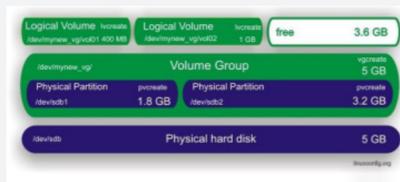


### LVM - schema



Example of an LVM layout.

#### LVM - terminology

<b>PE - physical extents</b>	The smallest unit of disk space that can be managed by LVM. It's set when you create a volume group. It determines the "granularity" of disk space allocation. All physical volumes in a volume group will use the same extent size.
<b>LE - logical extents</b>	The same as PE but for Logical Volumes
<b>PV - physical volume</b>	Disks, partitions, RAID volumes
<b>VG - volume group</b>	Group of one or more physical volumes
<b>LV - logical volume</b>	Entity that contains information, store on physical volumes, grouped within a volume group

#### LVM - commands and components

<b>lvm</b>	pseudo shell for LVM
<b>lvmconfig</b>	prints LVM configuration from /etc/lvm/lvm.conf

#### LVM - commands and components (cont)

<b>lvmdiskscan</b>	prints devices that could be used as physical volumes
<b>lvmdump</b>	dumps LVM configuration to tar file
<b>lvmetad</b>	caching service for LVM metadata

#### LVM - important files

<b>/etc/lvm/</b>	LVM configuration directory
<b>/etc/lvm/lvm.conf</b>	LVM configuration file
<b>/etc/lvm/archive/{backup}</b>	Backup and archive of LVM configuration
<b>/etc/lvm/profile</b>	A set of selected customizable configuration settings that can be used to achieve certain characteristics in various environments or uses. Normally, the name of the profile should reflect that environment or use. An LVM profile overrides existing configuration.

#### LVM - Physical Volume

<b>pvcreate</b>	create PV
<b>pvchange</b>	change attributes of PV
<b>pvs pvdisplay</b>	display information about PV
<b>pvmove</b>	moves extents from one PV to another
<b>pvremove</b>	removes a PV
<b>pvresize</b>	resizes a PV
<b>pvscan</b>	scans for changes in PV configuration and size

#### LVM - Volume Group

<b>vgs vgdisplay</b>	displays VG information
<b>vgmknodes</b>	create the special files for VG devices in /dev
<b>vgck</b>	check VG consistency
<b>vgcfs vgbackup</b>	make a backup of VG metadata
<b>vgcfs vgrestore</b>	restore VG metadata from backup
<b>vgrename</b>	rename a VG
<b>vgsplit</b>	move PVs to a new or existing VG
<b>vgconvert</b>	change VG metadata format

#### LVM - Volume Group (cont)

<b>vgimport</b>	register exported VG in the system
<b>vgexport</b>	unregister VG from the system
<b>vgimportclone</b>	import a VG from cloned PVs
<b>vgcreate</b>	create VG from one or more PVs
<b>vgremove</b>	remove VG
<b>vgreduce</b>	remove PVs from VG
<b>vgchange</b>	change VG attributes
<b>vgextend</b>	extend VG with one or more PVs
<b>vgmerge</b>	merge VGs
<b>vgscan</b>	scan VGs for changes in metadata and size

#### LVM - Logical Volume

<b>lvs lvdisplay</b>	print LV metadata
<b>lvchange</b>	change attributes of LV
<b>lvcreate</b>	create a new LV
<b>lvextend</b>	extend size of LV
<b>lvreduce</b>	reduce size of LV (offline!)
<b>lvresize</b>	change size of LV



LVM - Logical Volume (cont)	
<b>lvrename</b>	rename a LV
<b>lvconvert</b>	change layout of LVM
<b>lvscan</b>	scan LVs for metadata changes and size
<b>lvremove</b>	remove LV

LVM - create PV	
<b>pvccreate</b>	create a new PV from device
<code>/dev/&lt;disk_device&gt;</code>	device
<code>pvccreate /dev/sdb</code>	
To create a PV you need a usable storage device. Try <code>lvmdiskscan</code> command.	

LVM - create VG	
<b>vgcreate</b>	create a volume groups that consists of one or more PVs
<code>&lt;VG&gt;</code>	groups that consists of one or more PVs
<code>&lt;PV&gt;</code>	of one or more PVs
<code>&lt;PV&gt;</code>	
<code>vgcreate data_vg /dev/sdb /dev/sdc</code>	
To create a VG you need one or more PVs.	

LVM - create LV	
<b>lvcreate -n</b>	create a LV of a given size in a VG
<code>&lt;LV_NAME&gt;</code>	a given size in a VG
<code>L &lt;SIZE&gt;</code>	a VG
<code>&lt;VG&gt;</code>	
<code>lvcreate -n data_lv -L 100G data_vg</code>	
<b>lvcreate -n</b>	create a LV that fills 100% of free space in VG
<code>&lt;LV_NAME&gt;</code>	that fills 100% of free space in VG
<code>I 100%FREE</code>	of free space in VG
<code>&lt;VG&gt;</code>	in VG

LVM - create LV (cont)	
<code>lvcreate -n data_lv -l 100%FREE data_vg</code>	
To create an LV you need a volume group.	

