AERSP 497B/597E Presentation
1/15/2002, 4:15 pm

Introduction to Unix & Linux

Contents by
Prof. Remzi H. Arpaci

Presented by
Carlo Mozetic
[http://www.carlomozetic.net]

Overview

- Review and History of Unix and Linux
- Basic Unix Commands
- Networking in Linux
- Essential System Administration
- Installing Linux (if time permits)

History of Unix

- Unix was developed in the mid 70's for mini-computers and mainframes. It was the first true “multitasking” OS.
- Developed at Bell Labs (AT&T) by Ken Thompson and Dennis Ritchie.
- Unix can run on almost all computers (from mainframes to supercomputers to PC's to hand-held computers, different flavor for each one).
- The 'real' Unix (System V) from AT&T might cost you $1500 for one copy and will only run on limited hardware (if they still sell it!!).

History of Unix

- Current day Unixes are basically Unix “System V” clones that comply with the IEEE POSIX (Portable Operating System Interface) standard.
- They are all written in C and are thus very portable given the availability of robust C compilers on every known computer hardware since the early ‘70s.

History of Linux

- Linus Torvald began working on Linux as an OS course project in 1991, when he was a student at Univ. of Helsinki (Finland). [now works at Transmeta Corporation].
- The project that started as a hobby (a “hackers” OS), became a full-fledged OS when Linus posted the source code at a bulletin board asking people for suggestions and improvements, which received an overwhelming response!
- Linux is completely Free under GPL (GNU Public License).
- First stable release: Linux kernel v1.0 in March 1994.
- Stable kernel versions have even sub-version numbers (1.2, 1.4, 2.0, 2.2, 2.4). Experimental versions have odd numbers (1.1, 1.3, 2.1, 2.3, 2.5).

