Linux introduction and installation

In this week's lecture, we talk about the items in TCO 2. Each student will be required to perform a Linux installation. Planning for the installation is key, as well as deciding which type of Linux to use for either a workstation or server installation. Please use all items in this lecture and the accompanying lab assignment to provide research for your Project Paper as described in TCO 1. UNIX, the Precursor to Linux.

Back to Top UNIX is characterized as a multiuser, multitasking, stable, reliable, and portable OS. UNIX was developed at AT&T Bell Labs in 1969. Two programmers, Dennis Ritchie and Ken Thompson, had returned to Bell Labs after being loaded on a project named Multics with programmers from MIT and GE. The Multics project (1965–1969) was an attempt to write a dependable timesharing system in the days of batch processing on large mainframe computers. Bell Labs withdrew from the Multics project in 1969.

Upon returning to Bell Labs, Ritchie and Thompson began developing a game called Space Travel using the Multics language and a retired DEC PDP-7 minicomputer. Space Travel led to the development of a file system structure, paging, a command-line shell, and processes. In 1971, UNIX was chosen as the operating system name, a play on the name Multics. Ritchie wrote the C Language (1971–1973) to provide a more flexible programming environment for UNIX; earlier versions were written in the assembler and B languages, which had been written by Thompson.

In 1973, UNIX was selected as the operating system for AT&T's Switching Control Center System (SCCS). In 1983, the Class 5 Electronic Switching System (5ESS) running UNIX was announced by AT&T. UNIX was chosen for

its speed, flexibility, and reliability in switching voice and data across the telecommunications infrastructure.

In the early 1980s, UNIX was announced as the OS for the AT&T Computer Systems Division 3B computer product line. Today, AT&T UNIX System V Release 4 (SVR4) is one of two major UNIX releases. In the interim, Berkeley University wrote its own version of UNIX and subsequently became a major developer of UNIX distributions. To separate the university from its UNIX development project, Berkeley formed Berkeley Software Distribution (BSD) to distribute BSD-UNIX, the second major UNIX release. Other corporations adopting UNIX were Hewlett-Packard with HP-UX; Sun Microsystems with SunOS and Solaris; IBM with AIX; NCR with SVR4; DEC with Digital-UNIX; Santa Cruz Operations with SCO-UNIX; and Microsoft with XENIX (short lived).

In short, one can see that UNIX is indeed the grandfather of all operating systems in use today. Its high price, though, is a major hurdle for some organizations to overcome. In its place, several organizations have moved to using the Linux operating system.

The Introduction of Linux

Back to Top UNIX was traditionally developed for minicomputer, mainframe, and supercomputer environments, given their infrequent architectural changes. This fact makes UNIX support less costly to the vendor. However, high cost of supporting UNIX on the personal computer (PC) is prohibitive. In 1990, a student at the University of Helsinki, Linus Torvalds, decided to port an OS named Minix to the PC as a course project. Finding the project falling behind schedule, Linus sent a request for help out on the World Wide Web (the first ever). An enormous response was received from programmers

around the globe volunteering to help. At the suggestion of one such volunteer, Torvalds switched from porting Minix to the PC to porting a more popular OS named UNIX.

Today, Linux (named after Linus and UNIX, and pronounced " lin-uhks") is gaining popularity as an open-source operating system. Many organizations are using Linux as the operating system platform for their web servers, in combination with the Apache web server application. Linux is also used as a database server, DNS server, and DHCP. You are also able to create firewalls using Linux. It has gained rapid popularity in the server arena; however, it has not experienced the same growth as an end-user workstation operating system. Use of its command line interface (CLI) is still the preferred method of management by the system administrators, because the commands provide considerable flexibility and a significant number of options for management purposes.

General Public License and Linux

Back to Top Linus Torvalds wanted the entire world to be able to use Linux, without having to worry about costs. He decided to provide the kernel under the umbrella of the General Public License (GPL). Basically, the GPL states that anyone can use the kernel, and even modify it, but make all modifications available for public use and scrutiny. Thus, we have an operating system that joins other software processes in the world of open source.

Richard Stallman, an MIT graduate, started the open-source software movement in the 1980s. His idea was that software should be free and open in order to maximize development through the collective efforts of a wide

body of programmers and hobbyists alike. Stallman named his organization GNU, which stands for " GNU's not UNIX" (UNIX is a proprietary, costly OS). GNU wrote a large body of free, open-source applications. Torvalds and Stallman formed a partnership under which Linux distributions used the free GNU applications. Stallman later formed the Free Software Foundation (FSF) to fund the efforts of GNU. If you are ever up late at night browsing TV channels, look for a movie entitled Revolution OS to see the interesting story of Torvalds and Stallman.

Richard Stallman coined the term copyleft to position the ownership of Linux as the opposite of copyright under law. Copyleft implies that anyone who acquires a copy of Linux is free to add, delete, change, or modify the Linux OS in any manner he or she wishes as long as the modified source code is placed back into the public domain for others to enjoy. To legalize this concept, the GNU-General Public License (GPL) was drawn up. The GNU-GPL must accompany every distributed copy of Linux.

