

ceph-volume - Bug #44356

ceph-volume inventory: KeyError: 'ceph.cluster_name'

02/28/2020 06:59 PM - Sebastian Wagner

Status: New	% Done: 0%
Priority: High	
Assignee: Jan Fajerski	
Category:	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport:	Pull request ID:
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	

Description

```
Module 'cephadm' has failed: cephadm exited with an error code: 1, stderr:INFO:cephadm:/usr/bin/podman:stderr WARNING: The same type, major and minor should not be used for multiple devices.
INFO:cephadm:/usr/bin/podman:stderr WARNING: The same type, major and minor should not be used for multiple devices.
INFO:cephadm:/usr/bin/podman:stderr --> KeyError: 'ceph.cluster_name'
Traceback (most recent call last):
  File "<stdin>", line 3394, in <module>
  File "<stdin>", line 688, in _infer_fsid
  File "<stdin>", line 2202, in command_ceph_volume
  File "<stdin>", line 513, in call_throws
RuntimeError: Failed command: /usr/bin/podman run --rm --net=host --privileged --group-add=disk -e CONTAINER_IMAGE=registry.suse.de/suse/sle-15-sp2/update/products/ses7/update/cr/containers/ses/7/ceph/ceph:latest -e NODE_NAME=hses-nodel -v /var/run/ceph/002c389e-54fd-11ea-a99f-52540044d765:/var/run/ceph:z -v /var/log/ceph/002c389e-54fd-11ea-a99f-52540044d765:/var/log/ceph:z -v /var/lib/ceph/002c389e-54fd-11ea-a99f-52540044d765/crash:/var/lib/ceph/crash:z -v /dev:/dev -v /run/udev:/run/udev -v /sys:/sys -v /run/lvm:/run/lvm -v /run/lock/lvm:/run/lock/lvm --entrypoint /usr/sbin/ceph-volume registry.suse.de/suse/sle-15-sp2/update/products/ses7/update/cr/containers/ses/7/ceph/ceph:latest inventory --format=json

hses-nodel:~ # ceph -s
cluster:
  id: 002c389e-54fd-11ea-a99f-52540044d765
  health: HEALTH_ERR
    1 filesystem is offline
    1 filesystem is online with fewer MDS than max_mds
Module 'cephadm' has failed: cephadm exited with an error code: 1, stderr:INFO:cephadm:/usr/bin/podman:stderr WARNING: The same type, major and minor should not be used for multiple devices.
INFO:cephadm:/usr/bin/podman:stderr WARNING: The same type, major and minor should not be used for multiple devices.
INFO:cephadm:/usr/bin/podman:stderr --> KeyError: 'ceph.cluster_name'
Traceback (most recent call last):
  File "<stdin>", line 3394, in <module>
  File "<stdin>", line 688, in _infer_fsid
  File "<stdin>", line 2202, in command_ceph_volume
  File "<stdin>", line 513, in call_throws
RuntimeError: Failed command: /usr/bin/podman run --rm --net=host --privileged --group-add=disk -e CONTAINER_IMAGE=registry.suse.de/suse/sle-15-sp2/update/products/ses7/update/cr/containers/ses/7/
```

```
ceph/ceph:latest -e NODE_NAME=hses-node1 -v /var/run/ceph/002c389e-54fd-11ea-a99f-52540044d765:/va
r/run/ceph:z -v /var/log/ceph/002c389e-54fd-11ea-a99f-52540044d765:/var/log/ceph:z -v /var/lib/cep
h/002c389e-54fd-11ea-a99f-52540044d765/crash:/var/lib/ceph/crash:z -v /dev:/dev -v /run/udev:/run/
udev -v /sys:/sys -v /run/lvm:/run/lvm -v /run/lock/lvm:/run/lock/lvm --entrypoint /usr/sbin/ceph-
volume registry.suse.de/suse/sle-15-sp2/update/products/ses7/update/cr/containers/ses/7/ceph/ceph:
latest inventory --format=json
```

```
services:
  mon: 4 daemons, quorum hses-node1,hses-node2,hses-node3,hses-node4 (age 11m)
  mgr: hses-node1.jxzdin(active, since 13m), standbys: hses-node1.aogwfz, hses-node2.dlyvwy, hse
s-node3.delhzp, hses-node4.vgmgec
  mds: myfs:0
  osd: 8 osds: 8 up (since 6m), 8 in (since 6m)
  rgw: 4 daemons active (default.default.hses-node1.sirofe, default.default.hses-node2.kvzapg, d
efault.default.hses-node3.dhobfn, default.default.hses-node4.cuulnm)
```

task status:

```
data:
  pools: 6 pools, 168 pgs
  objects: 189 objects, 5.8 KiB
  usage: 8.2 GiB used, 152 GiB / 160 GiB avail
  pgs: 168 active+clean
```