History of Linux

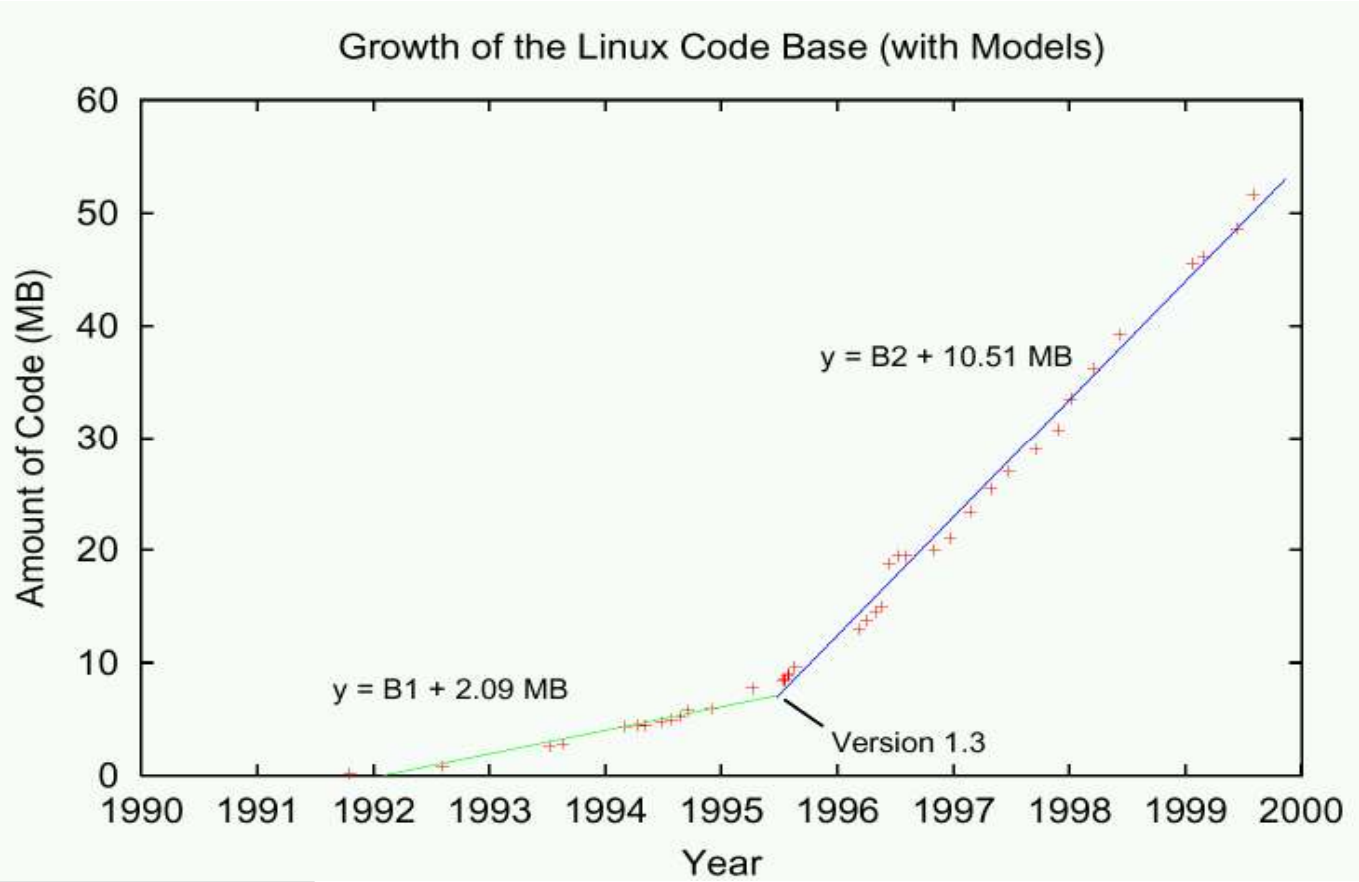

(Courtesy Prof. Remzi H. Arpaci-Dusseau)

http://www.cs.wisc.edu/~remzi/Linux/

Linux Statistics

- Official Linux Source code managed by Linus Torvalds is a whopping 107 MB today for the latest kernel v2.4.17!! (143 MB including documentation)
- Kernel v2.4.17 has 8214 unique C header/source files and 3.65 million lines of source code.
  - find . –type f –name “*.[ch]” –exec wc –l {} \; | awk ‘BEGIN {i=0} {i += $1} END {printf i “\n”}’
- About 50% of the source code consists of drivers!
- Currently well over 10 million people use Linux for their everyday needs.

Linux Copyrights

- Covered by GNU General Public License (GPL)
- People can sell and make a profit from free software, but they cannot restrict the right of others to use or distribute it.
- This is NOT public domain software
- This is NOT shareware
- Even a modified version of Linux is covered by GPL. (Read Eric Raymond's *"Cathedral and the Bazaar"* to learn more about the Open-Source movement).
- Open Source Software, means anyone can download it, modify it, and redistribute it.
- Instead of having teams of thousands of programmers (such as at Microsoft), people thru-out the world work on Linux (and other open software) as a team
- Linux is sort of notorious with Computer Security experts. Since the source is all available, hackers are able to break in sometimes. But this is often due to people running Linux who do not really understand network security. On the other hand, this also helps people remove bugs from Linux sooner than for other OSes.

What Other Operating Systems are There ?

- The infamous BSOD Win32 series: MS Windows 9x/NT/2000/XP.
- MacOS X (kernel based on the Mach kernel + BSD unix)
- BeOS (unfortunately, the company went bankrupt in December 2001!)
- IBM OS/2 (dead!)
- Various Proprietary Unixes: IBM AIX, HP-UX, SunOS, SGI IRIX, etc.
- Free POSIX OSes: FreeBSD, NetBSD, OpenBSD.

Who uses Linux ?

