African Virtual University

Applied Computer Science: CSI 4100

OPERATING SYSTEMS ADMINISTRATION

Ashenafi Kassahun

Foreword

The African Virtual University (AVU) is proud to participate in increasing access to education in African countries through the production of quality learning materials. We are also proud to contribute to global knowledge as our Open Educational Resources are mostly accessed from outside the African continent.

This module was developed as part of a diploma and degree program in Applied Computer Science, in collaboration with 18 African partner institutions from 16 countries. A total of 156 modules were developed or translated to ensure availability in English, French and Portuguese. These modules have also been made available as open education resources (OER) on oer.avu.org.

On behalf of the African Virtual University and our patron, our partner institutions, the African Development Bank, I invite you to use this module in your institution, for your own education, to share it as widely as possible and to participate actively in the AVU communities of practice of your interest. We are committed to be on the frontline of developing and sharing Open Educational Resources.

The African Virtual University (AVU) is a Pan African Intergovernmental Organization established by charter with the mandate of significantly increasing access to quality higher education and training through the innovative use of information communication technologies. A Charter, establishing the AVU as an Intergovernmental Organization, has been signed so far by nineteen (19) African Governments - Kenya, Senegal, Mauritania, Mali, Cote d'Ivoire, Tanzania, Mozambique, Democratic Republic of Congo, Benin, Ghana, Republic of Guinea, Burkina Faso, Niger, South Sudan, Sudan, The Gambia, Guinea-Bissau, Ethiopia and Cape Verde.

The following institutions participated in the Applied Computer Science Program: (1) Université d'Abomey Calavi in Benin; (2) Université de Ougagadougou in Burkina Faso; (3) Université Lumière de Bujumbura in Burundi; (4) Université de Douala in Cameroon; (5) Université de Nouakchott in Mauritania; (6) Université Gaston Berger in Senegal; (7) Université des Sciences, des Techniques et Technologies de Bamako in Mali (8) Ghana Institute of Management and Public Administration; (9) Kwame Nkrumah University of Science and Technology in Ghana; (10) Kenyatta University in Kenya; (11) Egerton University in Kenya; (12) Addis Ababa University in Ethiopia (13) University of Rwanda; (14) University of Dar es Salaam in Tanzania; (15) Universite Abdou Moumouni de Niamey in Niger; (16) Université Cheikh Anta Diop in Senegal; (17) Universidade Pedagógica in Mozambique; and (18) The University of the Gambia in The Gambia.

Bakary Diallo

The Rector

African Virtual University

Production Credits

Author

Ashenafi Kassahun

Peer Reviewer

Thadee Gatera

AVU - Academic Coordination

Dr. Marilena Cabral

Overall Coordinator Applied Computer Science Program

Prof Tim Mwololo Waema

Module Coordinator

Robert Oboko

Instructional Designers

Elizabeth Mbasu

Diana Tuel

Benta Ochola

Media Team

Sidney McGregor

Barry Savala

Edwin Kiprono

Kelvin Muriithi

Victor Oluoch Otieno

Michal Abigael Koyier

Mercy Tabi Ojwang

Josiah Mutsogu

Kefa Murimi

Gerisson Mulongo

Copyright Notice

Supported By

AVU Multinational Project II funded by the African Development Bank.

Table of Contents

Course Overview

Welcome to Operating Systems Administration

Operating systems are central to computing activities. An operating system is a program that acts as an intermediary between a user of a computer, application software and the computer hardware. Two primary aims of an operating systems are to manage resources (e.g. CPU time, memory) and to control users and software.

This course introduce the students to the administration of operating systems.It deals with user/group management, filesystem and disk management, backup, recovery and server virtualization, security, performance monitoring, applications management, and automating tasks on operating systems.

Prerequisites

The course is designed to acquaint students with operating system administration. It is believed that students already took the course "Introduction to Operating System" and a good understanding on architecture as well as the different data structures and programming paradigms used in computer software that have an impact on the structure and performance of an operating system.

Materials

The materials required to complete this course are:

- A personal computer (PC) loaded with operating system software, like Windows Server 2012 or Linux/Ubuntu
- A browser software
- An Internet connection

Course Goals

Upon completion of this course the learner should be able to ·

- Create and manage users and groups
- Setup a filesystem and manage disk
- Take a backup and perform recovery, design server virtualization
- Secure operating system and perform performance monitoring
- Manage applications and automate tasks

Units

Unit 0: Pre-assessment

This unit gives a high level review of the operating system. It discusses the definition, services, requirements, environments, installation and configuration of operating system.

Unit 1: Users and Group Administration

This unit is emphasizes how user, password and group is managed in operating system.

Unit 2: Filesystem and DiskManagement

This unit covers the major filesystemused by the operating system. It also covers how updating and maintain filesystem is done. Moreover, how disk quotas for users is allocated.

Unit 3: Backup, Recovery and Server Virtualization

This unit elaborates and describes the importance of backup, different backup methods, backup devices, data recovery and server virtualization.

Unit 4: Securing OS and Monitoring

The major goal of this unit is to enlighten students on how they can secure the computer system. This unit covers types of attacks, different types of security, how to secure the authentication process, how to protect against network intrusion, and performance monitoring.

Unit 5: Special Topics

This unit addresses concepts associated with application on operating systems, server hardware knowledge, and automating routine tasks.Assessment

1	Continuous assessment		Weight (%)
	1.1	Test I	10
	1.2	Test II	15
	1.3	Home take assignment	10
	1.4	Laboratory Project	25
2	Final Examination		40
Total			100

Schedule

Unit	Activities	Estimated time
Pre-assessment	Activity 0.1 – Definition of Operating System	20 Hours
	Activity 0.2 – Services of Operating System	
	Activity 0.3 – Requirements of Operating System	
	Activity 0.4 – Operating System Environments	
	Activity 0.5 – Installation and Configuration of Operating System	
Users and Group Administration	Activity 1.1 – User Management in Linux	15 Hours
	Activity 1.2 – Password Management	
	Activity 1.3 – Group Management in Linux	
Filesystem and Disk Management	Activity 1 – Filesystem	25 Hours
	Activity 2 – Updating and Maintaining Filesystem	
	Activity 3 – Disk Quotas	
Backup, Recovery and Server Virtualization	Activity 1 – Overview and importance of backups	15 Hours
	Activity 2 – Backup Methods and Backup Strategy	
	Activity 3 – Backup Devices and Backup Media	
	Activity 4 – Data Recovery	
	Activity 5 – Server Virtualization	

Securing OS and Monitoring	Activity 1 – Types of Attacks	20 Hours
	Activity 2 – Types of Security	
	Activity 3 – Securing the Authentication Process	
	Activity 4 – Protection Against Network Intrusion	
	Activity 5 – Performance Monitoring	
Special Topics	Activity 1 – Application on Windows and Linux OS	15 Hours
	Activity 2 – Server Hardware Knowledge (rack mount and blades)	
	Activity 3 – Automating Routine Tasks Using Scripts	

Readings and Other Resources

The readings and other resources in this course are:

Unit 0

Required readings and other resources:

- Modern Operating System, Andrew S. Tanenbaum, Prentice-Hall, Inc., 3rd Edition, 2008, Chapter 1
- A Silberschatz, Peter B Galvin, G Gagne: Operating System
- Concepts, 7th edition, Chapter 1 and 2
- Optional readings and other resources for Evolution, Structure and Services of Operating Systems:
- William Stallings, Operating Systems, 4th edition, 2002

Unit 1

Required readings and other resources:

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2nd Edition, 2002, Chapter 5

- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 6
- http://www.control-escape.com/linux/users-groups.html, Users and Groups, February 24, 2006
- http://www.tldp.org/HOWTO/Shadow-Password-HOWTO-2.html, Why shadow your passwd file?, February 24, 2006
- http://www.w3resource.com/linux-system-administration/user-management.php, User Management, February 24, 2016

Optional readings and other resources for Evolution, Structure and Services of Operating Systems:

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2nd Edition, 2002, Chapter 5

Unit 2

Required readings and other resources:

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2ndEdition, 2002, Chapter 6
- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 5

Unit 3

Required readings and other resources:

- http://www.pcmag.com/article2/0,2817,899676,00.asp; Backup Methods and Rotation Schemes; BY MATTHEW D. SARREL; FEBRUARY 21, 2003
- www.tandbergdata.com/us ;Tape Rotation BEST PRACTICES The Tandberg Data SMB Guide to Backup Data Protection Strategies for the Small-to-Medium Size Business p.10 Basic Backup Guide

Unit 4

- A. Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2ndEdition, 2002, Chapter 15
- B. LPI – Linux System Administration, Revision 1.0, 2000, Chapter 9
- C. [Moore] Moore, Sonia Marie (Ed.). "MS Windows NT Workstation 4.0 Resource Guide." 1995. http://www.microsoft.com/technet/archive/ntwrkstn/reskit/default.mspx.

Unit 5

Required readings and other resources:

- http://www.psychocats.net/ubuntu/installingsoftware, Installing Software in Ubuntu, February 24, 2016
- http://windows.microsoft.com/en-us/windows/install-program#1TC=windows-7, Install a Program, February 24, 2016
- http://www.unixmen.com/removing-applications-using-terminal-ubuntu/Removing Applications Using Terminal in Ubuntu, Enock Seth Nyamador ,February 24, 2016
- https://www.digitalocean.com/community/tutorials/how-to-schedule-routine-tasks-with-cron-and-anacron-on-a-vps, How To Schedule Routine Tasks With Cron and Anacron on a VPS, Justin Ellingwood, February 24, 2016
- https://www.raspberrypi.org/documentation/linux/usage/cron.md, Scheduling Tasks with Cron, February 24, 2016
- Brian Hilton, Michael Welch, Server Primer : Understanding the current state of the industry, Golisano Institute for Sustainability

Unit 0: Pre-assessment

Unit Introduction

This unit introduces definition, functionalities, system requirements, and types of operating system. As you were introduced in the course "Introduction to Operating Systems", the operating system is a system software. What services does an operating system provide? What does the system requirements of an operating system? What types of operating system exist?

This unit addresses the above questions providing the definition, services, system requirements and types of operating systems as a whole. Focus on terminologies used, functionalities of an operating systems has and types of operating systems.

Unit Objectives

Upon completion of this unit you should be able to:

i. Demonstrate the functionalities of operating system
ii. Determine the system requirements of operating system
iii. Determine the type of operating system to be used

KEY TERMS

CLI: Command line interface an operating system provides for user interaction by tying commands

GUI: Graphical user interface an operating system provides for user interaction by selecting commands from the screen.

Activity 1: Definition of Operating Systems

Introduction

System software are set of software that manage and operate components of the computer system. One of the system software is operating system. What is operating system?

What is Operating System?

Operating systems are system software that run on top of computer hardware. Thisdefinition needs to be observed from different perspectives of computer system, namely from application software and user hardware interaction.

Operating system can be defined as, system program that controls execution of application programs. It is responsible for the loading and execution of application programs. It has to make sure the availability of the required hardware and software requirements before executing the application program.

Operating system can also be defined as system program that acts as an interface between the hardware and its users. The operating system, according to this definition, has a responsibility to hide the complexities of the underlying hardware for the end user. The end user is not supposed to know the details of hardware components like CPU, memory, disk drives, etc.

Conclusion

As we have discussed the definitions of operating system so far, we can conclude that the operating system is found between user, application software, and hardware of the computer system. It facilitates the execution of application programs and acts as a middle man between hardware and users.

Assessment

Essay Type Questions

Instruction

Write short answers for the following questions

What is an operating system?

Answers

Operating system can be defined as, system program that controls execution of application programs. It is responsible for the loading and execution of application programs.

Activity 2: Services of Operating Systems

Introduction

Lots of issues have been discussed so far to understand the definition of an operating system. This activity looks at different operating system services provided by an operating system to users, processes and other systems.

Services of Operating System

Operating system provides several services to users, programs as well as the whole machine. As system software, it acts as a bridge between users and the hardware in accepting user requests and interpreting them into machine executable commands. Processes directly consult the operating system for their various resource requests as well. Some of the basic services provided by an operating system are listed below.

- User Interface
  i. All operating systems have a user interface (UI) they provide that allows users to interact with the machine and access its resources.
  ii. The UI provided can be:

  Command Line Interface (CLI): allows direct command entry by fetching commands from user and executing it. it provides text based interaction for the user where by users are forced to remember the label of each command. The commands range from a built-in to simple program names.

  Graphical User Interfaces (GUI): easy to use, user oriented interfaces where users interact with the system through icons. These are user friendly desktop metaphor interfaces. Users pass their commands through keyboard, mouse and monitor.

- System and utility function calls

  These are programming interfaces to the services provided by the operating system. Programs access these calls with the use of an Application Program Interface (API) They are written and used with a high-level programming language like C++ When implemented, each system call is associated with a number that is used for indexing. The API invokes the required system call in the OS kernel and sends back the status of the system call along with any return value. The API hides any detail of the OS from the caller.

- System calls can be about:

a. **Process control**

Creating, terminating and waiting for a process

b. **File management**

Create, read, write and delete files

c. **Communication management**

Creating connections, sending message, receiving message, sharing memory

d. **Information maintenance**

Getting/setting time, system data, file and process attributes

e. **Device management**

Access, read, write devices

Most common APIs are: WIN32 for Windows and POSIX for UNIX, LINUX and MAC OS

The following table shows the system calls for different activities in POSIX and WIN32

Process management		
POSIX	WIN32	Description
fork	CreateProcess	Clone current process
exec(ve)		Replace current process
waidpid	WaitForSingleObject	Wait for a child to terminate
exit	ExitProcess	Terminate current process and return status
File Management		
open	CreateFile	Open a file & return descriptor
close	CloseHandle	Close an open file
read	ReadFile	Read from file to buffer
write	WriteFile	Write from buffer to file
lseek	SetFilePointer	Move file pointer

stat	GetFileAttributesEx	Get status info
Directory and File system management		
mkdir	CreateDirectory	Create new directory
rmdir	RemoveDirectory	Remove empty directory
link	(none)	Create a directory entry
unlink	DeleteFile	Remove a directory entry
mount	(none)	Mount a file system
umount	(none)	Unmount a file system

Program Execution

The OS is capable of loading a program onto main memory and run the program

It ends the execution of a program normally, if no errors are found, or teminates abnormally otherwise

Input/Output Operations

Running programs mostly require input and output tasks.

The input/output can be either from/to file or I/O devices

It is the task of the OS to facilitate such kinds of requests coming from programs

File Management

The OS handles reading and writing files and directories

It also manages creation, deletion and searching of items in the files and directories

Access Permission of files and directories is also the service of an OS

Error Detection

Errors might occur in a system in different areas including CPU, memory, I/O devices or even user programs

The OS ensures correct and consistent operation of the system despite the errors occurred by taking the appropriate measurement on the detected errors

To do so, the OS requires to continuously inspect the system and identify possible errors

Communication

Processes found on the same computer or over a networked computers might need to communicate among themselves

This communication can be carried out either through message passing or shared memory concept

The OS facilitates such communication among processes through implementing different communication management mechanisms

Resource Allocation

A computer system consists of several types of resources

Some of these resources such as the CPU and main memory need special allocation code while others like the I/O devices might need a general request and release code to be accessed

It is the job of the OS to identify the type of resource required and run the appropriate code to access the resource

Especially, in a multi-tasking and multi-user environment, the OS should be able to assign resources for each of the concurrently running jobs

Accounting

The OS also keeps track of the number of users using the system,

It also maintains which users are using which kinds of resources and their consumption rate of the resource as well

System Protection

The OS controls all access to a systems' resource

It also enforces authentication and different access control mechanisms to keep the system free of undesired access attempts

Conclusion

Operating system provides number of services like user interface, program execution, system protection, accounting, input/output management, file management, error handling, etc.

Before we install operating systems, what information do we need to know about the machine? And its existing software? The next activity answers these questions.

Instruction

Write the answers of the following questions.

List at least four services of the operating system?

Answers

The following are services of the operating system: user interface, program execution, system protection, accounting, input/output management, file management, error handling, etc.

Activity 3: Requirements of Operating System

Introduction

How do we identify the hardware and software requirements of OS? In this activity, we are going to address these question.

Requirements Identification

As a software, operating systems do have hardware and software requirements. Before you perform installation of operating systems. You have to consider the following basic requirements:

1. First, identify what kind of microprocessor you have. Is it 32-bit? Or 64-bit? Modern operating systems comes with both flavors. Some are only 64-bit OSs. Pick the operating system with the appropriate microprocessor word length.
2. Second, make sure the main memory satisfies the minimum requirement of the operating system. In order to run several other applications after the operating system is bootstrapped, the memory required for the operating system mentioned by the vendor is the minimum.
3. Next, identify the hard disk space requirement of the operating system and make sure it is available on the machine. Just like the memory requirement, the operating system states the minimum requirement of the hard disk space.
4. Last, make sure the device drivers (software) of the peripheral devices are available. These software include device drivers for display, audio, network, etc…for your specific machine. If the operating system fails to identify a peripheral device connected to your machine, you cannot make you of that specific peripheral. To make sure it operates correctly, install the appropriate device driver for that specific machine after installing the operating system. Usually, modern operating systems have a pack of device drivers for known hardware. But, you have to take a precaution on taking a backup of the device drivers.

Conclusion

Operating systems have software and hardware requirements. Before you made an installation make sure that these requirements are met for a specific operating system based on its specific version.

What kind of operating system environments do we have? How can one interact with them? The next activity addresses such and other questions.

Assessment: Essay Type question

Instruction

Write the answers of the following questions.

- What are the hardware and software requirements of Windows Server 2012 operating system?
- What are the hardware and software requirements of Ubuntu 14.01 system?
- Does your machine satisfy the requirements of the two operating systems?
- What is the hardware specification of your machine?

Activity 4: Operating System Environments

Introduction

Interaction between the operating system and the user can be a major factor to classify operating system environments. Operating system can have a Command Line Interface (CLI) or Graphical User Interface (GUI). This activity addresses the features of these two operating system environments.

Command Line Interface (CLI)

CLI is an interface users use to access the services provided by the operating system. Operating system that use CLI support a single user at a time using such interfaces. CLI was the primary interaction mechanism where users are supposed to type commands to tell the operating system what they want. CLI do have interface which is also known as shell that accepts user commands and convert it to appropriate operating system functions. Example of such CLI operating system includes MS DOS, Apple DOS, UNIX systems, CP/M and Open VMS. A shell is a command line interpreter that acts as a primary interface between a user sitting at his terminal and the operating system. It is a process (running program) that reads commands. It is not part of the operating system but it makes heavy use of many operating system features.

How the shell works?

- When any user logs in, a shell is started up
- The shell displays a prompt which tells the user that it is ready to accept a command
- When the user enters a command, the shell creates a child process that runs the program for the corresponding command
- When the child process is running, the shell waits for it to terminate
- When the child process finishes, the shell displays the prompt and tries to read the next command

While running the shell, you will be prompted the following interface:

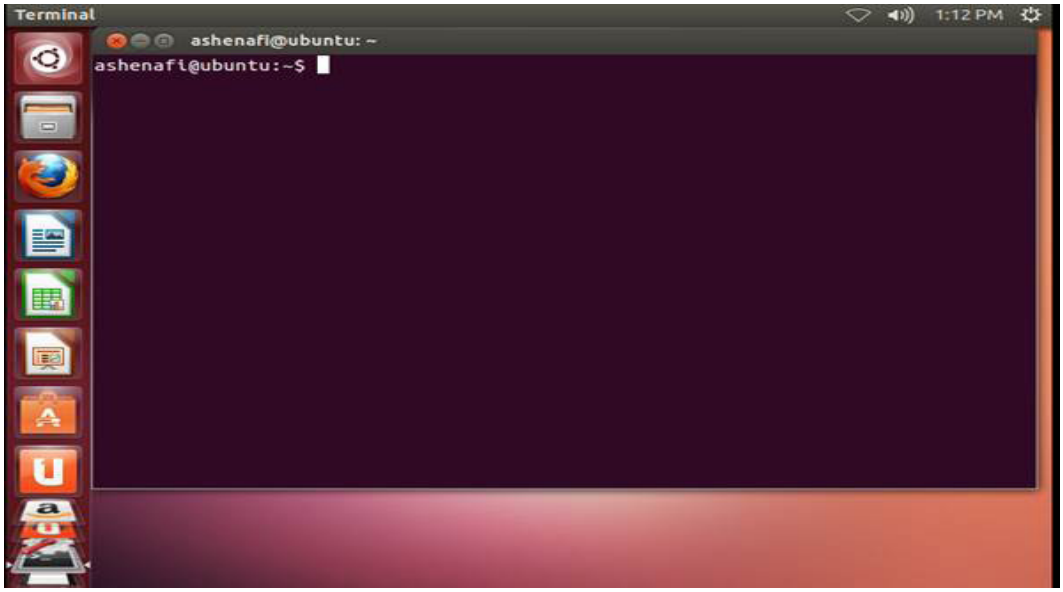

Figure 0.1: Shell of Ubuntu operating system

Figure 0.2: Shell of DOS operating system

Graphical User Interface (GUI)

GUI is an interface provided by the operating system for user interaction based on the graphics capability of the computer system. A GUI uses metaphors for objects familiar in real life, such as the desktop, the view through a window, or the physical layout in a building. GUI elements include such things as: windows, pull-down menus, buttons, scroll bars, icons, wizards, mouse, and more. Users do not need to remember the commands. Selection of commands does the job users are looking for. With the increasing use of multimedia as part of the GUI, sound, voice, motion video, and virtual reality interfaces seem likely to become part of the GUI for many applications. A system's GUI along with its input devices is sometimes referred to as its "look-and-feel".

Xerox Palo Alto Research Laboratory in the late 1970s introduces GUI. Apple used it in their first Macintosh computers. Later, Microsoft used many of the same ideas in their first version of the Windows operating system for IBM-compatible PCs. Now, almost all operating systems have GUI environment. Figure 0.3 and 0.4 show the GUI for Ubuntu and Windows operating systems.

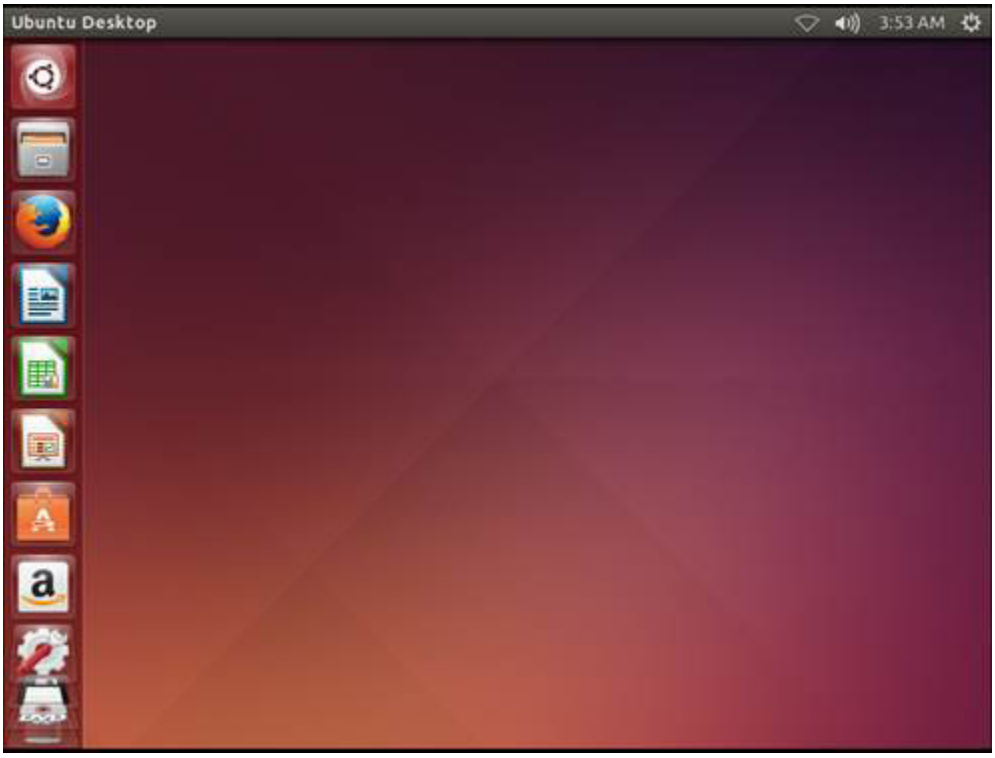

Figure 0.3: GUI of Ubuntu Operating System

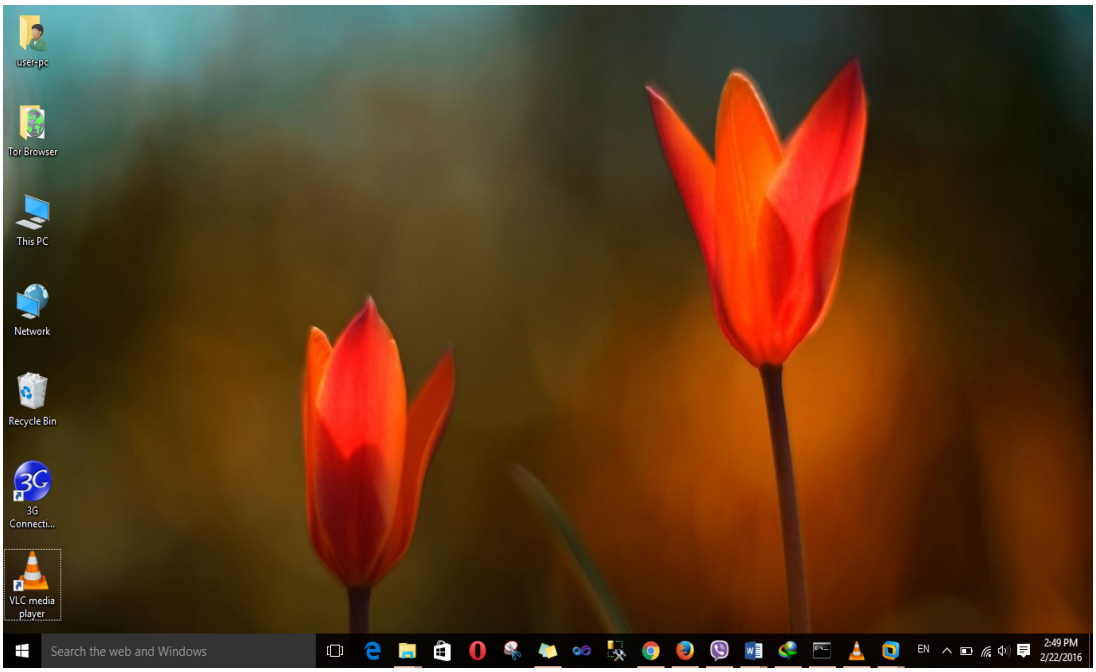

Figure 0.4: GUI of Windows OS

Conclusion

Operating systems have two ways of communication with the users: Command Line Interface (CLI) or Graphical User Interface (GUI). The earliest operating systems provide CLI for users to type in commands. Operating systems with GUI lets users to select their preference command than writing them.

Assessment: Essay Type Question

Instruction

Write your answers for the following questions

- How do we classify operating system environments based on their environment?
- List the GUI elements.
- What is the purpose of shell?

