

LVM(Logical Volume Manager) is a software layer that sits between the operating system and the physical storage. It allows you to create logical volumes that can be managed independently of the underlying physical storage. This means you can create logical volumes of any size, in any order, and on any physical storage configuration. You can also create logical volumes that span multiple physical volumes, allowing for increased flexibility and scalability.

LVM: Logical Volume Manager (LVM) is a software layer that sits between the operating system and the physical storage.

Physical Volume (PV)

A physical volume is typically a hard disk, though it may well just be a device that 'looks' like a hard disk (eg. a software raid device).
 You create a physical volume by using the `pvcreate` command on a partition (e.g. `/dev/sdb1, /dev/sdc1`).

Physical Extent (PE)

Each physical volume is divided into chunks of data, known as physical extents. These extents have the same size as the logical extents for the volume group.
 Example: `PE1 = 4M, VG = 65534, PE = 4M, VG = 65534 * 4M = 256G`.
 Example: `PE = 4M, VG = 8k-16GB`.

Volume group (VG)

The Volume Group is the highest level abstraction used within the LVM. It gathers together a collection of Logical Volumes and Physical Volumes into one administrative unit.
 Example: `PV (/dev/sdb1, /dev/sdc1) in VG (/dev/volum00) on /dev/sda`.

Logical Volume (LV)

The equivalent of a disk partition in a non-LVM system. The LV is visible as a standard block device; as such the LV can contain a file system (eg. `/home`).
 Example: `VG partition (/dev/volum00/LVOL1) on /dev/sda1`.

LVM is a software layer that sits between the operating system and the physical storage.

????????? partition(/dev/sdb1 /dev/sdc1) ??????????(/dev/vg).????????????????????? partition(/dev/vg/
lv1,/dev/vg/lv2).????????? partition ???.

????????? LVM2

????????????????????? LVM ???,?????????????????????????????.

```
[root@benjr ~]# cat /boot/config-kernel-version | grep -i lvm
#Multi-device support (RAID and LVM)
CONFIG_BLK_LVM=m
CONFIG_LVM_PROC_FS=m
```

?????????????????????????????????,????????????????????? LVM ???
????????????????????????????????!!????????????????????? Multi-device support (RAID and LVM) ?? Logical volume
manager(LVM) support ????

LVM2 ??? Device Mapper?????????

```
[root@benjr ~]# cat /boot/config-kernel-version | grep -i BLK_DEV_DM
CONFIG_BLK_DEV_DM=m
```

????????????????????????????,?? RedHat ?????????? LVM
LVM1:<ftp://sources.redhat.com/pub/lvm/current>

LVM2:<ftp://sources.redhat.com/pub/lvm2>

```
???? "/" (root)" ??? LVM ?????? LVM ?????????,????????????,????? initrd .
?????????
```

```
[root@benjr ~]# vi /etc/modules.conf
```

```
alias block-major-58 lvm-mod
alias char-major-109 lvm-mod
```

```
???? initrd, ?????????? initrd
```

note: ?? boot ?????????????? LVM ?.

Physical Volume (PV) ???

PV ??????? partition ?????? 3 ? PV,????? 3 ? partition (sdb1 , sdc1 ,sdd1) ,????????? type ?? 8e (LVM ??????)

```
[root@benjr ~]# fdisk -L
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            1         1217     9775521    8e  Linux LVM
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/sdc1            1         1217     9775521    8e  Linux LVM
```

```
Device Boot      Start         End      Blocks   Id  System
/dev/sdd1            1         1217     9775521    8e  Linux LVM
```

```
?? /etc/lvmtab ?????????????????? "#vgscan" ??????.
```

```
[root@benjr ~]# vgscan
```

```
vgscan -- reading all physical volumes (this may take a while...)
```

```
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
```

```
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

```
????????????? partition ? LVM ????
```

```
[root@benjr ~]# pvcreate /dev/sdb1 /dev/sdc1 /dev/sdd1
```

```
pvcreate -- physical volume "/dev/sdb1" successfully created
```

```
pvcreate -- physical volume "/dev/sdc1" successfully created
```

```
pvcreate -- physical volume "/dev/sdd1" successfully created
```

????????????,????????????