- Almost all types of computer users now use Linux
- Engineers and scientists use it for code development and simulation.
- System administrators.
- Network providers: networking is one of the real strengths of Linux (share files, remote logins, SAMBA, ...)
- Kernel hackers: lots of talented people on web for help .
- Multimedia authors : works with almost all sound & video cards. OpenGL has been ported. Even some Virtual Reality machines now use Linux. Very handy graphics tools called Gimp too.
- Antartica research stations
- Oceanography vessels
- Students

Linux Caveats

- Linux is just the kernel (i.e., the heart of the OS), not the OS itself.
- The OS consists of the kernel and the basic tools and utilities supporting the kernel, like the file manipulation and search commands, editors, compilers, etc.
- The kernel by itself is pretty useless…..it is like a brain without a body!
- Linux kernel + GNU utilities form the “Linux OS” as most people know it.
- e.g., RedHat Linux, Mandrake Linux, SuSe Linux, Debian Linux, Slackware Linux.

Linux Kernel

- Guts of the OS: middleman between user apps and system devices/hardware.
- A set of routines that are always in memory, and the other processes interact with it.
- Interrupt handler: Kernel puts one process to sleep and starts the next one.
- Can be either *Monolithic* or *Modular*:
  - Monolithic: all device drivers are in the kernel (as is BSD and System V)
  - Modular: most device drivers can be loaded and unloaded as modules at any time.
- Kernel can emulate math instructions if there is no math processor.
- Latest kernel series (2.4.xx) can support upto 16 processors and 64 GB of physical memory.
- Executables can use dynamically linked shared libraries (similar to SunOS). Saves space and makes efficient use of memory.

Linux Features

- Multi-tasking (more than one task/program can run simultaneously).
- Multi-user (more than one user can work simultaneously).
- Multi-processing (more than one processor is supported by the OS, SMP support is very stable).
- POSIX compliant…….behavior similar to traditional Unixes (look and feel).
- Runs on a variety of architectures (not just ix86 based PCs): Sparc, Alpha, PowerPC, MIPS, PalmPilot, ...
- An Embedded version of Linux exists for hand-held devices and real-time OS applications.

Linux Features

- Complete and very efficient networking implementation (TCP/IP, IPV6, PPP, SLIP, NFS, ...)
- Supports a wide range of filesystems:
  - Log-based filesystems (XFS, JFS, ReiserFS, ext3): favorites for servers.
  - Traditional Unix filesystems (minix, ext, ext2, UFS)
  - Network filesystems (NFS, Coda, AFS, DFS)
  - DOS/Windows filesystems (FAT16, FAT32, NTFS)
  - Other: UDF (DVDs), etc.
- Extensive Firewall support
- Standard ‘Unix' software written in C often just compiles and runs on Linux boxes as well as other machines effortlessly.
- Lots of easily and freely available public domain software.

Networking

- Networking features are atleast as good as any OS.
- Linux machines make excellent network servers, file servers, even for MS Windows machines.
- TCP/IP = transmission control protocol / internet protocol
- Web browsers, ssh, telnet, ftp, news, email, ...
- Most people use ethernet, since 10/100 Mbit/second cards are extremely inexpensive, but you can also use gigabit , ATM, ISDN, wireless, token ring, ...
- Linux based Beowulf clusters are very popular owing to the good networking infrastructure on Linux.
- PPP and SLIP are used with modems, for dial-up networking
- Apache web server is one of the most popular applications in the world and runs primarily on Linux and BSD.
- Samba lets a Linux machine act as a file and print server for MS Windows machines.

Basic Linux Commands

- Virtually every command and feature of Unix is in Linux (because of POSIX compliance).
- All the different versions of Unix (Solaris, System V, AIX, HPUX, ...) look mostly the same to 'users', but to System Administrators they can look quite different. We try to use as few flavors of Unix as possible.
- There is a large number of editors available in Linux (ed, vi/vim, elvis (GNU vi), emacs, xemacs, pico, jove, ...).
- Linux has a huge number of 'shells'. The shell is the program that reads and executes commands from the user. Modern shells do much more than that though: job control, input/output redirection, and a command language (i.e. batch files).
- The most common shells are:
  - *bash*: Bourne again shell, default, most popular, very functional.
  - *csh*: c shell
  - *tsch*: c shell with advanced features
  - *ksh*: Korn shell

