

## RADOS - Bug #37886

### Adding back the IOPS line for client and recovery IO in cluster logs

01/13/2019 03:37 AM - Vikhyat Umrao

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Normal	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b> Neha Ojha	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b>	<b>Affected Versions:</b>
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b> mimic luminous	<b>Component(RADOS):</b>
<b>Regression:</b> No	<b>Pull request ID:</b> 26105
<b>Severity:</b> 3 - minor	<b>Crash signature (v1):</b>
<b>Reviewed:</b>	<b>Crash signature (v2):</b>

**Description**

In luminous cluster logs, client and recovery IOPS log lines were removed it used to help a lot in RCA. These logs were present in jewel and before releases.

RHBZ - [https://bugzilla.redhat.com/show\\_bug.cgi?id=1664440](https://bugzilla.redhat.com/show_bug.cgi?id=1664440)

- Sample from a cluster which was under recovery.

```
018-12-10 03:23:09.149191 mon.0 192.168.124.15:6789/0 8668553 : cluster [INF] pgmap v50874961: 265
6 pgs: 1 active+undersized+degraded+remapped+backfilling, 1 active+remapped+backfill_toofull, 10 a
ctive+degraded+remapped+wait_backfill, 4 active+recovery_wait+undersized+degraded+remapped, 18 act
ive+undersized+degraded+remapped+wait_backfill, 17 active+recovery_wait+degraded+remapped, 150 act
ive+recovery_wait+degraded, 68 active+remapped+wait_backfill, 2387 active+clean; 211 TB data, 318
TB used, 75859 GB / 392 TB avail; 10845 kB/s rd, 10191 kB/s wr, 3507 op/s; 738922/406321200 object
s degraded (0.182%); 3751024/406321200 objects misplaced (0.923%); 67316 kB/s, 21 objects/s recove
ring
```

- Sample from a cluster which was HEALTH\_OK.

```
2018-12-29 03:14:22.348339 mon.0 192.166.124.23:6789/0 41474226 : cluster [INF] pgmap v39151502: 8
768 pgs: 8768 active+clean; 6658 GB data, 19621 GB used, 1196 TB / 1215 TB avail; 43564 B/s rd, 54
96 kB/s wr, 1675 op/s
2018-12-29 03:14:23.369658 mon.0 192.166.124.23:6789/0 41474227 : cluster [INF] pgmap v39151503: 8
768 pgs: 8768 active+clean; 6658 GB data, 19621 GB used, 1196 TB / 1215 TB avail; 1770 kB/s rd, 93
11 kB/s wr, 1528 op/s
2018-12-29 03:14:24.375619 mon.0 192.166.124.23:6789/0 41474228 : cluster [INF] pgmap v39151504: 8
768 pgs: 8768 active+clean; 6658 GB data, 19621 GB used, 1196 TB / 1215 TB avail; 1975 kB/s rd, 10
580 kB/s wr, 1618 op/s
```

**Related issues:**

Copied to RADOS - Backport #38107: mimic: Adding back the IOPS line for clien...	<b>Resolved</b>
Copied to RADOS - Backport #38108: luminous: Adding back the IOPS line for cl...	<b>Resolved</b>

#### History

#1 - 01/13/2019 03:41 AM - Vikhyat Umrao

- Description updated

#2 - 01/13/2019 03:41 AM - Vikhyat Umrao

- Description updated

### #3 - 01/13/2019 03:46 AM - Vikhyat Umrao

I did some search in the code and testing with vstart cluster and we need to bring back this commit.

