

CST8207 – Linux Operating Systems I

Logical Volume Manager

Overview

- Why is disk space management hard?
- What is LVM?
- What's it for?
- Behind the scenes

Disk Space Management

- Disk Space Management is difficult
- Why?
 - Hard drives were (looooong ago) small
 - And became larger and larger
 - Multiplied (and divided)
 - Needed to be backed up
 - The complexity invited data loss
 - Especially when backup tapes were used

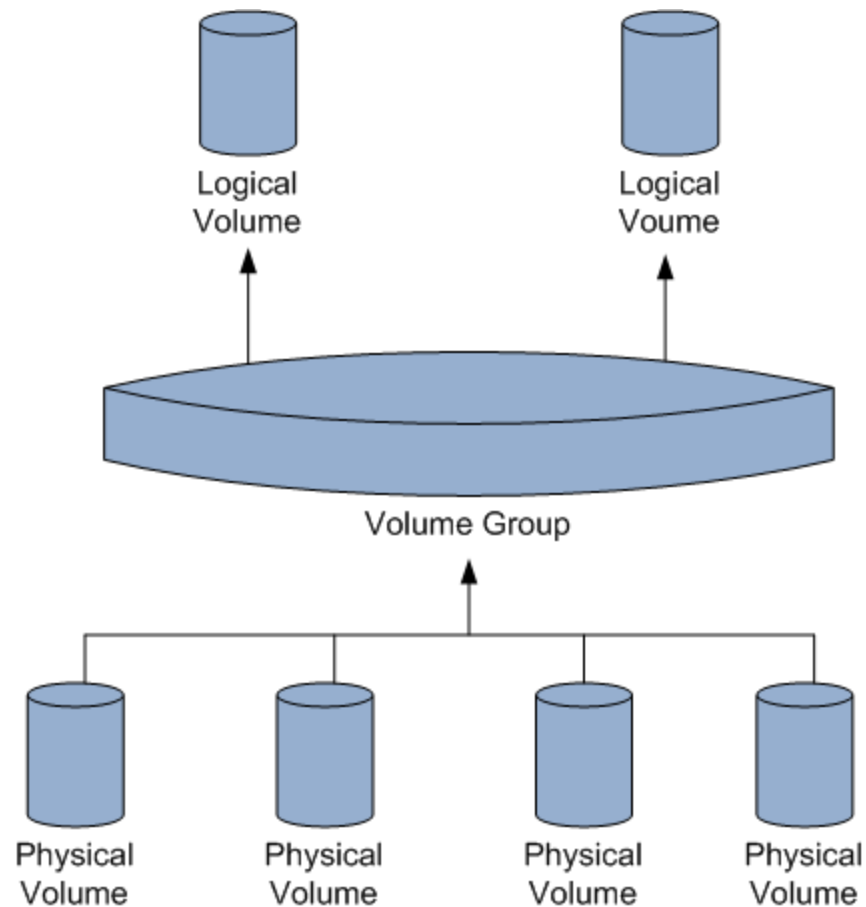
LVM Origins

- Probably inspired by HAL (Hardware Abstraction Layer)
 - HAL treats all different types of hardware the same
 - at least that's how it seems from an application's perspective
- As has been seen, file systems are not really as “virtual” as other aspects.
 - E.g. The first hard drive is sda, the next sdb, etc.

Virtual Disks

- If all disks are virtual, then they can be managed by software instead of hardware
- This will allow division and multiplication of hard drives without needing to be concerned about the actual hardware devices
- They can also be “named” instead of identified by device labels
- Virtual drives can be “virtually copied” making backups simple

LVM Logical Volume Components



LVM2 Capabilities

- Can create logical volumes, possibly spanning more than one physical disk, and give them any name.
- Can resize logical volumes.
- Can delete logical volumes.
- Can export logical volumes, in order to move them to a separate machine without changing the device names.
- Can create snapshots (a frozen version of a logical volume) in order to undo many changes in a single operation.

Glossary (<http://www.haifux.org/lectures/153>)

- Physical Volume (PV) - A disk (or any block device).
- UUID - A unique ID assigned to a physical volume by the LVM subsystem. The UUID is stored on the physical volume (as the last signature).
- Volume Group (VG) - A set of physical volumes.
- Physical Extent (PE) - The minimal chunk on a physical volume that is managed by LVM.
- Logical Volume (LV) - A "virtual" block device, that is created by LVM on some volume group. The logical volume is mapped to a list of physical extents, that may reside on any of the volumes in a single volume group.
- Snapshot - A (logical) freeze of a logical volume at some point in time. enables us to revert changes made to a logical volume since the snapshot was taken.

How LVM is used

- Setup with LVM2
 - need to choose some physical volumes
- Take these physical volumes, and place them in volume groups
- This causes LVM2 to assign a unique UUID for each physical volume
- Now start creating logical volumes
- Use these logical volumes like normal disks
 - Can create a file system and mount it
 - Can use them as raw devices, etc.
- Can add more physical volumes to existing volume groups, in order to add more capacity
- If the logical volumes become too full, the size can be increased
 - If there is a file system on such a logical volume, it will require resizing as well.

So... What happened when I installed Fedora?

- By default, LVM was used to set up the `/dev/sda` hard drive.
- Disk Druid set up `/dev/sda1` and `/dev/sda2`
- `/dev/sda1` was set as bootable and contains `/boot`
- To actually see where things are you can do the following:
 - `mount`
 - `lvdisplay`

My system

```
[sunger@localhost ~]$ mount
/dev/mapper/VolGroup-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sda1 on /boot type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
gvfs-fuse-daemon on /home/sunger/.gvfs type fuse.gvfs-fuse-daemon (rw,nosuid,nodev,user=sunger)
```

```
[root@localhost ~]# lvdisplay
--- Logical volume ---
LV Name                /dev/VolGroup/lv_root
VG Name                VolGroup
LV UUID                raiQHH-7RH3-fwsh-q14m-oscs-odyF-cftxA9
LV Write Access        read/write
LV Status              available
# open                 1
LV Size                13.80 GB
Current LE             3533
Segments              1
Allocation             inherit
Read ahead sectors    auto
 - currently set to   256
Block device          253:0
```

```
--- Logical volume ---
LV Name                /dev/VolGroup/lv_swap
VG Name                VolGroup
LV UUID                aBMwRw-51kD-30FZ-E1jr-XlSq-45eG-1nT44R
LV Write Access        read/write
LV Status              available
# open                 1
LV Size                1.00 GB
Current LE             256
Segments              1
Allocation             inherit
Read ahead sectors    auto
 - currently set to   256
Block device          253:1
```