Basic Linux Commands

- One of the neat things about Linux is if your machine is used for a limited number of things, you do not have to load the other features or utilities (so you can have a very lean OS). Most Linux distributions offer various categories of install (server, workstation, graphics/multimedia, development, etc.).
- Documentation: *man* and *info*.
  - e.g., *man bash* or *info bash*
  - ***Section 0 - Everything***
    Section 1 - Commands
    Section 2 - System Calls
    Section 3 - Library Calls
    Section 4 - Special Files
    Section 5 - File Formats and Conversions
    Section 6 - Games for Linux
    Section 7 - Macro Packages and Conventions
    Section 8 - System Management Commands
    Section 9 - Kernel Routines
  - E.g., *man 2 open* is different from *man open*

Linux Compilers

- You can get all the standard compilers for Linux (Fortran 77/90/95, Java, C, C++, Ada, Lisp, Pascal, ...)
- GNU C/C++ compiler (gcc/g++) is one of the best C compilers available for Linux, and it is free.
- Except for FORTRAN 90/95, there are free versions of compilers for almost every programming language in use today!
- The Java Development kit is available for Linux.
- GNU Make, RCS (revision control system) and CVS (Concurrent Versioning System) are included in most Linux distributions for efficient software development.

X Windows

- X is the standard graphical user interface for Unix.
- You can have multiple windows, and use a mouse for point and click apps.
- Look and feel of Linux machines is very professional and better than Windows machines, and highly customizable.
- For Linux, Xwindows is called Xfree86 (X11R6, XFree86 release 6) .
- It used to be a pain to get X working under Linux, since you had to configure the drivers manually, but now the new releases of Linux do this automatically
- Some standard X applications are:
  - xterm, xclock, xman, netscape, gnuplot, gimp
- There are several different X window managers, which give different look-and-feel:
  - KDE
  - GNOME
  - Enlightenment
  - Windowmaker
  - Other classic WMs: olvwm, twm, fvwm2 (classic MIT window manager, Tom's)

Interfacing Linux with Other OS

- Wine and WABI are windows emulators for Linux (limited success so far). Wine is GPL, WABI is commercial.
- DOSemu is a very stable MS-DOS emulator.Some of your partitions on your disk can be MS-DOS partitions. You can read MS-DOS floppies too.
- VMware is the best alternative, if you need to run both Linux and MS Windows. It is a commercial emulator that emulates the x86 hardware flawlessly, so any OS that can run on the x86 platform can be installed under VMware as a separate application!

