

## RADOS - Bug #40421

### osd: lost op?

06/19/2019 01:52 AM - Patrick Donnelly

<b>Status:</b>	New	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v15.0.0	<b>Spent time:</b>	0.00 hour
<b>Source:</b>	Q/A	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Component(RADOS):</b>	OSD
<b>Severity:</b>	3 - minor	<b>Pull request ID:</b>	

#### Description

Could use some help figuring out what happened here.

MDS got stuck in up:replay because it didn't get a reply to this message:

```
2019-06-15T21:19:48.822+0000 7fdf18a16700 1 -- [v2:172.21.15.135:6834/3199686756,v1:172.21.15.135:6835/3199686756] --> [v2:172.21.15.120:6808/12480,v1:172.21.15.120:6810/12480] -- osd_op(unknown.0.13:7 2.4 2:292cf221:::200.00000000:head [read 0~0] snapc 0=[] ondisk+read+known_if_redirected+full_force e25) v8 -- 0x55bec2f2f200 con 0x55bec2ef5b00
```

It does get a reply to an op just before that:

```
2019-06-15T21:19:48.822+0000 7fdf19217700 1 -- [v2:172.21.15.135:6834/3199686756,v1:172.21.15.135:6835/3199686756] --> [v2:172.21.15.120:6808/12480,v1:172.21.15.120:6810/12480] -- osd_op(unknown.0.13:5 2.6 2:654134d2:::mds0_openfiles.0:head [omap-get-header,omap-get-vals] snapc 0=[] ondisk+read+known_if_redirected+full_force e25) v8 -- 0x55bec20ead00 con 0x55bec2ef5b00
```

```
...
2019-06-15T21:19:48.998+0000 7fdf22a2a700 1 -- [v2:172.21.15.135:6834/3199686756,v1:172.21.15.135:6835/3199686756] <== osd.1 v2:172.21.15.120:6808/12480 1 ==== osd_op_reply(5 mds0_openfiles.0 [omap-get-header,omap-get-vals] v0'0 uv25 ondisk = 0) v8 ==== 202+0+8394 (crc 0 0 0) 0x55bec2cede40 con 0x55bec2ef5b00
```

^ The ops were received by the osd here:

```
2019-06-15T21:19:48.977+0000 7f35de67a700 1 -- [v2:172.21.15.120:6808/12480,v1:172.21.15.120:6810/12480] <== mds.0 v2:172.21.15.135:6834/3199686756 1 ==== osd_op(mds.0.13:5 2.6 2.4b2c82a6 (undecoded) ondisk+read+known_if_redirected+full_force e25) v8 ==== 263+0+16 (crc 0 0 0) 0x55fcf7172a00 con 0x55fcelc67600
```

```
...
2019-06-15T21:19:48.977+0000 7f35de67a700 1 -- [v2:172.21.15.120:6808/12480,v1:172.21.15.120:6810/12480] <== mds.0 v2:172.21.15.135:6834/3199686756 2 ==== osd_op(mds.0.13:7 2.4 2.844f3494 (undecoded) ondisk+read+known_if_redirected+full_force e25) v8 ==== 221+0+0 (crc 0 0 0) 0x55fcf7172700 con 0x55fcelc67600
```

From: /ceph/teuthology-archive/pdonnell-2019-06-15\_02:00:55-kcephfs-wip-pdonnell-testing-20190614.222049-distro-basic-smithi/4035802

/remote/smithi120/log/ceph-osd.1.log.1.gz

Looks like the message was just not ever processed?

## History

---

### #1 - 06/26/2019 09:10 PM - Greg Farnum

Has this recurred on master? What PRs were in that test branch?

### #2 - 06/28/2019 01:53 AM - Patrick Donnelly

Greg Farnum wrote:

Has this recurred on master? What PRs were in that test branch?

I haven't looked at a recent batch of tests to see if it has recurred. Here's the branch:

<https://github.com/ceph/ceph-ci/commits/wip-pdonnell-testing-20190614.222049>

None of the PRs should cause this IMO and many/all were merged already.

### #3 - 08/21/2019 08:01 PM - Greg Farnum

- *Priority changed from High to Normal*