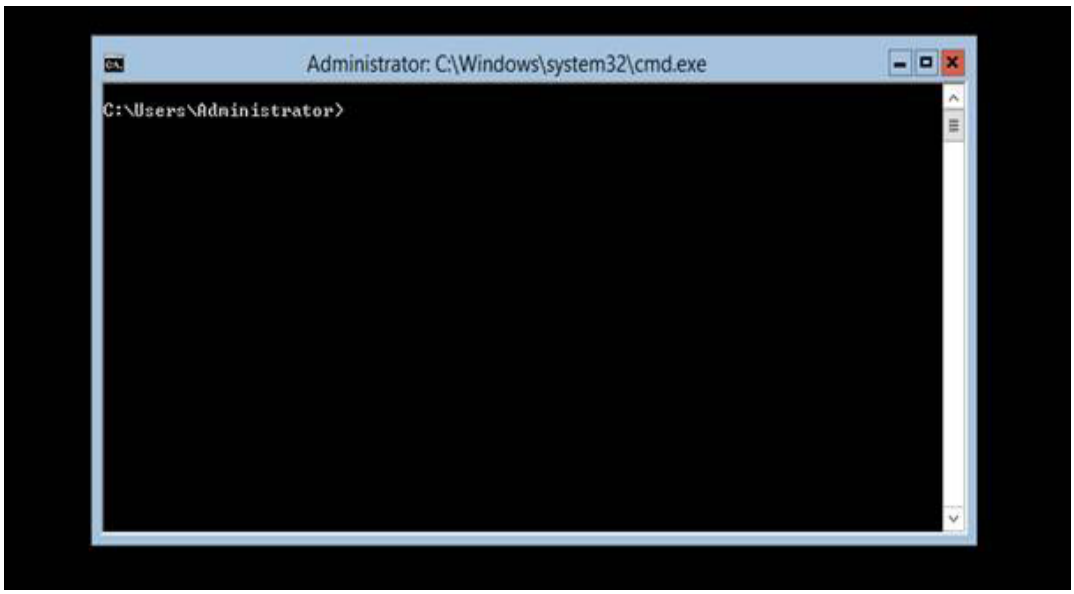

GUI: single dashboard server manager interface with icons. it has all graphical symbols including windows explorer, Internet explorer, control panel, etc.

There are also two other levels of installation known as features on demand and minimal server interface. The features on demand can be used for GUI or core installation by further reducing the footprint of the operating system through removing unused items from the operating system according to the purpose of the system. For instance, if the server is not needed to function a DNS (Domain Name Server), then this component is to be removed. The minimal server interface installation is like an intermediate between the GUI and the core installation in that it provides some components of the graphical user interface and some features of the core installation together.

virtualization strategy, if any, roles intended to be implemented on the server

Whether to add the server to production network or install it on a test network.

the 32-bit or 64-bit version

Windows server 2012 comes only as a 64 bit operating system. This decision had been made by Microsoft due to most major applications now being 64-bit and modern server configurations being supported on hardware that requires 64 bits

what setting will be set up in the server. Maximum hardware configurations

Minimum System/hardware requirements: though the actual requirements will vary based on the system's configuration and the applications and features installed, the

general requirements for the successful installation of windows server 2012 are:

- GHz 64-bit processor

- 512 MB RAM
- 32 GB available disk space
- DVD or USB flash drive
- Super VGA (800 x 600) or higher resolution monitor
- Keyboard and mouse (or other compatible pointing device)
- Internet access

The system requirement can be checked by running Microsoft planning and assessment toolkit also called MAPTOOLKIT. The tool can perform several important tasks for you including inventory, and virtualization.

Activity 5

6. Run MAPTOOLKIT 7.0 and run the inventory to check whether your machine meets the requirements to install Windows Server 2012
7. Set all requirements as reported on the inventory data to what is required by the server, if possible

Installation

Server installation can be either through upgrading or pure installation. Upgrade is possible only from Windows Server 2008 or 2008 R2 and 64 bits in both cases as Windows Server 2012 is also a 64 bit and supports an upgrade of same bit.

Pure installation

Before you install Windows Server 2012:

- Disconnect UPS devices
- Back up your servers
- Disable your virus protection software
- Configure Windows Firewall

Steps

- Insert the Windows Server 2012 DVD, and once you get the following message press Enter to boot from the setup.

- Wait for a while till the setup loads all necessary files (Depending on your machine, it will take couple of minutes)

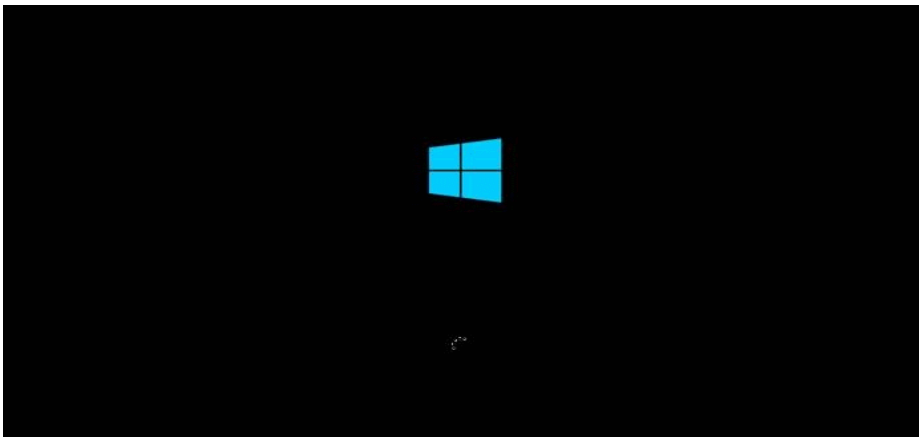

- Once the setup files are loaded, the setup will start with the following screen. You can change these to meet your needs but the default values should be fine for now

- Click Next to start installation and click install now on the coming window

- The installation options will be displayed once the setup started and you can select your preferred option. For now, select Windows Server 2012… (Server with GUI)

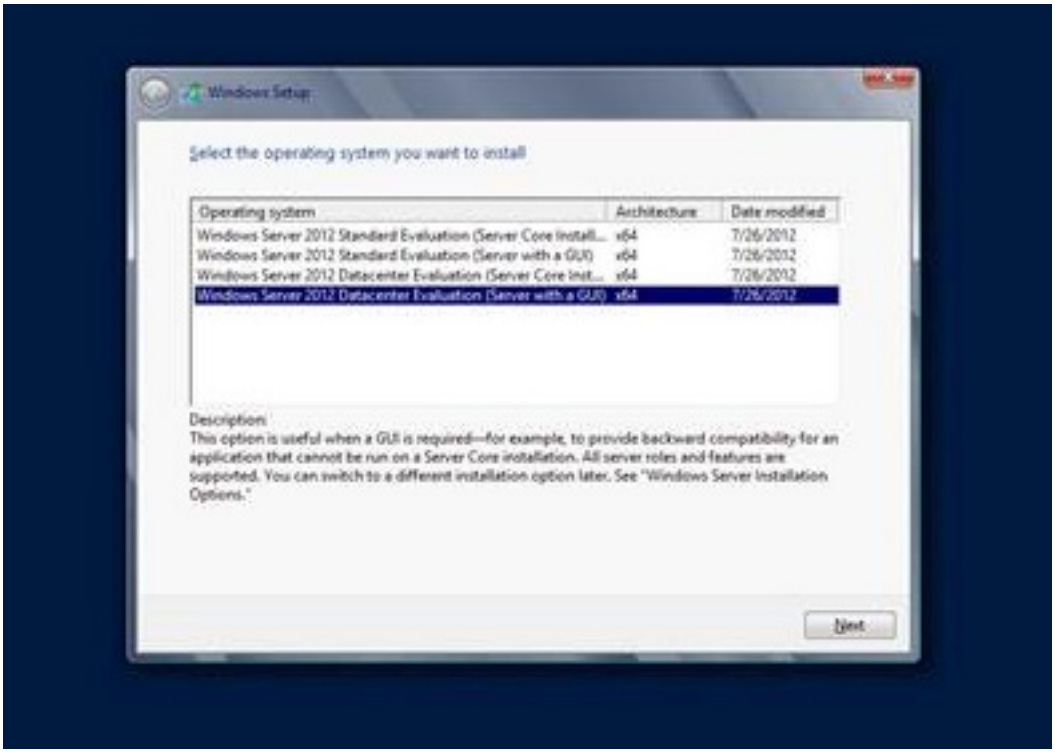

- In the next screen, Read the License terms, tick the "I accept the license terms" and click Next

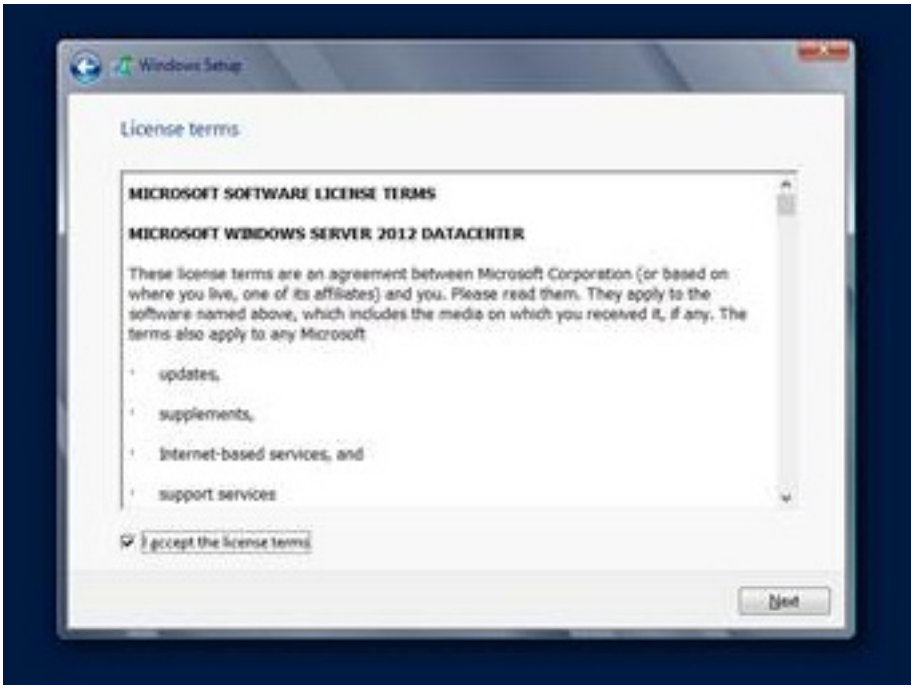

- Once you specify the drive or partition and click next on the above screen, setup will commence and takes a while to complete.

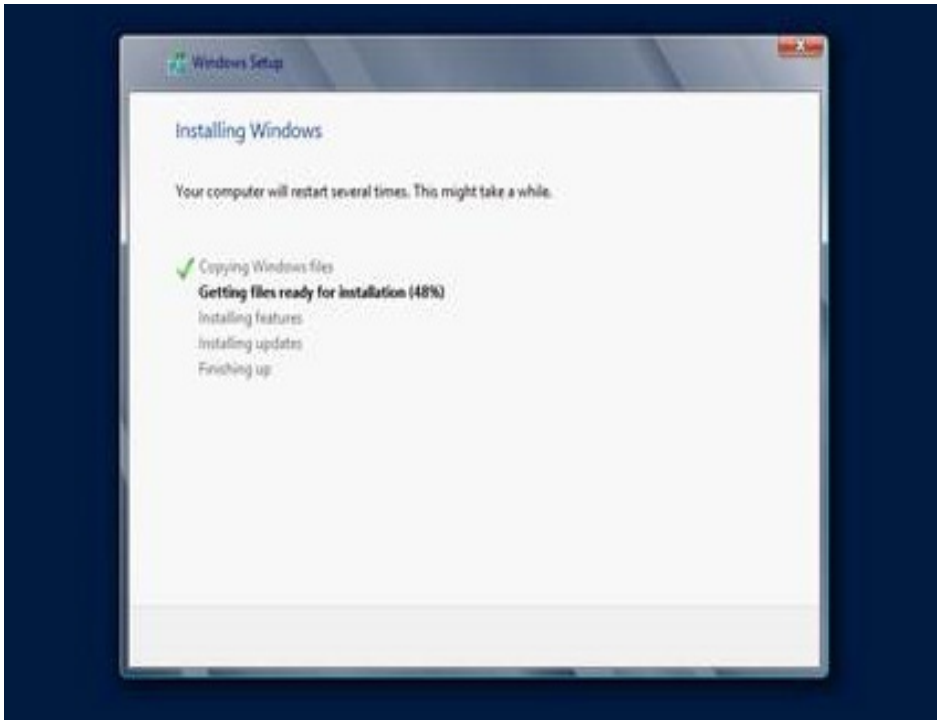

- On completion, the system will restart and the Windows Server 2012 installed will start running asking to setup password for the

Administrator user. this is a very important account an make sure to keep it strong and secure.

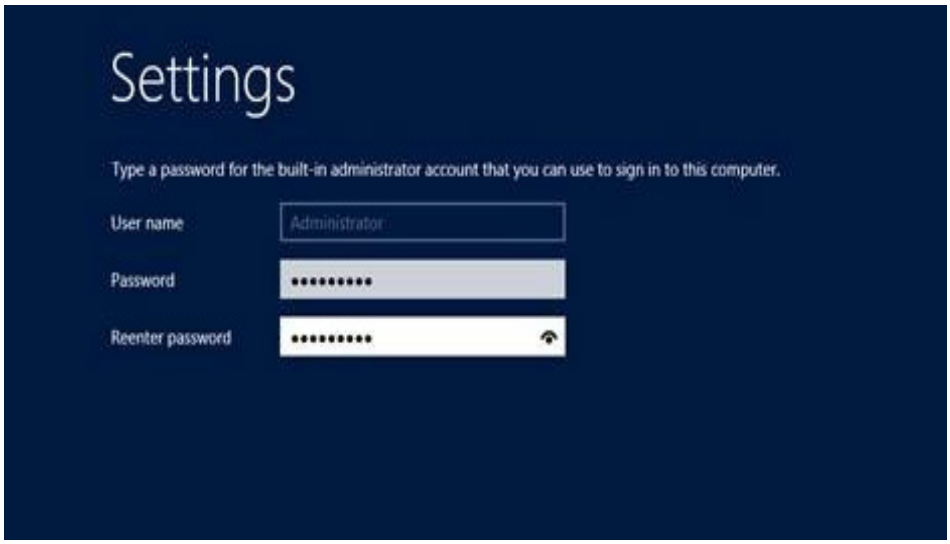

- Once the setup is done, the Administrator can login for the first time to the server with the password provided on the previous step

- You are now logged onto Windows Server 2012 with GUI

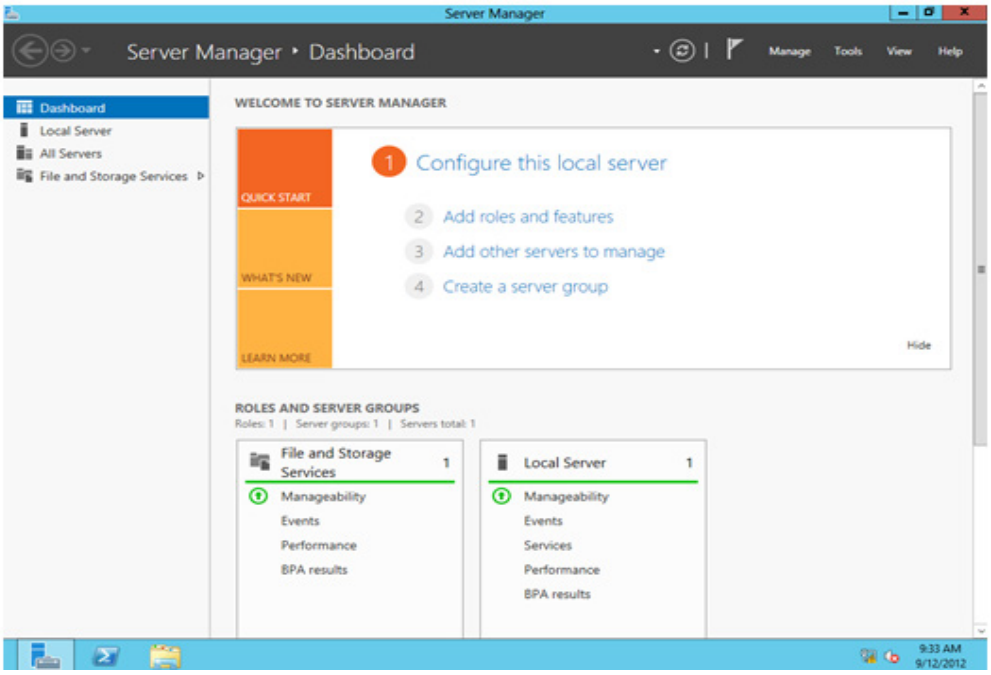

- The server manager page will be displayed by default once we are logged in. this is the place where we are going to do further administration tasks by configuring the server.

Configuration of the Server will be discussed in the coming chapters.

Note: the GUI of Windows Server 2012 is a bit different from the interfaces we are used to so it requires a bit exercise to get accustomed

Installing Ubuntu 14.04.1 LTS OS

Using a DVD

The easiest way to install Ubuntu is from a DVD. Here's what you need to do:

Put the Ubuntu DVD into the DVD-drive. Restart your computer. You should see a welcome screen prompting you to choose your language and giving you the option to install Ubuntu or try it from the DVD.

Using a USB drive?

Most modern computers can boot from USB. You should see a welcome screen prompting you to choose your language and giving you the option to install Ubuntu or try it from the CD, i.e., live CD.

If your computer doesn't automatically do so, you might need to press the F12 key to bring up the boot menu, but be careful not to hold it down - that can cause an error message.

Figure 6.32. Linux OS installation welcome page

Prepare to install Ubuntu

Plug your computer into a power source –recommended

You should also make sure you have enough space on your computer to install Ubuntu

We advise you to select Download updates while installing and Install this third-party software now

You should also stay connected to the Internet so you can get the latest updates while you install Ubuntu

If you're not connected to theInternet, we'll help you set up wireless at the next step

Figure 6.33. Screenshot captured during preparation for installation

Set up wireless

If you are not connected to the Internet, you will be asked to select a wireless network, if available. You are advised to connect during the installation so Ubuntu can ensure your machine is up to date. So, if you set up your wireless network at this point, it's worth then clicking the Back button to go back to the last screen (Preparing to install Ubuntu) and ticking the box marked 'Download updates while installing'.

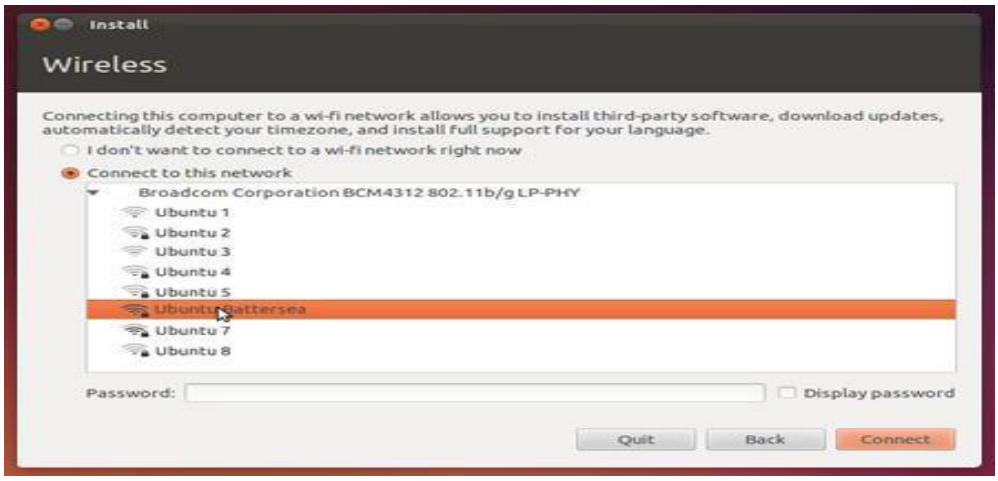

Figure 6.34. Connecting your system to a Wifi during installation

Allocate drive space

Use the checkboxes to choose whether you'd like to Install Ubuntu alongside another operating system, delete your existing operating system and replace it with Ubuntu, or — if you're an advanced user — choose the 'Something else'option

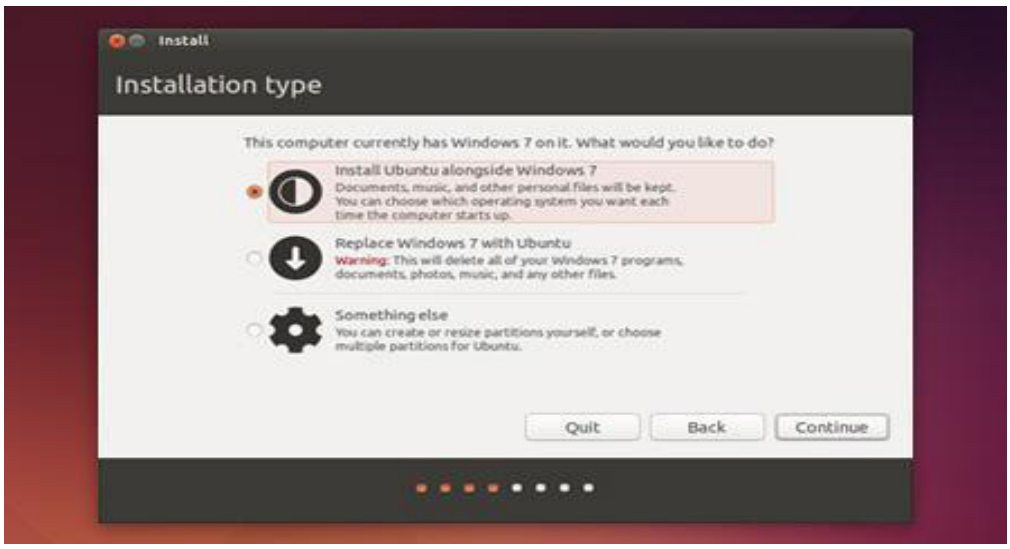

Figure 6.35. Select your preferred installation option for Ubuntu

Begin the installation

Depending on your previous selections, you can now verify that you have chosen the way in which you would like to install Ubuntu. The installation process will begin when you click the Install Now button.

Ubuntu needs about 4.5 GB to install, so add a few extra GB to allow for your files.

Figure 6.36. Allocating drive space for the Operating system to be installed

Select your location

If you are connected to the Internet, this should be done automatically. Check your location is correct and click'Forward' to proceed. If you're unsure of your time zone, type the name of the town you're in or click on the map and we'll help you find it.

TIP: If you're having problems connecting to the Internet, use the menu in the top-right-hand corner to select a network.

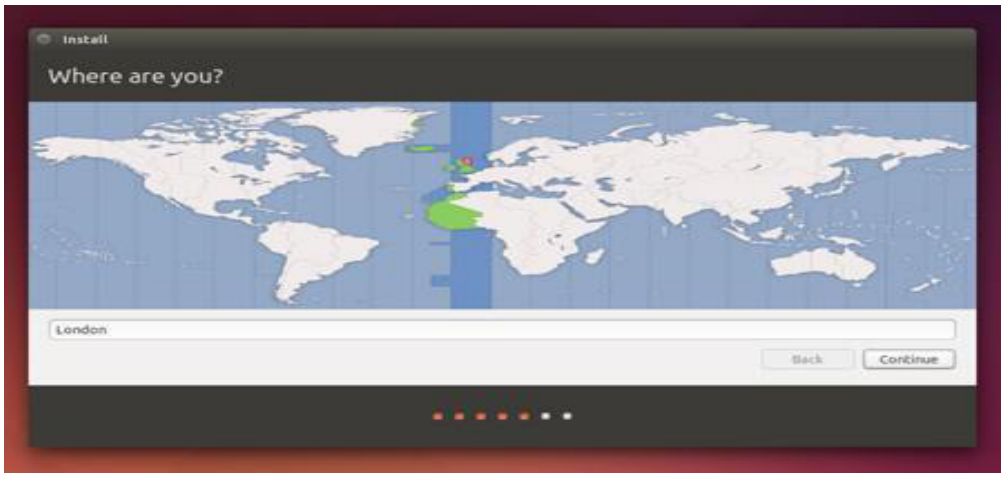

Figure 6.37. Screenshot showing your location selectionSelect your preferred keyboard layout

Click on the language option you need. If you're not sure, click the 'Detect Keyboard Layout' button for help

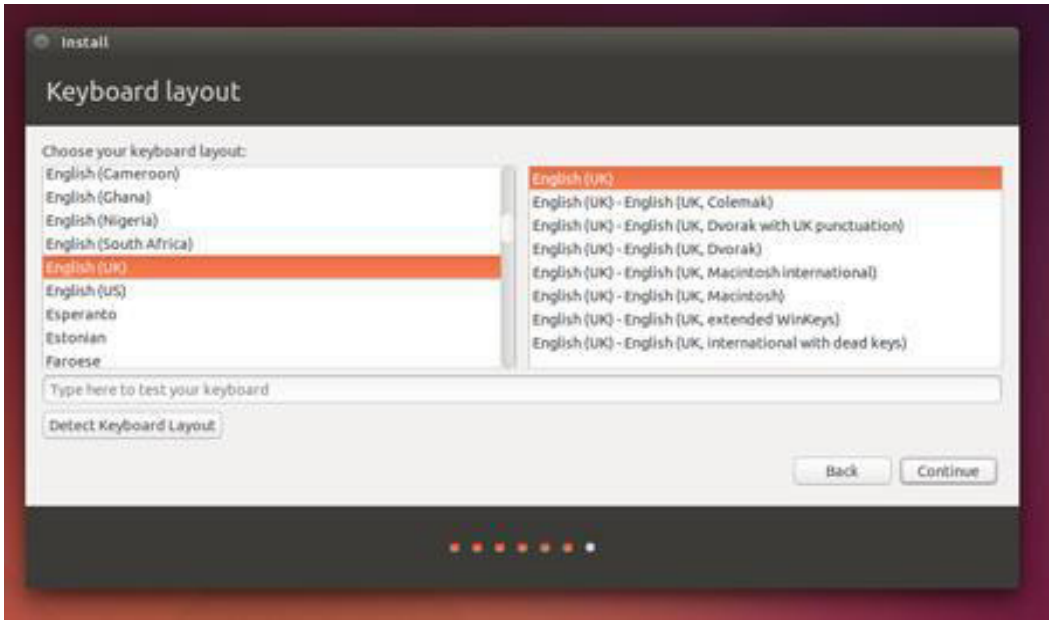

Figure 6.38. Menu to select preferred keyboard layout

Enter your login and password details

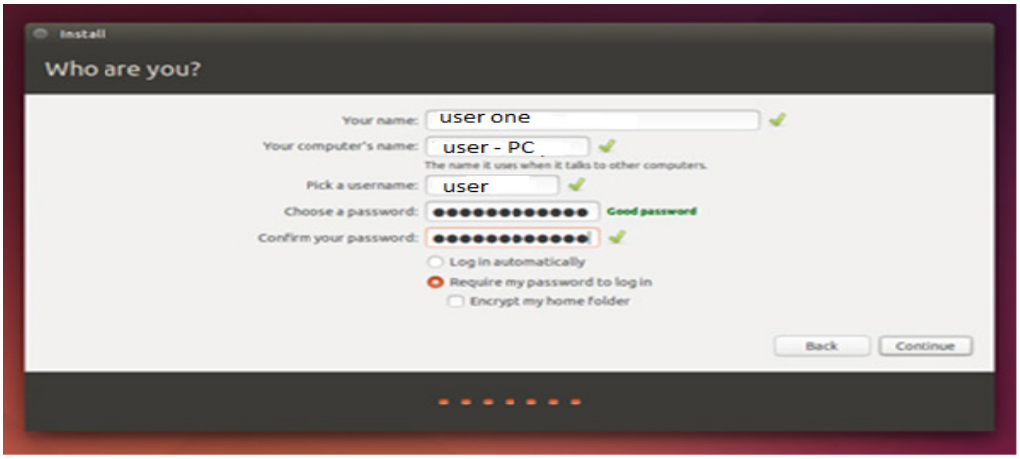

Figure 6.39. Menu to specify your credential information

Learn more about Ubuntu while the system installs…

Read the features of Ubuntu 14.04 while installing. A brief slideshow explaining about Ubuntu while the system installs. The setup will take something like 10-40 minutes, depending on your hardware and Internet connection, especially if you selected updates and third-party software earlier.