```
[root@benjr ~]# pvscan
PV /dev/sdb1 lvm2 [9.32GB]
PV /dev/sdc1 lvm2 [9.32GB]
PV /dev/sdd1 lvm2 [9.32GB]
Total:3[27.97GB] / in use:0[0GB] /in no VG:3[27.97 GB]
```

????pvdisplay???

Volume Group (VG) ???

???????????? 3 ? PV(sdb1,sdc1,sdd1) ????? VG(vg1)

```
[root@benjr ~]# vgcreate -s 32M vg1 /dev/sdb1 /dev/sdc1 /dev/sdd1
vgcreate -- INFO: using default physical extent size 32 MB
vgcreate -- INFO: maximum logical volume size is 2 Terabyte
vgcreate -- doing automatic backup of volume group "vg1"
vgcreate -- volume group "vg1" successfully created and activated
```

??PE??? VG ??,????????PE???????? -s (???????? 8k-16GB ???????),??? VG ????? vg1

```
[root@benjr ~]# vgscan
```

```
vgscan -- reading all physical volumes (this may take a while...)
```

```
vgscan -- found active volume group "vg1"
```

```
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
```

```
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

```
???? volume ?????
```

```
[root@benjr ~]# vgdisplay -v vg1
```

```
--- Volume group ---
```

```
VG Name vg1
```

```
VG Access read/write
```

```
.....?
```

```
--- No logical volumes defined in "vg1" ---
```

```
--- Physical volumes ---
```

```
PV Name (#) /dev/sdb1 (1)
```

```
PV Status available / allocatable
```

```
Total PE / Free PE 297 / 297
```

```
.....?
```

Logical Volume (LV) ???

LVM ??????????(mapping)??,1.Linear 2.Mirrored 3.Striped ??? Linear ?????????? PV(Physical Volume) ?????

?.???? VG(Volume Group) ??? PV(Physical Volume) ?????????? 2.Mirrored 3.Striped ??? RAID ?????,???????

? Soft RAID ? ??Mirrored ??? RAID 1 ,Striped ??? RAID 0.

1.Linear

???????????? sdb1,sdc1????? Linear LVM Disk,????????????,? sdb1 ?????????? sdc1 ??????.

2.Mirrored

???????????? sdb1,sdc1????? Mirrored LVM Disk,???????????????? sdb1 sdc1 ??????????.????? RAID 1.?????????????????
??.

3.Striped

???????????? sdb1,sdc1????? Striped LVM Disk,????????????????sdb1 ,sdc1 ?????.????? RAID 0.????????????????????????????

?????????????.????????????????????.

???????? linear ??

```
[root@benjr ~]# lvcreate -L 15G -n LVOL1 vg1
lvcreate -- doing automatic backup of "vg1
lvcreate -- logical volume "/dev/vg1/LVOL1" successfully created
```

note:????????,???????????? LVOL1 , LVOL2 ??? LV ????? LVOL1

```
[root@benjr ~]# lvscan
lvscan -- ACTIVE "/dev/vg1/LVOL1" [15 GB]
lvscan -- 1 logical volumes with 15 GB total in 1 volume group
lvscan -- 1 active logical volumes
```

???? LVOL1 ?????

```
[root@benjr ~]# lvdisplay /dev/vg1/LVOL1
--- Logical volume ---
LV Name /dev/vg1/LVOL1
VG Name vg1
LV Write Access read/write
.....?.....
```

??, ??????? LVM ???????????????????:

```
[root@benjr ~]# mkfs.ext3 /dev/vg1/LVOL1
mkfs.ext3 1.32 (09-Nov-2002)
Filesystem label=
.....?
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 25 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

?????????????????

```
[root@benjr ~]# mount /dev/vg1/LVOL1 /mnt
```

???? LV (Logical Volume)???????
 ??????????????LVM1,??? umount ????????????

```
[root@benjr ~]# umount /mnt/
```

?????????????? e2fsadm

```
[root@benjr ~]# e2fsadm -L +10G /dev/vg1/LVOL1
e2fsadm 1.32 (09-Nov-2002)
.....?
The filesystem on /dev/vg1/LVOL1 is now 6553600 blocks long.
e2fsadm -- ext2fs in logical volume /dev/vg1/LVOL1 successfully extended to 25 GB
```

?????? "/dev/vg1/LVOL1" ?? 25 GB ?(???)

```
[root@benjr ~]# lvscan
lvscan -- ACTIVE "/dev/vg1/LVOL1" [25 GB]
lvscan -- 1 logical volumes with 25 GB total in 1 volume group
lvscan -- 1 active logical volumes
```

????????????? "-" ??????

```
[root@benjr ~]# e2fsadm -L -10G /dev/vg1/LVOL1
```

?????????????????

```
[root@benjr ~]# mount /dev/vg1/LVOL1 /mnt
```


LVM2????????umount?????????. ????????

