

# Move DB to bigger storage

I.e: to move from 4GB partition to 30GB partition in case of spillover:

```
1 OSD(s) experiencing BlueFS spillover
```

```
osd.3 spilled over 985 MiB metadata from 'db' device (3.8 GiB used of 4.0 GiB) to slow device
```

```
lvdisplay
  --- Logical volume ---
    LV Path                /dev/ceph-66ad9fa6-0f65-4fc0-baeb-41089a509347/osd-
block-4a9e142e-ed4d-4890-a018-5651c36071a5
    LV Name                 osd-block-4a9e142e-ed4d-4890-a018-5651c36071a5

    --- Logical volume ---
    LV Path                /dev/ceph-ffd9911c-2d6c-43aa-b637-bec8a2092a30/osd-
db-ad0a2108-bcca-4f15-9747-29d86624b911
    LV Name                 osd-db-ad0a2108-bcca-4f15-9747-29d86624b911
    VG Name                 ceph-ffd9911c-2d6c-43aa-b637-bec8a2092a30

    LV Size                 4,00 GiB
    Current LE              1024
vgdisplay
  --- Volume group ---
    VG Name                 ceph-ffd9911c-2d6c-43aa-b637-bec8a2092a30
    VG Size                 4,00 GiB
    PE Size                 4,00 MiB
    Total PE                1025
    Alloc PE / Size         1024 / 4,00 GiB
    Free PE / Size          1 / 4,00 MiB

pvdisplay
  --- Physical volume ---
    PV Name                 /dev/nvme0n1p5
    VG Name                 ceph-ffd9911c-2d6c-43aa-b637-bec8a2092a30
    PV Size                 <4,01 GiB / not usable 2,00 MiB
    Allocatable             yes
    PE Size                 4,00 MiB
    Total PE                1025
    Free PE                 1
    Allocated PE            1024
    PV UUID                 pjYGgY-SS26-m8Sc-aWyQ-3zzY-f7Iz-XGoI90
```

## resize exsiting DB partition

Remove separate DB and migrate data back to main storage:

```
ceph osd set noout
systemctl stop ceph-osd@3.service

cat /var/lib/ceph/osd/ceph-3/fsid
lvdisplay

ceph-volume lvm migrate --osd-id 3 --osd-fsid 024c05b3-6e22-4df1-
b0af-5cb46725c5c8 --from db --target ceph-5582d170-
f77e-495c-93c6-791d9310872c/osd-block-024c05b3-6e22-4df1-b0af-5cb46725c5c8
--> Migrate to existing, Source: ['--devs-source',
'/var/lib/ceph/osd/ceph-0/block.db'] Target: /var/lib/ceph/osd/ceph-0/block
--> Migration successful.
```

Create new PV,VG,LV for DB and attach it:

```
pvccreate /dev/nvme0n1p6
vgcreate ceph-db-8gb /dev/nvme0n1p6
lvcreate -n db-8gb -l 100%FREE ceph-db-8gb

ceph-volume lvm new-db --osd-id 0 --osd-fsid 024c05b3-6e22-4df1-
b0af-5cb46725c5c8 --target ceph-db-8gb/db-8gb
--> Making new volume at /dev/ceph-db-8gb/db-8gb for OSD: 0
(/var/lib/ceph/osd/ceph-0)
Running command: /usr/bin/chown -h ceph:ceph
/var/lib/ceph/osd/ceph-0/block.db
Running command: /usr/bin/chown -R ceph:ceph /dev/dm-0
--> New volume attached.
```

```
ceph-volume lvm migrate --osd-id 0 --osd-fsid 024c05b3-6e22-4df1-
b0af-5cb46725c5c8 --from data db --target ceph-db-8gb/db-8gb
--> Migrate to existing, Source: ['--devs-source',
'/var/lib/ceph/osd/ceph-0/block'] Target: /var/lib/ceph/osd/ceph-0/block.db
--> Migration successful.
```

```
systemctl start ceph-osd@3.service
ceph osd unset noout
```

## OR: create new one

Create new 41GB (41984MB) partition (/dev/nvme0n1p6)

```
pvccreate /dev/nvme0n1p6
vgcreate ceph-db-40gb /dev/nvme0n1p6
lvcreate -n db-40gb -l 100%FREE ceph-db-40gb
```

```
ceph osd set noout
systemctl stop ceph-osd@3.service
```

```
cat /var/lib/ceph/osd/ceph-3/fsid
```

```
ceph-volume lvm migrate --osd-id 3 --osd-fsid 4a9e142e-ed4d-4890-a018-5651c36071a5 --from db --target ceph-db-40gb/db-40gb
```

```
systemctl start ceph-osd@3.service  
ceph osd unset noout
```

## Move DB disks

In case of replace (full copy by dd) NVMe drive where partition for OSD's DB is located, OSD service cannot start. Reason is that CEPH doesn't have permission to access device owned by root. To repair everything with one command:

```
ceph-volume lvm activate --all
```

## draft

```
$ systemctl stop ceph-osd@68  
$ ceph-bluestore-tool --path /var/lib/ceph/osd/ceph-68 --devs-source  
/var/lib/ceph/osd/ceph-68/block --dev-target  
/var/lib/ceph/osd/ceph-68/block.db bluefs-bdev-migrate  
$ systemctl start ceph-osd@68
```

From:

<https://niziak.spox.org/wiki/> - niziak.spox.org

Permanent link:

[https://niziak.spox.org/wiki/vm:proxmox:ceph:db:move\\_db\\_disks](https://niziak.spox.org/wiki/vm:proxmox:ceph:db:move_db_disks)

Last update: **2024/02/18 16:28**

