

## Ceph - Bug #20427

### osd: client IOPS drops to zero frequently

06/27/2017 10:17 AM - Alexey Sheplyakov

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Normal	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b>	<b>Affected Versions:</b> v10.2.7
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b> jewel,kraken	<b>Pull request ID:</b>
<b>Regression:</b> No	<b>Crash signature (v1):</b>
<b>Severity:</b> 3 - minor	<b>Crash signature (v2):</b>
<b>Reviewed:</b>	

**Description**

[From <http://www.spinics.net/lists/ceph-devel/msg37163.html>]

At Alibaba, we experienced unstable performance with Jewel on one production cluster, and we can easily reproduce it now with several small test clusters. One test cluster has 30 SSDs, and another test one has 120 SSDs, we are using filestore+async messenger on the backend and fio+librbd to test them. When this issue happens, client fio IOPS drops to zero (or close to zero) frequently during fio runs. And the durations of those drops were very short, about 1 second or so.

For the 30 SSDs test cluster, we use 135 client fio writing into 135 rbd images individually, each fio has only 1 job and rate limit is 3MB/s. On this fresh created test cluster, for all 135 client fio runs, during first 15 minutes or so, client IOPS were very stable and each OSD server's throughput was very stable as well. After 15 minutes and 360 GB data written, the test cluster entered an unstable state, client fio IOPS dropped to zero (or close) frequently and each OSD server's throughput became very spiky as well (from 500MB/s to less 1MB/s). We tried let all fio keeping writing for about 16 hours, cluster was still in this swing state.

This is very easily reproducible. I don't think it's caused by filestore folder splitting, since they were all done during the first 15 minutes. And also, OSD server mem/cpu/disk were far from saturated. One thing we noticed from perf counter is that op\_latency increased from 0.7 ms to >20 ms after entering this unstable state. Is this normal Jewel/filestore behavior? Anyone knows what causes it?

**Related issues:**

Copied to Ceph - Backport #20428: jewel: osd: client IOPS drops to zero frequ...	<b>Resolved</b>
Copied to Ceph - Backport #20443: kraken: osd: client IOPS drops to zero freq...	<b>Resolved</b>

#### History

##### #1 - 06/27/2017 10:19 AM - Alexey Sheplyakov

The bug has been fixed by <https://github.com/ceph/ceph/pull/15891>

##### #2 - 06/27/2017 10:19 AM - Alexey Sheplyakov

- Backport set to jewel,kraken

**#3 - 06/27/2017 10:45 AM - Nathan Cutler**

- Subject changed from client IOPS drops to zero frequently to osd: client IOPS drops to zero frequently

**#4 - 06/27/2017 11:08 AM - Alexey Sheplyakov**

- Status changed from New to Pending Backport

**#5 - 06/27/2017 11:13 AM - Alexey Sheplyakov**

- Copied to Backport #20428: jewel: osd: client IOPS drops to zero frequently added

**#6 - 06/28/2017 04:48 AM - Alexey Sheplyakov**

- Copied to Backport #20443: kraken: osd: client IOPS drops to zero frequently added

**#7 - 08/29/2017 07:16 AM - Nathan Cutler**

- Status changed from Pending Backport to Resolved