Linux Distributions

Back to Top The question of whether or not Linux is free is asked often. When we say Linux is free, this simply means that we are free to change the source code as long as we provide the code to the public. In the case of Red Hat Enterprise Linux, the Red Hat Corporation will provide you with the operating system for free, but you will need to purchase the support, payable as a monthly charge. In other words, companies subscribe to the Red Hat Linux Enterprise Server by paying a monthly fee in order to receive support and services. This support includes access to any bug fixes,

automated updates, unlimited 24/7 technical support, and product documentation. |

Linux can be downloaded for free from numerous websites. Downloading can take a long time and source code arrives without instructions or help desk support. You need to burn that image onto a CD and use the disc to install the operating system.

The Linux Kernel

Back to Top The UNIX kernel, the core operating system, is small compared to the entire operating system distribution made up of the shell, applications, and utilities. Its design makes it possible to write a kernel for specific architectures. The Linux kernel, when first introduced, was written for the Intel i386 processor. This feature makes Linux a portable operating system. It can be ported to many different hardware architectures. As long as a particular architecture has a C compiler, the kernel source code can be compiled into machine language for that particular architecture type. A small kernel (3. 7M lines of freely licensed source code) allows it to be compiled onto a wristwatch, PDA, or cell phone, as well as onto large minicomputers, mainframe computers, or supercomputers. Linux Installation

Back to Top Linux can be installed easily by inserting the DVD accompanying the textbook into a computer's DVD drive and then rebooting the computer. An installation screen will appear, requiring you to press Enter to begin the installation. A series of questions about your computer's hardware and network connection will walk you through the process. One key decision needs to be made during the installation process. This is about choosing the default partition configuration or the custom manual configuration and to

load specific Linux packages. If you choose manual, Disk Druid will provide a graphical interface for manually partitioning the hard drive. Six partitions (separate file systems) are required. File system names follow traditional conventions in the UNIX directory name structure. Before installing, list the type of video card, network interface card, sound card, monitor, mouse, and keyboard your system uses.

Include the monitor type and settings.

The one tool everyone should become familiar with is the hardware compatibility listing, or HCL. Many vendors of Linux have their own HCLs, and there are several online HCLs created for use by the open source community. One should note that neither the vendor HCLs nor any of the public HCLs guarantee performance. The vendor sites provide information that a particular device has worked with its software at one point in time or another. The public sites provide a broader range of testing results for several hardware and operating system combinations.

It is imperative that a thorough understanding of all hardware items be reviewed before installation. Items that must be considered are: memory type and amount, video card and monitor, central process unit, and media such as hard drives, CD, and DVD ROM devices.

After ensuring hardware compatibility, a personal computer must be equipped with a minimum configuration to host Linux. This configuration includes a 400 MHz Pentium II processor, 6. 9 GB of hard disk space, and 256 MB of RAM at a minimum for GUI. As we know, minimum configurations typically do not perform well in a business productivity setting.

To check on known Linux compatibility issues with various devices and computers, visit the site: http://www. linuxquestions. org/hcl/

Installing Linux in a business setting requires a faster method than loading each client from a DVD or set of CDs. This can be accomplished either by imaging hard drives or by copying installation files from the Linux installation CDs or DVD to a network server that can be accessed as an NFS (Network File System) directory. Another option is to access the installation files using a webpage address or FTP address from a remote server containing the installation files.

Linux Application Packages

Back to Top Application packages are user productivity tools requiring little or no formal training. Once the kernel is loaded, any number of applications can be installed that match the computing requirements for a user or user group. A popular Linux package named OpenOffice is a groupware package containing word processing, spreadsheet, and presentation packages similar to Microsoft Office that are beneficial for performing day-to-day tasks. A Web-hosting company will find the Apache server package beneficial for hosting webpages on a Linux machine. Linux users have choices to make when selecting application packages to load, update, and use. Packages that system administrators use to manage specific maintenance tasks are called utilities. Linux distributions contain a wide variety of applications and utilities that can be selected and loaded at any time. A list of applications can be found at http://freshrpms. net/.

Applications contain multiple files that must be copied to specific directories in the Linux file system. This process makes the installation of new

applications a labor-intensive process. Most Linux distributions provide a program to load, update, and remove application packages in a single step. Red Hat, Inc., developed the Red Hat Package Manager, or RPM. Today, it is called the Resource Package Manager or RPM Package Manager. The RPM system provides all the features needed to manage applications, including a database of installed packages arranged by version number for tracking purposes. Packages are provided in binary format (RPMS directory) and source code (SRPMS directory) in a depository. RPM is not a new idea. SVR4 UNIX provided the pkgadd command to install application packages.

The RPM package file format follows the official standard for Linux under the Linux Standards Base (LSB). Each package is made up of a bundle of related files and accompanying documentation. A package is labeled with four highly identifiable part names: (a) package name, (b) version number, (c) release number, and (d) architecture. Each part is separated with a dot or a dash. A package labeled as kernel-smp-2. 4. 18-3. i386. rpm is named kernel-smp. Its version is 2. 4. 18 in Release 3. The package is for an Intel 386 (i386) or compatible processor chip. All package names must end with the . rpm file identifier.