

Orchestrator - Bug #46606

cephadm: post-bootstrap monitoring deployment only works if the command "ceph mgr module enable prometheus" has already been issued

07/17/2020 07:26 PM - Nathan Cutler

Status:	New	% Done:	0%
Priority:	Normal		
Assignee:			
Category:	cephadm/monitoring		
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	39520
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			
Description			
Post-bootstrap monitoring deployment only works if the command "ceph mgr module enable prometheus" has already been issued			
Reported by Dmitri Savineau here: https://tracker.ceph.com/issues/46561#note-5			
deploying the monitoring after the bootstrap requires to run an extra ceph command to enable the prometheus mgr module (which is automatically done during the bootstrap) [1]			
[1] https://github.com/ceph/ceph/blob/master/src/cephadm/cephadm#L2877-L2879			

History

#1 - 07/17/2020 07:26 PM - Nathan Cutler

- Related to Bug #46561: cephadm: monitoring services adoption doesn't honor the container image added

#2 - 07/17/2020 07:26 PM - Nathan Cutler

- Subject changed from Post-bootstrap monitoring deployment only works if the command "ceph mgr module enable prometheus" has already been issued to cephadm: post-bootstrap monitoring deployment only works if the command "ceph mgr module enable prometheus" has already been issued

#3 - 07/17/2020 07:27 PM - Nathan Cutler

- Description updated

#4 - 07/17/2020 07:28 PM - Nathan Cutler

- Description updated

#5 - 07/20/2020 12:14 PM - Sebastian Wagner

- Category set to cephadm/monitoring

#6 - 02/11/2021 12:05 PM - Sebastian Wagner

- Priority changed from Normal to High

#7 - 02/17/2021 10:16 AM - Juan Miguel Olmo Martínez

- Assignee set to Sebastian Wagner

#8 - 03/10/2021 02:47 PM - Sebastian Wagner

- Status changed from New to Fix Under Review
- Pull request ID set to 39520

#9 - 04/14/2021 01:49 PM - Sebastian Wagner

- Status changed from Fix Under Review to New

#10 - 04/23/2021 01:57 PM - Sebastian Wagner

- Related to deleted (Bug #46561: cephadm: monitoring services adoption doesn't honor the container image)

#11 - 04/30/2021 04:46 PM - Sage Weil

A couple options:

- make the 'orch apply prometheus' fail if the mgr prometheus module isn't enabled. (maybe include a --force in case the user **really** wants to proceed?)
- make cephadm raise a health warning if there is a prometheus deployed but the prometheus module isn't enabled
- make 'orch apply prometheus' silently enable the prometheus module

#12 - 04/30/2021 07:42 PM - Sebastian Wagner

I'd **definitively** go for make 'orch apply prometheus' silently enable the prometheus module.

#13 - 05/02/2021 08:55 AM - Nathan Cutler

- make the 'orch apply prometheus' fail if the mgr prometheus module isn't enabled. (maybe include a --force in case the user **really** wants to proceed?)

This one is slightly problematic because there is not just "orch apply prometheus" with a prometheus-specific yaml blob, but also "orch apply" with a BIG yaml blob (with sections for various kinds of services/daemons).

Arguably, the "orch apply" command (with BIG yaml blob) should fail if any part of the yaml is not fulfillable. But that's not how the orchestrator works: the "orch apply" is fulfilled as a background task and when something goes wrong it's not always obvious to the user how to figure out what happened and why, since it typically involves conducting a post-mortem examination of the mgr logs.

To say it another way: "orch apply" is like a "moon shot". Everything has to be prepared in advance. Once the rocket is on its way up, there isn't any good way of aborting the mission.

(Caveat: this is just my impression as a casual user of "orch apply", not based on any deep knowledge of the code or even the design)

#14 - 05/10/2021 08:57 AM - Sebastian Wagner

- *Priority changed from High to Normal*

prio=normal, as this is not trivial to implement

#15 - 07/07/2021 09:40 AM - Sebastian Wagner

- *Assignee deleted (Sebastian Wagner)*