

## mgr - Bug #46743

### mgr/devicehealth: device\_health\_metrics pool gets created even without any OSDs in the cluster

07/29/2020 05:02 AM - Nathan Cutler

|   |                               |
|---|-------------------------------|
| <b>Status:</b> Resolved   | <b>% Done:</b> 0%             |
| <b>Priority:</b> Normal   |                               |
| <b>Assignee:</b> Sunny Kumar  |                               |
| <b>Category:</b>  |                               |
| <b>Target version:</b>  |                               |
| <b>Source:</b>  | <b>Reviewed:</b>              |
| <b>Tags:</b>  | <b>Affected Versions:</b>     |
| <b>Backport:</b> octopus  | <b>ceph-qa-suite:</b>         |
| <b>Regression:</b> No   | <b>Pull request ID:</b> 37085 |
| <b>Severity:</b> 3 - minor  | <b>Crash signature:</b>       |
| <b>Description</b>  |                               |
| <p>Currently, "cephadm bootstrap" appears to create a pool because "devicehealth", as an "always on" module, gets created when the first MGR is deployed.</p> <p>The pool actually gets created by mgr/devicehealth, not by cephadm - hence this bug is opened against mgr/devicehealth, even though - from the user's perspective - the problem happens when the "cephadm bootstrap" command is issued.</p> <p>Because mgr/devicehealth creates a pool before the cluster has any OSDs, the cluster enters HEALTH_WARN immediately after bootstrap:</p> <pre>master:~ # ceph -s cluster:   id:      fed46cbe-d157-11ea-901a-52540084b2ce   health: HEALTH_WARN           OSD count 0 &lt; osd_pool_default_size 3  services:   mon: 1 daemons, quorum nodel (age 67s)   mgr: nodel.ikkrrt(active, since 39s)   osd: 0 osds: 0 up, 0 in  data:   pools:  1 pools, 1 pgs   objects: 0 objects, 0 B   usage:  0 B used, 0 B / 0 B avail   pgs:    100.000% pgs unknown           1 unknown</pre> <p>The pool is:</p> <pre>master:~ # ceph osd pool ls device_health_metrics</pre> <p>It seems like the creation of this pool should be linked to deployment of the first OSD, not to the deployment of the first MON/MGR.</p> |                               |
| <b>Related issues:</b>  |                               |
| Copied to mgr - Backport #47739: octopus: mgr/devicehealth: device_health_met... <b>Resolved</b>  |                               |

## History

---

### #1 - 07/29/2020 05:33 AM - Nathan Cutler

Even when I immediately deploy OSDs after running "cephadm bootstrap", the cluster does not recover from the health warning very quickly. Here's what the health status looks like when deployment of the OSDs completes:

```
master:~ # ceph -s
cluster:
  id:          59f04650-d15c-11ea-bc40-52540017ea7a
  health: HEALTH_WARN
             Reduced data availability: 1 pg inactive, 1 pg peering

services:
  mon: 5 daemons, quorum master,node3,node1,node4,node2 (age 35s)
  mgr: master.oqkwii(active, since 88s), standbys: node2.rwzexl
  osd: 4 osds: 2 up (since 5s), 2 in (since 5s); 1 remapped pgs

data:
  pools: 1 pools, 1 pgs
  objects: 0 objects, 0 B
  usage: 2.0 GiB used, 14 GiB / 16 GiB avail
  pgs: 100.000% pgs not active
      1 creating+peering

progress:
  Rebalancing after osd.1 marked in (3s)
  [.....]
```

It just doesn't seem right for the cluster to bootstrap into a degraded state, and then have to rely on Ceph's "self-healing" capabilities to put it right.

### #2 - 07/29/2020 08:35 AM - Nathan Cutler

- Subject changed from Running "cephadm bootstrap" should not create any pools to Running "cephadm bootstrap" without "--apply-spec" should not create any pools

### #3 - 07/29/2020 10:13 AM - Sebastian Wagner

- Project changed from Orchestrator to mgr

- Subject changed from Running "cephadm bootstrap" without "--apply-spec" should not create any pools to mgr/devicehealth: device\_health\_metrics gets created even without any OSDs in the cluster

**#4 - 07/29/2020 11:43 AM - Nathan Cutler**

- Subject changed from *mgr/devicehealth: device\_health\_metrics gets created even without any OSDs in the cluster* to *mgr/devicehealth: device\_health\_metrics pool gets created even without any OSDs in the cluster*

**#5 - 07/29/2020 12:00 PM - Nathan Cutler**

- Description updated

**#6 - 07/29/2020 12:00 PM - Nathan Cutler**

- Description updated

**#7 - 08/07/2020 08:50 PM - Neha Ojha**

- Tags set to *low-hanging-fruit*

**#8 - 08/12/2020 09:08 PM - Neha Ojha**

- Assignee set to *Sunny Kumar*

**#9 - 08/13/2020 02:34 PM - Sunny Kumar**

- Status changed from *New* to *In Progress*

**#10 - 09/10/2020 03:16 PM - Sunny Kumar**

- Status changed from *In Progress* to *Fix Under Review*

- Pull request ID set to *37085*

**#11 - 09/26/2020 03:01 PM - Nathan Cutler**

- Backport set to *octopus*

**#12 - 09/26/2020 03:15 PM - Nathan Cutler**

- Tags deleted (*low-hanging-fruit*)

**#13 - 10/04/2020 05:51 AM - Kefu Chai**

- Status changed from *Fix Under Review* to *Pending Backport*

**#14 - 10/04/2020 07:43 AM - Nathan Cutler**

- Copied to Backport #47739: *octopus: mgr/devicehealth: device\_health\_metrics pool gets created even without any OSDs in the cluster added*

**#15 - 10/10/2020 08:59 AM - Nathan Cutler**

- Status changed from *Pending Backport* to *Resolved*

While running with `--resolve-parent`, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".