<https://github.com/ceph/ceph/commit/dcc4c52ee8a8c33fabd4d56a42b9667658ed730c>

- After backporting this patch in luminous branch.

```
tail -f out/cluster.mon.a.log
```

```
2019-01-13 06:58:37.966486 mgr.x client.4103 10.65.200.47:0/512226364 79 : cluster [INF] pgmap v90: 16 pgs: 16
  active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:39.967002 mgr.x client.4103 10.65.200.47:0/512226364 80 : cluster [INF] pgmap v91: 16 pgs: 16
  active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:41.099915 mon.a mon.0 10.65.200.47:40557/0 96 : cluster [WRN] Health check failed: 1 osds dow
n (OSD_DOWN)
2019-01-13 06:58:44.534572 mon.a mon.0 10.65.200.47:40557/0 100 : cluster [WRN] Health check failed: Reduced d
ata availability: 2 pgs inactive, 4 pgs peering (PG_AVAILABILITY)
2019-01-13 06:58:46.649394 mon.a mon.0 10.65.200.47:40557/0 101 : cluster [INF] Health check cleared: OSD_DOWN
(was: 1 osds down)
2019-01-13 06:58:46.666220 mon.a mon.0 10.65.200.47:40557/0 102 : cluster [WRN] Health check failed: Degraded
data redundancy: 21/63 objects degraded (33.333%), 6 pgs degraded (PG_DEGRADED)
2019-01-13 06:58:46.740718 mon.a mon.0 10.65.200.47:40557/0 103 : cluster [INF] osd.2 10.65.200.47:6809/5046 b
oot
2019-01-13 06:58:50.020969 mon.a mon.0 10.65.200.47:40557/0 106 : cluster [INF] Health check cleared: PG_AVAIL
ABILITY (was: Reduced data availability: 2 pgs inactive, 2 pgs peering)
2019-01-13 06:58:44.803209 osd.2 osd.2 10.65.200.47:6809/5046 1 : cluster [WRN] Monitor daemon marked osd.2 do
wn, but it is still running
2019-01-13 06:58:41.967553 mgr.x client.4103 10.65.200.47:0/512226364 81 : cluster [INF] pgmap v93: 16 pgs: 4
  stale+active+clean, 12 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:43.968192 mgr.x client.4103 10.65.200.47:0/512226364 82 : cluster [INF] pgmap v95: 16 pgs: 2
  stale+active+clean, 4 peering, 10 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:45.968905 mgr.x client.4103 10.65.200.47:0/512226364 83 : cluster [INF] pgmap v96: 16 pgs: 6
  active+undersized+degraded, 4 peering, 6 active+undersized; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
; 21/63 objects degraded (33.333%)
2019-01-13 06:58:47.969595 mgr.x client.4103 10.65.200.47:0/512226364 84 : cluster [INF] pgmap v99: 16 pgs: 6
  active+undersized+degraded, 4 peering, 6 active+undersized; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
; 21/63 objects degraded (33.333%)
2019-01-13 06:58:49.970277 mgr.x client.4103 10.65.200.47:0/512226364 85 : cluster [INF] pgmap v100: 16 pgs: 4
  active+clean, 5 active+undersized+degraded, 2 peering, 5 active+undersized; 2.19KiB data, 374GiB used, 1.39Ti
B / 1.76TiB avail; 17/63 objects degraded (26.984%)
2019-01-13 06:58:52.484662 mon.a mon.0 10.65.200.47:40557/0 107 : cluster [INF] Health check cleared: PG_DEGRA
DED (was: Degraded data redundancy: 17/63 objects degraded (26.984%), 5 pgs degraded)
2019-01-13 06:58:52.484708 mon.a mon.0 10.65.200.47:40557/0 108 : cluster [INF] Cluster is now healthy
2019-01-13 06:58:51.970935 mgr.x client.4103 10.65.200.47:0/512226364 86 : cluster [INF] pgmap v101: 16 pgs: 1
  4 active+clean, 2 peering; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:53.971574 mgr.x client.4103 10.65.200.47:0/512226364 87 : cluster [INF] pgmap v102: 16 pgs: 1
  6 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:55.972200 mgr.x client.4103 10.65.200.47:0/512226364 88 : cluster [INF] pgmap v103: 16 pgs: 1
  6 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:57.972782 mgr.x client.4103 10.65.200.47:0/512226364 89 : cluster [INF] pgmap v104: 16 pgs: 1
  6 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 06:58:59.973364 mgr.x client.4103 10.65.200.47:0/512226364 90 : cluster [INF] pgmap v105: 16 pgs: 1
  6 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail

2019-01-13 07:01:06.008118 mgr.x client.4103 10.65.200.47:0/512226364 153 : cluster [INF] pgmap v168: 16 pgs:
  16 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 07:01:08.008697 mgr.x client.4103 10.65.200.47:0/512226364 154 : cluster [INF] pgmap v169: 16 pgs:
  16 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
2019-01-13 07:01:10.009218 mgr.x client.4103 10.65.200.47:0/512226364 155 : cluster [INF] pgmap v170: 16 pgs:
  16 active+clean; 2.19KiB data, 374GiB used, 1.39TiB / 1.76TiB avail
```

