COMP421
Unix Environment for Programmers
Lecture 03: Installing Unix

Jeff Wiegley, Ph.D.
Computer Science

jeffw@csun.edu

09/12/2005

"Unix gives you just enough rope to hang yourself – and then a couple of more feet, just to be sure." –*Eric Allman*

Requirements

- Desktop, laptop or workstation with 386 or better CPU.
  - The Intel 386 CPU, did not have math Co-processing functions built in to the CPU but had a Memory Management Unit (MMU) making x86 hardware capable of running Unix efficiently.
  - Work has been done over the past several years to migrate the Linux kernel to devices which do not have a MMU (such as smaller embedded platforms.)
- bootable CDROM drive (early installation used 25 – −50 floppy discs.)
- The boot/install disc for your desired flavor of Unix.
- Enough disk space.
  - 2gig is minimal for a windowed environment.
  - 8Gig for system files is usually enough space in which to grow.
  - User's files take up an arbitrary amount space. Estimate your own needs.

Disk Partitions:

Hard drives are just a collection of sectors (usually 512 bytes per sector).

Windows (and many x86 operating systems) divide the hard disk into *partitions*

There are two kinds of partitions:

1. Primary (There can be up to four primary partitions defined)
   - One primary partitions can be of type "extended"
2. Logical (There can be any number of logical partitions and they are all stored in the "extended" primary partition.

This is x86 specific, not Unix specific. Windows uses this partitioning scheme as well but Solaris on Sparc hardware uses a different system.

Partition Purpose:

- Provides the ability to view a single drive as multiple logical, independent units.
- The BIOS boot sequence relies on partitions to locate and bring up the operating system.
- The drive has a Master Boot Record (MBR).
  - This is the **only** thing that the BIOS knows how to find and execute when the machine is powered on.
- Every primary partition has its own individual boot record.

Booting:

A computer is a very expensive brick until something starts executing.

This presents a chicken and egg problem:

- In order to start executing something a decision would have to be made about what to execute.
- In order to make a decision about what to execute the computer would already have to be running a program.

The computer solves this by executing the BIOS program by default.

The only thing the BIOS knows how to do is to load the 512 byte program stored in the master boot record and start executing it.

The master boot record that is written by a windows installation is also simple:

- search for the primary partition marked as active (aka bootable) and load the 512 byte program from that partition's boot record and transfer execution to it.

Reasonable Partitioning Example:

Let's say I want to have a computer that can run both Windows and Linux.

In general, the Linux file-system can't be mixed with the Windows file-system.

- What if both systems had a file with the same path?
- Can both operating systems understand the same file-system?

It is best to allocate some partitions to be used by Windows and others to be used by Linux.

Simplest Scheme:

Assuming that you had a 60gig drive and your purpose was to install Linux for COMP421 a minimal partition allocation would be:

Active?	*Partition*	*Primary?*	*Size*	*Type*	*Purpose*
Yes	1	Yes	43 Gig	NTFS	Windows XP
No	2	Yes	16 Gig	Linux (83)	/
No	3	Yes	1 Gig	Linux Swap (82)	swap space

"Swap" space is the area used as virtual memory to make a machine appear as though it has more physical memory. Linux can locate its swap contents as either a file or as a partition. Partitions are more efficient.[a]

[a]So yes, there is an even simpler scheme consisting of only two partitions if you move swap into the Linux file-system.

Better Scheme :

Assuming that you had a 60gig drive and your purpose was to install Linux for COMP421 and you wanted some flexibility:

Active?	*Partition*	*Primary?*	*Size*	*Type*	*Purpose*
Yes	1	Yes	43 Gig	NTFS	Windows XP
No	2	Yes	128 Meg	Linux (83)	/boot
No	3	Yes	8 Gig	Linux (83)	/
No	4	Yes	9 Gig	Extended	
No	5	No	1 Gig	Linux Swap (82)	swap space
No	6	No	8 Gig	Linux native (83)	/home

This scheme allows you to reinstall Linux over and over again without having to backup user files located in /home. As long as you don't re-format the contents of that partition you can simply mount it as /home again.[a][b]

[a] User and group ids will have to be the same on the new installation otherwise the ownership of the files in /home will have to be updated.

[b] This is the scheme the instructor defaults to.

I NEED WINDOWS!!!! :

If you edit the partition table with fdisk or any other editor and just change partition sizes you **will** destroy any existing data on the drive!!!!

"But I don't want to lose my windows data"

`Partition Magic` is a program that can resize NTFS partitions without destroying the contents.

- Its a commercial product so it costs money
- It works very well

`ntfsresize` is a free program for Linux that can resize NTFS partitions.

- Chicken/egg... ntfsresize runs under Linux... you need ntfsresize before you install Linux... argh.

Knoppix to the rescue!

Klaus Knopper is a brilliant person that developed the "Knoppix" bootable CD. (You can download a free ISO image from `http://www.knopper.net/knoppix-mirrors/index-en.html` Version 3.9 fits on a CD; version 4.0.1 requires a bootable DVD.

- Knoppix boots a machine entirely from the CDROM so does not require a hard drive. Any changes made to the operating system are lost when the machine is stopped.
- Knoppix has ntfsresize as one of its installed programs.
- Ubuntu has a similar bootable CD.
- Such CDs are typically called "Live-CDS".

Chicken/Egg solved because you can boot/run knoppix, use ntfsresize to shrink the windows partition down. Then you have room left to create Linux partitions in.

Windows Drive Letters:

Windows has this silly notion of drive "letters".

Windows is very, very fussy about what its "installation" drive is.

If you install to what Windows thinks is Drive "D" and then later windows thinks that same partition is drive "C" you're screwed. Basically if you do anything to alter driver letter assignments you're screwed.

Windows assigns drive letters to partitions in a particular order.

This is the main reason why the two partitioning schemes presented earlier have windows as the first partition.

When you had a single partition, windows assigned the drive letter "C" to it when the operating system was installed.

Leaving it the first partition guarantees that it will still get enumerated as "C" when more than one partition exists.

File-system types

Unix gets installed to a hard drive which holds the file-system. This file-system is where the first real choices come into play. The file-system "type" is an important decision and there are many to choose from:

- `minix`: file-system used by minix and very, very early Linux systems before `ext2`
- `ext`: A very early file-system type.
- `ext2`: A better, very stable and reliable file-system based on `ext`
- `ext3`: The next version, based on `ext2` it includes journaling support.[a]
- `ReiserFS`: Another journaling file-system that is faster than `ext3` but not supported by all distributions.
- `XFS`: yet another file-system type. Its claim to fame is increased performance for extremely large files. It can almost instantly delete a file of an size. (Other file-systems will have to take time to deallocate disk space
- `JFS`: IBM's journaling file-system ported to Linux.[b]

[a] journaling is a method that allows a file-system to recover from a crash with minimal damage to the file-system and maximum possibility of full and easy recovery.

[b] IBM is very helpful commercial player in the Linux world

File-system choice:

basically the file-system "type" describes the advanced data structures and disk layout used to represent and store files. Since it's basically a data structure it defines what the performance of file operations will be.

If in doubt, choose `ext3` as it is supported ubiquitously and has efficient, reliable performance characteristics.[a]

[a] Your instructor uses `ReiserFS` since it is a journaling system that appeared before `ext3` and he is a creature of habit.