

Linux Logical Volume Manager

What is Linux LVM?

The Linux Logical Volume manager (LVM) is a group of applications that provides the Administrator flexibility when working with mass storage. This includes the following:

- **Striping** The ability to write one file to multiple hard drives
- **Mirroring** The ability to copy verbatim information to a backup harddrive “on the fly.”
- **Dynamic Resizing** Resize partitions as necessary without downtime
- **Online Backups** Create “snapshots” of your filesystems for later restoration.

Why Linux LVM?

- Performance
 - Striping information across several hard drives increases I/O throughput
- Redundancy
 - Making an exact copy of one drive to another drive, or ‘‘mirroring’’ enables the server to continue running even after HDD failure
- Flexibility
 - Partitions are dynamic; can modify partition size as necessary.

Available Documentation

- My Website:
`www.metasource.us`
- LVM HOWTO:
`http://www.linuxdocs.org/`
- Sistina's Websight:
`http://www.sistina.com/products_lvm.htm`
- Meta Newsgroup:
`http://groups.google.com`

LVM Terms

- **Physical Volume**
 - The physical disk drive

- **Physical Extent**
 - Partion that is defined as a member of the Logical Volume Manager.

- **Volume Group**
 - Unit described as one or more Physical Extents grouped or “bundled” to create one volume (see diagram)



- It is possible (even likely) that a Volume Group has Physical Extents from several hard drives.

- **Logical Volume**




- Logical unit that is “usable” by the Administrator. The logical Volume `/dev/fill/in,` can have a filesystem placed on it and mounted for use:

```
mkfs.xfs /dev/fill/in/path/usr; mount  
/dev/fill/in/in/path/usr /usr
```


- The “meat” of the Logical Volume Manager



Volume Group



Volume Group






Installation Overview

- Partition Hard Drives
- Lay Down a Base System
- Fullfill Prerequisites (Kernel Compile, Userspace Utilities, etc.)
- Install LVM
- Create Physical Extents

- Define Volume Groups
- Create Logical Volumes
- Create Filesystems on Physical Volumes
- Mount Physical Volumes/Move Data
- Update Init Scripts, /etc/fstab

Beginning Configuration



Partition Hard Drives





- Drive 0 (/dev/hda)
 - Will use two drives
 - 1 4.3 Gig Drive
 - 1 2.5 Gig Drive
 - 5 Partitions:
 - /boot: 100 MB
 - swap 100 MB
 - root 1.5 GB (will become /usr in final config)
 - system: 600 MB
 - data: 200 MB

- Drive 1 (/dev/hdd)
 - 3 Partitions:
 - swap: 100 MB
 - system: 1600 MB
 - data: 200 MB
 - Note: Use of hdd keep drives on Separate IDE chains
 - Be sure to change the partition types to type “8e” (Linux LVM) when creating partitions targeted for LVM

Fullfill Prerequisites

- Compile “Logical volume manager (LVM) support” under the “Multi-device support (RAID and LVM)” section of the kernel compilation
- Get LVM tar.gz, RPM, apt-get, or emerge depending on your disrobution

Final Configuration



Create Physical Extents

- `pvcreate /dev/hda7 /dev/hda8`
- `pvcreate /dev/hdd1 /dev/hdd5 /dev/hdd6`
- This tells the LVM about the partitions available for use.

Create Volume Groups

- In this example, will create three Volume Groups:
 - System
 - Data
 - Swap
- `vgcreate system /dev/hda7 /dev/hdd5`





- Display a Volume Group
vgdisplay system
- vgcreate data /dev/hda8 /dev/hdd6
- Display a Volume Group
vgdisplay data
- vgcreate swap /dev/hdd1
- Display a Volume Group
vgdisplay swap

Create Logical Volumes

- Logical Volumes are the last step before we create a filesystem. Finally!!

- Basic command:
lvcreate -n var -L 1000M
- But, we would like to stripe the data for performance...
- # lvcreate -n var -i 2 -I 64 system -L 900M
- # lvcreate -n tmp -i 2 -I 64 system -L 200M
- # lvcreate -n usr system -L 1000M
Single block for right now, will stripe during data conversion transfer.
- # lvcreate -n home -i 2 -I 64 data -L 392M
- # lvcreate -n swap0 swap -L 100M

Final Configuration



Create Filesystems

- `# mkfs.xfs /dev/system/var`
- `# mkfs.xfs /dev/system/tmp`
- `# mkfs.xfs /dev/system/usr`
- `# mkfs.xfs /dev/data/home`
- `# mkswap /dev/swap/`

Mount Filesystems

- `# mkdir /mnt/tmp_var`
- `# mount /dev/system/var /mnt/tmp_var`

Update Init Scripts

In `/etc/init.d/rc.d/rc.local` (RedHat)
`# vgchange -a y`

Migrate Filesystems

- `# cd /var`
- `# find . -depth print — cpio pd /mnt/tmp_var`
- Repeat for the rest of the filesystems