All that's left is to restart your computer and start using Ubuntu. Once the procedure complete, you can continue using Ubuntu in the live session, but no changes will be persistent. You will have to reboot and select your operating system in the boot menu.

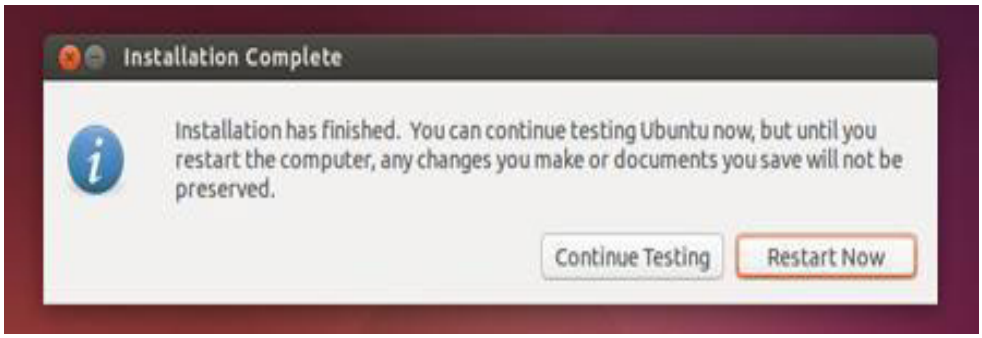

Figure 6.41. Message that pops up once installation has finishedConclusion

Installation of any operating system, especially of a server one, starts with good planning. The planning should not be of a complex task rather a simple procedure that helps to avoid any regrets that may arose after installation. Identifying system requirements, role of the server, type of server to be installed, level of installation are some of the basic planning activities done prior to installation. Windows Server 2012 comes with core or GUI installation the core being the default due to various security and efficiency reasons. The core installation requires us to know some basic commands to run on the power shell while the GUI will guide us through the installation as well as configuration tasks with menus and other graphical interfaces.

Upgrading of older Server Operating system is possible to Server 2012 only if the older is Windows Server 2008 or 2008 R2 and 64 bits.

Installation procedure of Ubuntu operating system for the server purpose takes the same procedure as installing Ubuntu for desktop purpose. The basic difference is made when we add services to the already installed desktop operating system. After installation of the Ubuntu operating system for the desktop, services like DHCP, FTP, Telnet, etc…are installed to make the operating system a server.

Assessment

1. Go through the GUI of Windows Server 2012
2. Practice some common tasks to be done through the interface like accessing start menu (charm), lock the computer, shutdown, etc.
3. What happen if we install Ubuntu operating system first and Windows operating system then on the same drive?
4. What is the importance of planning during server installation
5. Discuss some problems that may happen due to weak planning
6. Briefly discuss the different editions of Windows Server 2012

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

- Write your answer for the following questions
- What are the basic functionality of operating system?
- Why is it important to know the hardware and software requirements of operating system before installation?
- What kind of operating system does your machine use? What kind of environment does it have?
- What advantage does the accounting service of an Operating system has for users?
- What different kinds of Windows Server 2012 editions exist?
- How do we make Ubuntu a server operating system?

Grading Scheme

This unit's assessments worth a total of 20% from which

a. Activity Assessments: 5%

b. Unit Assessment I:3 %

c. Unit Assessment II (Lab Project): 12%

Answers

Rubrics for assessing activities

Assessment Criteria	Doesn't attempt both assessment or the answers are all wrong	Attempt both assessments and provide partial answers or partially correct answers	Attempt all assessments with full and correct answer

Grading Scales	Fail	Got half mark for each assessment	Score full mark for each assessment

Unit Readings and Other Resources

- Modern Operating System, Andrew S. Tanenbaum, Prentice-Hall, Inc., 3rd Edition, 2008, Chapter 1
- A Silberschatz, Peter B Galvin, G Gagne: Operating System
- Concepts, 7th edition, Chapter 1 and 2

Unit 1: Users and Group Administration

Unit Introduction

This unit deals about users and groups management. In a multi-user environment, management of users and groups is mandatory. Who is allowed to use the computer system with what level of permission? This is a question of user management. Creating, enabling, disabling, removing, etc… of users is possible. In case if there are two or more users that share similar permissions, it is wiser to create a group and add these users to the group. How can one do that? How do we managing users and groups?

Unit Objectives

Upon completion of this unit you should be able to:

- Create and manage user accounts
- Establish user level security
- Create and manage groups

KEY TERMS

User: an authorized person or process to use the computer system

Group: sets of users that can have a predefined permissions on system objects

accounts also identify individual users so that you have the ability to track what your users are doing. Setting up user accounts is one of the most visible jobs you'll have. Learning to do it efficiently will save you hours in the long run, and the confidence you'll exude from knowing it well will put you in good standing with your users.

Linux uses two or three files to maintain user and group information. The /etc/passwd file stores information about user accounts, and /etc/group stores information about groups. Linux systems also use a file called /etc/shadow to maintain passwords.

User Accounts

Linux puts a lot of power at your fingertips. That's the best reason to switch to Linux; it's also the most dangerous thing about the system. Linux controls how much power you can use on the computer based on your Login ID. It keeps a database of all users, and it keeps track of which user owns which files, and which users have permission to view, edit, and execute each file, folder or program. An ordinary user will not be able to do really dangerous things, like editing the user database, or deleting every file on the system. But right now you are logged in as root. You are not just an ordinary user, you are SuperUser. (SuperUser is a real Unix term, synonymous with root.) There are no restrictions on your power. You have the ability to crash the system and make it otherwise unusable in more ways than you can imagine. As a novice it is very easy to make your system completely unusable with a single erroneous command.

Because it is so dangerous to be logged in as root, you should never use this account unless you have to. The root account is meant to be used by the System Administrator to perform certain duties which can be destructive and therefore should only be performed by an expert. Some examples are emptying log files, mounting and unmounting file systems, installing or removing programs, and creating or deleting user accounts. If you are using Mandrake Linux, you will have a tool available to perform the most common administration tasks, even when logged in as a regular user. This is called the Mandrake Control Center, which you may find on your desktop or in the Configuration menu. It will ask you for the root password when you start it for security reasons. As a result of this handy tool, you may never need to actually log in as root.

Becoming SuperUser

No phone booth needed. The obvious way is simply to login as root. That may be the best way to do it if you plan on doing a bunch of system maintenance type stuff, but operating as root regularly is a bad idea, as you lose all the security protections that Linux provides. Logging in as the root user is generally discouraged and is in fact prohibited on some Linux systems by default. Fear not, there is a better way.

Try this:

```
[user]$ sudo ls
Password:*****
```

At the password prompt, type your password, not the root password. If it works, you will have just listed the current directory as superuser. If you got an error about not being in the sudoers file, see the section on configuring sudo.

Type this:

```
[user]$ su
Password:*****
[root]#
```

Just like that, you are SuperUser! A few cautions: Although you are now SuperUser, this is not a "login" shell, so your environment hasn't changed. The biggest way this will affect you is that some programs you normally run as root may appear to be missing. That's because your PATH environment variable, the list of places Linux looks for executables, does not contain /sbin or /usr/sbin. If you try to run a command like shutdown (see below) and it complains, try typing /sbin/shutdown instead. That should do it.

When you are finished with your maintenance tasks you should immediately change back to normal user mode:

```
[root]# exit

[user]$
```

Notice that while you are SuperUser, your command prompt looks different. An ordinary user is prompted with the dollar sign ($) while SuperUser gets a pound sign or hashmark (#). This makes it easy to tell which mode you are in. (This is true on most Linux distributions, but your prompts may be different, and they can be customized.)

The /etc/passwd File

The file /etc/passwd contains all information regarding the user (login, passwords, etc.). Only the superuser (root) must be able to change it. It is therefore necessary to change the rights of this file so that it can only be ready by the other users.

This file has a special format which makes it possible to mark each user, and each of its lines has the following format:

account_name : password : user_number : group_number : comment : directory : start_program

Seven fields are specified separated by the character :

- the account name of the user
- the password of the user (encoded, of course)
- the integer identifying the user for the operating system (UID=User ID, user identification)
- the integer identifying the group of the user (GID=Group ID, group identification)
- the comment in which the information on the user or simply its real name can be found
- the connection directory, which is directory which opens upon connection to the system
- the command is the one that is executed after connection to the system (often, this is the command interpreter)

Here is an example of a passwd file:

root:x:0:0:root:/root:/bin/bash
bin:x:1:1:bin:/bin:/bin/bash
daemon:x:2:2:daemon:/sbin:/bin/bash
news:x:9:13:News system:/etc/news:/bin/bash
uucp:x:10:14::/var/lib/uucp/taylor_config:/bin/bash
cquoi:x:500:100:Cool.....:/home/cquoi:/bin/bash

It is important to know that the passwords located in this file are encrypted. It is therefore useless to edit and replace the field password by directly typing the password, which would only cause the account to be blocked.

Once a user connects, the login program compares the password typed in by the user (after encrypting it) with the password stored in the passwd file. If they do not match, the connection can not be established.

To prohibit use, it is sufficient to replace the encrypted password by a star: "*".

Access to an account may be opened by leaving the field password open. Any person who wishes to connect via the account can then do so. To be able to modify the password of an account with the command passwd, you must either be the system administrator or the account owner (the system will than require that you enter the old password before asking you to enter the new password twice).

UID: (unique) identifier of each user account. Numbers between 0 and 999 are frequently reserved for the machine's own accounts. Numbers higher than 1000 are reserved for user accounts.

GID: group identifier. The default group (called group) has the number 50. This identifier is used in connection with access rights to the files. This question will not concern you if your system has more than one user group. (In that case, you must pay attention to the file/etc/group.

From the shell, it is possible to modify the command interpreter. To do so, use the commandchshorpasswd -s. Linux will look for program you have specified in the file /etc/shells. Only commands that are present in this file will be accepted and will replace the current value of the fieldstart_program. These restrictions do not apply to the superuser account. Make sure that the access rights to the file /etc/shells are the same as for the file /etc/passwdThe superuser may not necessarily be called root. To change this, just replace the root account name by the desired name. A privileged account is an account whose identifier (UID, User ID) is zero.

Shadowed Passwords

On a Linux system without the Shadow Suite installed, user information including passwords is stored in the /etc/passwd file. The password is stored in an encrypted format. If you ask a cryptography expert, however, he or she will tell you that the password is actually in an encoded rather than encrypted format because when using crypt(3), the text is set to null and the password is the key.

The algorithm used to encode the password field is technically referred to as a one way hash function. This is an algorithm that is easy to compute in one direction, but very difficult to calculate in the reverse direction. More about the actual algorithm used can be found in your crypt(3) manual page.

When a user picks or is assigned a password, it is encoded with a randomly generated value called the salt. This means that any particular password could be stored in 4096 different ways. The salt value is then stored with the encoded password.

When a user logs in and supplies a password, the salt is first retrieved from the stored encoded password. Then the supplied password is encoded with the salt value, and then compared with the encoded password. If there is a match, then the user is authenticated.

It is computationally difficult (but not impossible) to take a randomly encoded password and recover the original password. However, on any system with more than just a few users, at least some of the passwords will be common words (or simple variations of common words).

System crackers know all this, and will simply encrypt a dictionary of words and common passwords using all possible 4096 salt values. Then they will compare the encoded passwords in your /etc/passwd file with their database. Once they have found a match, they have the password for another account. This is referred to as a dictionary attack, and is one of the most common methods for gaining or expanding unauthorized access to a system.

If you think about it, an 8 character password encodes to 4096 * 13 character strings. So a dictionary of say 400,000 common words, names, passwords, and simple variations would easily fit on a 4GB hard drive. The attacker need only sort them, and then check for matches. Since a 4GB hard drive can be had for under $1000.00, this is well within the means of most system crackers.

Also, if a cracker obtains your /etc/passwd file first, they only need to encode the dictionary with the salt values actually contained in your /etc/passwd file. This method is usable by your average teenager with a couple of hundred spare Megabytes and a 486 class computer.

Even without lots of drive space, utilities like crack(1) can usually break at least a couple of passwords on a system with enough users (assuming the users of the system are allowed to pick their own passwords).

The /etc/passwd file also contains information like user ID's and group ID's that are used by many system programs. Therefore, the /etc/passwd file must remain world readable. If you were to change the/etc/passwd file so that nobody can read it, the first thing that you would notice is that the ls -l command now displays user ID's instead of names!

The Shadow Suite solves the problem by relocating the passwords to another file (usually /etc/shadow). The /etc/shadow file is set so that it cannot be read by just anyone. Only root will be able to read and write to the /etc/shadow file. Some programs (like xlock) don't need to be able to change passwords, they only need to be able to verify them. These programs can either be run suid root or you can set up a group shadow that is allowed read only access to the /etc/shadow file. Then the program can be run sgid shadow.

By moving the passwords to the /etc/shadow file, we are effectively keeping the attacker from having access to the encoded passwords with which to perform a dictionary attack.

The /etc/shadow file contains the following information:

username:passwd:last:may:must:warn:expire:disable:reserved

Where:

Username: The User Name

Passwd: The Encoded password

Last: Days since Jan 1, 1970 that password was last changed

May: Days before password may be changed

Must: Days after which password must be changed

Warn: Days before password is to expire that user is warned

Expire: Days after password expires that account is disabled

Disable: Days since Jan 1, 1970 that account is disabled

Reserved: A reserved field

An example might then be:

username:Npge08pfz4wuk:9479:0:10000::::

Adding and deleting users

To add a user:

sudo adduser username

```
user1@server:~$ sudo adduser newuser
Adding user `newuser' ...
Adding new group `newuser' (1003) ...
Adding new user `newuser' (1003) with group `newuser' ...
Creating home directory `/home/newuser' ...
Copying files from `/etc/skel' ...
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
Changing the user information for newuser
Enter the new value, or press ENTER for the default
        Full Name []: _
```

To delete a user (will retain the home directory):

Install the package:

sudo apt-get install unattended-upgrades

```
user1@server:~$ sudo apt-get install unattended-upgrades
Reading package lists... Done
Building dependency tree
Reading state information... Done
unattended-upgrades is already the newest version.
0 upgraded, 0 newly installed, 0 to remove and 2 not upgraded.
user1@server:~$ _
```

To disable a user:

sudo passwd -l username

```
user1@server:~$ sudo passwd -u newuser
passwd: password expiry information changed.
user1@server:~$ _
```

To enable a user:

sudo passwd -u username

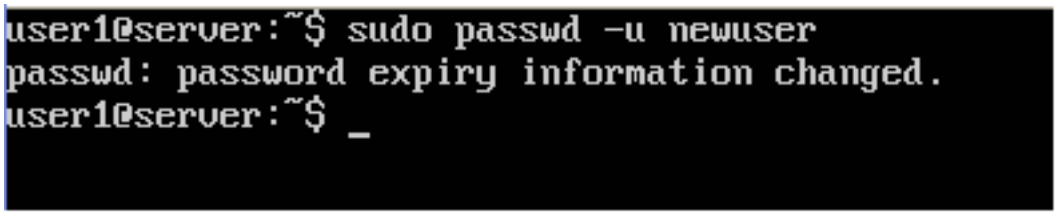

User profiles

By default, all new home directories can be accessed by everyone. You can enforce non-default access rights to new home directories by editing /etc/ adduser.conf:

Change “DIRMODE=0755”to “DIRMODE=0750”

The contents of this directory are modelled after the contents of /etc/skel.

Password policy

You can enforce strong user passwords by editing the password policy file /etc/pam.d/ common-password.

For example, if you want to enforce a password with minimum 8 characters and maximum 14 characters, edit the password line to look like this:

password required pam_unix.so nullok obscure min=8 max=14 md5

Password expiration

To see the password expiry value for a user, use the “chage” command:

The chage tool is for changing password expiration date.

sudo chage -l username

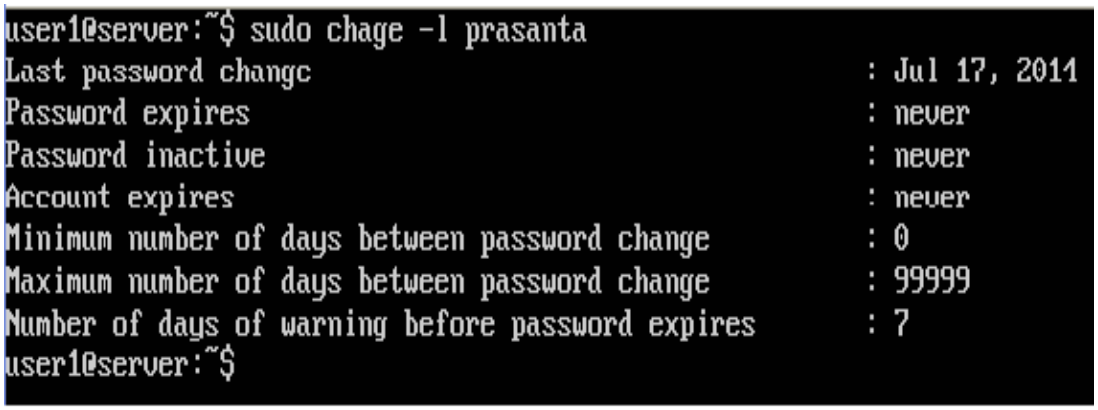

To make changes:

sudo chage username

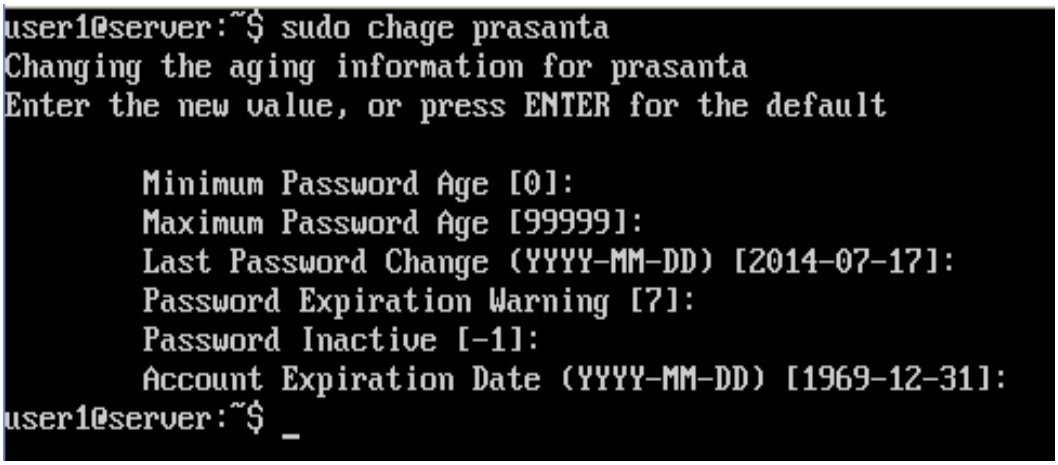

Disabling User Accounts

Deleting an account might result losing of user files and folders, rather it is better to deactivate the account. For the user whether it is deleted or disabled looks the same. It doesn't allow to login.

Manual Disabling

Disable an account is a way to make sure that the user's password has expired. To do this, modification of user entry on /etc/shadow file is possible. Dates in this /etc/shadow are displayed as the number of days since January 1, 1970. The third element in a record is the last modified date of the password, and the eighth field is the expiration date of the account. So you first want to change the user's password; the third field in the entry will then reflect today's date. Subtract one from that number, insert the new number immediately before the last colon, and save the file.

ashenafi:1ye1JK1VM$$MG9w14HdNwN9PF4dCEKrio:16788:0:99999:7:::

Subtracting 1 from 16788 yields 16787, so you'd change the entry to this:

ashenafi:1ye1JK1VM$$MG9w14HdNwN9PF4dCEKrio:16788:0:99999:7::16787:

chage: To Disabling an Account

One canuse the chage (change aging) command to modify the user's password expiration date toany previous daye. chage allows you to input this as the number of days since January 1, 1970, or in theYYYY−MM−DD format as follows:

 # chage −E 2016−2−23 ashenafi

If the date has already passed, the account will be disabled.And enabled later using the same procedure.

When the user attempts to log in, s/he will see the following message:

Your account has expired; please contact your system administrator

Conclusion

Authorization of user to use a system needs created first. In this activity we have seen how to add, modify, disable/block, and delete user accounts.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. Add a user called Patrick.
2. Add a user called Daniel specifying C shell.

Activity 2: Password Management

Introduction

Password for a user account is mandatory in Linux systems. How is it defined? What precautions shall a user take? We will deal with the issues related to password management in this activity.

Password

Users can set their passwords using the passwd command. Administrators usually force users to change passwords regularly for the sake of security. Usage of this tool is demonstrated in the previous activity. The only user who is able to change the password without providing the old password is the root. Otherwise, the old password must be provided to change it to a new one. NULL passwords are not allowed in Linux.

Choosing Passwords

Passwords are sensitive. You have to make sure passwords are not easily guessable. Much of the operating systems do have password policy that sets minimum-length of password, enforce non-dictionary passwords and the like.

Consider the following while assigning password for user accounts to make it more secured:

- Don't dictionary words or names.
- Use uppercase and lowercase letters and digits.
- Include special symbols.

Not only must a password be well chosen, it must be easily remembered. It is especially attention must be given when you choose the root password. You should also choose the password policies with care. The following list has few password policies to ensure the strength of the password:

- Maintain updated dictionaries of disallowed words.
- Enforce the use of non-alphanumeric.
- Do not ever have guest accounts (anonymous ftp may be an exception).
- Enforce a minimum password length.
- If someone has to use your system, give them an account with a password.

Conclusion

A user is identified if s/he provides the correct user name/password combination. It ensures security in the computer system. Passwords must be strong enough to be easily guessed. In this activity we addressed the mechanism how to make passwords strong.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. Add a password for user Patrick.
2. Force Patrick to change his password at next login.

Activity 3: Group Management in Linux

Introduction

Groups are an integral part of Linux security. When you create a user, automatically Linux creates a group by the same name. Upon login, the user assumes the group identity specified in the /etc/passwd file. In Linux, a user's default group ID, all users are by default put in a single group. As an administrators you are able to create groups, associating users for specific purposes such as projects that require certain users to have access to the same set of files.

The /etc/group File

The file /etc/group contains a list of the users who belong to the different groups. As a matter of fact, whenever a large number of users may have access to the system, they are frequently placed in different groups, each of which has it own access rights to the files and directories.

It has different fields that are separated by ":"

groupe_name : special_field : group_number : member1, member2

The special field is frequently blank. The group number is the number which makes the link between the /etc/group and /etc/passwdfiles.

Here is an example of a /etc/group file:

root:x:0:root
bin:x:1:root,bin,daemon
daemon:x:2:
tty:x:5:
disk:x:6:
lp:x:7:
wwwadmin:x:8:
kmem:x:9:
wheel:x:10:
mail:x:12:cyrus
news:x:13:news

- When the ls command is used with the option -l, the group number is displayed with the number of the user to whom the file (or the directory) belongs. This unique number corresponds to a unique group name (often 8 characters max.).
- The same user can appear in several groups. When he connects to the system he belongs to a group specified in the /etc/passwd (in the GID field). He can change this using the newgrpcommand. The file access rights are then defined.
- File protections must prevent the modification of files by non-privileged users.
- To add a group, the administrator can change the /etc/group file using a text editor. He can also use the addgroup or groupadd command (not always present). In the first instance, he will only have to add the line(s) relating to the groups. For example, the line:
- admin : : 56 : ccm

- To add a user to a group, just edit the file /etc/group and add the name at the end of the line by separating the names of the members by a comma.
- To delete a group, edit the /etc/group file and delete the corresponding line. Please note, do not forget to change the the numbers (GID) of the deleted group in the /etc/passwd file, if users belonged to it. It is also important to search the files and directories of this group to change this (otherwise, the files and directories may become inaccessible).

Creating and deleting groups

To create a group:

sudo addgroup groupname

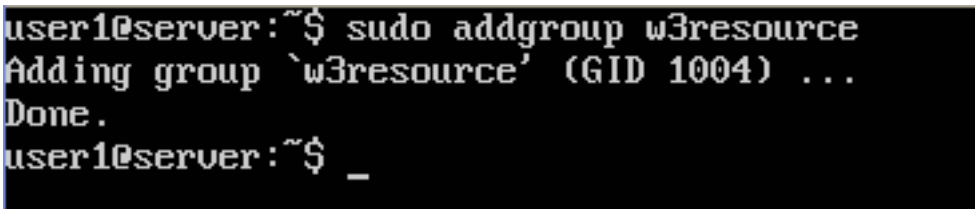

Creating and deleting groups

To create a group:

sudo addgroup groupname

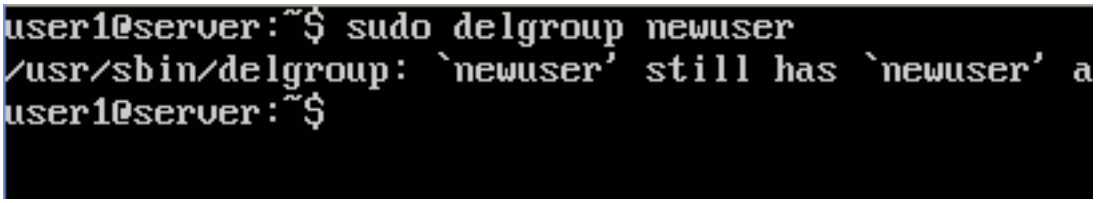

Just as users can be added manually, with a command–line utility, groups can also be created in any of these ways. The different methods are explained below so that you can choose the method you prefer. (Once you've created a new group, you presumably want to add users to that group.

Changing Group Membership

A user has one primary group it belongs but several supplementary group. A user may only act as a member of one group at a time. For instance, a user might need to create files that carry a certain group membership, so that other members of the proper group can read those files. Every user has a main group association, specified in the user's /etc/passwd file entry. Upon logging in, the user plays as a member of that group. The newgrp command provides a means for users to temporarily change the group associated with their actions. To use newgrp, type the command and the new group name. For instance, to become a member of the project1 group, type:

$ newgrp – project1

If the user ashenafi has used newgrp to join the group project1, when he creates a new file, such as a project document, it will show up in a long listing (ls –l) with group ownership by the new group:

–rw–r––r–– 1 ashenafi project1 10332 Feb 23 16:07 proj_doc

Modifying Groups

You can modify groups using the groupmod command. You can modify the group name, the group ID number, and group members. You might need to change a group name or ID. Modifying group members is a redundant task of system administration. All appropriate system files, like /etc/group, are modified. The syntax for groupmod is similar to that of groupadd:

groupmod [–g GID [–o]] [–n group_name] group

–g GID: The numeric value of the group's ID. This value must be unique. As with the groupadd command, group IDs must be nonnegative and should not be between 0 and 999 unless the group is a systemgroup.

–n group–name: The new name of the group.

group: The old name of the group.

Conclusion

In a multiuser environment having too much users with the same need of access may arise. Managing every user will be very difficult and unlikely. Groups are used in such kind of environment. Group may have several users with the same privileges on particular object. In this activity, we covered how to create, modify, and delete groups.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. Add a group called avuproject1.
2. Add a user called Fekade with UID of 1100 to the group avuprojet1.

