

mgr - Bug #40871

osd status reports old crush location after osd moves

07/22/2019 07:48 PM - Noah Watkins

Status:	Resolved	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:	Kefu Chai		
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	luminous, mimic, nautilus	Pull request ID:	30448
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

**Description**

Scenario:

Move an OSD disk from host=worker1 to a new node (host=worker0) and on that new node we update the crush location (e.g. host=worker0). After the OSD starts up `ceph osd status` reports the old location, but `ceph osd tree` reports the new location.

```
[nwatkins@smash rook]$ kubectl -n rook-ceph exec -it rook-ceph-tools-7cf4cc7568-kz4q6 ceph osd status
```

id	host	used	avail	wr ops	wr data	rd ops	rd data	state
0	worker1	1027M	8188M	0	0	0	0	exists,up
1	worker0	1027M	8188M	0	0	0	0	exists,up

```
[nwatkins@smash rook]$ kubectl -n rook-ceph exec -it rook-ceph-tools-7cf4cc7568-kz4q6 ceph osd tree
```

ID	CLASS	WEIGHT	TYPE	NAME	STATUS	REWEIGHT	PRI-AFF
-1		0.01758	root	default			
-5		0.01758	host	worker0			
0	hdd	0.00879		osd.0	up	1.00000	1.00000
1	hdd	0.00879		osd.1	up	1.00000	1.00000

And then after restarting the manager, `ceph osd status` also starts reporting the correct location

```
[nwatkins@smash rook]$ kubectl -n rook-ceph exec -it rook-ceph-tools-7cf4cc7568-kz4q6 ceph osd status
```

id	host	used	avail	wr ops	wr data	rd ops	rd data	state
0	worker0	1027M	8188M	0	0	0	0	exists,up
1	worker0	1027M	8188M	0	0	0	0	exists,up

related: <http://tracker.ceph.com/issues/40011>

<b>Related issues:</b>	
Related to mgr - Bug #40011: ceph -s shows wrong number of pools when pool wa...	Resolved

Copied to mgr - Backport #41982: luminous: osd status reports old crush locat...	<b>Rejected</b>
Copied to mgr - Backport #41983: nautilus: osd status reports old crush locat...	<b>Resolved</b>
Copied to mgr - Backport #41984: mimic: osd status reports old crush location...	<b>Rejected</b>
Copied to mgr - Bug #44446: osd status reports old crush location after osd m...	<b>Resolved</b>

## History

### #1 - 07/22/2019 07:50 PM - Noah Watkins

[https://bugzilla.redhat.com/show\\_bug.cgi?id=1724428](https://bugzilla.redhat.com/show_bug.cgi?id=1724428)

### #2 - 08/22/2019 10:21 AM - Sebastian Wagner

- Related to Bug #40011: ceph -s shows wrong number of pools when pool was deleted added

### #3 - 09/12/2019 02:31 PM - Kefu Chai

- Assignee set to Kefu Chai

### #4 - 09/18/2019 04:06 AM - Kefu Chai

- Status changed from New to Fix Under Review

- Pull request ID set to 30448

### #5 - 09/18/2019 04:07 AM - Kefu Chai

- Backport set to mimic, nautilus

### #6 - 09/22/2019 03:46 AM - Kefu Chai

- Status changed from Fix Under Review to Pending Backport

- Backport changed from mimic, nautilus to luminous, mimic, nautilus

### #7 - 09/23/2019 09:06 AM - Nathan Cutler

- Copied to Backport #41982: luminous: osd status reports old crush location after osd moves added

### #8 - 09/23/2019 09:06 AM - Nathan Cutler

- Copied to Backport #41983: nautilus: osd status reports old crush location after osd moves added

### #9 - 09/23/2019 09:06 AM - Nathan Cutler

- Copied to Backport #41984: mimic: osd status reports old crush location after osd moves added

### #10 - 03/05/2020 03:33 PM - Kefu Chai

- Copied to Bug #44446: osd status reports old crush location after osd moves added

### #11 - 01/27/2021 08:06 PM - 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".