

Orchestrator - Bug #54142

quincy cephadm-purge-cluster needs work

02/04/2022 06:35 PM - Tim Wilkinson

Status:	Resolved	% Done:	0%
Priority:	Normal		
Assignee:	Redouane Kachach Elhichou		
Category:	cephadm		
Target version:			
Source:		Affected Versions:	v17.0.0
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

For the sake of tracking ...

The purge process in quincy is not yet ready for prime time in this early stage. The preflight & purge playbooks were used but ultimately I went through the manual steps I've used previously when the process fails somewhere ...

```
# Clean all hosts excluding bootstrap
cephadm_in_host=$(ls /var/lib/ceph/$fsid/cephadm*)
python3 $cephadm_in_host rm-cluster --fsid $fsid --force
systemctl stop ceph.target
systemctl disable ceph.target
rm -f /etc/systemd/system/ceph.target
systemctl daemon-reload
systemctl reset-failed
rm -rf /var/log/ceph/*
rm -rf /var/lib/ceph/*

# clean bootstrap
cephadm_in_host=$(ls /var/lib/ceph/$fsid/cephadm*)
python3 $cephadm_in_host rm-cluster --fsid $fsid --force
#cephadm rm-cluster --fsid $fsid --force
systemctl stop ceph.target
systemctl disable ceph.target
rm -f /etc/systemd/system/ceph.target
systemctl daemon-reload
systemctl reset-failed
rm -rf /etc/ceph/*
rm -rf /var/log/ceph/*
rm -rf /var/lib/ceph/*

# on OSD nodes
declare -a devList=(/dev/nvme0n1 /dev/nvme1n1 /dev/sdc /dev/sdd /dev/sde /dev/sdf /dev/sdg /dev/sdh /dev/sdi /dev/sdj /dev/sdk /dev/sdl /dev/sdm /dev/sdn /dev/sdo /dev/sdp /dev/sdq /dev/sdr /dev/sds /dev/sdt /dev/sdu /dev/sdv /dev/sdw /dev/sdx /dev/sdy /dev/sdz /dev/sdaa /dev/sdab /dev/sdac /dev/sdad /dev/sdae /dev/sdaf /dev/sdag /dev/sdah /dev/sdai /dev/sdaj /dev/sdak /dev/sdal)
for device in ${devList[@]}; do
    echo $device
    sgdisk --zap-all $device
done
for fsid in `systemctl list-units ceph*.target |grep target|grep -v services|awk '{print$NF}'`; do
    echo $fsid
```

```
/perfl/tim/tools/svc-clean.sh $fsid
done
for fsid in `ls /etc/systemd/system/ceph-*.target |cut -c 26- |cut -d. -f1` ; do
  echo $fsid
  /perfl/tim/tools/svc-clean.sh $fsid
done
for i in `lsblk -ro NAME |grep ceph` ; do
  echo $i
  dmsetup remove -f $i
done
```

... but that was insufficient. Subsequent Pacific deployments would fail due to remnant pods still running and holding onto ports, etc. Those had to be searched out and stopped. A couple of purge output examples are included FWIW.

Related issues:

Related to Orchestrator - Bug #54018: Suspicious behavior when deleting a clu...	Resolved
Related to Orchestrator - Feature #53815: cephadm rm-cluster should delete lo...	Resolved
Related to Orchestrator - Bug #53010: cephadm rm-cluster does not clean up /v...	Resolved

History

#1 - 02/04/2022 07:21 PM - Vikhyat Umrao

- Category changed from orchestrator to cephadm

#2 - 04/13/2022 11:45 AM - Redouane Kachach Elhichou

- Related to Bug #54018: Suspicious behavior when deleting a cluster (by running cephadm rm-cluster) added

#3 - 04/13/2022 11:46 AM - Redouane Kachach Elhichou

- Related to Feature #53815: cephadm rm-cluster should delete log files added

#4 - 04/13/2022 11:46 AM - Redouane Kachach Elhichou

- Related to Bug #53010: cephadm rm-cluster does not clean up /var/run/ceph added

#5 - 04/13/2022 11:49 AM - Redouane Kachach Elhichou

Logs and other issues were fixed as part of the related BUGs. But I'm not sure about the OSDs part.

#6 - 04/21/2022 09:18 AM - Redouane Kachach Elhichou

- Assignee set to Redouane Kachach Elhichou

#7 - 04/21/2022 09:18 AM - Redouane Kachach Elhichou

- Status changed from New to In Progress

#8 - 05/12/2022 07:41 PM - Tim Wilkinson

I was able to return to quincy deployments and purges using cephadm-17.2.0-0.el8.noarch and have had no problems running the preflight/purge/preflight/bootstrap procedure. There was no need to manually prepare any previously used devices or search & destroy remnant pods.

My only comment would be /var/run/ceph is not wiped and as such older cluster fsid's remain ...

```
root@f22-h01-000-6048r:~
# ll /var/{run,lib,log}/ceph
/var/lib/ceph:
total 8.0K
drwxr-x---  3 ceph ceph  50 Apr 18 22:36 .
drwxr-xr-x. 37 root root 4.0K May 12 13:09 ..
drwx----- 30 ceph ceph 4.0K May 12 13:08 c56d7946-d1f2-11ec-8d0b-000af7995d6c

/var/log/ceph:
total 4.0M
```

```
drwxrws--T 3 ceph ceph 69 Apr 18 22:36 .
drwxr-xr-x. 11 root root 4.0K May 12 12:54 ..
drwxrwx--- 2 ceph ceph 4.0K May 12 13:08 c56d7946-d1f2-11ec-8d0b-000af7995d6c
-rw-r--r-- 1 root ceph 4.0M May 12 19:36 cephadm.log
```

/var/run/ceph:

```
total 0
drwxrwx--- 4 root root 80 May 12 13:00 .
drwxr-xr-x 38 root root 2.2K May 12 13:08 ..
drwxrwx--- 2 ceph ceph 540 May 11 23:00 aa8ec022-calc-11ec-a5a0-000af7995d6c # old deployment
drwxrwx--- 2 ceph ceph 540 May 12 13:08 c56d7946-d1f2-11ec-8d0b-000af7995d6c
```

root@f22-h01-000-6048r:~

ll /var/run/ceph/aa8ec022-calc-11ec-a5a0-000af7995d6c

```
total 0
drwxrwx--- 2 ceph ceph 540 May 11 23:00 .
drwxrwx--- 4 root root 80 May 12 13:00 ..
srwxr-xr-x 1 ceph ceph 0 May 11 23:00 ceph-client.rgw.rgws.f22-h01-000-6048r.nbegha.7.94591718512160.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.107.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.114.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.122.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.12.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.130.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.137.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:45 ceph-osd.145.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.153.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.161.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.169.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.177.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.185.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.20.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:46 ceph-osd.28.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.35.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.42.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.49.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.57.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.5.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.66.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:47 ceph-osd.74.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:48 ceph-osd.82.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:48 ceph-osd.90.asok
srwxr-xr-x 1 ceph ceph 0 May 2 13:48 ceph-osd.97.asok
```

#9 - 09/28/2022 10:28 AM - Redouane Kachach Elhichou

- Status changed from In Progress to Resolved

I'm not able to reproduce these issues with the code on the main branch anymore. Please, feel free to re-open if you think the related BUG is still valid.

Files

220127-1930_cephadm-purge_f28-h28-000-r630	35.2 KB	02/04/2022	Tim Wilkinson
220127-1942_cephadm-purge_f28-h28-000-r630	32.8 KB	02/04/2022	Tim Wilkinson