

## Orchestrator - Bug #45393

### Containerized osd config must be updated when adding/removing mons

05/06/2020 06:41 AM - Tim Serong

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> High	
<b>Assignee:</b> Tim Serong	
<b>Category:</b> cephadm	
<b>Target version:</b> v15.2.4	
<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> Yes	<b>Pull request ID:</b> 34922
<b>Severity:</b> 3 - minor	<b>Crash signature:</b>
<b>Description</b>	
Try this:  - bootstrap a cluster (1 mon, 1 mgr) - add a bunch of osds (ceph orch apply osd --all-available-devices) - add some more mons and mgrs (ceph orch apply mon 3 ; ceph orch apply mgr 3)  At this point, ceph.conf as seen by the osds (/var/lib/ceph/\$FSID/osd.\$ID/config) still only lists the first mon. If you restart any osds, and that first mon is down for some reason, the osd's can't join the cluster, because they don't know the other mons exist. They'll just sit there logging "monclient(hunting): authenticate timed out after 300" every five minutes until that first mon comes back.  I can think of two potential ways to address this:  1) Have cephadm mon apply go out and update every single osd config file with the new list of mons. This would of course not work completely if some osd hosts were down at the time. Also it might take a while... 2) Have the osds update their own config file automatically based on current monmaps.  This probably also needs to go into troubleshooting docs (check mon_host in each containerized osd's individual config file)	
<b>Related issues:</b>	
Related to Orchestrator - Feature #45378: cephadm: manage /etc/ceph/ceph.conf	<b>Resolved</b>

#### History

##### #1 - 05/06/2020 06:41 AM - Tim Serong

- Related to Feature #45378: cephadm: manage /etc/ceph/ceph.conf added

##### #2 - 05/06/2020 07:42 AM - Sebastian Wagner

- Priority changed from Normal to High

- Regression changed from No to Yes

This was fixed in <https://github.com/ceph/ceph/pull/33855> . Looks like we have to figure out, what went wrong here.

##### #3 - 05/06/2020 08:36 AM - Tim Serong

- Assignee set to Tim Serong

Thanks for the pointer, I'll try to figure out what's going on, seeing as I'm the one who hit this :-)

##### #4 - 05/06/2020 08:58 AM - Tim Serong

A quick grep of my logs shows it reconfiguring the mons and mgrs, but not the osds.

**#5 - 05/06/2020 11:23 AM - Tim Serong**

- *Status changed from New to Fix Under Review*
- *Pull request ID set to 34922*

It's always the little things...

**#6 - 05/11/2020 10:43 AM - Tim Serong**

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

**#7 - 06/02/2020 01:08 PM - Sebastian Wagner**

- *Status changed from Pending Backport to Resolved*
- *Target version set to v15.2.4*