

ceph-volume - Bug #41374

journal size can't be overridden with --journal-size when using --journal-devices in lvm batch mode

08/21/2019 01:07 PM - Guillaume Abrioux

Status: Resolved	% Done: 0%
Priority: Normal	
Assignee:	
Category:	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport: octopus,nautilus	Pull request ID: 36847
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	

Description

trying to deploy filestore OSDs with 'ceph-volume lvm batch' ends up like following when trying to override what there is in ceph.conf for the journal size (not sure whether this is valid for other params):

seen in : ceph version 14.2.2 (4f8fa0a0024755aae7d95567c63f11d6862d55be) nautilus (stable)

```
-bash-4.2# cat /etc/ceph/ceph.conf
# Please do not change this file directly since it is managed by Ansible and will be overwritten
[global]
cluster network = 192.168.40.0/24
fsid = b295b2c9-241d-4399-a3b2-3dc00b5235f3
mon host = [v2:192.168.39.10:3300,v1:192.168.39.10:6789],[v2:192.168.39.11:3300,v1:192.168.39.11:6789],[v2:192.168.39.12:3300,v1:192.168.39.12:6789]
osd_pool_default_size = 1
public network = 192.168.39.0/24
[osd]
osd journal size = 100
osd mkfs options xfs = -f -i size=2048
osd mkfs type = xfs
osd mount options xfs = noatime,largeio,inode64,swaloc
```

```
-bash-4.2# docker run --rm -ti --ulimit nofile=1024:4096 --privileged=true --net=host --pid=host --ipc=host -e CEPH_VOLUME_DEBUG=1 -v /dev:/dev -v /var/lib/ceph:/var/lib/ceph -v /etc/ceph:/etc/ceph -v /var/run:/var/run --entrypoint=ceph-volume docker.io/guits/ceph:ceph_volume_legacy_devices-nautilus-centos-7-x86_64 lvm batch --journal-size 4096 --no-systemd --filestore /dev/sda /dev/sdb --journal-devices /dev/sdc
```

Total OSDs: 2

Solid State VG:

Targets:	journal	Total size:	49.00 GB
Total LVs:	2	Size per LV:	4.00 GB
Devices:	/dev/sdc		

Type	Path	LV Size	% of device
------	------	---------	-------------

```

e
-----
---
[data]          /dev/sda                49.00 GB      100%
[journal]       vg: lv/vg                4.00 GB       8%
-----
---
[data]          /dev/sdb                49.00 GB      100%
[journal]       vg: lv/vg                4.00 GB       8%
--> The above OSDs would be created if the operation continues
--> do you want to proceed? (yes/no) yes
Running command: /usr/sbin/vgcreate -s 1G --force --yes ceph-journals-a128f13e-a2fa-44a2-8b39-b76
6bd7calf5 /dev/sdc
stdout: Physical volume "/dev/sdc" successfully created.
stdout: Volume group "ceph-journals-a128f13e-a2fa-44a2-8b39-b766bd7calf5" successfully created
--> Refusing to continue with configured size for journal
Traceback (most recent call last):
  File "/usr/sbin/ceph-volume", line 9, in <module>;
    load_entry_point('ceph-volume==1.0.0', 'console_scripts', 'ceph-volume')()
  File "/usr/lib/python2.7/site-packages/ceph_volume/main.py", line 38, in init
    self.main(self.argv)
  File "/usr/lib/python2.7/site-packages/ceph_volume/decorators.py", line 59, in newfunc
    return f(*a, **kw)
  File "/usr/lib/python2.7/site-packages/ceph_volume/main.py", line 148, in main
    terminal.dispatch(self.mapper, subcommand_args)
  File "/usr/lib/python2.7/site-packages/ceph_volume/terminal.py", line 206, in dispatch
    instance.main()
  File "/usr/lib/python2.7/site-packages/ceph_volume/devices/lvm/main.py", line 40, in main
    terminal.dispatch(self.mapper, self.argv)
  File "/usr/lib/python2.7/site-packages/ceph_volume/terminal.py", line 206, in dispatch
    instance.main()
  File "/usr/lib/python2.7/site-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
  File "/usr/lib/python2.7/site-packages/ceph_volume/devices/lvm/batch.py", line 325, in main
    self.execute()
  File "/usr/lib/python2.7/site-packages/ceph_volume/devices/lvm/batch.py", line 288, in execute
    self.strategy.execute()
  File "/usr/lib/python2.7/site-packages/ceph_volume/devices/lvm/strategies/filestore.py", line 3
55, in execute
    journal_size = prepare.get_journal_size(lv_format=True)
  File "/usr/lib/python2.7/site-packages/ceph_volume/util/prepare.py", line 67, in get_journal_si
ze
    raise RuntimeError('journal sizes must be larger than 2GB, detected: %s' % journal_size)
RuntimeError: journal sizes must be larger than 2GB, detected: 100.00 MB
-bash-4.2#

```

looks like the "--journal-size 4096" passed to the CLI doesn't override as expected the value present in ceph.conf (osd journal size = 100)

by the way, it should output a nice error message instead of throwing a python trace.