```
[root@benjr ~]# lvextend -L +10G /dev/vg1/LVOL1
```

```
[root@benjr ~]# ext2online -L +10G /dev/vg1/LVOL1
```

???? VG(Volume Group)??????

???? VG ??????,????????????,?????? VG ???.??? 1 ? partition ,????????? type ?? 8e(LVM ??????),?????????????
partition ? LVM ?????.

```
[root@benjr ~]# pvcreate /dev/sde1
```

```
pvcreate -- physical volume "/dev/sdb1" successfully created
```

? /dev/sde1 ?? vg1

```
[root@benjr ~]# vgextend vg1 /dev/sde1
```

```
vgextend -- INFO: maximum logical volume size is 2 Terabyte
```

```
vgextend -- doing automatic backup of volume group "vg1"
```

```
vgextend -- volume group "vg1" successfully extended
```

```
[root@benjr ~]# vgdisplay vg1
```

```
--- Volume group ---
```

```
VG Name vg1
```

```
VG Access read/write
```

```
VG Status available/resizable
```

```
VG # 0
```

```
MAX LV 256
```

```
Cur LV 1
```

```
Open LV 0
```

```
MAX LV Size 2 TB
```

```
Max PV 256
```

```
Cur PV 4
```

```
Act PV 4
```

```
VG Size 37.12 GB
```

```
.....?.....
```

???? VG Size ???,????????? #e2fsadm ?????? LV ?????

?? LV,VG

??? LV???? umount

```
[root@benjr ~]# umount /mnt
```

???????? LV ?

```
[root@benjr ~]# lvremove /dev/vg1/LVOL1
```

??? VG???? deactivate

```
[root@benjr ~]# vgchange -a n volume  
vgchange -- volume group "vg1" successfully deactivated
```

???????? VG ?

```
[root@benjr ~]# vgremove vg1  
vgremove -- volume group "vg1" successfully removed
```

??????

???LVM?????,?????????.??? manpage

I. PV ????

pvchange pvdata pvm pvscan

pvcreate pvdisplay pvmove pv.sh

II. VG ????

vgcfgbackup vgcreate vgmerge vgscan

vgcfgrestore vgdisplay vgmknodes vgscan.static

vgchange vgexport vgreduce vgsplit

vgchange.static vgextend vgremove vgwrapper

vgck vgimport vgrename

III.LV ????

lvchange lvdisplay lvmchange lvmsadc lvreduce lvrename

lvcreate lvextend lvmdiskscan lvmsar lvremove lvscan

LVM mapping modes (linear/striped) ???

???? LV Properties ??? LVM ??????????(mapping)??,1.Linear 2.Mirrored 3.Striped ??? Linear ????????????

Linear ,?????????? PV(Physical Volume) ??????.???? VG(Volume Group) ???PV(Physical Volume) ??????????

? 2.Mirrored 3.Striped ??? RAID ?????,????????? Soft RAID ? ??Mirrored ??? RAID 1 ,Striped ??? RAID 0.

Striped:

Striped(RAID 0) 2 PV(Physical Volume) VG(Volume Group) PV(Physical Volume) LV(Logical Volume) PV(Physical Volume).

```
[root@benjr ~]# lvcreate -i 3 -l 4 -L 15G -n striped vg1
```

```
lvcreate -- doing automatic backup of "vg1"
```

```
lvcreate -- logical volume "/dev/vg1/striped" successfully created
```

note: striped 2 striped 2, -i 3 3 3 (/dev/sdb1, /dev/sdc1, /dev/sdd1) -l 4 stripesize (2*n, n=2 to 9).

Mirrored:

```
[root@benjr ~]# lvcreate -L 15G -m1 -n mirrored vg1 --corelog
```

```
lvcreate -- doing automatic backup of "vg1"
```

```
lvcreate -- logical volume "/dev/vg1/mirrored" successfully created
```

note: mirrored 2 mirrored 2, --corelog PV(Physical Volume) Mirrored Log, Log mirror.

LV striped, mirrored

```
[root@benjr ~]# lvscan
```

```
lvscan -- ACTIVE "/dev/vg1/striped" [15 GB] inherit
```

```
lvscan -- ACTIVE "/dev/vg1/mirrored" [15 GB] inherit
```

```
lvscan -- 2 logical volumes with 30 GB total in 1 volume group
```

```
lvscan -- 2 active logical volumes
```

striped, mirrored.