UNIT SUMMARY

As a multiuser OS, Linux relies upon user accounts and groups to maintain security and keep the system usable to all its users. Over the years, manual methods, text–based tools, and GUI tool shave been developed to manage these accounts and groups. However you do it, though, it's important that you understand how Linux handles its accounts and groups—the relationship between usernames and UIDs or group names and GIDs; where passwords are stored; and how users are assigned to groups. Understanding these topics will allow you to effectively manage your user base, whether it's just one or two people or thousands. These subjects are also critically important for securing the files and processes your system uses. We discuss disk management in the next unit.

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

Write your answer for the following questions

1. Add a user called Tessema using /home/Tessema as the home directory.
2. Modify Tessema to use the Bash shell.
3. Modify Tessema to give him new UID 1112.
4. Set the expiration date of Tessema to 28 Feb 2016.
5. Lock Tessema's account.
6. Unlock Tessemas's account.
7. As root, use chage –l to show the status of Tessema's password protection.

   Change the password aging to:

   Maximum number of days: 7

   Minimum number of days: 2

   Warning number: 7

   Log in as tessema and try to change the password.
8. Create a new group called wamp. Modify tessema to be a member of group wamp; do not modify tessema's default group.

9. In one command, add a new user nesredien with full name Nesredien Sulieman, userID 319, and membership in the supplementary group wamp (leave the user's default group as its default value). Set a password for nesredien and log out.

Grading Scheme

This unit's assessments worth a total of 20% from which

- Activity Assessments: 5%
- Unit Assessment I:3 %
- Unit Assessment II (Lab Project): 12%

Answers

Rubrics for assessing activities

Assessment Criteria	Doesn't attempt both assessment or the answers are all wrong	Attempt both assessments and provide partial answers or partially correct answers	Attempt all assessments with full and correct answer
Grading Scales	Fail	Got half mark for each assessment	Score full mark for each assessment

Unit Readings and Other Resources

Required readings and other resources:

- Linux System Administration, Vicki Stanfield, Roderick W. Smith,Sybex., 2nd Edition, 2002, Chapter 5
- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 6
- http://www.control-escape.com/linux/users-groups.html, Users and Groups, February 24, 2006
- http://www.tldp.org/HOWTO/Shadow-Password-HOWTO-2.html, Why shadow your passwd file?, February 24, 2006

- http://www.w3resource.com/linux-system-administration/user-management.php, User Management, February 24, 2016

Optional readings and other resources for Evolution, Structure and Services of Operating Systems:

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2nd Edition, 2002, Chapter 5
- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 6

Unit 2: File System and Disk Management

Unit Introduction

This unit deals about file systems, disk management and quotas. One of Linux's strengths as an operating system is its support for a wide variety of filesystems. Not only does Linux support its own native filesystem, the Extended Filesystem (ext2fs or ext2, ext3, and ext4), but it also supports a wide variety of filesystems used by other Unix family OSs, Windows, and others. Since, Linux supports these filesystem which makes it effective OS in a multi−boot environment. One can use Linux's extensive filesystem support to read removable media and even hard disks from any computers.

As a normal user or system administrator, the low−level data structures that organize a filesystem inits first meaning are largely invisible to you—but their effects are important nonetheless, as demonstrated throughout this unit.

Accessing of a filesystem is possible if you accustomed to the tools used to manage the file system. You must distinguish how Linux mounts and unmounts filesystems—that is, how the OS incorporates a removable disk or hard disk partition onto its existing directory structure.

Linux file systems implement the disk quota mechanism. Users can be allocated disk quotas on specific file systems and can be restricted by number of disk blocks and/or number of inodes. How is disk quotas are assigned for users?

Unit Objectives

Upon completion of this unit you should be able to:

- Demonstrate the different file system Linux use and support
- Manage file system of Linux
- Set, manage, and view disk quotas

KEY TERMS

File system: a disk− or network−based data structure used for storing filesor hierarchy of directories and files on a disk

Disk quota: amount of disk space allocated for a user

Activity 1: File System

Introduction

The file system defines how disks are structured. Individual disks are hidden from the casual user. All access to a file is through the directory structure, and file permissions are used to control user access to the system. Different disks can be formatted with different file systems and must be mounted into the directory structure to be accessed. A disk or partition can be mounted into any directory, which is then referred to as the mount point for that disk.

The Linux file system is an inverted tree structure with a single, top-level directory called the root directory (/). All file and device access in a Linux system is through the file system structure. Different disks (fixed and removable) and partitions on the same disk are mapped into the directory hierarchy using the mount command.

The superblock contains information about the file system installed on the disk (or disk partition). This includes a list of the first few free blocks on the disk and of the first few free inodes. These lists are sufficiently large to meet the immediate needs of the operating system. Whenever these lists run down, the system will automatically replenish them by scanning the free block list on the disk or the inode list as required. Inodes define the files in the file system. Each file has one allocated inode containing all of the file information except the file's name. The filename is stored in the directory containing the file. The inode contains a list of the first ten blocks allocated to the file. Subsequent block numbers are kept in a separate block pointed to by the inode. This is called the indirect block list. To handle larger files, the inode also has a double-indirect block pointer and a triple-indirect block pointer.

File System Types

File systems have a type associated with them. The types of file systems supported by a given Linux implementation are compiled into the kernel, allowing the builder to configure the file systems that will actually be used. This saves memory space by only including code for file systems that are needed. System manufacturers can add their own proprietary file system types into the kernel. In practice, most systems only implement the most common file systems.

A file system type defines how the disk data is structured. Linux supports many different file system types (typically twelve or more). The ext2fs and ext3 are standards across all Linux distributions. All Linux systems can include support for Microsoft's FAT, VFAT, FAT32, and NTFS for use with DOS/Windows systems as well as OS/2 and NFS.

Making a File System

The mkfs command is used to create a file system on a disk or partition. Making a file system is similar to formatting, as shown in the following example:

1. Format a hard disk as an ext2 disk:

```
# mkfs -t ext2 /dev/hda5 102400
```

2. Format a flash memory as an ext2 disk:

```
# mke2fs /dev/sda1
```

The examples use -t to override the file system default type. The command requires theraw disk device name (/dev/hda5, /dev/fd0). Additional parameters will vary dependingon the file system type. The command creates a superblock and an inode list. The inodelist is fixed in size and cannot be extended, but you can override the number of inodesinitially created.

If you want to change the number of inodes, you must know the physical size of the disk(in physical blocks) and convert this to the number of logical blocks. Typically, a diskblock is 512 bytes and a logical block is 1,024 blocks, so a 64-MB disk actually has64-KB logical blocks (65,536) and 128-KB physical blocks (131,072). Use the defaultnumber of inodes unless you expect to create a lot of very small files or a few very bigones. A larger block size (2,048) will give better disk performance but will waste morespace due to partially filled blocks at the end of files. Smaller block sizes (512) use lessspace, but disk performance is worse.

You do not have to make the file system the same size as the disk. Any space not allocatedto a file system can be allocated to another, separate file system or can be left empty forlater use.Once a file system has been made, it cannot be altered without destroying all of the dataon the file system. Increasing the number of inodes requires backing up the entire disk,making a new file system with more inodes, and restoring the backed-up data.

Mounting a File System

A Linux file system can only be accessed if it is mounted on the file system hierarchy (tree structure). The directory used to mount the file system is called the mount point. Any directory can act as a mount point for any type of file system, and mounted file systemscan contain other mounted file systems. Since mounting a file system will hide the previous contents of the mount point directory, it is customary to choose an empty directory for the mount point. The /mnt directory is provided as a temporary mount point.

The mount command is used to add any disk containing a file system to the tree structure.

```
# mount -t ext2 /dev/hda2 /home
# ls /home
lost+found
```

File systems that are no longer needed can be removed from the root file structure using the umount command. File systems with open files or containing other mount points cannot be unmounted. The fuser command can be used to determine which processes (and users) are using the file system and can be used to kill these processes if necessary.

```
# umount /home

# ls /home
```

As long as the kernel understands the format (type) of the file system, it can be mounted onto the hierarchy. An MS-DOS file system mounted in this way is accessed as though it was a normal Linux file system. Linux uses file system type msdos (or vfat for Windows). The mount mechanism is also the way network disk systems, such as NFS or RFS, are accessed.

The mount command, without any parameters, lists the mounted file systems. The df command provides a different view, showing percentages used and available. The mount command maintains the /etc/mtab file with information concerning the currently mounted files. Corrupting this file will upset the mount and umount commands.

File System Configuration Files

The /etc/fstab file is maintained by the system administrator and is used to define the standard mountable disks on the local system. The system startup and the commands use the information in this file to mount file systems automatically.

Entries in /etc/fstab consist of one line per file system, with space- or tab-separated fields of the following form:

Block device	mountpoint	Fstype	mount options		dump	fsck order

```
# cat /etc/fstab
```

/dev/hda5	/	ext2	defaults	1	1	
/dev/hda1	/boot	ext2	defaults	1	2	
/dev/hda6	swap	swap	defaults	0	0	
/dev/fd0	/mnt/floppy	ext2	noauto		0	0
/dev/cdrom	/mnt/cdrom	iso9660	noauto,ro		0	0
none	/proc	proc	defaults	0	0	
/dev/hdb5	/daj	ext2	defaults	0	3	

Use a hyphen for fields that do not require the following values: no raw device, no fsck pass number, or no mount options. Lines starting with a # symbol are ignored and can be used for comments.

The fields in /etc/fstab are listed in the following table.

Field	Example
block device	/dev/hda5
mount point	/
file system type	ext2
mount options	auto, rw
dump order	1 (*yes*)
fsck order	1 (*first*)

Free Disk Space

Use the df command (disk free) to report on free disk space. Older versions would report size in 512-byte blocks, with a -k switch to give size in kilobytes (now the default); a –m switch gives size in megabytes. Use the -t switch to get a report on partitions of a specific file system type. The df command only reports on mounted file systems.

When first sizing disks for a new system, df is invaluable for working out file system sizes; on a stable running system, df can be used to monitor file system space usage. Run df as an hourly cron job and filter the output (in a shell script) to look for file systems getting within 10% or 5% of capacity. Some Linux systems provide administration utilities that do this automatically.

```
# df -k
```

File system	1024-blocks	Used	Available	Capacity	Mounted on
/dev/hda5	642009	446269	162579	73%	/
/dev/hda1	16554	593151	06	4%	/boot
/dev/hdb5	1189119	1311	27670	0%	/daj

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

a. Create an ext2 file system on a 200-MB disk hda5.

b. Mount this file system on /usr.

Activity 2: Updating and Maintaining Filesystems

Introduction

If you have a new disk, you must break itinto partitions and create filesystems on those partitions. Only then will you be able to mount thefilesystems. These steps are necessary when adding a disk and when replacing one, but theprecise details of these operations differ. There's also the issue of filesystem maintenance. In somesituations—particularly after a system crash or power failure—Linux must check its filesystems forintegrity. You may need to supervise this process, so it's important to understand what goes onduring such a check.

Disk Partitioning

If you've bought a new disk, your first task once you've connected it to your computer is to partition the disk. This procedure carves the disk into smaller chunks so that you can share the disk acrossmultiple OSs, or subdivide the space used on a single OS to protect files on one partition should another develop problems. You can also use assorted partitioning tools after installation, to change your configuration or to add a new disk.

Tip:

If you want to change an existing partition configuration, one of the best tools available is PartitionMagic, from PowerQuest (http://www.powerquest.com/). This commercial package allows you to add, delete, move, resize, and copy FAT, HPFS, NTFS, ext2fs, and Linux swap partitions, without damaging their contents. Theopen source GNU Parted (http://www.gnu.org/software/parted/parted.html) provides some of PartitionMagic's functionality. ReiserFS, XFS, ext2fs, andext3fs all include their own partition–resizing utilities, but they aren't as easy to use or as flexible as PartitionMagic.

Linux Disk–Partitioning Software

Linux's main partitioning tool is called fdisk (for fixed disk). It's named after the DOS FDISK utility but works quite differently. To use fdisk, type its name followed by the device file you want to modify, such as /dev/sda or /dev/hdb, thus:

```
# fdisk /dev/hdb
```

Warning:

Every x86 OS has its own disk–partitioning software. Linux's fdisk is unusually flexible, and so can produce partitions that other OSs don't like. As a general rule of thumb, you should use each OS's partitioning tools to create its own partitions.

Alternatively, you can use a more OS–neutral tool, such as PartitionMagic, to do the job for all OSs. On modern disks, you'll likely be told that the number of cylinders exceeds 1024. You can safely ignore this warning. Once fdisk is running, you see only a prompt that reads Command (m for help):.You type single–character commands at this prompt in order to accomplish various tasks. You can type m or ? to see what these commands are.

It's generally a good idea to start any fdisk session with a p command to display the current contents of the disk. This allows you to verify that you're modifying the correct disk, and gives you the partition numbers for partitions you might want to delete. You also need this information in planning where to put new partitions.

Once you've seen what (if anything) already exists on the disk, you can proceed to delete, add, and otherwise modify the partition table using fdisk's commands. Consider the following fdisk exchange:

Command (m for help): p

Disk /dev/hdb: 255 heads, 63 sectors, 1216 cylinders

Units = cylinders of 16065 * 512 bytes

Device Boot	Start	End	Blocks	Id	System
/dev/hdb1		257	1216	7711200	5 Extended
/dev/hdb2		1	192	1542208+	fb Unknown
/dev/hdb3		193	256	514080 17	Hidden HPFS/NTFS
/dev/hdb5		257	516	2088418+	6 FAT16
/dev/hdb6		517	717	1614501	7 HPFS/NTFS

Command (m for help): n

Command action

l logical (5 or over)

p primary partition (1–4)

l

First cylinder (718–1216, default 718): 718

Last cylinder or +size or +sizeM or +sizeK (718–1216, default 1216): +400M

In this situation, the initial configuration included five partitions, and the n command added a new one. fdisk gave the option of creating a logical or primary partition. The x86 partitioning scheme originally provided for only four partitions per disk, which soon became inadequate. The workaround was to use one of the original four primary partitions as a placeholder for a potentially large number of logical partitions. The "placeholder" primary partition is then referred to as an extended partition. In Linux, the primary partitions use numbers from 1 to 4; the logical partitions are numbered 5 and up. Linux doesn't care whether its partitions are primary or logical, so we recommend using mostly or exclusively logical partitions for Linux. This reserves primary partitions for OSs that do need them, such as DOS, Windows, and FreeBSD.

Linux's fdisk lets you specify partition sizes either in terms of an ending cylinder number or in bytes, kilobytes, or megabytes. The preceding example specified a 400MB partition starting at cylinder 718.

Tip:

For the final partition on a disk, enter the size by specifying an ending cylinder number that corresponds to the maximum available. This practice minimizes the amount of unused disk space.

By default, fdisk creates partitions that use the type code 0x83 (Linux native). Such partitions are suitable for holding Linux's ext2fs or any of the journaling filesystems available for Linux. If you want to create a Linux swap partition or a partition to be used in another OS, however, you must change its type code. You do this with the t command, which prompts you for a hexadecimal code. If you don't know the code, type L at this point for a list. (You can enter a code that's not on the list if you like, but fdisk won't be able to identify the associated OS if you do so.) You can use this feature to convert a partition created with another tool for use by Linux.

When you're done editing the partition table, look it over with the p command; then verify that everything's okay with the v command. Chances are that v will report your disk has some number of unallocated sectors. This is normal and reflects sectors lost to the standard PC method of accessing the disk. You should write down the partition numbers and your intended uses for them, so that you don't forget these details. Once you're satisfied with your new partitioning scheme, type w to commit the changes to disk and exit.

Conclusion

Filesystem creation and maintenance involves several tools, including the fdisk tool for partition creation, mkfs and its helper programs for filesystem creation, and fsck and its helper programs for filesystem integrity checking. Understanding how to use these tools is critically important for

Upgrading and maintaining your system. Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

- Why is it important to partition disk?
- Why is it recommended to use third party softwares like Partition Magic than fdisk tool?

Activity 3: Disk Quotas

Introduction

Linux file systems implement the disk quota mechanism. Users can be allocated disk quotas on specific file systems and can be restricted by number of disk blocks and/or number of inodes.

Disk Quota Management

Quotas have two kinds of limits: hard and soft. The soft limit can be exceeded for a certain grace period; failure to reduce the allocated space before the timer expires will trigger the quota. Hard limits apply immediately. Users exceeding their quota on a disk will be advised that they have run out of space. This is often confused with a full disk, as the system does not differentiate between the two error conditions.

For the user in question, the disk is full. The disk quota mechanism can be switched on and off, as required, once it is enabled for a file system and can be applied by user or group. Users not given disk quotas are not included in the system. It is advisable to apply disk quotas to some users but not to others. Especially note that any files created by a quota user when quota is not switched on will not be counted by the quota system.

To set up quotas on a file system, first create the quota files (one or the other, or both) atthat file system's root level:

```
# touch /home/quota.user
# touch /home/quota.group
```

Edit a user or group's quota properties with edquota. (Adjust the EDITOR systemenvironment variable as necessary to enable your preferred editor.)

```
# edquota amsale
```

Turn on quotas for a file system with the quotaon command:

```
# quotaon -v /home
```

Quotas can be turned on for all systems with quotas enabled using the -a switch:

```
# quotaon -a
```

Turn off quotas for a file system with the quotaoff command:

```
# quotaoff /home
```

A user can examine his or her own quota status using quota, but only root can look at the quotas of others. Unless a user is over quota, the quota command will return no output.For a more detailed output of quota information, a user would run quota -v.

```
$ quota -v
```

Disk quotas for amsale (uid 1002):

Filesystem	usage	quota	limit	grace	files	quota	limit	grace
/usr/homeA	0	5120	6144		0	0	0	
/usr/homeB	4612	5120	6144		305	0	0	

For each file system where a quota has been defined, the current user will receive a line of output with the following fields:

Filesystem quotas	Mountpoint of the file system with
Usage	Amount of blocks used
Quota	Number of blocks allowed (soft)
Limit	Blocks allowed (hard)
Grace	Applicable only if over quota
Files	Current number of files used
Quota	Number of blocks allowed (soft)
Limit	Blocks allowed (hard)
Grace	Applicable only if over quota

Conclusion

In a multiuser environment, disk quotas are required on file systems. Disk quotas limit the user to use excess amount of disk than allocated for him/her. This activity showed how disk quotas are managed in Linux environment.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. Modify the following line in /etc/fstab to allow for user and group quotas:

   /dev/hda1 /home ext2 defaults 1 2

2. Examine the following example output from edquota -u tempuser: Quotas for user tempuser:

   /dev/hda2:

   blocks in use: 6502, limits (soft = 8000, hard = 10000)

   inodes in use: 814, limits (soft = 2000, hard = 2500)

Write the command to duplicate this quota for every user in the tools group.

UNIT SUMMARY

The file system defines how disks are structured. Individual disks are hidden from the casual user. All access to a file is through the directory structure, and file permissions are used to control user access to the system. Different disks can be formatted with different file systems and must be mounted into the directory structure to be accessed. A disk or partition can be mounted into any directory, which is then referred to as the mount point for that disk.

Disk quota mechanism should be implemented in systems where several users are available. Linux implements disk quota to limit and manage disk usage of users. This unit addressed disk quota and file system management in Linux.

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

Write your answer for the following questions

1. Create an ext2 file system on a 150-MB disk hda6.
2. Mount this file system on /usr/lib. Create this directory if it does not already exist.
3. Create a ext3 file system on 100-MB disk slice hdb2 and mount on / home.
4. Unmount all three file systems; notice that you must unmount /usr/lib before you can unmount /usr.

Grading Scheme

This unit's assessments worth a total of 20% from which

a. Activity Assessments: 5%

b. Unit Assessment I:3 %

c. Unit Assessment II (Lab Project): 12%

Answers.

Rubrics for assessing activities

Assessment Criteria	Doesn't attempt both assessment or the answers are all wrong	Attempt both assessments and provide partial answers or partially correct answers	Attempt all assessments with full and correct answer
Grading Scales	Fail	Got half mark for each assessment	Score full mark for each assessment

Unit Readings and Other Resources

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2ndEdition, 2002, Chapter 6
- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 5

Unit 3: Backup, Recovery and Server Virtualization

Backup and Recovery Systems and methodologies

- Overview and importance of backup
- Backup methods and backup strategy
- Backup devices and media
- Data Recovery methods
- Server Virtualization

Introduction

1. What will you do if something went wrong in the system you administer?
2. How disastrous can a data loss be?
3. What sorts of disaster might strike your valuable data?
4. How can a system rebuild itself and give non-interruptable service to its users?
5. Is it possible to run several services on a single physical system?

Today, most businesses are being conducted through electronic communications and companies are becoming more and more dependent on IT to carry out their day to day activities. More and more corporations connect their computer systems to communication systems, such as LANs, and the Internet, thereby increasing the vulnerability of their data to attacks, among which only few corporations are aware of the data loss or damage risks and invest in backup and recovery techniques. Recovery techniques are means by which a system restores to its normal state after a disruptive event occurred. The event might be something huge, like a hard disk or memory failure, or something small, like malfunctioning software caused by a computer virus. Having a well-planned backup and disaster recovery technique plays a vital role in the sustainability of a business and assures uninterrupted service provision by the organization. In this unit, we will discuss the importance of a backup plan, backup methods and techniques, devices to use for backup, backup timing and also advantage of recovery and techniques to use.

Virtualization is another important issue to be addressed in this unit. We will mainly discuss server virtualization and its importance.

Unit Objectives

Upon completion of this unit you should be able to:

- Define backups and the importance of backups
- Identify the fundamentals of backup
- Discuss backup strategies and scheduling
- Explain the basics of backup devices and media
- Define system recovery
- Define virtualization
- Explain server virtualization
- Discuss three approaches to server virtualization

KEY TERMS

Backup: is the process of transferring data from your company's primary computer system to a separate storage device, such as a tape drive

Disaster: A situation which makes a system loss or damage its resource

Recovery: A process of reviving from a disaster

Archive:Data maintained permanently

Dump: Copy or Duplicate of data

Virtualization: an abstraction of computer resources that can be accessed in a consistent way before and after abstraction through virtualization

Activity 1. Overview and importance of backups

Introduction

There are instances where everyone who has used a computer system has experienced losing important data in some kind of accident. Many reasons can be identified why data get lost: it might be due to carelessness or natural disaster. The purpose of a backup is to make a copy of data, which is unlikely to be lost or destroyed by the same act as the original and a crucial function of system administration is to backup file systems. Backups safeguard against data loss, damage, or corruption. In this activity we will begin by defining what a backup is, how it works and why it is required

Overview

A system can encounter critical data loss disaster through various ways. Human error is among the leading causes of data loss. For example, mistakes as simple as typing "rm *" at the UNIX command prompt can have disastrous results. In addition, software failure, hardware failure, computer viruses, ill-minded employees, and natural disasters such as fire and flood can cause a catastrophic system failure.

Backup is duplicate information of a system to restore it to its normal functionality in case of failure or data loss disaster. This backup is used when the original copy is lost or cannot be retrieved anymore. Taking backups of all data seems like the obvious solution to overcome the data loss, though many organizations gave very low priority to it and lack well planned and adequate backup practices due to deficiency of knowledge as well as budget. It is very difficult and in fact sometimes impossible to function normally during a crisis. Thus, it is important to think about data backups before a disaster strikes. The backup copy is usually retained over a period of time and can be placed on any media such as disk or tape.

Some of the importance of maintaining backup includes:

- Recover data in the event of system failure
- Allow users to recover accidentally (or deliberately) corrupt or deleted files
- Used to retain data while systems are upgraded
- Backups are also used for transferring data between non-networked machines

Backups are crucial to provide an alternate storage media for important data.When to perform backups is an important question to be addressed by system administrators. Ideally, we need to backup our system as often as possible or even daily. However, in order to answer this question, one needs to consider the importance of the data to be backedup along with its rate of change,availabilty of resource on which to backup and also the required frequency of the backup which leads to a backup plan. Having a backup plan helps to specify all tasks and activities needs to be done before, during and after a backup. In reality, all computer systems fail now and then causing loss or corruption of some data. Having a regular and efficientbackup plan enables the system administrator to recover the system data to the state it was in when the last backup took place. Unless backup data is used to recover a system, it is as useless as not doing a backup.

- **Which Files Should Be Backed Up?**

Should we backup OS Binaries, Applications, Configuration files or User and log files? In general, backing up everything a system has is useful and also easy to manage. However, backing up system files everyday has no relevance, as it does not change frequently, except wasting of time and resource. Backup of only user files is not also enough to recover the system after a disaster happens. Thus, those resources which changes frequently including log files as well as configuration information needs to be maintained as backup.

- **How frequent should Backups Be Performed?**

In general, to assure backups maintain as much of the system as possible, the gap between backups should be kept minimal. Also, the system administrator shall be able to determine the level of acceptance of data loss and delay tolerance of the business to determine the backup frequency.

- **Where should a backup be stored?**

Backups are certainly necessary, but where you keep your data is also important decision .There are two backup options we can choose from. These are onsite and offsite backup options.

Onsite backup is a process of taking backups and storing important data on a periodic basis on local storage devices, such as hard drives, DVDs, magnetic tapes, or CDs. This method gives us immediate access to data, it is less expensive and more it does not require any Internet connection to be accessed. But the main limitation of this backup option is its equal vulnerability for disasters and we may end up losing both the original data as well as backed up data if a catastrophic situation happens. Moreover, the backup media can also be stolen.

Offsite backup is a process of storing important data on a remote server, usually via the Internet, although it can also be done via direct access. This is a method of using external backup services which can be set to automatically backup specific areas of your computer as often as you please. The advantages of these method are:

- Possibility of accessing backup data from any location, via Internet or FTP
- High degree of data preservation even in a catastrophic damages
- Sharing of backup data among different remote locations

In general, which backup option to choose depends on one's personal preference. However, Backups shall be maintained in a different building or room to that of the system in order to secure the backup. It is also advised to use a lockable fireproof safe for the backup devices to avoid any kind of loss or corruption vulnerability.

- **When Should Backups Take Place?**

Ideally, backup should be taken after the regular business hours of an organization is over or when employee demands on the network are at a minimum. This off-hours time frame is called the "backup window." However, for a very demanding business that runs 24x7, the backup window might be minimum, in which case, the system administrator shall identify the time during which the normal execution of the business is least affected or devise a means by which copying the system data quickly can be achieved which can be later backed up to a backup media without affecting the users' of the system.

Conclusion

One of the important task expected from a system administrator is carrying out backups as often as possible. Backup is a process of keeping data redundantly to recover a loss in case a disastrous incident happens on the actual system. A well planned and efficient backup plan helps the administrator to identify which files to backup, how frequent to backup, where to keep the backup and also when to perform the backup procedure?

Assessment

- Define the term backup
- What advantages does a backup bring to a system?

Activity 2. Backup Methods and Backup Strategy

Introduction

- What kinds of backup techniques shall we use to take a backup?
- Is it possible to reuse a backup media for another backup?

In this activity, we will try to address these and more questions by going through the various backup schemes and media rotation algorithms.

Backup Methods

Clearly, making backups of every file is an infeasible activity as it requires lots of storage and also time-consuming process. There are two kinds of backup

1. Full backup: copies every file, the system files, the software files, and the data files on a source medium to a backup medium every time a backup is taken regardless of any change made to the data itself. This has an advantage of restoring everything to its previous state if a problem occurs. But it is very resource intensive.
2. Incremental backup: is a partial backup method which includes only those files that have changed since the last backup operation of any kind
3. Differential backup: this is a special type of incremental backup, often referred as cumulative backup scheme and includes all files changed since the last full backup, whether they have been changed since the last backup operation or not

Incremental or differential backup defines different levels of backup and copies files according to the level of the backup. A level 0 backup copies everything.