#### #4 - 01/13/2019 03:48 AM - Vikhyat Umrao

- Subject changed from Luminous - adding back the IOPS line for client and recovery IO in cluster logs to Adding back the IOPS line for client and recovery IO in cluster logs

#### #5 - 01/14/2019 06:01 PM - Vikhyat Umrao

- Description updated

#### #6 - 01/21/2019 06:24 PM - Vikhyat Umrao

- Backport set to mimic luminous

#### #7 - 01/23/2019 09:34 PM - Vikhyat Umrao

- Assignee changed from Vikhyat Umrao to Neha Ojha

#### #8 - 01/24/2019 03:49 AM - Neha Ojha

- Status changed from In Progress to Fix Under Review

- Pull request ID set to 26105

#### #9 - 01/26/2019 12:14 AM - Vikhyat Umrao

Hi Neha,

As discussed I did some testing in luminous branch after adding this patch and changes looks great and working as expected as they should. Below are my test results.

Patch in luminous branch:

```
$ git diff
diff --git a/src/common/options.cc b/src/common/options.cc
index 231a765..1c474f3 100644
--- a/src/common/options.cc
+++ b/src/common/options.cc
@@ -444,7 +444,7 @@ std::vector<Option> get_global_options() {
     .set_description(""),

     Option("mon_cluster_log_file_level", Option::TYPE_STR, Option::LEVEL_ADVANCED)
-    .set_default("info")
+    .set_default("debug")
     .set_description(""),

     Option("mon_cluster_log_to_graylog", Option::TYPE_STR, Option::LEVEL_ADVANCED)
diff --git a/src/mgr/DaemonServer.cc b/src/mgr/DaemonServer.cc
index 87c399b..7799478 100644
--- a/src/mgr/DaemonServer.cc
+++ b/src/mgr/DaemonServer.cc
@@ -1457,6 +1457,7 @@ void DaemonServer::send_report()
     pg_map.get_health_checks(g_ceph_context, osdmap,
                             &m->health_checks);

+
+     dout(10) << m->health_checks.checks.size() << " health checks"
+         << endl;
+     dout(20) << "health checks:\n";
@@ -1464,6 +1465,9 @@ void DaemonServer::send_report()
     jf.dump_object("health_checks", m->health_checks);
     jf.flush(*_dout);
     *_dout << endl;
+     if (osdmap.require_osd_release >= CEPH_RELEASE_LUMINOUS) {
+         clog->debug() << "pgmap v" << pg_map.version << ": " << pg_map;
+     }
+ });
+ });
```

- After patching cluster logs in vstart cluster.

```
$ tail -f out/cluster.mon.a.log
2019-01-26 01:34:59.247787 mgr.x client.4102 10.65.200.47:0/3098712987 58 : cluster [DBG] pgmap v70: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:01.248397 mgr.x client.4102 10.65.200.47:0/3098712987 59 : cluster [DBG] pgmap v71: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:03.248929 mgr.x client.4102 10.65.200.47:0/3098712987 60 : cluster [DBG] pgmap v72: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:05.249529 mgr.x client.4102 10.65.200.47:0/3098712987 61 : cluster [DBG] pgmap v73: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:07.250140 mgr.x client.4102 10.65.200.47:0/3098712987 62 : cluster [DBG] pgmap v74: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:09.250657 mgr.x client.4102 10.65.200.47:0/3098712987 63 : cluster [DBG] pgmap v75: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:11.251176 mgr.x client.4102 10.65.200.47:0/3098712987 64 : cluster [DBG] pgmap v76: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:13.251754 mgr.x client.4102 10.65.200.47:0/3098712987 65 : cluster [DBG] pgmap v77: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:15.252305 mgr.x client.4102 10.65.200.47:0/3098712987 66 : cluster [DBG] pgmap v78: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:17.252795 mgr.x client.4102 10.65.200.47:0/3098712987 67 : cluster [DBG] pgmap v79: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:19.253354 mgr.x client.4102 10.65.200.47:0/3098712987 68 : cluster [DBG] pgmap v80: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:21.253884 mgr.x client.4102 10.65.200.47:0/3098712987 69 : cluster [DBG] pgmap v81: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:23.254424 mgr.x client.4102 10.65.200.47:0/3098712987 70 : cluster [DBG] pgmap v82: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:25.254956 mgr.x client.4102 10.65.200.47:0/3098712987 71 : cluster [DBG] pgmap v83: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 01:35:27.255491 mgr.x client.4102 10.65.200.47:0/3098712987 72 : cluster [DBG] pgmap v84: 16 pgs: 1
6 active+clean; 2.19KiB data, 476GiB used, 1.29TiB / 1.76TiB avail
```