Related issues:

Duplicated by Orchestrator - Bug #45604: mgr/cephadm: Failed to create an OSD

[Duplicate](#)

History

#1 - 03/06/2020 10:57 PM - Sage Weil

- Project changed from Orchestrator to ceph-volume

- Subject changed from cephadm: ceph-volume inventory: KeyError: 'ceph.cluster_name' to ceph-volume inventory: KeyError: 'ceph.cluster_name'

#2 - 03/17/2020 04:14 PM - Jan Fajerski

Can we get a stack trace from ceph-volume for this? Setting env CEPH_VOLUME_DEBUG=true would do the trick.

#3 - 03/17/2020 04:14 PM - Jan Fajerski

- Assignee set to Jan Fajerski

#4 - 03/17/2020 04:51 PM - Jan Fajerski

Just from looking at the code path, I'd say there was an lv, that had a ceph.osd_id tag set but no ceph.cluster tag.

Is this reproducible and if so how? I can tighten some tests to avoid the situation described above, but that just a good guess.

#5 - 05/19/2020 10:47 AM - Sebastian Wagner

- Duplicated by Bug #45604: mgr/cephadm: Failed to create an OSD added

#6 - 09/14/2020 01:25 PM - Jan Fajerski

lvs output from an incident that looks like this.

```
sesnode2:~ # lvs -o +lv_tags
  LV                               VG                               Attr      LSize Poo
l Origin Data% Meta% Move Log Cpy%Sync Convert LV Tags
  osd-data-0e6e3192-9933-480d-9784-f5103d9e1a1e ceph-7e3c615c-447f-4008-966a-1e7d4c7628b1 -wi-a----- 3.49t
                                     ceph.cluster_fsid=null,ceph.osd_fsid=null,ceph.osd_id=null,ce
ph.type=null
  osd-data-8576aff5-39e2-4120-a886-ddc92382fead ceph-7e3c615c-447f-4008-966a-1e7d4c7628b1 -wi-a----- 3.49t
                                     ceph.cluster_fsid=null,ceph.osd_fsid=null,ceph.osd_id=null,ce
ph.type=null
  osd-data-90de784d-0f6f-4228-94b5-74e92d5a81c0 ceph-7e3c615c-447f-4008-966a-1e7d4c7628b1 -wi-a----- 3.49t
```

```

ceph.cluster_fsid=null,ceph.osd_fsid=null,ceph.osd_id=null,ce
ph.type=null
  osd-data-bdcd5eb1-954c-4647-899a-8cf22c0d8172 ceph-7e3c615c-447f-4008-966a-1e7d4c7628b1 -wi-a----- 3.49t
      ceph.cluster_fsid=null,ceph.osd_fsid=null,ceph.osd_id=null,ce
ph.type=null
sesnode2:~ # lsblk
NAME                                                                                               MAJ:MIN R
M  SIZE RO TYPE MOUNTPOINT
sda                                                                                               8:0
0 223.6G 0 disk
ââsda1                                                                                             8:1
0 512M 0 part /boot/efi
ââsda2                                                                                             8:2
0 4G 0 part [SWAP]
ââsda3                                                                                             8:3
0 219.1G 0 part /
nvme5n1                                                                                           259:0
0 698.7G 0 disk
nvme7n1                                                                                           259:1
0 2.9T 0 disk
nvme2n1                                                                                           259:2
0 14T 0 disk
nvme4n1                                                                                           259:3
0 698.7G 0 disk
nvme3n1                                                                                           259:4
0 2.9T 0 disk
nvme6n1                                                                                           259:5
0 14T 0 disk
nvme0n1                                                                                           259:6
0 14T 0 disk
ââceph--7e3c615c--447f--4008--966a--1e7d4c7628b1-osd--data--bdcd5eb1--954c--4647--899a--8cf22c0d8172 254:0
0 3.5T 0 lvm
ââceph--7e3c615c--447f--4008--966a--1e7d4c7628b1-osd--data--0e6e3192--9933--480d--9784--f5103d9e1a1e 254:1
0 3.5T 0 lvm
ââceph--7e3c615c--447f--4008--966a--1e7d4c7628b1-osd--data--8576aff5--39e2--4120--a886--ddc92382fead 254:2
0 3.5T 0 lvm
ââceph--7e3c615c--447f--4008--966a--1e7d4c7628b1-osd--data--90de784d--0f6f--4228--94b5--74e92d5a81c0 254:3
0 3.5T 0 lvm
nvme1n1                                                                                           259:7
0 14T 0 disk

```