A level 1 backup copies everything, which has changed since the last level 0 dump. A level 2 backup copies everything which has changed since the last level 1 dump or level 0 dump and so on. (i.e. all levels which are lower). This method saves a great deal of time, storage and backup effort though with a restore and recovery overhead as all backup tapes till the system crashes will be required to restore the system during a problem.

Three important factors have to be considered while choosing one of these backup methods. These are:

- Capacity or size of the media to be used for backup
- The period of time or window available for your backup
- The level of urgency experienced when a file restoration is necessary

Backup Strategy

After the appropriate backup method is selected, we also need to have a backup strategy which picks the most appropriate rotation scheme for our organization needs, which can reduce media costs and extend the longevity of our tapes while ensuring that every file is protected. A backup strategy is a planned approach to data protection and data recovery. There are several backup algorithms to choose from and put the data on the backup media among which Volume/Calendar, Grandfather/Father/Son, and Tower of Hanoi are the most commonly used backup algorithms.

1. Volume/Calendar Backup

This strategy defines to take a full system backup once a month, incremental backups for files that change often once a week and daily incremental backups of files that have changed since the last daily backup. For instance, a system administrator may perform level 0 backup one Sunday a month, weekly incremental backups all the remaining Sundays' of the month and daily incremental backups Monday through Saturday. This would count the total backup media to eight, one monthly media, one weekly media and six daily medias. Table 4.1 below depicts the backup scheme defined by the Volume/Calendar backup strategy

Table 1. The Volume/Calendar Backup Algorithm

Date	Sun	Mon	Tue	Wed	Thur	Fri	Sat
Week 1 Tape	A	B	C	D	E	F	G

Backup Level	0	5	5	5	5	5	5
Week 2,3,4 Tapes	H						
Backup Level	3	5	5	5	5	5	5

When a complete data loss happens and restore is required in a system that uses a volume/ calendar backup algorithm,

- first restore from the most recent full backup,
- Then restore from the most recent weekly backup,
- Finally, recover from each daily backup since the weekly backup

An advantage of this algorithm is the minimum media it requires for the backup as backup media are immediately reused for the next backup which also brings a major problem to the algorithm. In such cases, every daily backup overwrites last week's same day backup (every Monday backup overwrites last Monday's backup). So what would happen if one of the backup media fails during the second backup date? If such a scenario happens, the system won't be able to recover all its data since the media failed while the overwriting was being performed.

2. Grandfather/Father/Son(GFS) Backup

This strategy is similar to the volume/calendar strategy except GFS incorporates a one-month archive in the backup scheme which overcomes the problem of overwriting a tape before completing a more recent backup of the file system.GFS is knows as one of the simplest, most effective , most commonly used rotation methods for making and keeping backup copies of your data. It has three basic sets of backups being the most common setup. The Grandfather backups are the monthly full backups (Level 0 Dump) that are kept for, say, 12 months before rotating on a first-in-first-out (FIFO) basis. In between the monthly backups are weekly incremental (Level 3 Dump) backups (Father) that could be retained for 48 weeks before rotating. To save duplicating backups one Father backup per month is generally promoted to Grandfather status, rather than generating a new full backup. The Son backups are daily incremental (Level 5 Dump) backups that hang off the Father for that week.

The following table shows how GFS backup algorithm is used.

Table 2. the GFS Backup Algorithm

Date	Sun	Mon	Tue	Wed	Thur	Fri	Sat
Week 1 Tape	A	B	C	D	E	F	G
Backup Level	0	5	5	5	5	5	3
Week 2Tape	H						I
Backup Level	5						3
Week 3 Tape	H						J
Backup Level	5						3
Week 4 Tape	H						K
Backup Level	5						3
Week 5 Tape	L						M
Backup Level	0						3
Week 6 Tape	H						G
Backup Level	5						3
Week 7 Tape	H						I
Backup Level	5						3
Week 8 Tape	H						J

Backup Level	5						3
Week 9 Tape	A						
Backup Level	0						

In this scheme,

- The monthly full backup is placed into a permanent storage to provide a one-month archive
- Each weekly backup should also be placed in storage.
- The second monthly full backup, should use new media.
- When the third monthly backup is due, the first month's full backup media should be reused.
- The weekly backups are archived in a similar manner.

Overall, this algorithm sets a requirement of two monthly backup media sets, five weekly backup media sets, and six daily backup media sets requiring a total of 13 sets of media to implement it with a one-month archive of information.

During recovery from complete data loss,

- first restore the most recent level 0 backup tape
- Next, restore from the most recent of the level 3 backups, if that backup was written after the level 0 backup.
- When the level 3 backup has been restored, the operator would restore from each of the level 5 backups written after the level 3 backup.

The advantage of this scheme is the increased data survivability it provides than the volume/calendar scheme. It also gives great granularity for the most recent backups which is a desired property during recovery. It is considered as a very useful backup scheme for archiving purposes. However, the biggest disadvantage of this method is the more number of media required and that media isn't guaranteed to be replaced before it goes bad.

3. Tower of Hanoi(ToH) Backup

This scheme is based on the solution to the "Tower of Hanoi" puzzle involving transferring discs between poles. Imagine you have N tapes that can each store a single backup. Using ToH you would use the first tape every other day (Day 1,3,5 etc.). Then you would use tape 2 on every fourth day starting on day 2 (So 2,6,10 etc.). The general rule being: Use tape N

every 2 N days starting on day N1.This method requires you to perform a full backup on five tapes labeled A, B, C, D, and E.

Tape A is used every other backup session, tape B every 4 sessions, tape C every 8 sessions, tape D every 16 sessions, and tape E every 32 sessions, or once a month. These backups can be performed daily, weekly, or at whatever intervals you determine. This method is more complex than any of the other schemes described here but gives archives that are more evenly distributed. The table below shows how ToH works with five tapes

Table 3. ToH Backup Algorithm.

Date	Sun	Mon	Tue	Wed	Thur	Fri	Sat
Week 1 Tape	E	A	B	A	C	A	B
Backup Level	0	5	4	5	3	5	4
Week 2Tape	A	D	A	B	A	C	A
Backup Level	5	1	5	4	5	3	5
Week 3 Tape	B	A	E	A	B	A	C
Backup Level	4	5	0	5	4	5	3
Week 4 Tape	A	B	A	D	A	B	A
Backup Level	5	4	5	1	5	4	5
Week 5 Tape	C	A	B	A	E	A	B
Backup Level	3	5	4	5	0	5	4
Week 6 Tape	A	C	A	B	A	D	A
Backup Level	5	3	5	4	5	1	5

Week 7 Tape	B	A	C	A	B	A	E
Backup Level	4	5	3	5	4	5	0
Week 8 Tape	A	B	A	C	A	B	A
Backup Level	5	4	5	3	5	4	5
Week 9 Tape	D	A	B				
Backup Level	1	5	4				

As you can see, the minimal tapes (lower lettered tapes) get used far more frequently than higher lettered tapes which depicts that this scheme won't maintain daily backups for a week like the other tape rotation schemes, but it does give you exponential retention periods. For example, in the five tape rotation scheme, Tape A is used 29 times (every other day). Tape B is used 15 times (every four days). Tape C is used 7 times (with seven days in between). Tapes D and E are each used 4 times. This means, using a very few number of tapes, all the daily and monthly backups are obtained making it a very cost effective scheme, which is the main advantage of the algorithm. Another advantage of ToH is the easy file restoration procedure required during recovery because all backups are full and there is no need to shuffle through partial backups.

The disadvantage of this algorithm, however, is that it is very complex in that our backup software shall automate tape rotation and not supported by all backup applications. Another problem with the ToH is that it wears out lower numbered tapes since the frequency at which they are being used is very much larger than the higher numbered tapes. It also needs a large enough backup window to accommodate daily full backups.

To recover from complete data loss,

- first restore from the most recent level 0 backup
- Then restore from the level 1 backup if that backup was written after the level 0 backup
- Next, restore consecutively from the most recent level 3 and 4 backups if both were written after the level 0 backup
- Finally, restore each of the level 5 backups that were written after the level 0 backup

Conclusion

It is very well known by all organizations that taking a backup is a very important task that needs to be performed regularly. The question lies on which backup method is suitable for the business's need from the full, incremental or differential ones. The choice made dictates other parameters to consider as well such as backup time, media cost and recovery time needed. Full backup is known to take large backup time and large tape space but the recovery time is minimal. In the differential backup, the backup time and media might be large but has moderate restore time requirement. Incremental backups are the fastest backup schemes with a very long restoration time.

Whether a full, differential or incremental backup is chosen, how the backup shall be performed is another issue we need to address. There are various backup strategy to choose from to achieve this goal. Volume/Calendar, Grandfather/Father/Son and the Tower of Hanoi are the most commonly used backup strategy. But, there is no one single backup scheme that is clearly superior as each scheme has inherent advantages and disadvantages. It is ultimately up to the system administrator to choose the scheme that is the best fit for the organization's needs.

Assessment

1. What are the main differences between Full backup, Differential backup and Incremental backup methods
2. Discuss the advantages and disadvantages of the three backup strategy discussed.

Activity 3. Backup Devices and Backup Media

Introduction

Once the data to be backed up is identified and a backup scheme is decided, what follows is selecting the appropriate backup device that meets your requirements. This activity provides some considerations for planning the total number of tapes your backup strategy requires and the tape costs associated with that strategy.

Backup Device

There are various storage devices on which data can be maintained permanently. These storage devices include Tapes, Hard disks, Optical disks, Remote Backup service, Flash Memory, etc. There are different desired properties a backup device must exhibit. These are:

- User ability to write data to the device.
- Media capable of storing the data for long periods.
- Support of standard system interconnects
- Support of reasonable input/output throughput

Let's discuss briefly each of the media listed above.

1. **Magnetic Tape**: these are probably the most commonly used secondary storage device due to its inexpensivity, moderate performance, standardized data format and ease of use. The device used to write or read data to/from a magnetic tape is known as tape drive. Tapes have sequential access scheme. There are several kinds of tape storages such as Cartridge Tape Drive, 8-mm Tape Drive, Digital Audio Tape Drive, Linear Tape Open, Digital Linear Tape, and Jukebox/Stacker Systems. Major problem with all magnetic tapes is quality degradation from time and usage. Tapes should be replaced on a regular basis dependent upon usage (ideally no more than 10 writes) and time (discard tapes over 2 years old). Magnetic tapes are also susceptible to magnetic fields like power cables, motors, magnets, etc so care should be taken not to expose them to such environments.

a. **Cartridge Tape Drive**

Cartridge tape drives store between 10 Mb and several Gb of data on a small tape cartridge. Most cartridge tape systems use SCSI interconnections to the host system. These devices support data transfer rates up to 5 Mb per second. The actual transfer rate from the tape drive memory to the tape media is typically about 500 Kb per second.

b. **8-mm Tape Drive**

These tape drives are also small and fast, and use relatively inexpensive tape media. The 8-mm media can hold between 2 and 100 GB of data, depending on the drive model and type of tape in use. The 8-mm drives use the SCSI bus as the system interconnection. Low-density 8-mm drives can store 2.2 Gb of information on tape. And transfer data to the tape at 250 Kb per second. High-density 8-mm drives can store up to 80 GB of information on a tape at a 16 MB/second. "low" end, the 8-mm drives, do not use data compression to store the data on tape. "high" end, advanced intelligent tape drives, incorporate compression hardware and improved recording techniques to increase the amount of information that can be stored on the tape.

c. **Digital Audio Tape Drive, DAT**

These are small, fast, and use relatively inexpensive tape media. Typical DAT media can hold between 2 and 40 GB of data. Although manufacturers of DAT devices have announced the end-of-life for these products, they will remain in use for many years.

The various densities available on DAT drives are due to data compression. A standard DAT drive can write 2 Gb of data to a tape. By using various data compression algorithms, and various lengths of tape, manufacturers have produced drives that can store between 2 and 40 GB of data on a tape. DAT drives use SCSI bus interconnections to the host system, and typically offer 3 MB/second throughput.

d. Linear Tape Open, LTO

A consortium of Hewlett Packard, IBM, and Seagate developed the LTO technology. LTO encompasses two formats: the Ultrium, a high-capacity solution, and Accelis format, a fast-access format. The two formats use different tape drives, and tape cartridges. LTO Ultrium drives can store up to 100 Gb of data on a single tape cartridge at 16 Mb/second.

e. Digital Linear Tape, DLT

DLT backup devices are relatively new on the backup market. These tape devices offer huge data storage capabilities, high transfer rates, and small yet expensive media. These drives can store up to 110 Gb of data on a single tape cartridge and has a transfer rates of 11 Mb/second on high-end Super-DLT drives, making them very attractive at sites with large on-line storage systems.

f. Stacker/ JukeBox Systems

Stackers are sequential tape systems where tapes are stacked in a hopper, and the tape drive starts by loading the tape at the bottom of the stack. When the tape is full, it is ejected, and the next tape is loaded from the stack and so on. Many stackers do not have the capability to load a specific tape in the drive. Instead, these stackers simply cycle (sequentially) through the tapes until the last tape is reached. At this point they can either start the cycle over again or wait for a new group of tapes to be loaded into the hopper. Unlike stackers, jukebox systems employ multiple tape drives, and special "robotic" hardware to load and unload the tapes. Jukebox systems require special software to control the robotics. The software keeps track of the content of each tape and builds an index to allow the user to quickly load the correct tape on demand. Each tape is "labeled" with a barcode label and the mechanism contains a label reader that keeps track of what tape is in the drive. Many commercially available backup software packages allow the use of jukebox systems to permit backup automation.

1. **Hard Disk**: is a popular storage device that is used to store large amount of data. Disks are orders of magnitude faster than tape devices, and therefore offer a solution to one of the backup problems on large-scale systems. A hard disk is a magnetic disk that generally resides inside the computer and data is stored on to the hard disk as a collection of bytes. Hard disks provide relatively fast retrieval because they rotate constantly at high speed. However, disks are more expensive and consume large system resource than magnetic tapes.

a. **Redundant Array of Inexpensive Disks, RAID Disk Arrays**

RAID is a backup storage implementation of hard disks which enables the system to make mirror image copies of all data on backup disk drives. RAID disk arrays also allow data striping for high-speed data access. RAID allows storage of data across multiple drives to maximize storage space, while protecting data if one or more of those drives fails .In this scheme, the primary copy of the data is stored on one disk, and the backup copy of the data is stored on another disk. However, both disks are housed in a single box. This makes the task of moving one drive off-site much more complicated. There are a number of different RAID types indicated by numbers from which the two most commonly used types are RAID 1 and RAID 5. RAID 1 has a minimum of two disks and if something goes wrong on a single drive, the other is able to happily take over along in its place. RAID 5 requires a minimum of three disks and files are spread across the disks along with a parity block which allows the drives to rebuild files if a disk develops an error or fails.

b. **High-Density Removable Media Backups**

A relatively recent addition to the backup market is the high-density removable media drive. These devices are capable of recording 100 Mb to 2 Gb of data on a removable medium that resembles a floppy diskette. Many of these devices employ a parallel port interface. A few of them offer SCSI interfaces, allowing them to be connected to the external SCSI port on a workstation. Examples of these devices include the Iomega ZIP and JAZ drives, and the Imation Superdisk drives.

1. **Optical Disk**: is another alternative storage drive to be considered while taking backups. Optical disks store data on reflective discs read by a moving laser head that can also write data onto rewritable discs. CD, DVD and Blu-ray disks are examples of such devices. Storage capacity varies greatly among the available optical media, from 682 megabytes on CDs, to as much as 9.4 gigabytes on DVDs, to up to 50 gigabytes on Blu-ray discs—none of which can rival the storage capacity of hard disks, solid state media, or digital tapes.CD-ROM is used as a long time archive of information as it offers excellent data survivabilitybecause data cannot be accidentally erased or inherit corruptions. Another advantage to the CD-ROM is the availability of reliable data transportability between systems as CD-ROM's adhere to industry standardized data formats. However, using CD-ROM as a backup device has a disadvantage of time intensivity operation during CD creation and setup. Moreover, as CD-ROM are non-rewritable discs, rolling backups is not possible, so they might not contain the most recent version of primary files

2. **Remote Backup Service**: this is a method of backing up your data using a cloud storage or Internet. Here, our data are automatically copied to a remote storage systems and make it available from any Internet-connected computer with the appropriate access rights. This service is given by companies termed service providers. The service providers offer a system for the backup, storage, and recovery of computer files to their clients. Remote Backup Services can protect our data against some of the worst-case scenarios, such as natural disasters or critical failures of local devices due to malware. Additionally, cloud services give you anytime access to data and applications anywhere you have an Internet connection, with no need for us to invest in networks, servers, and other hardware. This is achieved by purchasing more or less cloud service as needed, and the service provider transparently manages

1. **Flash Memory**: these devices are known for their very large storage capacity, durability, robustness, rewriteable and inexpensive. Drives are typically "plug & play". However, they can be a medium for propagation of viruses.

In general, which media to use as a backup device depends on:

- Storage capacity required
- Portability of device
- Speed of data transfer
- Speed of access (time it takes for device to find a specific piece of stored data

Conclusion

Media used for backup needs to be selected according to the data being backed up and the property of the backup device. There are several backup devices to choose from before performing your backup. This includes: magnetic tapes, hard disks, optical disks and flash memory. Data can also be backed up remotely, using online services. Each of these media have their own advantage and disadvantage so care has to be taken while choosing.

Assessment

1. Explain the term backup media
2. Among the backup devices discussed here, pick the one you think is the best one and justify your answer.
3. What are the disadvantages of using remote backup services?
4. What are the required properties of any backup device?

Activity 4. Data Recovery

Introduction

Once a system losses its valuable data, the main concern will be is there a way out? How can the system rebuild itself from such disaster? Shall the system administrator sit hands crossed when such a situation happens? What is there to follow so that user's won't notice the data loss that happened to the system? This activity goes through data recovery schemes.

Definition

As we have discussed, backup is a crucial activity one needs to do so that redundant copy of original data is maintained in different locations. However, a backed up data comes to importance when we need to restore our system to its previous, healthy condition after a disastrous event happens. There are various reasons as to why or how a disastrous event happens on a system. These are:

- Hardware malfunction
- Human error
- Software malfunction
- Malicious Software/Viruses
- Natural disaster

Recovery is, thus, a process of restoring a system and get the system working again. Just like the need to plan for a backup, recovery also requires a well-planned strategy. Many organizations using ICT are required 24/7 to run and provide their services. Failure to do so could have financial consequences as well as customer dissatisfaction. To overcome these, policies are required including:

- Location of backup
- Security of backup
- Appropriate backup for continuous service e.g. RAID

Thus, a Recovery Plan tries to ensure that businesses can:

- Resume trading quickly by ensuring ICT services are back online as soon as possible
- Give customers confidence in the reliability of the business
- Retain customers by offering full service satisfaction

There are various methods that can be used for data restoration and system recovery. To mention some

1. Use of backups: By performing regular backups, data can be quickly restored, and business activities can be quickly resumed, after a data loss situation.
2. Anti-virus software: Anti-virus software, when used properly, can provide constant protection against virus corruption
3. Disaster recovery/business continuity plans: These plans can provide the necessary procedures to help an organization resume vital operations and return to normal business functions as quickly as possible following a disaster situation.
4. Mirrored site: is a fully redundant facility that has real-time information mirroring and is identical to the primary site in its all technical aspects.
5. Commercial file recovery software: File recovery software can provide the necessary tools to recover lost files

Conclusion

Once a problem happens on a system's data, restoring the system shall commence immediately. To do this, we need to have a well-planned recovery strategy to follow. A system can be recovered using backup data, installing and managing anti-viruses, using a mirrored site or implementing file recovery software. Which recovery method to choose also varies on the degree of the data loss encountered, reliability of the recovery scheme and the cost of recovery?

Assessment

1. Can you think of any other data recovery method aside from the ones discussed?

Activity 5. Server Virtualization

Introduction

Recent computer systems have become increasingly large and very complex making it difficult to build efficient and optimal infrastructure. Server virtualization represents a base technology to address this problem. In this activity, we will try to define virtualization, importance of virtualization, and virtualization approaches.

Definition and Background

The term virtualization broadly describes the separation of a resource or request for a service from the underlying physical delivery of that service. It is the creation of a non-actual version of a computing resource, such as a hardware platform, operating system, a storage device or network resources.

Before the concept of virtualization came to existence, single OS runs on a single machine and there was high coupling between hardware and software of the system. Due to this, execution conflicts often happen when multiple applications are run on a single machine. This traditional method also exhibits weak resource utilization and inflexibility making it costly.

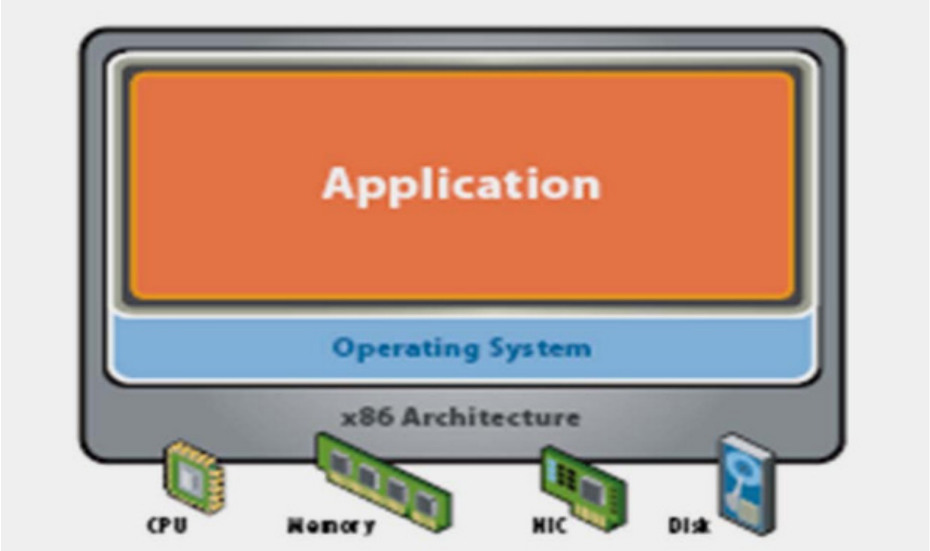

Figure 0.1. Traditional Computing Model

Virtualization is a means by which multiple operating system instances run concurrently on a single computer. The operating system that runs on the bare metal is referred as Host OS and the ones installed on the Host OS are referred as guest OS. Each guest OS is managed by a Virtual Machine Monitor (VMM), also known as a hypervisor. The virtualization system controls the guests's CPU, Memory and Storage usage.

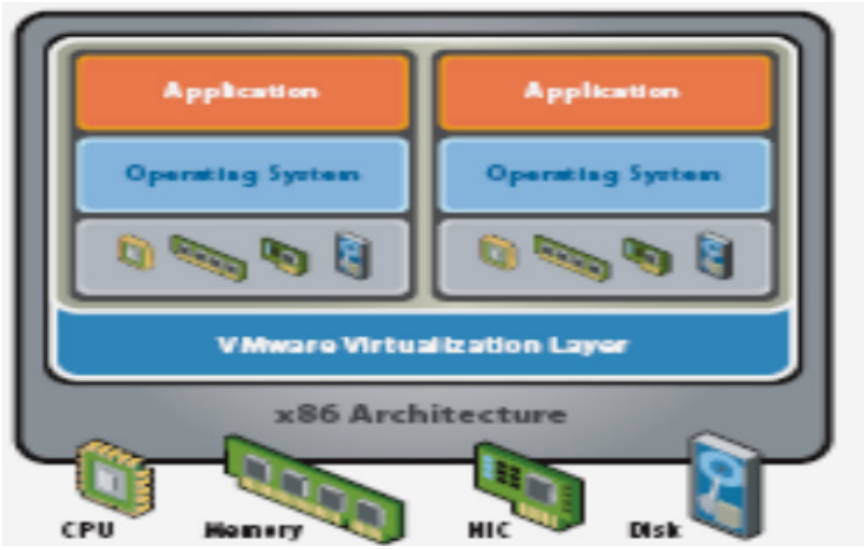

Figure 0.2. The Virtual Computing Model

Some of the advantages of virtualization are:

- Hardware independence of operating system and application software
- Provision of a virtual machine to any system
- Management of OS and applications as a single unit through encapsulation
- Simplified administration
- Increassed hardware utilization and less hardware cost
- Improved security
- Scalability and Flexibilty

There are various virtualization levels based on the application of the virtualization. These include Application Virtualization, Desktop virtualization, OS virtualization, Server Virtualization, Network Virtualization and Storage virtualization.

1. Server virtualization was first implemented more than 30 years ago by IBM as a way to logically partition mainframe computers into separate virtual machines. These partitions allowed mainframes to "multitask": run multiple applications and processes at the same time. Since mainframes were expensive resources at the time, they were designed for partitioning as a way to fully leverage the investment. By 1980/90 servers virtualization adoption initiated a reduction due to inexpensive x86 hardware platforms and Windows/Linux adopted as server OSs

Server virtualization enables flexible construction of multiple virtual servers with no hardware limitations on one physical server. Services running on a server are interacting with the resources available in the server hardware (CPU, memory, network card, etc). By placing a virtualization layer on top of the hardware, the available resources become invisible to the service and can even be shared. An operating system can be operated and applications can also run on each virtual servers similar to a physical server. The physical server can be partitioned using a software or hardware to construct the virtual server. If the physical server is divided by hardware, problems of one virtual server won't affect the remaining virtual servers. If the physical server is divided by a software, system resources such as CPU, memory and I/O devices can easily be assigned to the virtual servers. Server virtualization involves virtualizing CPU, memory and I/O devices of a system.

The CPU virtualization is accomplished by assigning a virtual CPU to the domain with correspondence of the real CPU changed to virtual CPU according to the CPU scheduler. Memory virtualization is performed by selecting a specific size of the physical memory and assigning this memory to each domain. I/O devices virtualization is accomplished through the virtual device model.

This enables services to be decoupled from the hardware and allow multiple operating systems or applications to run in their own secure container on a single server, while still appearing to be running in their own physical environment.

For example, we could run both Windows and Linux-based operating systems on the same physical machine and still receives an experience that is equivalent in its speed and presentation to having a dedicated piece of hardware.

There are three main advantages that a server virtualization brings to a system. These are:

1. Reduced capital expenditure due to efficient utilization of all hardware and less number of hardware required
2. Reduced operating expenditure due to less power supply, increased automation and fewer physical servers to be maintained
3. Increased flexibility and responsiveness as new services can be easily and quickly provisioned with no additional hardware setup

Virtualization Approaches

Three evolutions of server virtualization had been observed corresponding to the technology's development over time and each step added efficiency, security and resilience.

Single Server Virtualization : A dedicated server with a hypervisor supporting one or more dedicated services. The hypervisor is the hardware or software that separates the existing services from the underlying hardware, letting each application think it is running on its own physical machine. Although this represents virtualization in its most basic form, it already renders a large amount of hardware obsolete because a modern server can support the equivalent of about 60 virtual machines.

Multiple Servers, Common Instruction Set, Manual Management: Instruction Set is a term that describes the architecture defining the data types, instructions, registers, addressing modes and memory architecture, amongst other things, in a physical system. Here the hypervisor can allow the virtual machines to move from one physical server to another, thus creating a manually operated virtualized server farms and allows the services sitting on top of the hypervisor to be moved from one physical server to another. This can be beneficial if more flexible control utilization is required, for instance when the CPU usage is over a certain threshold, as it enables services to be moved to another, less utilized machine to maintain full performance of the service. The maximum available resource for a single service is bound to the available physical resources on a specific server. Movement of services is done manually, leaving the server/service relationship dedicated. This repositioning of services requires the operating system and application to be installed on a shared central storage system.