Interfacing Linux with Other OS

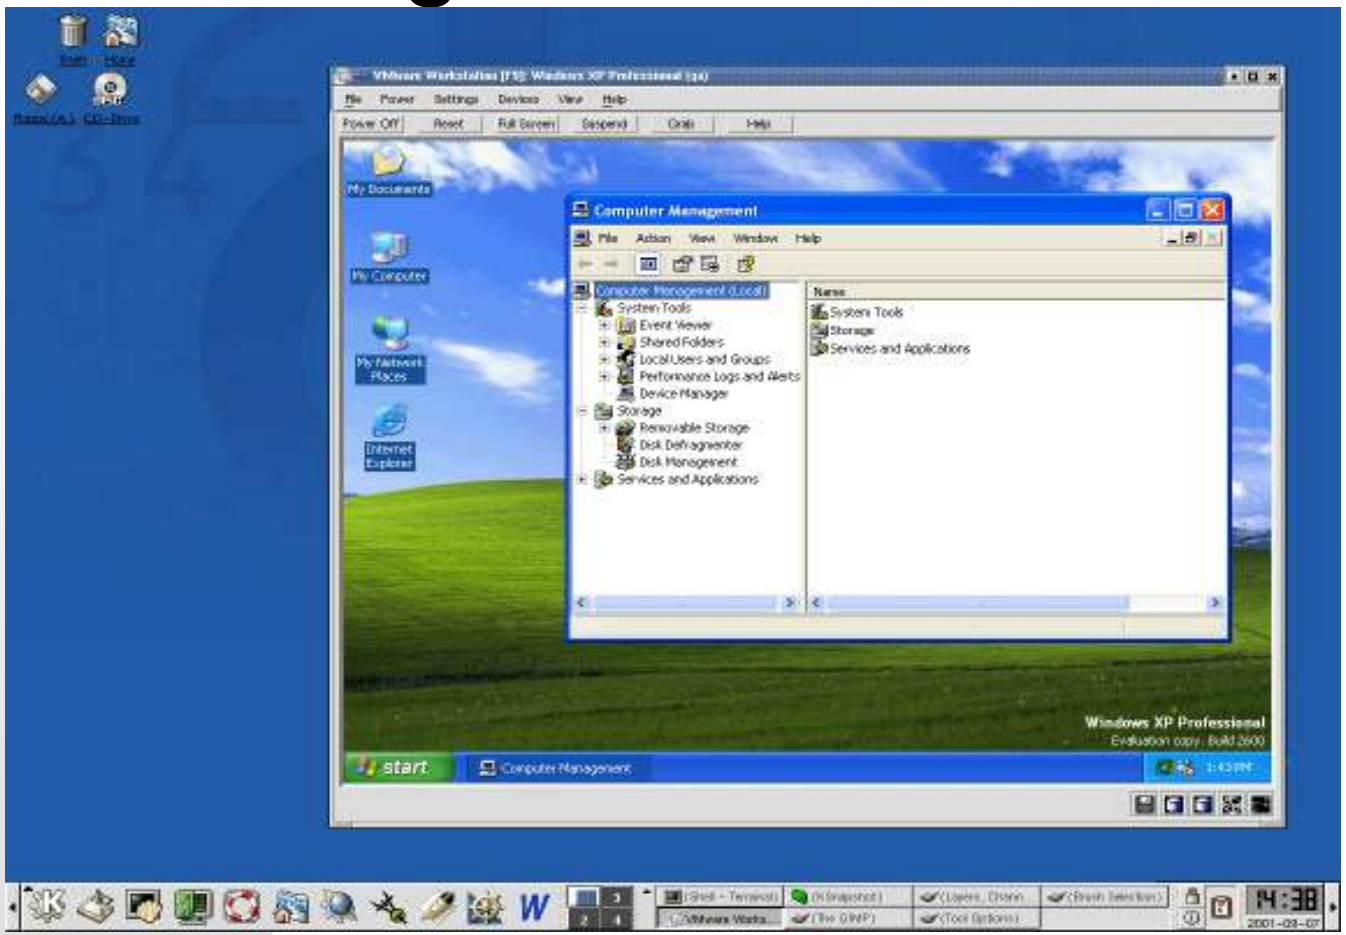

VMWare: Windows XP under Linux

Text and Word Processing

- On PC's and Mac's it is common to have WYSIWYG
- On Unix machines, text processing is the norm, i.e. 'typesetting'
- In this approach, the user puts in the content but also the commands to format the content (like writing html code).
- troff: Bell labs text processing
- groff, GNU troff, for man pages
- TeX, extremely popular
- LaTeX (the Linux book was written in LaTeX)
  - LyX is a WYSIWYG app for LaTeX
- However, Word Processing applications (like MS Office) are available on Linux (e.g.Sun's Star Office , a complete suite of applications, including a spreadsheet, a word processor, HTML editor, presentation manager, ...). This is FREE.
- Corel WordPerfect works on Linux also.
- KOffice that comes with KDE is also getting popular.

Linux Commands Tutorial

- *ls:* file listing.
- *cat, echo*
- *mkdir, rmdir*: create directories
- *chmod*
- *find*
- *sed*: stream editor
- *awk*
- *vi*
- etc.