LVM - Snapshots	
<b>lvcreate -L &lt;SIZE&gt;</b>	Create a snapshot of an LV
<code>-s -n &lt;NAME_OF_SNAPSHOT&gt;</code>	of an LV
<code>&lt;PATH_TO_TH_TO_LV&gt;</code>	
<code>lvcreate -L 4G -s -n data_lv-snapshot /dev/mapper/data--vg--data--lv</code>	
<b>lvremove &lt;PATH_TO_SNAPSHOT&gt;</b>	Remove snapshot
<code>lvremove /dev/mapper/data--vg--data--lv</code>	
<b>lvextend -L &lt;SIZE&gt;</b>	Extend snapshot
<code>&lt;PATH_TO_LV&gt;</code>	snapshot
<code>lvextend -L +10G /dev/mapper/data--vg--data--lv</code>	
<b>lvconvert --merge</b>	Restore data from snapshot
<code>&lt;PATH_TO_SNAPSHOT&gt;</code>	from snapshot
<code>lvconvert --merge /dev/mapper/data--vg--data--lv</code>	

LVM - Snapshots (cont)	
modify <b>snapshot_autoextend_threshold</b> and <b>snapshot_autoextend_percent</b>	Enable automatic extension of snapshots by adjusting the values
Snapshots contain differences from the point a given LV has been created - not the real data. If the size of a snapshot is exceeded it becomes useless and there is no way to restore data from it. Old state of an LV can be restored from a snapshot, but please remember to unmount the filesystem first.	

Useful storage commands (Bonus)	
<b>lsblk -o name, mountpoint, label, size, uuid</b>	List block devices with useful information
<b>multipathd show maps status</b>	Print path status of multipath devices
<b>lvs -o +devices</b>	Print LVs, VGs and device paths
<b>vgs -o +lv_size, lv_name</b>	Print information about LVs, VGs, sizes and attributes
<b>mkfs.ext4 &lt;FS&gt;</b>	Format a filesystem for EXT4
<b>mkfs.xfs &lt;FS&gt;</b>	Format a filesystem for XFS

Useful storage commands (Bonus) (cont)	
<b>mount &lt;PATH_TO_LV&gt;</b>	Mount a filesystem in a given location
<code>&lt;MOUNTPOINT&gt;</code>	given location
<b>mount -a</b>	Mount all filesystems listed in /etc/fstab
Useful in many situations when you need to manipulate storage devices and layout.	

LVM - extend PV	
<b>rescan-scsi-bus.sh   echo 1 &gt; /sys/block/&lt;DEVICE&gt;/device/rescan</b>	Rescan underlying storage to detect change in disk size
<code>rescan-scsi-bus.sh   echo 1 &gt; /sys/block/sdb/device/rescan</code>	
<b>partprobe</b>	Detect changes in partition size (if necessary)
<b>pvresize &lt;PV&gt;</b>	Resize PV to the maximum possible size of partition or disk/LUN
<code>pvresize /dev/sdb</code>	

To resize a PV you'll need to know the underlying storage and know how to rescan its size. All steps can be performed online.



### LVM - Extend VG

**vgextend** Extend existing  
<VG> VG with one or  
<PV> more PVs  
<PV>

```
vgextend data_vg /dev/sdb
/dev/sdc
```

To extend a VG you need a VG and one or more PVs that are not assigned to any PVs yet.

### LVM - extend LV

**lvextend -L** extend logical  
<SIZE> volume by given  
/dev/m- size (add 10GB to  
apper<- existing size)  
PATH\_T-  
O\_LV>

```
lvextend -L +10G /dev/mapper/-
data--vg_data--lv
```

**lvextend -r** extend logical  
-L <SIZE> volume and its  
/dev/m- filesystem by  
apper<- given size (add  
PATH\_T- 10GB to existing  
O\_LV> size)

```
lvextend -r -L +10G /dev/mapp-
er/data--vg_data--lv
```

To extend an LV you need its path. Use **df** command. In most cases you'll also want to extend the underlying filesystem, so use **-r** option to do it.

### Extend filesystem (bonus)

**umount <FS>** Unmount  
filesystem

```
umount /data
```

**e2fsck -f** Check and fix  
<PATH\_TO\_- potential errors  
LV>

### Extend filesystem (bonus) (cont)

```
e2fsck -f /dev/mapper/data--vg_-
data--lv
```

**resize2fs <PA-** Reduce size  
**TH\_TO\_LV>** of filesystem  
<SIZE>

```
resize2fs /dev/mapper/data--v-
g_data--lv 10G
```

**lvreduce -L <SI-** Reduce size  
**ZE> <PATH\_-** of LV  
**TO\_LV>**

```
lvreduce -L 10G /dev/mapper/-
data--vg_data--lv
```

**e2fsck -f <PA-** Check for  
**TH\_TO\_LV>** errors once  
more

```
e2fsck -f /dev/mapper/data--vg_-
data--lv
```

**mount <FS>** Mount  
filesystem

```
mount /data
```

Remember to make sure the size of both: FS (filesystem) and LV (logical volume) are the same and there are no errors. Reducing the size can ONLY be done OFFLINE. The filesystem must be unmounted. Not all filesystems can be reduced at all.