Multiple Servers, Diverse Hosts and Automated Management: In this case,the virtual infrastructure is extended across a number of underlying server machines (which can have different Instruction Sets), but the intelligent management software is smart enough to provision new services on-demand or to move services to different parts of the server farm to maintain performance.

These automation rules can be based on a number of variables, for example Resource Usage (CPU, Memory,etc), server maintenance hours, or backup. Automated management can be extended by use of a Self Service Portal, where end-users are able to request a service, which is then initiated. Currently, this service is available using Intel hardware in combination with VMWare software such as Citrix XenSource, Parallels, Microsoft Hyper-V and Opensource Xen, KVMt..

There are three approaches to server virtualization.

Full Virtualization : this is a 1st Generation offering of x86/x64 server virtualization which has a dynamic binary translation. The emulation layer talks to an operating system which talks to the computer hardware. The guest OS doesn't see that it is used in an emulated environment though all of the hardware is emulated including the CPU. Two popular open source emulators are QEMU and Bochs. Full virtualization offers two advantages to a system. These are:

- Isolation of VMs from the host OS and from each other as well
- Controled access by individual VM to system resources, preventing an unstable VM from impacting system performance
- Providing total VM portability by emulating a consistent set of system hardware and letting VMs to transparently move between hosts with dissimilar hardware without any problems.

However, the main problem with this approach is the high performance cost. Hardware emulation comes with a performance price as the hypervisor translates instructions between the emulated hardware and the actual system device drivers.

Paravirtualization: this was developed as a means to overcome the emulation requirement of privileged instructions from virtual machines. With paravirtualization, virtualization application programming interfaces (APIs) and drivers are loaded into the kernel of guest operating systems. This allows the guest operating systems to run while fully aware of the virtualization architecture and thus run kernel-level operations at Ring 1. The end result is that privileged instruction translation is not necessary. The architectural differences between paravirtualization and full virtualization exist between the VM and the Virtual Machine Monitor. Paravirtualization requires the existence of paravirtualization device drivers in the guest VM, the guest VM's OS, the VMM, and the hypervisor. By including paravirtualization APIs within the guest OS kernel, the guest is fully aware of how to process privileged instructions; thus, privileged instruction translation by the VMM is no longer necessary. Furthermore, paravirtualized device drivers such as for network and storage devices are written to communicate with the VMM and hypervisor drivers. Hence, the VMM does not have to present a legacy device to the guest OS and then translate its instructions for the physical host operating system or hypervisor.

Removing the heavy emulation requirements from the VMM reduces its workload to merely isolating and coordinating individual VM access to the physical host's hardware resources. Paravirtualization also offers hardware access. With appropriate device drivers in its kernel, the guest OS is now capable of directly communicating with the system hardware. Note that this doesn't mean that the VM has direct access to all system hardware.

In most instances, some system hardware will be available, while other hardware devices will appear as generic representations, as determined by the paravirtualization drivers within the VM. To determine which elements of hardware are paravirtualized and which are available for direct access, consult with the prospective virtualization software vendor.A disadvantage of this approach is the need to have a modified operating system that includes this specific API, and for certain operating systems (mainly Windows), this is an important disadvantage because that kind of API is not available.

Hardware-assisted virtualization: Hardware-assisted virtualization is very likely to emerge as the standard for server virtualization well into the future. While the first-generation hardware that supports hardware-assisted virtualization offers better CPU performance and improved virtual machine isolation, future enhancements promise to extend both performance and isolation on the hardware level. The key to isolation and memory performance lies in dedicating hardware space to virtual machines. This will come in the form of dedicated address space that is assignable to each VM. CPUs that support hardware-assisted virtualization are fully aware of the presence of the server virtualization stack. With hardware-assisted virtualization enabled via the system's Complementary Metal Oxide Semiconductor (CMOS) setup, the system will automatically reserve physical address space exclusively for virtual machines. This provides true isolation of virtual machine resources. Also note the existence of a device I/O pass-through bus in the virtualization stack. This is significant because virtual machines can use this bus to access high I/O devices such as disk and network directly instead of through emulated hardware resources. However, the pass-through bus, also known as the VMBus, is part of the VMM/hypervisor architecture for hypervisors designed to support hardware-assisted virtualization. Keep in mind that while the pass-through bus can provide a clear data path to physical hardware resources, all control information is processed by the VMM, which prevents one VM from taking full control of a hardware resource. The advantage of this approach is it allows to run unmodified OSs (so legacy OS can be run without problems). However, it also has speed and flexibility problems as the unmodified OS does not know it is running in a virtualized environment and so, it can't take advantage of any of the virtualization features

Conclusion

Virtualization is a very broad term in a computing system which refers to abstraction of computing resources. Virtualization brings several advantages to a system including resource utilization, operation efficiency and performance enhancements.

Depending on the virtualization implementation, Virtualization can be application level, desktop level, OS level, server level, storage level or network level. Server virtualization is a process of running several OS and applications on a single physical server enhancing hardware utilization and minimize infrastructure request. There are three common approaches to server virtualization. These are full virtualization, para virtualization, and hardware-assisted virtualization.

Assessment

1. Define the term virtualization
2. What are the limitations of traditional computing model?
3. What advantages does virtualization bring to a system?

UNIT SUMMARY

This unit tried to address backup and recovery systems and server virtualization in brief. It talked about backup, backup methods as well as backup algorithms. It also elaborated on storage devices and media to be used for backup. Data recovery and the various options to be used as a recovery strategy had been touched. The discussion continued to concept of virtualization and in particular focused on server virtualization; what it is, it's importance, it's evolutions and finaly approaches to server virtualization.

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

Write your answer for the following questions

1. What is one of the fastest and easiest ways to cut your backup time?
2. What is the importance of having a backup plan?
3. Explain the points to be addressed before taking a backup?
4. What are the disadvantages of using remote backup services?
5. What are the required properties of any backup device?

6. Which of the presented recovery methods do you think is best? Why?
7. Explain the main differences between the three server virtualization approaches
8. What are the three evolutions to server virtualization?

Unit readings

- http://www.pcmag.com/article2/0,2817,899676,00.asp; Backup Methods and Rotation Schemes; BY MATTHEW D. SARREL; FEBRUARY21, 2003
- www.tandbergdata.com/us ; Tape Rotation BEST PRACTICES The Tandberg Data SMB Guide to Backup Data Protection Strategies for the Small-to-Medium Size Business p.10 Basic Backup Guide

Unit 4: Securing OS and Monitoring

Unit Introduction

This unit deals about securing operating system environments and performance monitoring. There is no such thing as a secure system. The closest thing to a secure system is a standalone system locked in a room that only one person has access to. While this system is close to being secure, it is not very practical. For a system to be useful in today's business world, it almost always needs to be attached to a network of some kind. While we can never achieve a 100% secure system, we can take steps to make things as secure as possible. How do we achieve security for the operating system?

Performance of the operating system needs to be monitored to conduct any performance tuning. Performance issues related to the following will still keep coming up and you have to proactively troubleshoot.

- Storage issues and bottlenecks
- CPU load
- Memory bottlenecks
- Network issues

This unit, also introduces tools that can be used to monitor and analyze system and application performance, and points out the situations in which each tool is most useful. The data collected by each tool can reveal bottlenecks or other system problems that contribute to less-than-optimal performance.

Unit Objectives

Upon completion of this unit you should be able to:

- Setup host security
- Identify attack types and manage security solution
- Demonstrate tools that are used for performance monitoring
- Explain performance monitoring of a computing system
- Elaborate the advantage of performance monitoring and performance monitoring tools and types

KEY TERMS

Performance: Act of doing something

Monitoring: A process of assessing and observing an activity or action

Tuning: Refining or adjusting an action

Activity 1: Types of Attacks

Introduction

There are so many types of security attacks that it is impossible to list all the variations. We explain some of the most common here, with the emphasis on clarifying how the attack is made and what vulnerability it attempts to exploit. It is important to say that there is no such thing as absolute security from these attacks. Even if there were, some enterprising cracker would soon find something else to exploit. This is true not only of Linux but of every operating system out there: all Unix versions, all Microsoft Windows versions, VMS, and any other operating system you can think of.

The goal is to maintain the highest possible level of security on your system by keeping up with new attacks as they hit. Watch your distribution's home site for updates that fix security exploits, and install them immediately. Each major distribution of Linux has a security–announce list where users and staff post news of new security exploits and other security–related news. Be sure to subscribe to the appropriate security list to stay ahead of would–be intruders.

Trojan Horse

You probably remember the story of the Trojan horse from Homer's Iliad. The Greeks sent a giant wooden horse as a present to end their ten–year war with Troy. When it was within the gates of Troy, Greek soldiers emerged from inside and defeated the unsuspecting Trojans.

In security, then, a Trojan horse is a program or a script that is hidden inside an authorized program. This code, when executed, causes an undesirable effect. Many of these are aimed at Microsoft machines, targeting their mailers. For example, you may remember the "Love Bug" Trojan horse from May of 2000. It appeared to be an e–mailed love letter but, when opened, it would cause problems, such as sending itself to everybody on your e–mail address book, erasing your files, and downloading another Trojan horse program that would steal your passwords.

This particular Trojan horse affected only Microsoft Windows machines; while Trojan horses are rare on Linux machines, they are not unheard of. In January of 1999, a Trojan horse was found in the util−linux−2.9g.tar.gz on some Web sites. This type of attack is difficult to protect against. As a result, many Linux distributions have begun adding checksums or Pretty Good Privacy (PGP) signatures to the packages they make available for download. Make use of these verification techniques if they are available to you.

Pretty Good Privacy (PGP)

Pretty Good Privacy is a software package that provides both asymmetric data encryption, which is usually used for e−mail, symmetric encryption for other files, and secret key management. PGP is generally used to ensure that the e−mail you receive is as its author intended. Mail sent through the Internet can easily be intercepted, changed, or created as if it were from someone else entirely. With PGP, you can digitally sign your e−mail, generating a mathematical value based on your e−mail's content as encrypted using your private key. If the recipient of e−mail has your public key, the PGP software makes the same calculation, thereby determining that the message has not been altered. Since only you have the private key that encrypted the hash value that was successfully decrypted with your public key, only you could have digitally signed the e−mail. Packages made available for download may be signed the same way.

The difficulty comes with obtaining the public key in such a way as to be sure that the person who claims the key actually is the one who generated it. The most secure way is to meet face to face with the person, exchanging identification and public keys at the same time. Retrieving a signed software package from a Web site is not as secure. Typically you obtain the developer's key from a web site or possibly via e−mail. Still, this is a step in the right direction. For more information on PGP, see its international home page at http://www.pgpi.org/.There is a GNU version of PGP available that is gaining popularity in the Linux world. GNUpg, also referred to as GPG, is compatible with the existing PGP. For more information on GNUpg or to download a copy to try, go to http://www.gnupg.org/.

Checksum

A checksum is another type of file signature that is usually used for validation when files are transferred by means that are considered unreliable, for instance from disk to tape. It is valuable as a security check as well. The checksum is computed from the binary bytes of the file, creating a "fingerprint" of sorts. When the file is transferred, a new checksum is computed and compared with the saved fingerprint. This method is not foolproof. A knowledgeable cracker will simply pad the file to obtain the appropriate checksum value. Like other methods, however, it is better than nothing. For more information about creating a checksum, see the chksum man or info page.

Back Door

A back door is code added to a program to allow unauthorized access to the host system. Take, for example, a program that has been altered to e–mail information about the person running the program to some mail account whenever the program is run. Very few users look at the size of binary files to see whether the program has been changed. Fewer still check beyond that if the sizes match. Even worse is that some programmers–for–hire include back doors in the programs they write as a matter of course. Once there is any negative interaction between the programmer and the client, the programmer exploits the back door and gets revenge. In this case, the only way to find the offensive code is to wade through the source code. Typically, if the code works, it isn't looked at until it requires some updating, so the programmer has plenty of time to play spy.

There is little if any defense against a back door if it is inserted by someone doing software development for your company. If you believe that a coder has inserted a back door in a product, you can have another programmer look at the code, but beyond that, you have to depend upon you rnormal security watches to catch his or her entry onto the system. With a system that uses a package manager like the Red Hat Package Manager (RPM) or Debian's Package Management System, you can run a command to verify that a package as a whole has not been changed and output codes indicating any changes to individual files. Using this information, you can determine whether files on your system and their attributes have been changed and how.

Trusted Host

One of the most talked–about vulnerabilities in Linux involves the trusted host files. Trusted host files allow users to log in without a password, because their logins and host machine names are included in the appropriate files: /etc/hosts.equiv and .rhosts in the user's home directory. The commands that rely on these files are referred to as the r–commands and include rlogin, rsh, rcp, and rexec. Most Linux system administrators simply disable these commands and replace them with Secure Shell (SSH).

SSH is a client/server package that encrypts all outgoing data from either direction, thus creating an encrypted tunnel. The client side sends a request to the server (sshd), and the server sends back a public host key and a public server key. The client attempts to verify the public host key by locating it in either /etc/ssh_known_hosts or the .ssh/known_hosts file in the user's home directory on the client. If the client is able to verify the public host key, it generates a 256–bit random number, encrypts it with the two keys from the server, and sends it to the server. This generated number is then used by both the server and the client to derive the session key, which is used to encrypt any further communication between the two during this session. If the process fails to generate a session key, access by the client is denied. This is a somewhat simplified description, but it illustrates the basic process.

Red Hat 7.3 includes a free version of SSH that was developed by the OpenBSD Project. Called OpenSSH, the package includes the following: the ssh program, to be used in place of Telnet, rsh,and rlogin; sftp and scp, to be used in place of rcp and ftp; and the server−side sshd program and its supporting utilities.

Buffer Overflow

A buffer overflow occurs when a program written to read data is not properly restricted to the intended area of memory. This is similar to when your cellular phone picks up someone else's conversation. You don't really know whose conversation you are hearing, but you might hear something that the speaker didn't intend to reveal to anyone outside the target audience. A buffer overflow causes a program to read data that was not intended for it. It is quite difficult for even skilled programmers to take advantage of this, but there are programs available on the Internet that were written explicitly for this purpose. Typically, programs prone to overflow buffers are reported on the main security lists for your distribution.

Scanning or Sniffing

Network scanning involves the probing of an IP address or some of its ports to determine which services are available on that system. Sniffing is the process by which data packets are intercepted as they are sent over the network. Specifically, crackers look for passwords and the like. There are packages available that help you to detect whether your system is being attacked in this way and even to fool the sniffer.

Spoofing

Spoofing is the process of masquerading as a known user. This can be in the form of e−mail that appears to the receiving system to have come from a known user or data packets that look like they came from a trusted machine. The details of how this is accomplished are beyond the scope of this book, but there is a way to minimize or possibly even stop it completely.

Denial of Service

In a denial−of−service attack, the attacker tries to make some resource too busy to answer legitimate requests for services. In recent years, several large companies, including Amazon.com, Yahoo!, and CNN.com, have had service interruptions due to denial−of−service attacks.

Here's an example. Bad Guy and his buddy have two computers on the Internet. Each uses scanning techniques to find computers on the Internet that they can spoof through. Each finds several vulnerable computers they can crack. At a pre−agreed upon time, Bad Guy and his friend cause all the compromised computers to ping the target site, thereby increasing the load and hiding his and his buddy's identities. The target site can't handle all of the false traffic and its regular load too, so everything grinds to a halt. This is especially effective with the HTTP daemon, which attempts to confirm the address of any system that makes a request upon it. The HTTP daemon thus waits until its request to the spoofed IP times out before being "convinced" that it is a bad address. In this case, Bad Guy and his friend send so many data packets with spoofed IP addresses that the server is too busy trying to confirm IP addresses to be able to do any Web servicing.

Detecting and preventing denial–of–service attacks is difficult since the attacks come from several sources at once. To minimize the effects of the attacks, thoroughly screen incoming network packets and deny packets using firewall or routing software. You can write firewall rules that will limit the number of packets of a given type from a specific address at a time. This might not completely prevent delays in service, but it will reduce the load that the attack creates on your system, preventing all but the most intense attacks from crippling your system.

Accurate logging at the Web or mail server can assist law enforcement in catching the cracker; however, this is complicated in many cases by the fact that crackers rarely launch this type of attack from their own machines. Instead, the cracker launches the attack via other compromised machines on the Internet, thereby implicating only the owners of the exploited machines.

Password Cracking

Password cracking is the acquisition by whatever means of a valid user password. This is a real threat to any system but particularly to systems whose administrators don't teach their users the importance of selecting a hard–to–guess password or don't force users to change their passwords fairly often. We'll discuss the actual authentication process in the upcoming "Securing the Authentication Process" section.

For now let's talk about how to select a secure password. Passwords should not be standard dictionary words like misery or supersecret. Passwords should have a mixture of numbers, letters, and symbols like #@$%&. Use capital letters and lowercase letters. A good password might beR#t34%Ka. Linux uses a tool called crack lib to evaluate the security of a password. Regular users cannot change their passwords to anything based on a word in the cracklib dictionary, although the superuser may do so. This restriction is accomplished by using Pluggable Authentication Modules, which we'll discuss in this unit.

Social Attacks

Social attacks are performed by an insider or by the manipulation of an insider. These attacks areas simple as walking up behind a user who is logging in and watching to see the password as it is keyed in. In another form of social attack, an individual pretends to be a system administrator and obtains information—such as a user's password—that provides unauthorized access. These types of attacks are best dealt with by making your users security–conscious. Teach them to be alert for such situations, and make sure to introduce new members of your system administration staff to as many users as possible. Instruct them never to give their password over the phone or in e–mail. Also impress upon your users the importance of not using family member names or other easily guessed words as passwords.

Physical Attacks

If an intruder gains physical access to your system, s/he may try to hack it, interfere with or intercept network traffic, or even physically destroy it. Linux in particular has one physical security vulnerability that you need to avoid. Anyone with a boot disk can boot the average system into single–user mode and have root access to it! Obviously this is not something you want to have happen. The vulnerability is with respect to the boot loader. You must change the default configuration to safeguard the system.

The first step is to assign a password that must be typed in when the system is rebooted into single–user mode. Without the password, the machine doesn't initialize, so no breach is possible.

Under GRUB

To do so on a system using GRUB, you must run the password command from within the/etc/grub.conf file by adding the following line:

```
password ––md5 ENCRYPTEDPASSWD
```

After having created PASSWORD by using the md5crypt command, at the grub prompt:

```
grub> md5crypt
PASSWORD: **********
Encrypted: ENCRYPTEDPASSWD
```

Under LILO

To do so on a system using LILO, add the restricted and password lines to /etc/lilo.conf as indicated below:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
prompt
message=/boot/message
linear
default=linux
append="aha152x=0x140,12,7,0,0"
restricted
password=#94hd*1@
image=/boot/vmlinuz–2.2.16–22
```

```
label=linux

read-only

root=/dev/hda1
```

The password line is self-explanatory. The restricted line prevents the password from being required at every boot, which enables automatic reboots while still requiring the password before kernel parameters, such as single, are accepted from the console. This is just what you want: You want your system running the default kernel and you don't want intruders to change that. Remember, if you include the prompt line to accept console input, add a timeout value so the system automatically reboots after a crash. Omitting the timeout option removes the ability to boot the machine unattended.

If you are using LILO or GRUB without the --md5 flag to establish a boot loader password, you need to secure the configuration file to prevent an intruder from viewing the password, which is in clear text. First use the chmod command to set the file so that only the root user has read/write privileges.

```
# chmod 600 /etc/lilo.conf

or

# chmod 600 /etc/grub.conf
```

Next use the chattr command to change the file attributes so that the immutable bit is set on the file.This prevents the file from being modified, linked to, removed, or renamed except by the superuser.

```
# chattr +i /etc/lilo.conf

or

# chattr +i /etc/grub.conf
```

Finally, you'll need to ensure that the machine's BIOS is set to prevent booting by flash stick or CD/DVD-ROM, since there are bootable images that would allow the intruder to circumvent the bootloader hurdle. Finally, password-protect your BIOS settings. Now you are secure against someone sitting down and rebooting the system unless they know the password to either the BIOS or the bootloader. An intruder who can open the computer's case can still steal the hard disk or reset the BIOS to gain entry to the computer, though. The best protection against this is to keep the computer in a physically secure area.

Conclusion

In this activity we discussed types of attacks on the computer system along with solutions. As a system administrator, one has to identify the treats against the computer system and more specifically the operating system.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

- What kind of attack is buffer overflow?
- Why do people attack servers using denial of service?
- What solutions do you suggest for Trojan horse attack?

Activity 2: Types of Security

Introduction

Any operating system that is going to be on a shared or public network needs to be secured. The default installations are getting better about not leaving systems open to attack, but there are still several things that must be done to prepare for the exposure associated with a network environment, especially if the system will ever come in contact with the Internet. As the Internet expands, the frequency of attacks on systems increases as well.

Types of Security

Of course, a computer's security needs are relative to the importance of the data it contains. Home computers that are used primarily for learning, data processing, and sending e–mail need much less security than do bank systems, which are heavily targeted by crackers. Your system's security needs will probably fall somewhere in between. While a home user might not want the intrusive measures higher security requires, a data communications firm would probably be willing to jump through those hoops. Security requires a lot of extra effort by users as well as the system administrator, so determining just how much security you need is important. The job of securing your system, called risk assessment, has two parts: assessing vulnerability and assessing threat.

Assessing vulnerability means assessing the loss that the company would suffer if your security were breached. Ask yourself where your systems are vulnerable. Which information might be profitable to a thief? Which information might generate a profit if taken by disgruntled employees? Is the mission–critical data available to naive users who might inadvertently damage or delete it? Is the mission–critical data available to all employees or only to those who need it?

How much downtime would spent fixing a hacked system cost the company in man–hours, lost profit, etc.? These are the questions to ask before developing the security plan.

Assessing threat means assessing how many people out there are really going to attempt to break in. Be too cautious. If you are on the Internet, the threat of a cracker on the hunt is a real issue. Also ask yourself whether you have competitors or other groups who might seek to damage your company's reputation by stealing data or taking the system down. Try to determine the threat. Of course, you must never completely discount any of the above possibilities, but the answers to these questions can help you to focus your efforts where the threat is greatest.

Once the security system is in place, an auditing plan should be written to ensure that the security that you've set up remains effective. It is better to use a formal auditing checklist than to "shoot from the hip" since repetitive tasks like security evaluation are all too easy to rush through. A checklist forces you to be thorough.

Although no system is ever completely safe from intrusion, there are ways to minimize the possibility of a break–in or other security breach. For this discussion, we'll divide security into the areas that are under the system administrator's direct control:

1. User–based security basically answers two questions: "Is this user who s/he claims to be?" and "What resources should be available to this user at this time?" Linux uses Pluggable Authentication Modules (PAM) to secure the system from intrusion by unauthorized users.
2. Port security is designed to protect network ports from unauthorized hosts and networks. This type of security is handled largely by the kernel but is also affected by IP firewall administration, commonly referred to as IP chains or IP tables. Port security is important since any open port is an invitation to a cracker. Port security also helps to control outgoing information from the users on your systems.
3. You must also restrict network access to system resources and services based on the requesting host. Quite often you know that the users on www.gooddomain.com should have access to a given resource while users on www.baddomain.com should not. This is the job of host–based security. Under Linux, host–based security is handled in any of several ways, including TCP Wrappers, xinetd, and individual daemon security options.
4. You should restrict physical access to your systems to prevent tampering. Also use the boot loader securing measures that we mentioned in the "Physical Attacks" section.
5. A last type of security is the assignment of permissions to exclude certain users from having access to specific files or devices. This is the most commonly discussed method of securing your system and the best understood. Unfortunately, by itself it is insufficient to defend your system adequately.

Conclusion

In this activity, we discussed the types of security that can be applied in operating systems. Some solutions address number of treats. One of the major security types is authentication. Do we need to secure the authentication process itself? The next activity will address this question.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

- Who do you identify potential attacks on your system?
- What kind of security is port security? Which treats are handled by port security?

Activity 3: Securing the Authentication Process

Introduction

The authentication process is a mystery to most of your users. Generally users don't worry about how they are authenticated as long as they gain access to their files. They type their login and password combinations and wait while the magic happens. If they are authenticated, they are content; if the authentication fails, they call you for help. As a system administrator, you need to know how this process works.

First, Linux uses two different systems for authenticating user passwords: DES (the Data Encryption Standard) and MD5 (message digest algorithm, version 5). The DES algorithm can deal with MD5passwords, but the MD5 algorithm can't understand DES passwords. DES passwords are encoded using the Federal Data Encryption Standard algorithm. MD5 uses the RSA Data Security, Inc.'sMD5 message digest algorithm. By default on most Linux systems, MD5 is used now. Passwords encrypted with the MD5 hash are longer than those encrypted with the DES hash and also begin with the characters 1. A shorter password that does not begin with a $ is typically a DES password. In heterogeneous environments, all systems normally use the same method, usually DES, for compatibility.

Hashing Passwords

On a Linux system the passwords, which are often referred to as encrypted, are actually encoded or hashed. This means that they are encoded in such a way that they generally cannot be decrypted. Figure 1.1 illustrates the process.

Creating Password

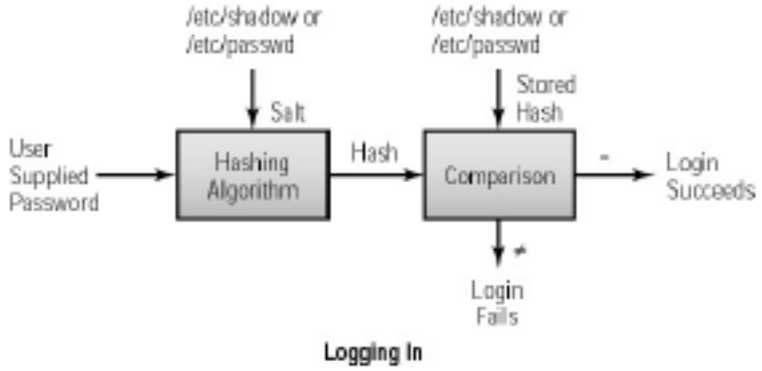

Figure 4.1: How Linux passwords are hashed

When users change their passwords, the variable−length password that they enter and a two−character salt that the system generates are run through a modified DES algorithm or MD5 algorithm to generate a hash. Although the hash is often referred to as an encrypted password, it actually is encoded instead. The salt is a sort of monkey wrench that changes the hash in any one of 4096 ways. Under Linux, the DES result is a 13−character one−way hash, a string of characters that is almost impossible to regenerate into its plain text form. The first two characters represent the salt, and the remaining 11 represent the hashed password. For MD5, the password is 32 characters long and begins with 1. The characters between the second and third $ represent the salt, while the rest is the hashed password.

Linux systems do not "decrypt" the stored passwords at all during authentication. Instead, the password that a user supplies to a DES system at login and the two−character salt hash taken from the /etc/passwd file are run through the same DES algorithm, generating a one−way hash. In the case of MD5 passwords, as with the DES passwords, the salt is removed and passed with the login password into the appropriate algorithm. The result is compared to the hash stored in /etc/passwd. If they are the same, it is assumed the proper password was supplied. In old−school Unix, the user was considered to be authenticated if this encrypted password was identical to the second field of the user's entry in the /etc/passwd file. The potential for intruders masquerading as authenticated users was simply ignored. An /etc/passwd file has the following format:

username:passwd:UID:GID:full_name:directory:shell

The /etc/passwd file on a Linux system must be world readable for the other user information that is stored in this file to be available to the programs that need it. Unfortunately, this makes the hashes available as well. While you cannot decrypt a hash, it is possible to compare a list of generated hashes with the ones stored in /etc/passwd to find a match, which can be exploited. Crackers simply encode a dictionary of words and common passwords using all possible 4096 salt values.

They then compare the hashed passwords in your /etc/passwd file with those in their list. Once they have found a match, they have the password for an account on your system. This is referred to as a dictionary attack and is one of the most common methods for gaining unauthorized access to a system. The attack is significantly easier if intruders can obtain a copy of your /etc/passwd file since they then only have to encode their dictionaries with the salts that exist in the /etc/passwd file. To reduce this vulnerability, Linux adopted the practice of shadowing the passwords.

Pluggable Authentication Modules (PAM)

At first, most privileges were granted based on password authentication or membership in a specific group. That was not a bad beginning. It soon became obvious, however, that it was insufficient to assume that a user's identity is known simply because the login password is correct and rely on that authentication as the sole means of determining what access should be granted. New programs began including their own authentication methods instead of assuming that the login authentication was sufficient. Each authentication method required the login, FTP, and other programs to be rewritten to support it.

Now authentication is more flexible and is handled under Linux by Pluggable Authentication Modules (PAM). The different major distributions began including PAM in the following versions: Red Hat Linux 5.0, Debian Linux 2.1, Caldera Linux 1.3, and SuSE Linux 6.2.

PAM is a flexible set of library modules that allow a system administrator to configure each individual application with its own authentication mechanism—a defined subset of the available library modules. Although some operating systems look for their PAM loadable object files in/usr/lib/security, Red Hat chose to follow the Linux Filesystem Standard by locating them in/lib/security. A listing of this directory includes a number of files with names like pam_access.so,pam_console. So, pam_cracklib.so, and so on. These files are the shared libraries that support the various capabilities of PAM.

Any program that uses PAM modules instead of some internal security method is called a PAM client. PAM uses four basic modules to authenticate users, manage user accounts, manage passwords, and manage actual login sessions. The authentication module is responsible for both verifying the user's credentials and giving the okay to proceed. The account management module provides the capability to determine if the current user's account is valid.

This module checks for password or account expiration and verifies that the login meets any restrictions like access hours or maximum number of users. The password module provides functionality to change a user's authentication token or password. A session management module provides functionality to set up and terminate login sessions and performs any setup tasks such as mounting required directories. These modules are referenced in the configuration file, which we'll look at below.In order to work with PAM, applications are developed to be independent of any particular authentication scheme. This means that the program is no longer concerned with which authentication scheme it will use. Instead, it will have "authentication modules" attached to it at runtime in order to work.

The Linux administrator can change the authentication scheme for a PAM–aware application at any time, without rewriting the application code, by editing a PAM configuration file, which may be application–specific or used system–wide. The modules referenced within the configuration file implement the authentication scheme; so to change schemes, you edit the file to include different modules. If no existing modules meet your security needs, you may write new modules and include them instead. New pluggable modules are being developed all the time, so it pays to watch the PAM Modules site: http://www.kernel.org/pub/linux/libs/pam/modules.html to see what is already out there.

Conclusion

One of the major issues while one secures the operating system environment is making sure that the authentication process itself is secured. Among the issues to be handled hashing passwords is mandatory. PAM is provided in Linux operating system, which activity deals with too.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. How does linux authenticates users for login?
2. What is the use of PAM? How can a system administrator make use of it?

Activity 4: Protection against Network Intrusion

Introduction

One of the biggest tasks with regard to security is to protect the network against intrusion. Any time you have a network, especially one that's connected to the Internet, there is potential for intrusion. A number of things can be done to minimize the danger, but no method is foolproof. Usually several methods are used in combination to maximize the chance of catching or stopping an intruder. Among these are firewalls, TCP Wrappers, and post–break–in analysis to enhance prevention (learning from your mistakes)

Firewalls

The purpose of a firewall is twofold. Firewalls keep would–be intruders out and keep local system users in.

If a company wants to protect proprietary information, it can block access to the system from any IP address not within the company. This is usually more desirable than taking the machine off the Internet entirely, since it allows remote users within the company to retain access. If that same company wanted to prevent its users from having access to certain outside Web sites, it could set up a firewall, which would restrict certain IP addresses from being able to send packets through the firewall to the specified sites on the Internet. The packet itself has two parts: header and data. The header contains the IP address of the sender, the intended recipient, the packet size, and other relevant information. As a result, the firewall need not waste time on the data portion of the packet but instead inspects the header to determine what to do with the packet.

IP Chains

The IP Chains package makes managing the kernel's inherent firewall capabilities simpler. An IP chain is a set of rules to be considered when a data packet is attempting to pass through the firewall. The default IP Chains configuration includes three permanent chains: input, output, and forward. The input chain governs incoming packets, the forward chain manages packets destined for another host, and outgoing packets are subjected to the rules of the output chain. You may add and edit other chains as needed.

IP Tables

The IP Tables package, more often called Netfilter, is a bit more sophisticated than IP Chains. IP Chains establish a stateless firewall that doesn't concern itself with the ordering of packets received. There is no mechanism for noticing that a packet arrived out of order or that a packet that was expected wasn't received. The IP Tables package is a stateful firewall, which adds tracking of open connections in a state table. There are three default tables: filter, nat, and mangle.

The default filter table includes these permanent chains: INPUT, OUTPUT, and FORWARD. Building on our IP Chains explanation, the use of these chains is rather intuitive except for a couple of differences: whereas all incoming datagrams are sent to the IP Chains input chain, only datagrams destined for the localhost are sent to the Netfilter INPUT chain; datagrams destined for other hosts are sent to the FORWARD chain. Another key difference is that the IP Chains output chain sees all outgoing packets regardless of where they originated. Netfilter's OUTPUT chain only looks at packets originating in the localhost destination, ignoring those datagrams being routed from a different host.

For masquerading, you'll need to use the nat table. This table has two chains: PREROUTING and POSTROUTING. Use –t nat to specify that you want to affect the nat table.

Proxy Server

While a packet–filtering firewall handles filtering at the network level, a proxy server works at the application level. The filtering firewall knows nothing of which application sent the packet—only the IP address of the machine that sent it and the destination IP—but the proxy server receives data directly from an application that is set up to talk to that specific proxy on that particular port. For example, Netscape may be configured to send all requests to a proxy server that will get the requested page and return it to the Netscape application. The proxy server may also maintain a cache of the last pages visited and check each page against that cache before going onto the Internet to retrieve the requested page.

The proxy server can reside on your firewall machine, or on a different machine inside the firewall, or connected to the Internet and the local area network without the benefit of a firewall. It is usually preferable to have as little outside the firewall as possible, so putting the proxy server on a different machine that is protected by the firewall is a good idea.

TCP Wrappers

TCP Wrappers is a host–based security layer that applies to daemons started by the inetd daemon. The tcpd executable is a wrapper program that checks the /etc/hosts.allow and /etc/hosts.deny files to determine whether or not the specified service should be run. When tcpd is started by inetd, it reads the /etc/hosts.allow file and then /etc/hosts.deny. It grants or denies access based on the rules in these two files.

The /etc/hosts.allow file is a configuration file for the tcpd program. This file contains rules that describe which hosts are allowed to access which services on the localhost. The format is as follows:

