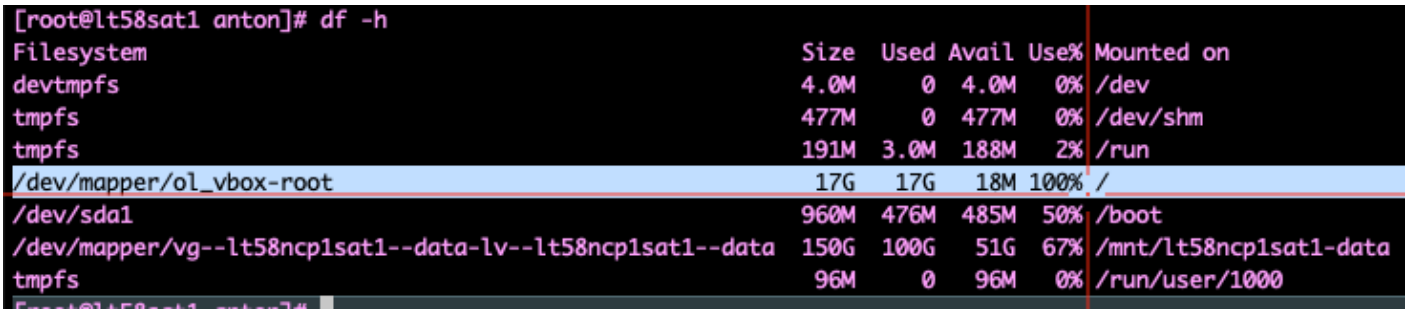
LVM - extend the volume

To extend the volume, we have to attach new physical (virtual) disk on hypervisor.

Preparations

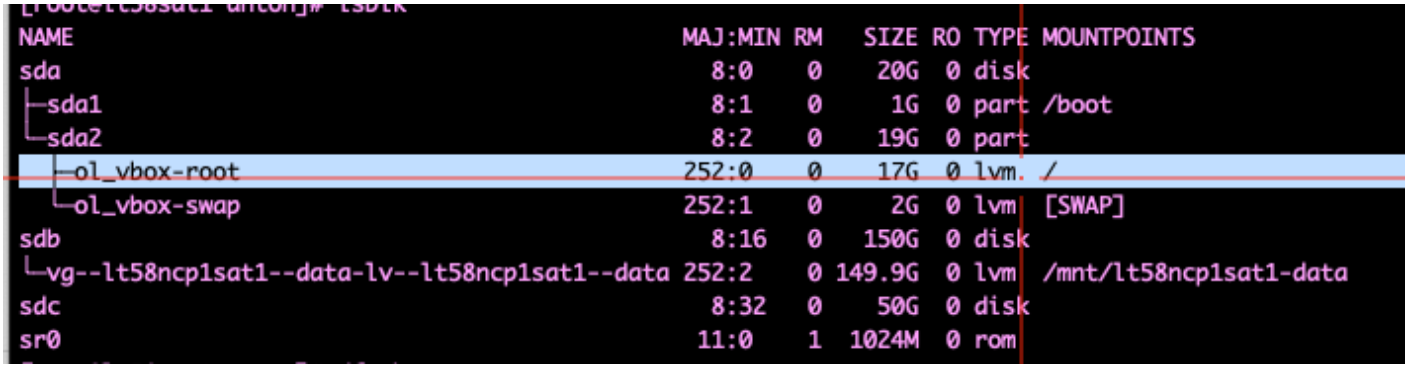
We need to understand which disk need to be extended: In my case it is a root which is full.

```
df -h
```



and verify it is LVM

```
lsblk
```



After adding the disk, it should appear in the list:

```
lsblk
```

```
[root@lt58sat1 anton]# lsblk
NAME                                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINTS
sda                                             8:0    0    20G  0 disk
├─sda1                                          8:1    0     1G  0 part /boot
└─sda2                                          8:2    0    19G  0 part
  ├─ol_vbox-root                              252:0    0    17G  0 lvm  /
  └─ol_vbox-swap                              252:1    0     2G  0 lvm  [SWAP]
sdb                                             8:16   0   150G  0 disk
└─vg--lt58ncp1sat1--data-lv--lt58ncp1sat1--data 252:2    0 149.9G  0 lvm  /mnt/lt58ncp1sat1-data
sdc                                             8:32   0    50G  0 disk
sr0                                            11:0    1  1024M  0 rom
```