- Testing slow request with vstart cluster to verify if all looks good changing log level to DBG.

```
$ bin/ceph daemon mon.a config show | grep mon_cluster_log_file_level
"mon_cluster_log_file_level": "debug",
```

```
$ bin/ceph daemon osd.0 config show | grep osd_op_complaint
"osd_op_complaint_time": "1.000000",
```

```
bin/ceph osd pool create scbench 8
bin/rados bench -p scbench 300 write
```

<.....>

```
2019-01-26 05:25:20.266225 min lat: 0.639216 max lat: 11.8141 avg lat: 5.87682
  sec Cur ops   started finished avg MB/s  cur MB/s last lat (s)  avg lat (s)
  300    16     823     807   10.7588    0      -    5.87682
  301    16     823     807   10.723    0      -    5.87682
  302    16     823     807   10.6875    0      -    5.87682
Total time run:           302.776360
Total writes made:       824
Write size:              4194304
Object size:            4194304
Bandwidth (MB/sec):     10.8859
Stddev Bandwidth:       10.0805
Max bandwidth (MB/sec): 56
Min bandwidth (MB/sec): 0
Average IOPS:           2
Stddev IOPS:            2
Max IOPS:               14
Min IOPS:               0
Average Latency (s):    5.87837
Stddev Latency (s):    2.50552
Max latency (s):        11.8141
Min latency (s):        0.427808
Cleaning up (deleting benchmark objects)
Removed 824 objects
Clean up completed and total clean up time :2.755849
```