```
service_list: host_list: [optional_command]
```

The /etc/hosts.deny file uses the same format as /etc/hosts.allow but serves the opposite purpose. Lines in this file determine which hosts will be denied access to the specified service. The /etc/hosts.allow and /etc/hosts.deny files work in conjunction. If the /etc/hosts.allow file listed the rule ALL:ALL, allowing open access to all services, but the /etc/hosts.deny file included lines that prohibited some hosts from using some services, /etc/hosts.deny would override the open access for only the hosts and services specified. The deny file is the dominant file. An open system would include the ALL:ALL entry in the /etc/hosts.allow file and nothing in the /etc/hosts.deny file, granting access to everyone and denying it to no one. A completely closed system would have nothing in /etc/hosts.allow and ALL:ALL in /etc/hosts.deny, thereby denying access to everyone.

Conclusion

This activity assesses the different measures that can be set to protect the system from attacks which are coming from network. Firewalls, IP chains, IP tables, proxy servers and TCP wrappers has been discussed.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. Apart from the ones listed as protection mechanisms of network intrusions, list potential protecting mechanisms.
2. How does proxy servers keep network intruders from the system?
3. On which machine, do you suggest, firewall should be installed?

Activity 5: Performance Monitoring

Introduction

Monitoring and performance tuning are essential parts of system administration. Performance monitoring is the process of finding bottlenecks in a system and tuning the operating system to eliminate these bottlenecks. Monitoring is carried out to ensure resources are running smoothly and troubleshoot problems as they occur. And then the performance of resources is tuned to achieve optimal performance based on the current system resources and traffic load. Though it is believed that performance tuning can be achieved by setting some parameters in the kernel, this is not the case.

In this activity, we will try to discuss performance monitoring of a computing system and issues related with it.

Overview

Monitoring your system is essential in order to be able to improve its performance, locate the source of a problem and take more targeted corrective actions. Performance tuning is about achieving balance between the different sub-systems of an OS including:

- CPU
- Memory

- I/O
- Network

These sub-systems are all highly dependent on each other and any one of them with high utilization can easily cause problems in the other. Thus, performance monitoring of a system revolves around monitoring the performances of the sub-systems and tuning their issues. Administrators monitor a system for several reasons such as

- Establishing a baseline of a system performance
- Monitor current system performance and compare it with the baseline
- Study trends to detect problems before they become critical

The general advantages brought through performance monitoring are:

- debug processes and systems
- effectively manage system resources
- making system decisions and
- evaluating and examining systems

Unless performance is monitored, we can't tell success from failure and if we can't recognise failure, we can't correct it. thus, way has to be there that can be used to measure performance of a system. Performance can be measured through resources spent in computing plus resources spent in waiting for data or some other actions. But, before measuring performance of a system and conclude it has a performance issue, we need to understand what is expected of the system. What kind of performance should be expected and what do those numbers look like? The only way to establish this is to create a baseline. Statistics must be available for a system under acceptable performance so it can be compared later against unacceptable performance.

One way of performance measurement is through use of performance counter. Performance counters can be implemented with special registers in the hardware to measure events and then we can investigate the value in this register, and use that as part of our performance tuning. Performance counter method is implemented by manually inserting code to start, read, and stop counter. This method measures exactly what we want and anywhere we want. It also can measure communication and computation duration. For instance, we can wrap our DMA get and DMA put methods by timer calls, and measure the actual work by timer calls. However, counter method requires manual changes to code. Examples are cycles, number of instructions, number of loads, number of stores, etc. Another performance measurement is dynamic profiling. Dynamic profiling can in turn be event based or time based where in the first one, interrupt is executed when an event counter reaches a threshold and in the later one, interrupt is executed every t second.

Dynamic profiling does not require that you know where problem might be and supports multiple languages and programming models making it efficient for appropriate sampling frequencies. Other useful performance measurements can be processor utilization, instructions per cycle, bandwidth consumed, etc.

Performance Monitoring Tools

Monitoring tools provide plenty of critical metrics which give us information on performance and other OS activities. The monitoring tools also help us to identify the cause of performance bottlenecks of a system. Most OSs have monitoring tools as part of the base installation or add-ons that can be installed easily. Efficient CPU utilization is an important metric to measure system performance and most performance monitoring tools compute CPU utilization from user time, system time, idle time and I/O wait where:

- User Time The time percentage a CPU spends executing process threads in the user space.
- System Time – The time percentage the CPU spends executing kernel threads and interrupts.
- I/O wait– The time percentage a CPU spends idle because all process threads are blocked waiting for IO requests to complete.
- Idle time – The percentage of time a processor spends in a completely idle state.

The performance of a CPU can be determined from it's utilization, amount of Context Switch performed and the Run Queue. In general, a system is expected not to have a run queue of more than three threads per processor, full utilization of CPU (65% – 70% User Time, 30% - 35% System Time, 0% - 5% Idle Time), and less frequent context switches.

There are several performance monitoring toos that can measure these statistics. These tools are categorized into two broad classes: Real-time Monitoring tools and Log-based Monitoring tools.

Real time performance monitoring tools: monitor and capture the current system state and provide up to date performance information or statistics. Such kinds of performance monitoring tools base on system calls that are built into the operating system to extract the performance readings due to the fact that system calls are built into the operating system and the significant impact they bring to the system performance. Task Manager, Performance Monitor (Perfmon), Process Monitor, etc are examples of Real time performance monitoring tools for Windows OS.

Task Manager is a very well-known tool for monitoring processes on the Windows operating system. It shows all applications and their state, all processes and some of their most frequently used performance measures, and some general system performance measurements.

Among the tabs found on the Windows Task Manager, Two of them provide operating system performance data to the user. The Processes tab, and the Performance tab. The Processes tab shows the current memory and percentage of CPU usage of every process running on the computer as well as the total CPU and memory usage of the system. While this is not a lot of information, it is a very good first indicator when a process is taking too much of the CPU or has a memory leak. It is easy to use and even allows sorting by name, user name, CPU usage, or memory usage. The Performance tab also provides a top level view of the system state in terms of CPU and memory usage.

Figure 4.1 and 4.2 below depicts the Task Manager's two performance data display tabs.

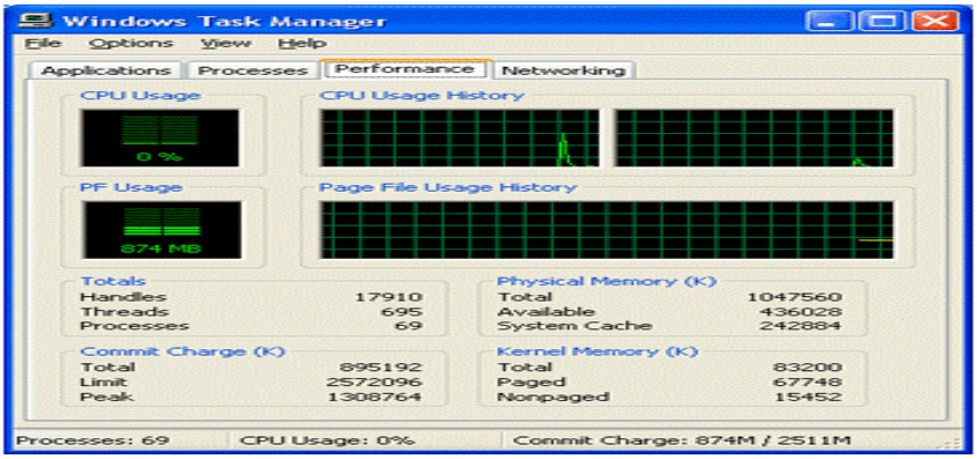

Figure 4.1. Performance tab of Windows Task Manager

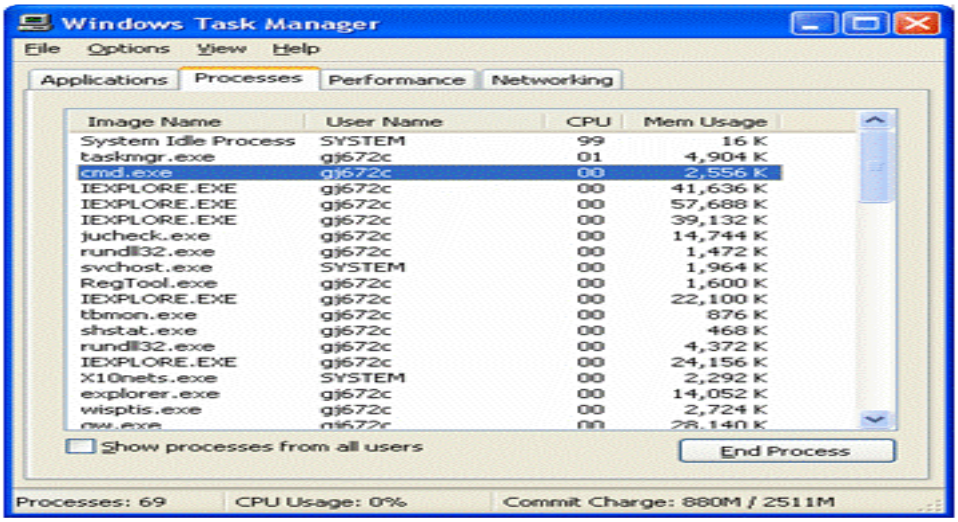

Figure 4.2. Processes tab of Windows Task Manager

Some examples of real time performance monitoring tools used in Unix OS include Process status(ps), top, Xosview and Treeps.

Process Status (ps) is the most widely used performance monitoring tool for Unix based systems. The ps command provides detailed information and performance statistics about running processes such as the names and process IDs as well as the current CPU usage. The ps command is one of the oldest performance measurement tools for Unix systems that is still used widely today, and it has a large number of options and parameters that it can display of which the most useful for process performance monitoring include:

- pcpu - Percentage of the CPU that the process is currently using
- etime - Elapsed time since the process started
- time - CPU time used by the process

Log-based performance monitoring tools: record system performance information for post-processing and analysis the record to find trends in the system performance. These tools have the ability to thoroughly evaluate resource utilization of a system. In addition, they are also used to produce and maintain detailed records of system performance measurements which can be, in turn,used for a number of purposes other than simple data analysis. Since measurements have to be written permanently to the disk, high overhead operation is inevitable. And also measurements that occur very frequently can quickly fill up the available disk space. These two are the main disadvanatges of log-based performance monitoring tools.

Examples of log-based mmonitoring tools in Windows OS are: Performance Logs and Alerts, Event Log Service and Event Viewer, and Performance Data Log Service (PerfLog).

Performance Data Log Service (PerfLog) is a new tool that logs performance data from system counters to a log file and produces comma or tab separated log files. It lets the user select which performance counters to log and the rate to log them at.

PerfLog automatically logs directly to comma or tab separated files for easy import into almost any program used to view logs. It also allows for a number of useful scheduling and automation options. It can be set to start monitoring on system boot up and can start prior to a user logging on. It can also be set up to use a base file name and append a date or automatically incremented number on the end to avoid user interaction or produce logs at a set frequency. Or can also be set to start a new log file based on a set frequency in hours, days, or months or based on a file size in kilobytes or megabytes. In addition, Performance Data Log Service has a number of preset counters that can be selected by choosing a desired level of detail for monitoring. This allows novices at performance logging to select a reasonable data set for logging and provides a good starting point for all users when selecting which counters should be recorded.

Examples of Unix log-based monitoring tools are: System Activity Reporter (sar), Cpustat, Vmstat, and mpstat. Vmstat, for instance is a general system monitoring utility found in Unix operating system which provides a good low-overhead view of system performance. The utility runs in two modes: average and sample mode. The sample mode will measure statistics over a specified interval. This mode is the most useful when understanding performance under a sustained load. The following example demonstrates vmstat running at 1 second intervals.