Let's note which file system in use: 'XFS' or 'ext4', will be needed later..

```
lsblk -f
```

```
[root@lt58sat1 anton]# lsblk -f
NAME                                          FSTYPE      FSVER    LABEL UUID                                   FSAVAIL FSUSE% MOUNTPOINTS
sda
├─sda1                                        xfs                      03636341-0b72-4bea-9c7b-da2bf8850754    484.9M    49% /boot
└─sda2                                        LVM2_member LVM2 001     lNSByB-LaV5-PwbI-J7uK-tyer-c7TL-vg30PB
  ├─ol_vbox-root                              xfs                      8e68b127-555a-47ab-aa26-a50ece118e91     17.1M   100% /
  └─ol_vbox-swap                              swap        1            9fd91766-82f0-49e4-9281-03f20efb6632                [SWAP]
sdb                                           LVM2_member LVM2 001     eFwsQp-rLnt-WRGw-eYrG-h2Y6-qc9x-NdISvc
└─vg--lt58ncp1sat1--data-lv--lt58ncp1sat1--data xfs                      79f14e77-d9bf-4145-ace1-b10cdaa0c702     50.3G    66% /mnt/lt58ncp1sat1-data
sdc
sr0
```

We are good to go, prepare for activities. Replace your values for variables `dev` and `fs` with correct ones. Subtract `0.1` from partition size:

```
export dev="/dev/sdc"
```

Disk need to be marked as a "Physical Volume" to be able to join the "volume group" and list:

```
pvcreate ${dev}
pvscan
```

```
[root@lt58sat1 anton]# pvcreate /dev/sdc
  Physical volume "/dev/sdc" successfully created.
[root@lt58sat1 anton]# pvscan
  PV /dev/sda2   VG ol_vbox               lvm2 [<19.00 GiB / 0    free]
  PV /dev/sdb    VG vg-lt58ncp1sat1-data  lvm2 [<150.00 GiB / 96.00 MiB free]
  PV /dev/sdc                             lvm2 [50.00 GiB]
  Total: 3 [218.99 GiB] / in use: 2 [168.99 GiB] / in no VG: 1 [50.00 GiB]
```

First, let's define, which 'volume group' need to be extended

```
vgscan
```

```
[root@lt58sat1 anton]# vgscan
  Found volume group "ol_vbox" using metadata type lvm2
  Found volume group "vg-lt58ncp1sat1-data" using metadata type lvm2
```

In my case, it is `ol_vbox`. Let's define it:

```
export vg="ol_vbox"
```

Let add "new Physical Volume" to the "Virtual Group":

```
echo ${vg}
echo ${dev}
vgextend ${vg} ${dev}
```

```
[root@lt58sat1 anton]# vgextend ${vg} ${dev}
  Volume group "ol_vbox" successfully extended
```