- We have slow requests and IOPS lines both a slow requests comes at WARN level.

```
2019-01-26 05:25:21.529681 osd.1 osd.1 10.65.200.47:6805/27875 1655 : cluster [WRN] slow request 4.880542 seconds old, received at 2019-01-26 05:25:16.649059: osd_repop(client.24143.0:817 3.1 e18/15 3:935e6664::benchmark_data_buildnode_29015_object816:head v 18'109) currently journaled_completion_queued
2019-01-26 05:25:21.529688 osd.1 osd.1 10.65.200.47:6805/27875 1656 : cluster [WRN] slow request 4.720721 seconds old, received at 2019-01-26 05:25:16.808880: osd_repop(client.24143.0:818 3.1 e18/15 3:853d05b1::benchmark_data_buildnode_29015_object817:head v 18'110) currently journaled_completion_queued
2019-01-26 05:25:21.529694 osd.1 osd.1 10.65.200.47:6805/27875 1657 : cluster [WRN] slow request 4.716851 seconds old, received at 2019-01-26 05:25:16.812750: osd_repop(client.24143.0:819 3.5 e18/15 3:a9ec099e::benchmark_data_buildnode_29015_object818:head v 18'106) currently journaled_completion_queued
2019-01-26 05:25:22.493178 osd.1 osd.1 10.65.200.47:6805/27875 1658 : cluster [WRN] 19 slow requests, 5 included below; oldest blocked for > 6.443075 secs
2019-01-26 05:25:22.493183 osd.1 osd.1 10.65.200.47:6805/27875 1659 : cluster [WRN] slow request 5.677283 seconds old, received at 2019-01-26 05:25:16.815800: osd_repop(client.24143.0:820 3.6 e18/15 3:632d8071::benchmark_data_buildnode_29015_object819:head v 18'100) currently journaled_completion_queued
2019-01-26 05:25:22.493187 osd.1 osd.1 10.65.200.47:6805/27875 1660 : cluster [WRN] slow request 6.439246 seconds old, received at 2019-01-26 05:25:16.053836: osd_op(client.24143.0:806 3.4 3:360c75f3::benchmark_data_buildnode_29015_object805:head [set-alloc-hint object_size 4194304 write_size 4194304,write 0~4194304] snapc 0=[] ondisk+write+known_if_redirected e18) currently commit_sent
2019-01-26 05:25:22.493191 osd.1 osd.1 10.65.200.47:6805/27875 1661 : cluster [WRN] slow request 5.432530 seconds old, received at 2019-01-26 05:25:17.060552: osd_op(client.24143.0:822 3.4 3:323cebeb::benchmark_data_buildnode_29015_object821:head [set-alloc-hint object_size 4194304 write_size 4194304,write 0~4194304] snapc 0=[] ondisk+write+known_if_redirected e18) currently waiting for subops from 0,2
2019-01-26 05:25:22.493194 osd.1 osd.1 10.65.200.47:6805/27875 1662 : cluster [WRN] slow request 6.433768 seconds old, received at 2019-01-26 05:25:16.059315: osd_repop(client.24143.0:807 3.5 e18/15 3:af33e4d1::benchmark_data_buildnode_29015_object806:head v 18'105) currently commit_sent
2019-01-26 05:25:22.493202 osd.1 osd.1 10.65.200.47:6805/27875 1663 : cluster [WRN] slow request 5.427220 seconds old, received at 2019-01-26 05:25:17.065863: osd_repop(client.24143.0:821 3.6 e18/15 3:7c8e3751::benchmark_data_buildnode_29015_object820:head v 18'101) currently started
```

```
2019-01-26 05:25:21.752638 mgr.x client.4101 10.65.200.47:0/1554384750 539 : cluster [DBG] pgmap v555: 24 pgs: 24 active+clean; 3.17GiB data, 506GiB used, 1.26TiB / 1.76TiB avail; 10.0MiB/s wr, 2op/s
2019-01-26 05:25:23.753368 mgr.x client.4101 10.65.200.47:0/1554384750 540 : cluster [DBG] pgmap v556: 24 pgs: 24 active+clean; 3.16GiB data, 507GiB used, 1.26TiB / 1.76TiB avail; 9.33MiB/s wr, 2op/s
```

```
2019-01-26 05:25:43.759827 mgr.x client.4101 10.65.200.47:0/1554384750 550 : cluster [DBG] pgmap v566: 24 pgs: 24 active+clean; 2.19KiB data, 477GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 05:25:45.760387 mgr.x client.4101 10.65.200.47:0/1554384750 551 : cluster [DBG] pgmap v567: 24 pgs: 24 active+clean; 2.19KiB data, 477GiB used, 1.29TiB / 1.76TiB avail
2019-01-26 05:25:47.761015 mgr.x client.4101 10.65.200.47:0/1554384750 552 : cluster [DBG] pgmap v568: 24 pgs: 24 active+clean; 2.19KiB data, 477GiB used, 1.29TiB / 1.76TiB avail
```

- Then tested setting it back to INFO and it went back to old state no IOPS line.

```
$ bin/ceph daemon mon.a config show | grep mon_cluster_log_file_level
  "mon_cluster_log_file_level": "info",
```

```
$ bin/ceph daemon osd.0 config show | grep osd_op_complaint
  "osd_op_complaint_time": "1.000000",
```

Cluster log goes back to info

```
2019-01-26 05:34:16.030528 mon.a mon.0 10.65.200.47:40557/0 85 : cluster [INF] daemon mds.a assigned to filesystem cephfs_a as rank 0 (now has 1 ranks)
2019-01-26 05:34:16.420411 mon.a mon.0 10.65.200.47:40557/0 88 : cluster [INF] daemon mds.a is now active in filesystem cephfs_a as rank
```

Then again ran rados bench and this time only slow requests because I have complaint set as 1 but no IOPS logs

```
 47      16      136      120  10.2116      0      -      5.66609
 48      16      136      120  9.99885      0      -      5.66609
 49      16      136      120  9.79479      0      -      5.66609
 50      16      136      120  9.59889      0      -      5.66609
 51      13      137      124  9.72437      2.66667      11.2601      5.80878
Total time run:          51.358409
Total writes made:      137
Write size:             4194304
Object size:           4194304
```

Bandwidth (MB/sec): 10.6701  
Stddev Bandwidth: 9.86784  
Max bandwidth (MB/sec): 48  
Min bandwidth (MB/sec): 0  
Average IOPS: 2  
Stddev IOPS: 2  
Max IOPS: 12  
Min IOPS: 0  
Average Latency(s): 5.98196  
Stddev Latency(s): 2.9864  
Max latency(s): 12.1471  
Min latency(s): 0.428529