```
[root@benjr ~]# hdparm -t /dev/vg1/mirrored
```

```
Timing buffered disk reads: 182 MB in 3.01 seconds = 60.49 MB/sec
```

```
HDIO_DRIVE_CMD(null) (wait for flush complete) failed: Inappropriate ioctl for device
```

```
[root@benjr ~]# hdparm -t /dev/vg1/striped
```

```
/dev/vg1/striped:
```

```
Timing buffered disk reads: 702 MB in 3.01 seconds = 233.54 MB/sec
```

```
HDIO_DRIVE_CMD(null) (wait for flush complete) failed: Inappropriate ioctl for device
```

Snapshots ??

??? snapshot ,??? snapshot ??????????,????????????? LV(Logical Volume) ??,???? LV(Logical Volume) ?????
 LV(Logical Volume) ???,?????????,snapshot ??????????(copy on write),? LV(logical Volume) ???? snapshot
 ?.LV(Logical Volume) ? snapshot ?????????????,??? LV(Logical Volume) ? snapshot ?????????????,?? snapshot ?
 ??????????.??? LV(Logical Volume) ?????????????,????????????? snapshot ,???? LV(logical Volume) ??????????.
 <?? snapshot ?????????????????? snapshot(?),????? snapshot ???
 ??????????.LVM 1 sanpshot ????????? LVM2 ? snapshot ???????.

?? snapshot LV

??? snapshot LV(logical Volume) ????????????????? VG(Volume Group),????????? LV ??? 15-20 %.?????????????
 /dev/vg1/stripped ?? 15 GB???,????????? 3G ? snapshot ???,????????????????,????? LV ??? read only ???,????????????
 LV.

```
[root@benjr ~]# mount -o ro /dev/vg1/LVOL1 /databse
```

```
???? /dev/vg1/stripped ?????????????,?????
```

```
[root@benjr ~]# mount -o remount,ro /database
```

```
[root@benjr ~]# lvcreate -s -L 3G -n sanpshot /dev/vg1/stripped
```

```
lvcreate -- WARNING: the snapshot must be disabled if it gets full
lvcreate -- INFO: using default snapshot chunk size of 64 KB for "/dev/vg1/sanpshot"
lvcreate -- doing automatic backup of "vg1"
lvcreate -- logical volume "/dev/vg1/sanpshot" successfully created
```

```
????????? LV ??? Read/Write ??.
```

```
[root@benjr ~]# mount -o remount,rw /database
```

???? snapshot

```
?????????,????? read only ???.
```

```
[root@benjr ~]# mkdir /backup
```

```
[root@benjr ~]# mount -o ro /dev/vg1/snapshot /backup
```

??????????????

```
[root@benjr ~]# tar -cf /dev/st0 /backup
```

Snapshot ???

?????????? snapshot ?????? (partition) ??????????????. snapshot ?????????? /dev/vg1/stripped ?????(?,?,?)
?,????????????? /dev/vg1/snapshot ??,????????????????????????????????!!!!

```
[root@benjr ~]# ls /database
```

```
total 13
-rw-r--r-- 1 root root 1 Apr 2 00:28 a.test.file
drwxr-xr-x 2 root root 12288 Apr 2 00:28 lost+found
```

```
[root@benjr ~]# ls /backup
```

```
total 13
-rw-r--r-- 1 root root 1 Apr 2 00:28 a.test.file
drwxr-xr-x 2 root root 12288 Apr 2 00:28 lost+found
```

?????? /database ?? /backup????????,???? /database/a.test.file ??.

```
[root@benjr ~]# rm /database/a.test.file
```

```
[root@benjr ~]# ls /backup
```

```
total 13
-rw-r--r-- 1 root root 1 Apr 2 00:28 a.test.file
drwxr-xr-x 2 root root 12288 Apr 2 00:28 lost+found
```

?????????? a.test.file ?????? /dev/vg1/snapshot ??,??? snapshot ??????????????.

note:

????? a.test.file ?,???? inode ??????. ?????? a.test.fil ??? chunk(64KB) ??? 'dirty' ????????? 'snapshot' volume. ?
????? snapshot ??????????,??? PE(chunks,blocks) ??????.????????? snapshot ??????,??? snapshot??????
(deactivated).

?? snapshot

???????? snapshot ??

```
[root@benjr ~]# umount /backup
```

```
[root@benjr ~]# lvremove /dev/vg1/snapshot  
lvremove -- do you really want to remove "/dev/vg1/snapshot"? [y/n]: y  
lvremove -- doing automatic backup of volume group "vg1"  
lvremove -- logical volume "/dev/vg1/snapshot" successfully removed
```

??? LVM ???????????