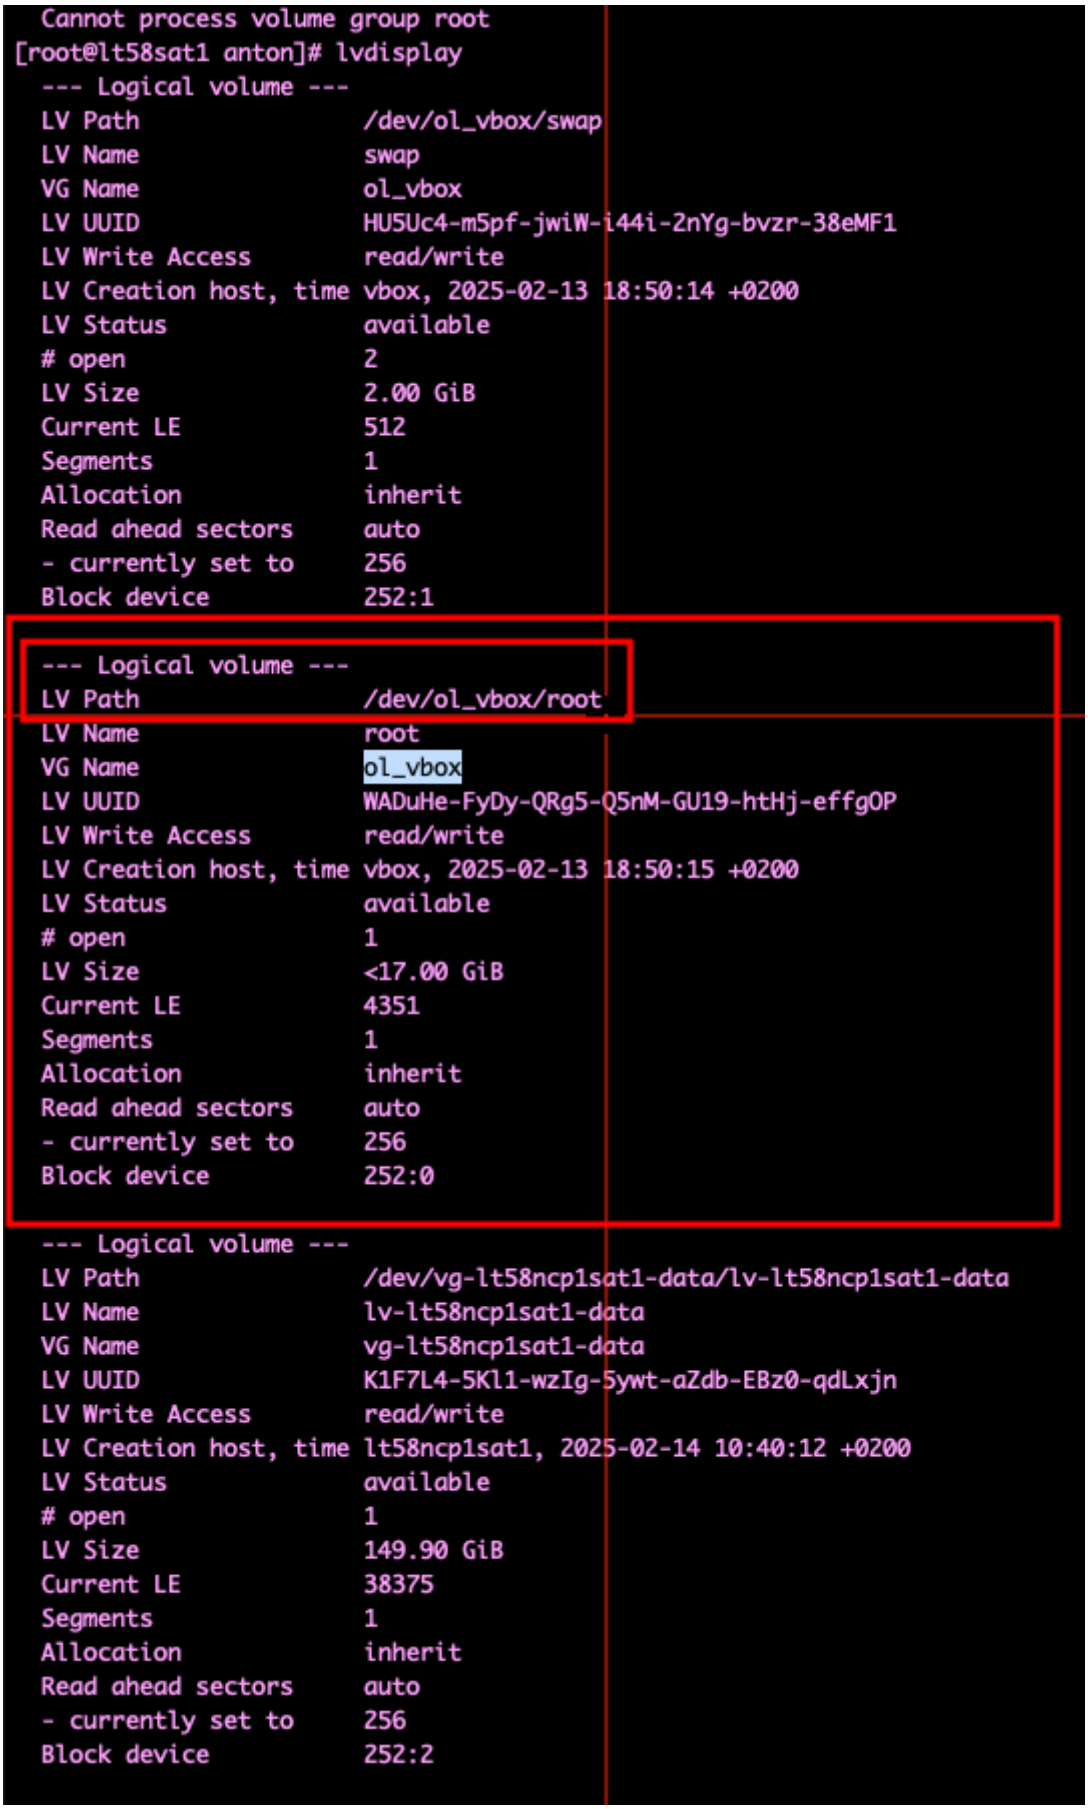
List Physical Volumes, new disk should be shown

```
pvscan
```

```
[root@lt58sat1 anton]# pvscan
  PV /dev/sda2   VG ol_vbox                lvm2 [<19.00 GiB / 0    free]
  PV /dev/sdc    VG ol_vbox                lvm2 [<50.00 GiB / <50.00 GiB free]
  PV /dev/sdb    VG vg-lt58ncp1sat1-data   lvm2 [<150.00 GiB / 96.00 MiB free]
  Total: 3 [<218.99 GiB] / in use: 3 [<218.99 GiB] / in no VG: 0 [0   ]
```

