What is a hard drive?

- Read/write head, spinning platters
- Billions or trillions of 1's and 0's
- Nonvolatile memory
  - Means the data stays if the machine is turned off

Usually split into “partitions”

- Partition doesn't mean wall in this case...
- Regions of the disk
  - More a “room” than a “wall”
- Put 1 file system on each partition
  - File system is the set of rules for how you arrange data on the partition.
- Usually can't be resized without losing data

The /dev directory

- Files that aren't normal files
  - They refer to a physical device, not data
- hda, hdb are the first two IDE drives
- sda, sdb, etc are SCSI or SATA drives
- A number after that references a partition
  - hda1, hda2, hda3, etc are partitions on hda

/dev example

- sdb3
  - S is SCSI or SATA drive
  - D means disk
  - B means 2nd disk
  - 3 means third partition
- 3rd partition of 2nd SCSI/SATA disk

Basic questions

- When is a hard drive not a hard drive?
  - When there's a “disconnect” between software and hardware
- Why would you do this?
  - Flexibility of installation
  - Protection from hardware failures

RAID

- Redundant
  - Data is stored in multiple places
- Array
  - Multiple disks acting as one
- Inexpensive
  - At least inexpensive relative to losing your data
- Disks
  - Could work with any nonvolatile storage

If hard drives were cakes

- You have 4 sheet cakes and 40 people
- Split them up evenly
- cut each one into 40 pieces?
- Make a 2x2 grid, and pretend it's a big cake
- Cut the big cake into 40 pieces
- Now do it with partitions rather than cakes

Hardware vs Software RAID

- Hardware
  - Expensive, computer CPU is free
  - All drives must be the same type
- Software
  - Built into Linux (and some other OS's)
  - Doesn't require special card
  - Can mix drives (IDE, SATA, SCSI, etc)
  - Slightly slower

Types (“levels”)of RAID

- Require at least 2 drives
- RAID 0, striping
  - Fastest
  - As storage-efficient as a single drive
  - Not actually redundant. Any drive failure kills the array
- RAID 1, mirroring
  - Good performance
  - Redundant. Can survive failure of one drive
  - Only half the total space is usable

Parity

- 3 + 5 = 8
- 3 + ? = 8
  - If you “lose” the 5
- ? = 8 – 3
- ? = 5
  - It can be recovered

Higher RAID levels

- Level 4
  - Rarely used
  - Requires 3 or more drives
  - 1 drive for parity
  - Only “lose” 1 drive's worth of storage
- Level 5, stripe with parity
  - Still requires 3 drives
  - Parity data is spread across all drives
  - Much more common.
  - Available in software RAID

/dev again

- Software RAID creates more /dev files
  - Acts the same regardless of RAID levels
- Usually /dev/md0, /dev/md1, etc
  - Referencing the “big cake”
- Install a file system on the array as if it was a partition.
  - mkfs /dev/hda1 for a partition
  - mkfs /dev/md0 for an array

Installing the Operating System

- Types of data
  - Boot files
    - Necessary to start the system
  - Operating system files
    - Used by the OS itself
  - Applications
    - Word processor, web browser, etc.
  - Temporary data / scratch pad
  - Data files
    - Music, papers, photos, etc.

The traditional (Microsoft) way

- One giant partition
- Advantages
  - Very simple
  - Can't guess wrong on the sizes
- Disadvantages
  - No separation between operating system, applications, and data
  - Usually can't reinstall the operating system without losing all data

The slightly-more-complicated (Linux) way

- /boot
  - Small partition
  - Just enough to get the main partition running
- Swap
  - Another partition used as “emergency memory”
  - Lets you run more programs, but the hard drive is slower than main memory
- / (root)
  - The rest of the file system

The more-complicated way

- /boot and swap as before
- A much smaller / (root) filesystem
- Attach other filesystems to it
  - /tmp (temporary file storage)
  - /var (log files)
  - /home (users' home directories)
- A “runaway logfile” won't fill the whole directory tree

Other options

- Can also create a partition for a particular application
  - Runaway process is completely contained
  - Can use an optimized file system

More letters: L V M

- Logical Volume Manager
- Without LVM
  - Hard disk
  - Partition
  - File system
  - Exactly 1 file system & mount point per partition

With LVM

- Hard disk
- Partition
- Assign partitions to LVM
- Logical volumes
  - Show up as more device files
  - /dev/lvm/<name>
- File system

Advantage of LVM

- Flexibility
  - Logical volumes (LV's) can be resized
  - LV's can span several physical volumes
    - A partition on a new drive can be added to an existing logical volume
  - You can specify which Physical Volume (or volumes) will hold a Logical Volume

What we can do

- Can create many file systems on a physical volume
- Can spread a file system across several physical volumes (or drives)
- Newer filesystems can be resized without losing data
  - ReiserFS, xfs, jfs can
  - Ext2, ext3, NTFS, FAT32, etc. can not

Common Implementation

- Two similar-sized drives
- Create RAID 1 mirror on most of it
  - Possibly keep /boot and swap separate
- Assign the mirror to LVM
- Create logical volumes of any size
  - All Logical volumes will survive a drive failure
- Mount them on your system

What we have

- Protection from hardware failures
  - Because the volumes are on a RAID mirror
- Changeable sizes
  - Because the file systems are logical volumes
- Space can be added later as needed
  - Likely would add drives as mirrored pairs