

Orchestrator - Bug #45594

cephadm: weight of a replaced OSD is 0

05/19/2020 07:42 AM - Kiefer Chang

Status:	Resolved	% Done:	0%
Priority:	Normal		
Assignee:	Joshua Schmid		
Category:	cephadm/osd		
Target version:	v15.2.5		
Source:		Reviewed:	
Tags:		Affected Versions:	v15.2.2, v16.0.0
Backport:	octopus	ceph-qa-suite:	
Regression:	No	Pull request ID:	36151
Severity:	3 - minor	Crash signature:	
Description			
Not sure if this is intended.			
After deleting an OSD with `--replace` flag and create new OSD on it, OSD's WEIGHT is 0.			
<pre>—# bin/ceph osd tree *** DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH *** 2020-05-19T07:38:24.802+0000 7f88013fb700 -1 WARNING: all dangerous and experimental features are enabled. 2020-05-19T07:38:24.822+0000 7f88013fb700 -1 WARNING: all dangerous and experimental features are enabled. ID CLASS WEIGHT TYPE NAME STATUS REWEIGHT PRI-AFF -1 0.32983 root default -3 0.29576 host ceph-master 0 ssd 0.09859 osd.0 up 1.00000 1.00000 1 ssd 0.09859 osd.1 up 1.00000 1.00000 2 ssd 0.09859 osd.2 up 1.00000 1.00000 -7 0.03407 host osd0 3 hdd 0 osd.3 up 1.00000 1.00000 <----- 4 hdd 0.01949 osd.4 up 1.00000 1.00000 5 hdd 0.01459 osd.5 up 1.00000 1.00000</pre>			

History

#1 - 05/19/2020 01:58 PM - Sebastian Wagner

- Assignee set to Joshua Schmid

#2 - 05/20/2020 09:55 AM - Joshua Schmid

- Status changed from New to In Progress

This is certainly not intended.. I'll investigate

#3 - 06/16/2020 10:49 AM - Joshua Schmid

The initial weight is never restored after `draining` the OSDs.

We can save the initial weight/reweight and reset it after the osd is empty and has a `destroyed` flag assigned. This will prevent the OSD from receiving new PGs

#4 - 07/22/2020 10:59 AM - Sebastian Wagner

- Category changed from *cephadm* to *cephadm/osd*

#5 - 07/31/2020 08:33 AM - Joshua Schmid

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

- Pull request ID set to 36151

#6 - 08/17/2020 12:52 PM - Sebastian Wagner

- Status changed from *Fix Under Review* to *Resolved*

- Target version set to *v15.2.5*