Now it is a time to extend Logical Volume. Let's observe current ones first

```
lvdisplay
```



In my case, "Volume Group" is linked to "Logical Volume" has path "/dev/ol_vbox/root". Set up extension size.

```
export lv="/dev/ol_vbox/root"
export size="49.9G"
```

and extend "Logical Volume"

```
lvextend -L +${size} ${lv}
```

```
[root@lt58sat1 anton]# lvextend -L +${size} ${lv}
  Rounding size to boundary between physical extents: 49.90 GiB.
  Size of logical volume ol_vbox/root changed from <17.00 GiB (4351 extents) to <66.90 GiB (17126 extents).
  Logical volume ol_vbox/root successfully resized.
```

Depending on which file system is in use, note current partition size, expand it and check size again:

```
df -hT
# df -hT | grep -v overlay | sort

# for "ext4"
resize2fs ${lv}

# for "xfs"
xfs_growfs ${lv}

df -h
```

```
[root@lt58sat1 anton]# df -h /
Filesystem               Size  Used Avail Use% Mounted on
/dev/mapper/ol_vbox-root   17G   17G   18M 100% /
[root@lt58sat1 anton]# xfs_growfs ${lv}
meta-data=/dev/mapper/ol_vbox-root isize=512    agcount=4, agsize=1113856 blks
         =                       sectsz=512   attr=2, projid32bit=1
         =                       crc=1        finobt=1, sparse=1, rmapbt=0
         =                       reflink=1    bigtime=1 inobtcount=1 nrext64=0
         =                       exchange=0
data     =                       bsize=4096   blocks=4455424, imaxpct=25
         =                       sunit=0      swidth=0 blks
naming   =version 2              bsize=4096   ascii-ci=0, ftype=1, parent=0
log      =internal log           bsize=4096   blocks=16384, version=2
         =                       sectsz=512   sunit=0 blks, lazy-count=1
realtime =none                   extsz=4096   blocks=0, rtextents=0
data blocks changed from 4455424 to 17537024
[root@lt58sat1 anton]# df -h /
Filesystem               Size  Used Avail Use% Mounted on
/dev/mapper/ol_vbox-root   67G   18G   50G  26% /
```

Q.E.D.

Storage management using LVM is a logical process, once understood.

Revision #2
Created 15 March 2026 11:21:37 by Anton
Updated 6 August 2026 15:28:55 by Anton