Cleaning up (deleting benchmark objects)  
Removed 137 objects  
Clean up completed and total clean up time :0.502795

2019-01-26 05:37:05.485409 osd.0 osd.0 10.65.200.47:6801/30883 300 : cluster [WRN] slow request 6.822828 seconds old, received at 2019-01-26 05:36:58.662494: osd\_repop(client.4147.0:132 3.2 e17/14) currently commit\_queue\_d\_for\_journal\_write  
2019-01-26 05:37:05.485412 osd.0 osd.0 10.65.200.47:6801/30883 301 : cluster [WRN] slow request 6.818966 seconds old, received at 2019-01-26 05:36:58.666356: osd\_repop(client.4147.0:133 3.2 e17/14 3:4d7ca680::benchmark\_data\_dell-per320-4.gsslab.pnq.redh\_31808\_object132:head v 17'26) currently commit\_queued\_for\_journal\_write  
2019-01-26 05:37:05.485418 osd.0 osd.0 10.65.200.47:6801/30883 302 : cluster [WRN] slow request 6.553486 seconds old, received at 2019-01-26 05:36:58.931836: osd\_op(client.4147.0:136 3.6 3:70f47ab9::benchmark\_data\_dell-per320-4.gsslab.pnq.redh\_31808\_object135:head [set-alloc-hint object\_size 4194304 write\_size 4194304,write 0~4 194304] snapc 0=[] ondisk+write+known\_if\_redirected e17) currently sub\_op\_commit\_rec from 1  
2019-01-26 05:37:03.851504 osd.1 osd.1 10.65.200.47:6805/31117 262 : cluster [WRN] 8 slow requests, 1 included below; oldest blocked for > 10.196165 secs  
2019-01-26 05:37:03.851509 osd.1 osd.1 10.65.200.47:6805/31117 263 : cluster [WRN] slow request 4.921332 seconds old, received at 2019-01-26 05:36:58.930139: osd\_repop\_reply(client.4147.0:119 3.2 e17/14) currently queued\_for\_pg  
2019-01-26 05:37:04.814854 osd.1 osd.1 10.65.200.47:6805/31117 264 : cluster [WRN] 8 slow requests, 1 included below; oldest blocked for > 11.159502 secs  
2019-01-26 05:37:04.814859 osd.1 osd.1 10.65.200.47:6805/31117 265 : cluster [WRN] slow request 8.943792 seconds old, received at 2019-01-26 05:36:55.871016: osd\_op(client.4147.0:127 3.0 3:0d4c372c::benchmark\_data\_dell-per320-4.gsslab.pnq.redh\_31808\_object126:head [set-alloc-hint object\_size 4194304 write\_size 4194304,write 0~4 194304] snapc 0=[] ondisk+write+known\_if\_redirected e17) currently op\_applied  
2019-01-26 05:37:13.664460 mon.a mon.0 10.65.200.47:40557/0 100 : cluster [INF] Health check cleared: POOL\_APP\_NOT\_ENABLED (was: application not enabled on 1 pool(s))  
2019-01-26 05:37:13.664524 mon.a mon.0 10.65.200.47:40557/0 101 : cluster [INF] Cluster is now healthy

**#10 - 01/26/2019 12:18 AM - Vikhyat Umrao**

- Status changed from Fix Under Review to Pending Backport

Once it is merged in master I can backport it to mimic and luminous.

**#11 - 01/26/2019 04:32 PM - Neha Ojha**

- Status changed from Pending Backport to 7

**#12 - 01/27/2019 08:44 PM - Neha Ojha**

- Status changed from 7 to Pending Backport

**#13 - 01/30/2019 12:56 PM - Nathan Cutler**

- Copied to Backport #38107: mimic: Adding back the IOPS line for client and recovery IO in cluster logs added

**#14 - 01/30/2019 12:56 PM - Nathan Cutler**

- Copied to Backport #38108: luminous: Adding back the IOPS line for client and recovery IO in cluster logs added

**#15 - 02/04/2019 08:47 PM - Yuri Weinstein**

merged <https://github.com/ceph/ceph/pull/26208>

**#16 - 02/05/2019 04:56 PM - Nathan Cutler**

- Status changed from Pending Backport to Resolved