Related issues:

Copied to ceph-volume - Backport #47283: nautilus: journal size can't be over...

Resolved

Copied to ceph-volume - Backport #47284: octopus: journal size can't be overr...

Resolved

History

#1 - 10/31/2019 06:45 PM - Alfredo Deza

- Status changed from New to Can't reproduce

I tried replicating this issue without success:

```
[vagrant@node3 ~]$ sudo ceph-volume lvm batch --journal-size 2096 --filestore /dev/sdae
```

Total OSDs: 1

Type	Path	LV Size	% of device
[data]	/dev/sdae	6.95 GB	77%
[journal]	/dev/sdae	2.05 GB	22%

--> The above OSDs would be created if the operation continues

--> do you want to proceed? (yes/no) n

--> aborting OSD provisioning for /dev/sdae

```
[vagrant@node3 ~]$ sudo ceph-volume lvm batch --journal-size 200 --filestore /dev/sdae
```

Total OSDs: 1

Type	Path	LV Size	% of device
[data]	/dev/sdae	8.80 GB	97%
[journal]	/dev/sdae	200.00 MB	2%

--> The above OSDs would be created if the operation continues

--> do you want to proceed? (yes/no) no

--> aborting OSD provisioning for /dev/sdae

Trying again with default from ceph.conf of 100mb (complains correctly):

```
[vagrant@node3 ~]$ sudo ceph-volume lvm batch --filestore /dev/sdae
```

--> Refusing to continue with configured size for journal

--> RuntimeError: journal sizes must be larger than 2GB, detected: 100.00 MB

```
[vagrant@node3 ~]$ cat /etc/ceph/ceph.conf
```

[global]

fsid = c9c407a8-fccd-4812-9013-9caea855abe8

mon_initial_members = node2

mon_host = 192.168.111.101

auth_cluster_required = cephx

auth_service_required = cephx

auth_client_required = cephx

[osd]

osd journal size = 100

osd mkfs options xfs = -f -i size=2048

osd mkfs type = xfs

osd mount options xfs = noatime,largeio,inode64,swallock

Version:

```
[vagrant@node3 ~]$ ceph --version
```

ceph version 14.2.4 (75f4de193b3ea58512f204623e6c5a16e6c1e1ba) nautilus (stable)

Are you still seeing this problem Guillaume?

#2 - 11/04/2019 08:45 AM - Guillaume Abrioux

I was seeing this in 14.2.2, didn't have a chance to test with a newer release.

If you didn't reproduce using 14.2.4, I guess it means it has been fixed meantime and we can close this issue.

#3 - 03/03/2020 04:04 PM - Dimitri Savineau

This issue is still present in master/octopus (15.1.0-1575-g8034044)

```
# grep 'osd journal size' /etc/ceph/ceph.conf
osd journal size = 1024
```

```
# ceph-volume --cluster ceph lvm batch --filestore --yes --journal-size 4096 /dev/sda /dev/sdb --journal-devic
es /dev/sdc
```

```
Running command: /sbin/vgcreate --force --yes ceph-journals-b7f7def1-7e86-448c-920d-d0e8c6e754a6 /dev/sdc
stdout: Physical volume "/dev/sdc" successfully created.
```

```
stdout: Volume group "ceph-journals-b7f7def1-7e86-448c-920d-d0e8c6e754a6" successfully created
```

```
--> Refusing to continue with configured size for journal
```

```
--> RuntimeError: journal sizes must be larger than 2GB, detected: 1024.00 MB
```

#4 - 03/03/2020 04:22 PM - Guillaume Abrioux

additional info:

it seems to happen only when using --journal-devices

#5 - 03/03/2020 04:30 PM - Sébastien Han

- Status changed from Can't reproduce to New

#6 - 03/03/2020 04:35 PM - Guillaume Abrioux

- Subject changed from *ceph-volume CLI doesn't allow to override ceph.conf param to journal size can't be overridden with --journal-size when using --journal-devices in lvm batch mode*

#7 - 09/02/2020 01:34 PM - Jan Fajerski

- Status changed from New to Pending Backport

- Backport set to octopus,nautilus

- Pull request ID set to 36847

#8 - 09/03/2020 08:55 AM - Nathan Cutler

- Copied to Backport #47283: *nautilus: journal size can't be overridden with --journal-size when using --journal-devices in lvm batch mode added*

#9 - 09/03/2020 08:55 AM - Nathan Cutler

- Copied to Backport #47284: *octopus: journal size can't be overridden with --journal-size when using --journal-devices in lvm batch mode added*

#10 - 10/03/2020 02:07 AM - 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".