 PM - Adam Emerson

osdc: Fix race condition with tick_event and shutdown

- Clear the tick_event whether it was in the timer queue or not.
- Make sure we don't schedule a new tick_event if someone calls shutdown while tick() is running
- Get rid of some assertions that aren't relevant

Fixes #14256

History

#1 - 01/07/2016 03:38 PM - John Spray

Hmm, looks like Objecter is making the assumption that if tick_event is set then it must also be in ceph_timer::events. That's not the case, because if we happen to enter shutdown just as the event is getting called, it will have been removed from events (done before the callback), but tick_event will still be set (it's cleared or reset during the callback).

Introduced by:

```
commit ecf2bebe99b43735b930406cb1fedf51283a62f0
Author: Adam C. Emerson <aemerson@redhat.com>
Date: Mon Sep 14 12:19:58 2015 -0400
```

```
time: Update OSDC for C++11 Time
```

```
Signed-off-by: Adam C. Emerson <aemerson@redhat.com>
```

#2 - 01/07/2016 06:11 PM - Adam Emerson

- Category changed from 47 to 46
- Assignee set to Adam Emerson

#3 - 01/07/2016 06:12 PM - Adam Emerson

- Status changed from New to In Progress

#4 - 01/07/2016 06:56 PM - Adam Emerson

- File bugfix.patch added
- Status changed from In Progress to Testing

This patch should fix it. I'll run make check and push.

#5 - 01/07/2016 07:29 PM - Adam Emerson

- Status changed from Testing to Pending Upstream

Pushed upstream as:

```
commit b308791290b48d1142eb8c222086ffe7509e3449
Author: Adam C. Emerson <aemerson@redhat.com>
Date: Thu Jan 7 14:15:34 2016 -0500
```

```
osdc: Fix race condition with tick_event and shutdown
```

- Clear the tick_event whether it was in the timer queue or not.
- Make sure we don't schedule a new tick_event if someone calls shutdown while tick() is running
- Make tick_event atomic so we can check it without a lock/while only holding a read lock

```
Fixes #14256
```

#6 - 01/07/2016 09:24 PM - John Spray

Link to the pull request for convenience:

<https://github.com/ceph/ceph/pull/7151>

#7 - 01/07/2016 09:25 PM - John Spray

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

I don't know if we ever used "Pending upstream" before, usually when a PR is outstanding we use "Needs review"

#8 - 01/13/2016 02:24 AM - Sage Weil

- Status changed from *Need Review* to *Resolved*

#9 - 07/13/2016 12:25 AM - Greg Farnum

- Category changed from 46 to *Correctness/Safety*

- Component(FS) *MDS, osdc added*

Files

bugfix.patch	2.43 KB	01/07/2016	Adam Emerson
--------------	---------	------------	--------------