

## Orchestrator - Bug #46558

### cephadm: paths attribute ignored for db\_devices/wal\_devices via OSD spec

07/15/2020 06:20 PM - Dimitri Savineau

<b>Status:</b>	Pending Backport	<b>% Done:</b>	0%
<b>Priority:</b>	High		
<b>Assignee:</b>	Juan Miguel Olmo Martínez		
<b>Category:</b>	cephadm/osd		
<b>Target version:</b>	v17.0.0		
<b>Source:</b>	Community (dev)	<b>Affected Versions:</b>	v16.0.0
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>		<b>Pull request ID:</b>	39415
<b>Regression:</b>	No	<b>Crash signature (v1):</b>	
<b>Severity:</b>	2 - major	<b>Crash signature (v2):</b>	
<b>Reviewed:</b>			

#### Description

When creating an OSD spec for using dedicated devices for either DB and/or WAL bluestore devices, we can't use the paths drive group attribute [1] because it seems currently ignored but works for data\_devices.

```
# cat osds.yml
service_type: osd
service_id: foo
placement:
  label: osds
data_devices:
  paths:
    - /dev/sdc
db_devices:
  paths:
    - /dev/sdd
encrypted: false
# cephadm shell -m osds.yml -- ceph orch apply osd -i /mnt/osds.yml
INFO:cephadm:Inferring fsid b6dc7042-c6ac-11ea-974b-fa163e8d447e
INFO:cephadm:Inferring config /var/lib/ceph/b6dc7042-c6ac-11ea-974b-fa163e8d447e/mon.ofgnapinv-1/c
onfig
INFO:cephadm:Using recent ceph image docker.io/ceph/daemon-base:latest-master
WARNING: The same type, major and minor should not be used for multiple devices.
Scheduled osd.foo update...
```

As a result, there's an OSD created on /dev/sdc only without the bluestore DB on /dev/sdd

```
# cephadm shell -- ceph-volume lvm list
INFO:cephadm:Inferring fsid b6dc7042-c6ac-11ea-974b-fa163e8d447e
INFO:cephadm:Inferring config /var/lib/ceph/b6dc7042-c6ac-11ea-974b-fa163e8d447e/mon.ofgnapinv-1/c
onfig
INFO:cephadm:Using recent ceph image docker.io/ceph/daemon-base:latest-master
WARNING: The same type, major and minor should not be used for multiple devices.
WARNING: The same type, major and minor should not be used for multiple devices.
WARNING: The same type, major and minor should not be used for multiple devices.
```

```
===== osd.5 =====
```

```
[block] /dev/ceph-bd41b3ba-23a0-4e03-bb1c-5e17f89136a1/osd-block-fde2284b-1f59-48ab-995f-7
06a3e49c6bf
```

```
block device /dev/ceph-bd41b3ba-23a0-4e03-bb1c-5e17f89136a1/osd-block-fde2284b-1f59-48ab-995f-706a3e49c6bf
block uuid avv5vs-2gZg-how1-kYnq-PAGL-bsfL-qAHg8m
cephx lockbox secret
cluster fsid b6dc7042-c6ac-11ea-974b-fa163e8d447e
cluster name ceph
crush device class None
encrypted 0
osd fsid fde2284b-1f59-48ab-995f-706a3e49c6bf
osd id 5
osdspec affinity foo
type block
vdo 0
devices /dev/sdc
```

[1] [https://github.com/ceph/ceph/blob/master/src/python-common/ceph/deployment/drive\\_group.py#L27](https://github.com/ceph/ceph/blob/master/src/python-common/ceph/deployment/drive_group.py#L27)

#### Related issues:

Related to Orchestrator - Bug #44738: drivegroups/cephadm: db_devices don't g...	<b>Won't Fix</b>
Related to Orchestrator - Bug #49191: cephadm: service_type: osd: Failed to a...	<b>Duplicate</b>
Related to Orchestrator - Bug #46687: MGR_MODULE_ERROR: Module 'cephadm' has ...	<b>New</b>

#### History

##### #1 - 07/22/2020 10:39 AM - Sebastian Wagner

- Category changed from *cephadm* to *cephadm/osd*

##### #2 - 08/24/2020 04:31 AM - Saputro Aryulianto

Dimitri Savineau wrote:

When creating an OSD spec for using dedicated devices for either DB and/or WAL bluestore devices, we can't use the paths drive group attribute [1] because it seems currently ignored but works for data\_devices.

[...]

As a result, there's an OSD created on /dev/sdc only without the bluestore DB on /dev/sdd

[...]

[1] [https://github.com/ceph/ceph/blob/master/src/python-common/ceph/deployment/drive\\_group.py#L27](https://github.com/ceph/ceph/blob/master/src/python-common/ceph/deployment/drive_group.py#L27)

Hi,  
there is any update for this issue? since create wal devices by paths attribute is very common and important

**#3 - 08/24/2020 01:22 PM - Dimitri Savineau**

It looks like it has been decided to ignore this at the moment

<https://github.com/ceph/ceph/pull/36543>

**#4 - 08/24/2020 01:48 PM - Sebastian Wagner**

- Status changed from *New* to *Resolved*
- Target version changed from *v16.0.0* to *v15.2.5*
- Pull request ID set to *36543*

**#5 - 08/24/2020 01:53 PM - Dimitri Savineau**

@Sebastien : does that mean we need another tracker for implementing this feature or it won't happen at all ?

I thought PR 36543 was only a temporary solution until something was implemented.

**#6 - 09/11/2020 11:32 AM - Joshua Schmid**

DriveGroups are supposed to `describe` a state/layout without explicitly pointing to disk identifiers.

If there is a specific example where the existing filters are not sufficient to depict the desired cluster layout we can certainly talk about extending the functionality of the drivegroups.

`paths` for `data_devices` are mainly for testing purposes and should not be used in production clusters unless absolutely unavoidable(for whatever reason).

**#7 - 01/28/2021 03:51 PM - Juan Miguel Olmo Martínez**

- Status changed from *Resolved* to *In Progress*
- Assignee set to *Juan Miguel Olmo Martínez*
- Priority changed from *Normal* to *High*
- Target version changed from *v15.2.5* to *v17.0.0*
- Affected Versions deleted (*v15.2.4*)

**#8 - 02/05/2021 02:30 PM - Sebastian Wagner**

- Related to Bug #44738: *drivegroups/cephadm: db\_devices don't get applied correctly when using "paths" added*

**#9 - 02/05/2021 02:31 PM - Sebastian Wagner**

- Related to Bug #49191: *cephadm: service\_type: osd: Failed to apply: "NoneType" object has no attribute "paths" added*

**#10 - 02/17/2021 11:53 AM - Juan Miguel Olmo Martínez**

- Related to Bug #46687: *MGR\_MODULE\_ERROR: Module 'cephadm' has failed: No filters applied added*

**#11 - 03/10/2021 02:51 PM - Sebastian Wagner**

- Status changed from *In Progress* to *Pending Backport*
- Pull request ID changed from *36543* to *39415*