```
procs -----------memory---------- ---swap-- -----io---- --system-- ----cpu----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in    cs us sy id wa
 0  0 104300  16800  95328  72200    0    0     5    26    7    14  4  1 95  0
 0  0 104300  16800  95328  72200    0    0     0    24 1021    64  1  1 98  0
 0  0 104300  16800  95328  72200    0    0     0     0 1009    59  1  1 98  0
```

Figure 4.3. Output displayed when vmstat utility runs

Field	Description
r	The amount of threads in the run queue. These are threads that are runnable, but the CPU is not available to execute them.
b	This is the number of processes blocked and waiting on IO requests to finish.
in	This is the number of interrupts being processed.
cs	This is the number of context switches currently happening on the system.
us	This is the percentage of user CPU utilization.
sys	This is the percentage of kernel and interrupts utilization.
wa	This is the percentage of idle processor time due to the fact that ALL runnable threads are blocked waiting on IO.
id	This is the percentage of time that the CPU is completely idle.

Conclusion

Performance monitoring of operating system is essential for debugging processes and systems, effectively manage system resources, making system decisions, and evaluating and examining systems. Performance of a system is monitored by taking various metrics and comparing these metrics against a pre-determined or specified statistics the main metrics considered being the CPU utilization.

There are two primary types of performance monitoring tools: those for real time performance monitoring and those for logging performance measurements. Though both kinds of monitoring tools are used to assess system's performance, their difference lies in their performance identification or scheduling scheme.

Assessment

1. What are the main reasons why performance monitoring is needed
2. Can performance monitoring assure effective resource utilization? Why or why not?
3. Among the performance monitoring methods discussed, which one is superior to the other? How?
4. What kind of performance related information can we get through Windows Task Manager? What kinds of tuning decisions can we make accordingly?

UNIT SUMMARY

Security is such a big topic that there is always more to be said. We have included the specific tools that we use or have had recommended to us. You should search the Internet and talk to other system administrators to find the specific products that meet the needs of your situation. Test everything. Be very thorough and always watch your distribution's homepage for security advisories and updated packages.

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

Write your answer for the following questions

1. Install the ipchains package if it is not already installed.
2. Set the default forwarding policy to DENY. Mount this file system on /usr/lib.
3. Filter out all packets coming from www.linux.org.
4. Test access to www.linux.org using ping and HTTP.
5. Remove the filter you placed on www.linux.org.

Grading Scheme

This unit's assessments worth a total of 20% from which

a. Activity Assessments: 5%

a. Unit Assessment I:3 %

a. Unit Assessment II (Lab Project): 12%

Answers

Rubrics for assessing activities

Assessment Criteria	Doesn't attempt both assessment or the answers are all wrong	Attempt both assessments and provide partial answers or partially correct answers	Attempt all assessments with full and correct answer
Grading Scales	Fail	Got half mark for each assessment	Score full mark for each assessment

Unit Readings and Other Resources

- Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2ndEdition, 2002, Chapter 15
- LPI – Linux System Administration, Revision 1.0, 2000, Chapter 9
- [Moore] Moore, Sonia Marie (Ed.). "MS Windows NT Workstation 4.0 Resource Guide." 1995. http://www.microsoft.com/technet/archive/ntwrkstn/reskit/default.mspx.

Unit 5: Special Topics

Unit Introduction

This unit deals with number of topics including application software on Windows and Linux operating systems, server hardware knowledge, and automating routine tasks using scripts. How application software is installed and removed will be addressed on Windows and Linux. Moreover, server hardware knowledge including server hardware, components, and management is addressed.

Unit Objectives

Upon completion of this unit you should be able to:

- Install and remove application software on Windows and Linux
- Schedule tasks to automatically execute
- Manage server racks and blades
- Define a server and distinguish between desktop PC and server
- Differentiate between server platforms

KEY TERMS

Performance: Act of doing something

Monitoring: A process of assessing and observing an activity or action

Tuning: Refining or adjusting an action

Server: Powerful system with redundant processor, Memory, I/O drives and network components

Server Farm: Collection of servers

Data Center: Place where server components are housesd

Disk Drive: Device that is able to read and write on disk.

Memory: A component which is able to store program and data.

Activity 1: Application on Windows and Linux OS

Introduction

Operating systems are responsible for application software execution. Prior to launching, application software needs to be installed. In case the user doesn't need the application anymore, removal of application can be done.

Installing Application on Windows

Adding of a program to your machine depends on where you put the installation files. Programs can be installed from a CD or DVD, from the Internet, or from a network.

To install a program from a CD or DVD

Insert the disc into your computer, and then follow the instructions on your screen. " " If you're prompted for an administrator password or confirmation, type the password or provide confirmation.

Many programs installed from CDs or DVDs open an installation wizard for the program automatically. In these cases, the AutoPlay dialog box appears and you can choose to run the wizard.

If a program doesn't begin installation automatically, check the information that came with the program. This information will likely provide instructions for installing the program manually. If you can't access the information, you can also browse through the disc and open the program setup file, usually called "Setup.exe" or "Install.exe".

To install a program from the Internet

1. In your web browser, click the link to the program.
2. Do one of the following:

- To install the program immediately, click "Open" or "Run", and then follow the instructions on your screen. " " If you're prompted for an administrator password or confirmation, type the password or provide confirmation.
- To install the program later, click "Save", and then download the installation file to your computer. When you're ready to install the program, double-click the file, and then follow the instructions on your screen. This is a safer option because you can scan the installation file for viruses before you proceed.

To install a program from a network

If you connect your computer to a network that has programs that you can add, you can install programs from Control Panel.

1. Open Get Programs by clicking the Startbutton , clicking Control Panel, clicking Programs, clicking Programs and Features, and then, in the left pane, clicking Install a program from the network.
2. Click a program in the list, and then click Install.
3. Follow the instructions on your screen. If you're prompted for an administrator password or confirmation, type the password or provide confirmation.

Removing Application from Windows 10

You can uninstall programs on the Settings page. To repair a program, use Control Panel.

1. On the Start menu select Settings.
2. In Settings, select System>Apps & features.
3. Select the program, and then select Uninstall.
4. Follow the directions on the screen.

Installing Application on Ubuntu

Most Windows users who migrate to Ubuntu end up confused about software installation. They go to a website, download a .tar.gz file, double-click it, and don't see a Next-Next-Next-Finish wizard.

Rather than leaving it up to the user to track down installer files and keep applications updated, Ubuntu (like many other Linux distributions) has a software package management system that provides a searchable database of easily installable applications (like an online shopping cart but the software is cost-free), which it will download and install for you with a few clicks.

Ubuntu Software Center

The best place to start with this package management process is to use a simple interface for it called Ubuntu Software Center.

You can most likely find Ubuntu Software Center on the left side of your screen.

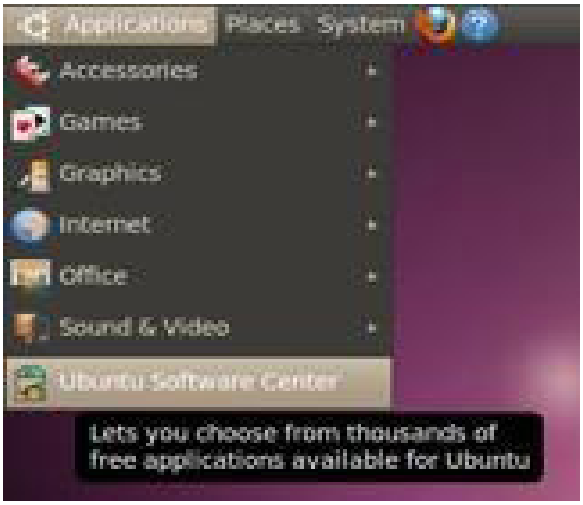

Go to Applications and select Ubuntu Software Center. If you already know what software you're looking for, you can begin typing the name of it in the top-right corner to begin the filtering process. If you don't know, you can also browse by category.

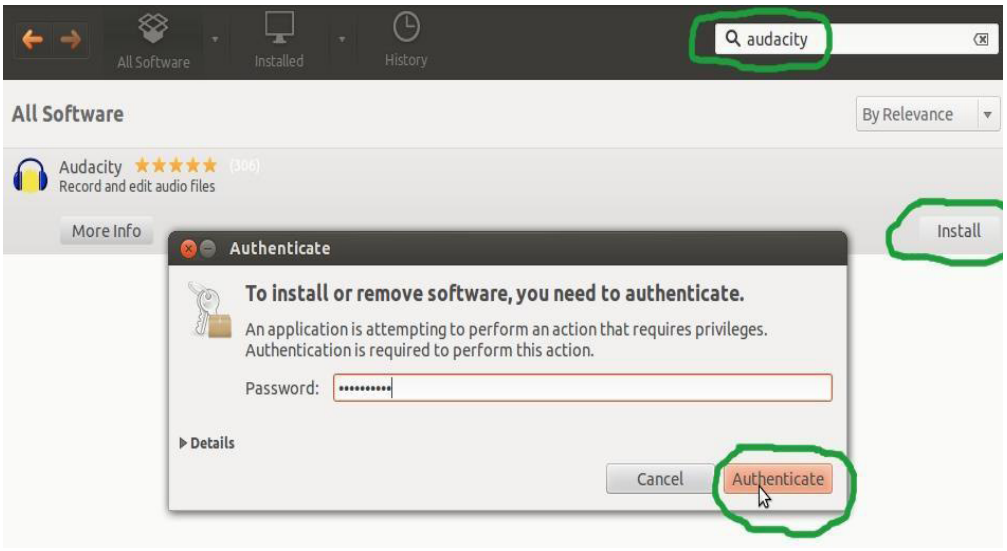

In this case, let's say you're looking for an audio recording and editing program called Audacity. So after finding it, click Install and get prompted to authenticate with your password.

Then wait for Audacity to be downloaded and installed. Depending on the size of the software you're installing and the speed of your internet connection, this could take a few seconds or several minutes to complete.

Once it's finished installing, it's ready for you to use.

Manual download: .deb

Most of the time, if you need software, you can use the Software Center to install it from the online repositories. Sometimes, software is not available in the repositories, and you have to go to a website to download it. If you are able to download a file with a .deb extension, this is the software package format Ubuntu prefers.


Download the latest version o...
www.skype.com/intl/en-us/get-skype/on-your-computer/linux/
English (Amer
skype
Features
Get Skype
Prices
Accessories
Busi
Skype 2.2 Beta for Linux
Test drive the new beta of Skype for Linux with improved audio and video quality.
Download now |
Select your linux distribution
Ubuntu 10.4+ 32-bit
Ubuntu 10.4+ 64-bit

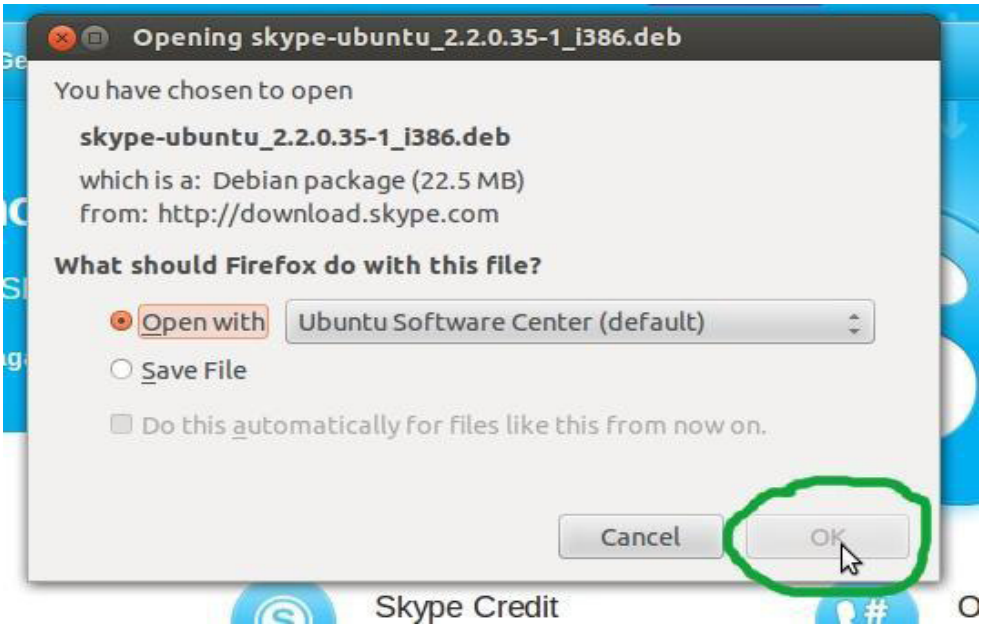
Opening skype-ubuntu_2.2.0.35-1_i386.deb
You have chosen to open
skype-ubuntu_2.2.0.35-1_i386.deb
which is a: Debian package (22.5 MB)
from: http://download.skype.com
What should Firefox do with this file?
Open with
Ubuntu Software Center (default)
Save File
Do this automatically for files like this from now on.
Cancel
OK
Skype Credit

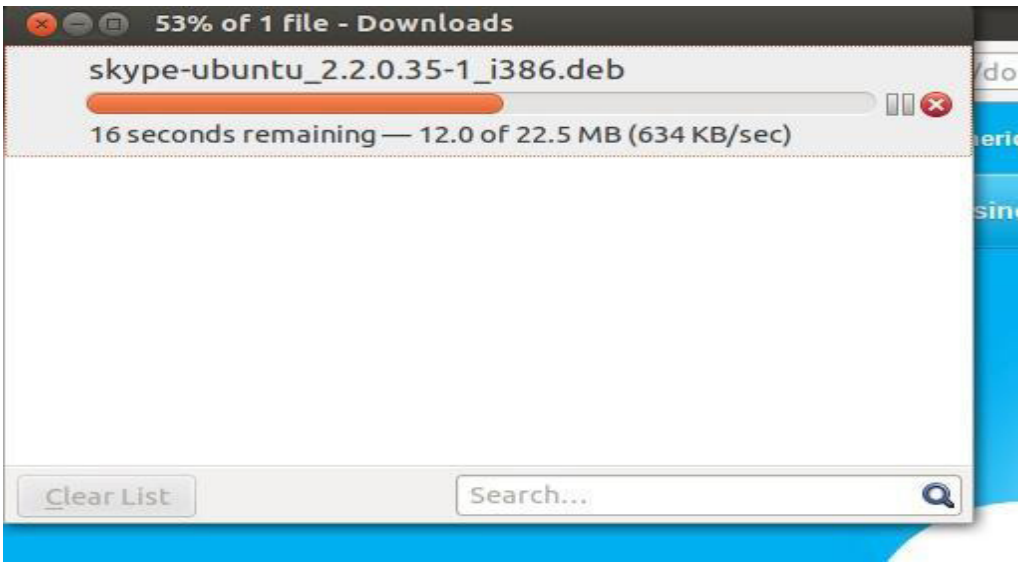

After you download the .deb, it should run in the Ubuntu Software Center by default. You can run it directly in that program (you can also download the .deb file and double-click it to launch the installer program).

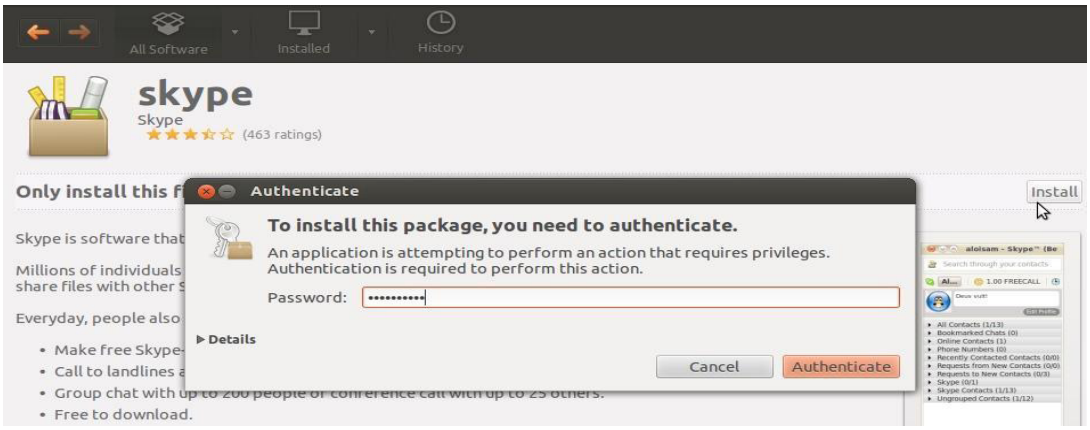

After Ubuntu Software Center loads, click Install to install the .deb file. You'll be prompted for your password. Go ahead and enter it if you trust the source you downloaded the file from.

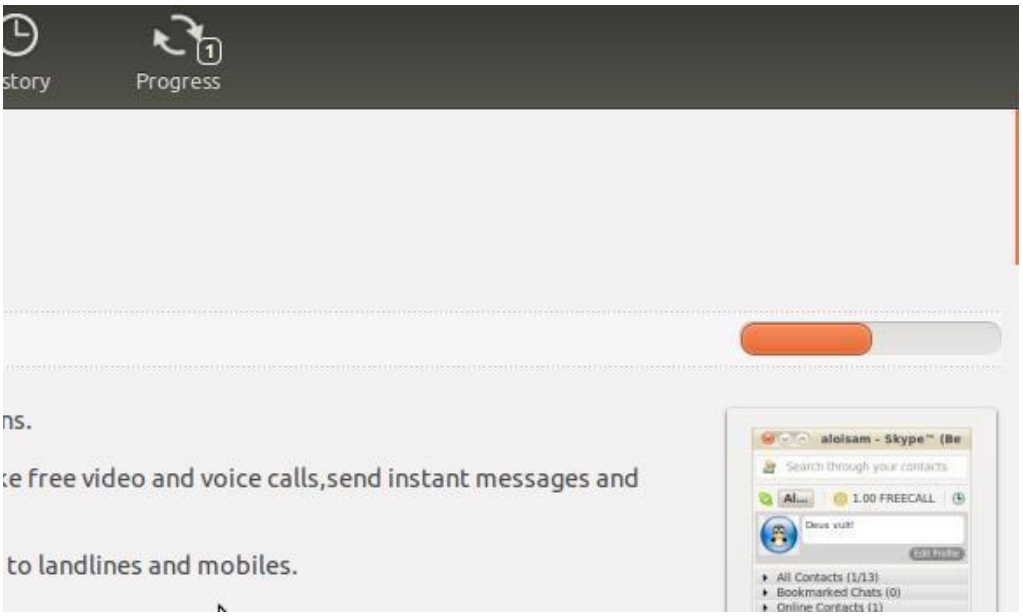

Pretty soon your application should be ready to use.

Last resorts: .rpm and .tar.gz

The preferred way to install software in Ubuntu is to use the package manager, which you can access through Ubuntu Software Center. As we've seen with Skype, sometimes you can also find a .deb for software not in the repositories. But what if you can't find a .deb?

.rpm

If you can't find a .deb, you can try a .rpm. These files are packaged for other Linux distributions (usually Fedora or Mandriva), but there is an application called alien(which you can install using Synaptic) that allows you (most of the time) to convert .rpm files to .deb. Read more about this process.

.tar.gz

As a last resort, you can download a .tar.gz file. The .tar.gz file extension indicates the file is a compressed set of files and folders (the compressed files you see in Windows usually have a .zip extension). If you see the .tar.gz, it could be compressed files that have a precompiled binary file, or it could be compressed files that have the source code allowing you to compile the application from source.

Uninstalling Application on Ubuntu

You are a Linux user (Ubuntu). You always install applications from the terminal using apt-get install. Even though you can add and remove applications using GUI (Ubuntu Software Center), it is always said that the command line is power. After installing a whole lot of applications both small and large in size. You now need disk space for other applications or you no longer want to see such applications on your system.

Most Ubuntu users, beginners specifically can install from the terminal after searching for how to install it. It becomes very difficult removing it. I also went through this.

Open Terminal (ctrl + alt + t)

Type

```
$ dpkg --list
```

(this displays all installed packages at a go)

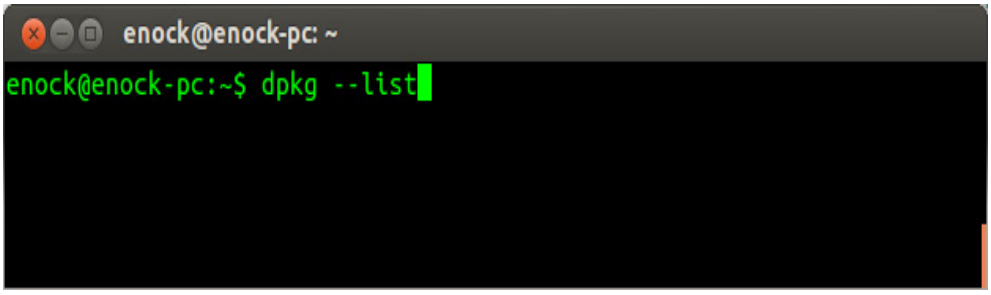

OR

```
$ dpkg --list | less
```

(to easily navigate using the user keys)

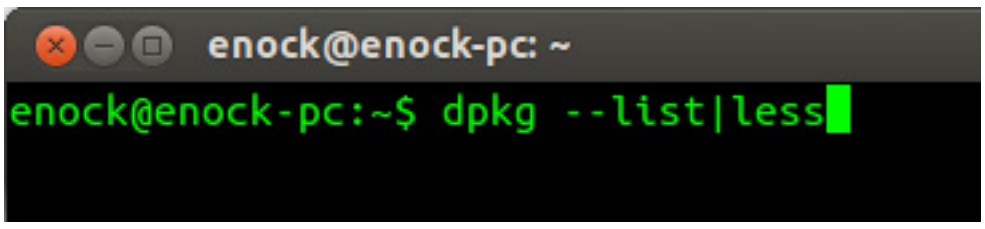

OR

If you know the package name, you can pipe it with grep command to locate it using the syntax below:

$ dpkg --list|grep -i ‘packagename’

Example for VLC player.

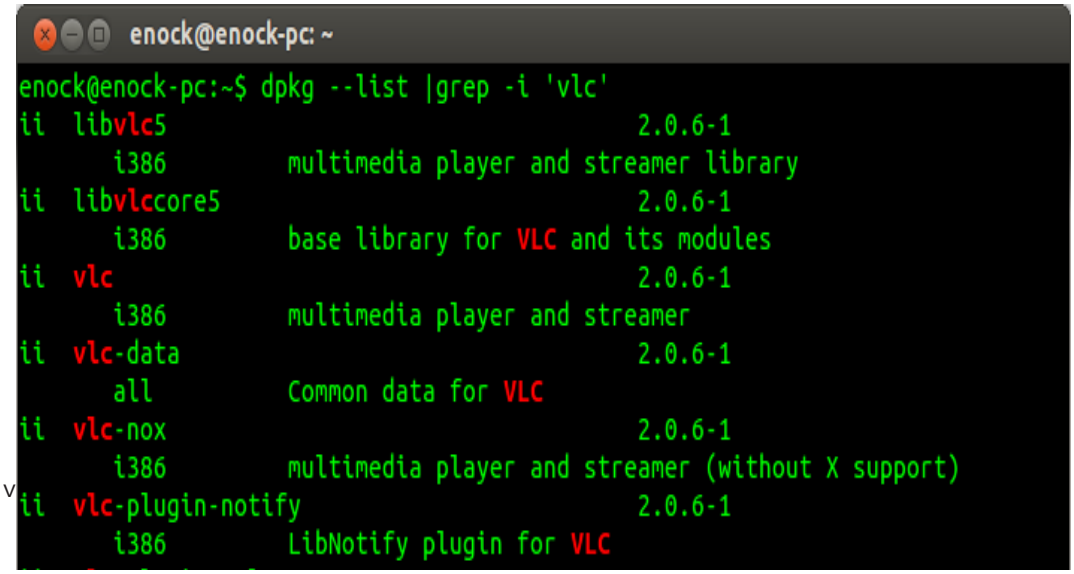

Locate the name of the package to be removed. Here am going to usesox.

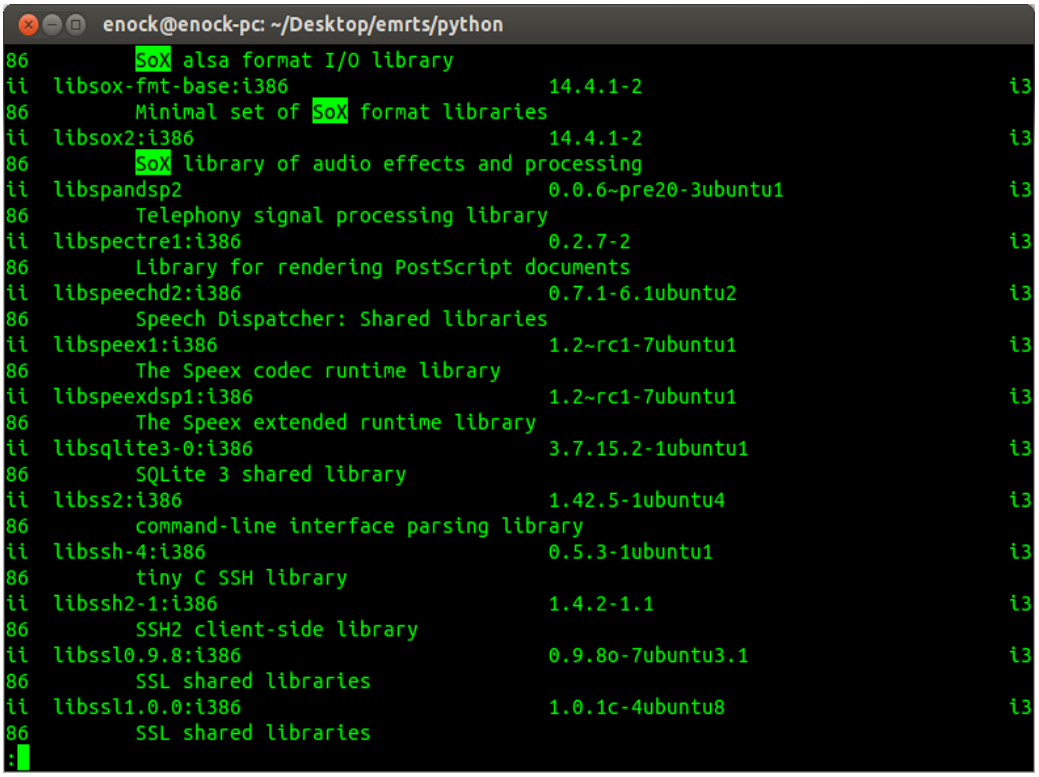

4. To remove only a package use:

$ sudo apt-get remove <packagename>

For example to remove package called sox, enter:

$ sudo apt-get remove sox

Type y for Yes to uninstall.

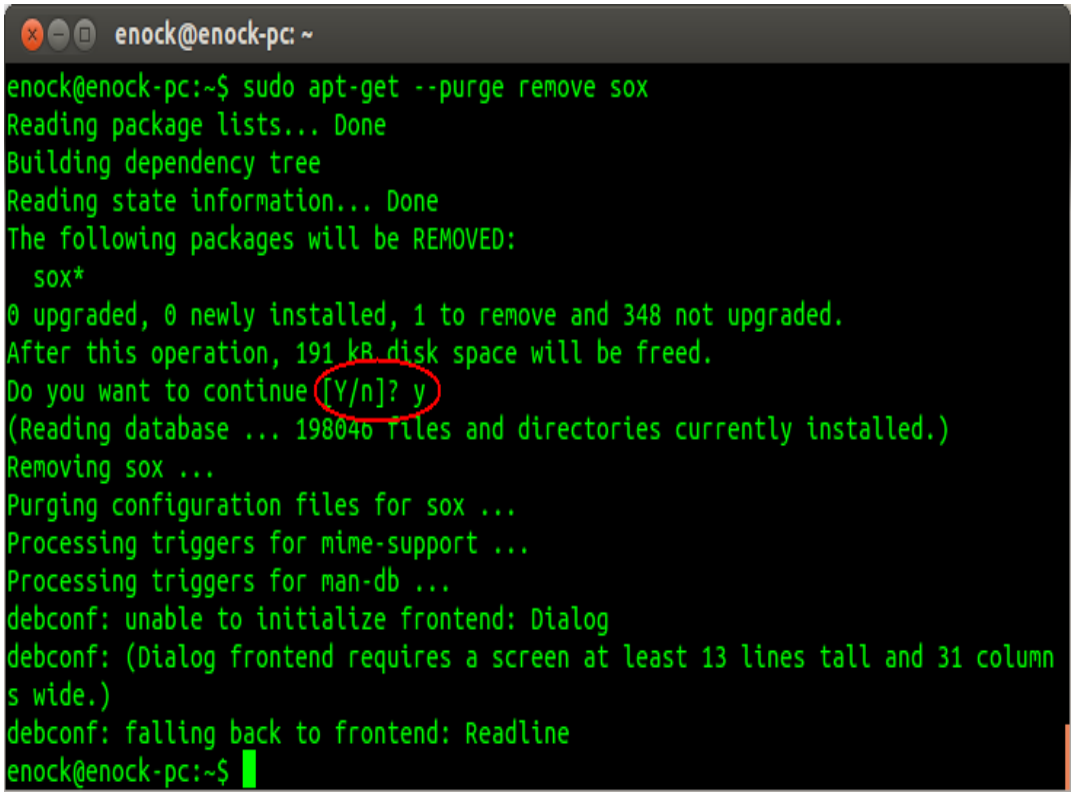

To remove a package with it's configuration files, enter:

$ sudo apt-get --purge remove <packagename>

For example removing a package called sox and all configuration files:

$ sudo apt-get --purge remove sox

Conclusion

In this activity we discussed how to install and remove application on Windows and Linux operating systems. There are different ways to install operating system i.e., it can be using graphical user interface tools or command based.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

- Download and install vlc player from the Internet on Windows operating system.
- Use the vlc player you installed and play a video,
- Uninstall vlc player.
- Install openshh server on Ubuntu operating system using the Ubuntu Software Center.

Activity 2: Server Hardware Knowledge (rack mount and blades)

Introduction

Today, most businesses are shifting from paper-based to digital processing and information management. This puts a larger requirement to the computing industry where large, efficient and optimal devices as well as software must be provided to clients. One of the major resource required being a server, which is a high performance, high capacity machine; the computer server industry is in the midst of major change stimulated by increasing demand for data processing and storage. We will explore the general feature of servers and move our discussion to server hardware knowledge.

Server Hardware

A computer server is a hardware device connected to a network or software whose purpose is to manage networked resources.

Traditionally, computer server hardware has been dedicated to manage a single functional purpose. The dedicated server functions include: application servers, file servers, web servers, mail servers, print servers, database servers, and many more. Several servers are typically required to enable a computer to properly interact with other network clients due to this dedicated nature of a server.

A collection of servers is referred to as a server farm or server cluster and the facility used to house the server farm and associated components is referred to as a data center. To keep pace with increased server space, the traditional data center has evolved to include cooling equipment, network equipment, and storage equipment.

Even though, the server hardware can range widely in size, performance, cost, capability, and environmental impact, generally, they will contain similar hardware components

Typically, server components include:

- External enclosure or case
- Central processing unit (CPU)
- 1-4 CPU sockets
- Main motherboard
- Memory
- Storage (hard drives, solid state drive (SSD))
- Input / Output adaptors
- Fans
- Power supplies
- and may include a small screen

Figure 5.1 and 5.2 below depicts the front and back view of the x3650 IBM server respectively.

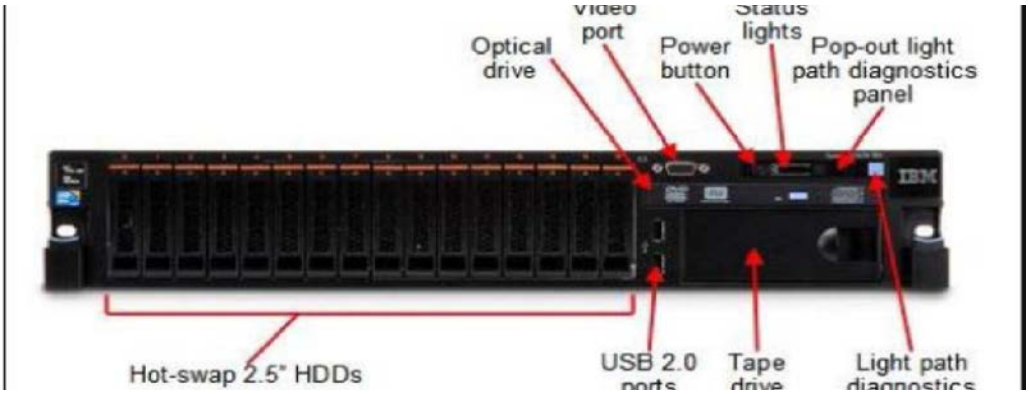

Figure 5.1. IBM's X3650 system front view

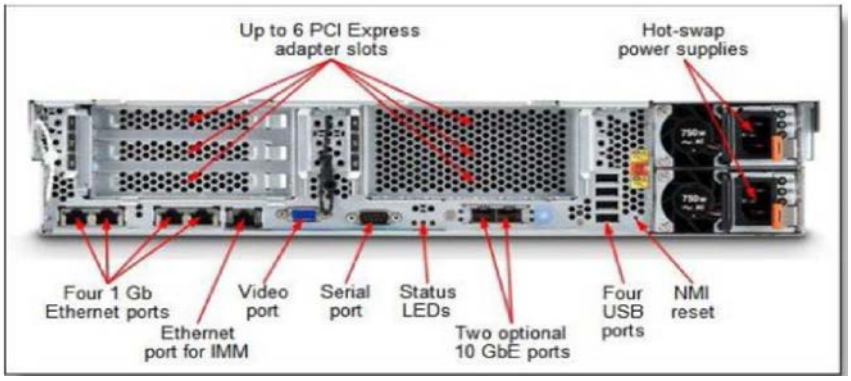

Figure 5.2. IBM's X3650 system back view

In many ways server hardware is similar to desktop PC hardware as both have same basic components such as memory, CPUs, and power supplies. However, in server hardware, the individual components tend to be more advanced than the components that are found in desktop computers. Servers may also make use of redundant and / or hot swappable components.

Server form factors

Servers are designed to run for long periods without interruption, which makes hardware reliability and durability extremely important. To be able and accomplish this requirement server will have hardware redundancy such as: Power supplies, Fans and Hard drives. Servers typically contain a higher number of computer fans or water cooling systems to remove heat. Factors to be considered before choosing particular server platform are what we call them server form factors. These factors can be:

- Storage capacity
- Processing power
- Available interface
- Network connectivity

Considering these factors, there are several different platforms of servers, but the three most common are tower servers, rack servers and blade servers.

Tower Servers

Tower servers look a lot like PCs. Each tower server is a standalone machine that is built into an upright case. Tower servers are used mostly in smaller businesses with 25 or fewer employees and that require only one or two servers. For instance, if a business needs a file server to support a growing network of users, installing a tower server is an ideal first step, especially when cost is a priority. Larger datacenters typically avoid the use of tower servers because of the amount of physical space that they consume and because they tend to be noisy. Another disadvantage to using tower servers is that the cabling can get messy. But, tower servers are also known for their easy setup as they are almost similar to a pc. Examples include the PowerEdge (T320, T420 and T620) range from Dell.

Figure 5.3. PowerEdge T110 II Tower server

Rack Servers

As the name implies, rack servers are servers that are mounted within a rack. The rack is of a uniform width and servers are mounted to the rack using screws. Each rack can accommodate multiple servers and the servers are typically stacked on top of each other. These servers are used in medium sized businesses where more than two servers are needed. Rack servers can be stacked on top of one another and accessed by the rail system, which is similar to a cabinet drawer reducing the footprint the server leaves in the datacenter.

A rack will need specific installation hardware, but this gives us the flexibility to design just the server platform our business needs now, and more importantly, that can be easily expanded at a later date. Rack mount components follow a form factor that is referred to as a rack unit. A standard rack mount server is referred to as a 1U server meaning that it is 1 rack unit in size. A 2U server consumes two rack units of space within the rack.

Some vendors also offer 4U and ½U servers. The larger form factors are usually used when the server needs to be able to accommodate a large amount of storage. Examples of rack servers include the Dell PowerEdge R320, R420 and R520.

Figure 5.4. Dell PowerEdge Rack Servers

Blade Servers

Like rack servers, blade servers also adhere to a standard size and mount inside a special "rack". In the case of a blade server however, the rack is known as a chassis. Blade servers move many of the components that you would typically see on a tower or rack server into what is called the chassis - a mount into which the blades fit. Blade servers tend to be vendor proprietary. You can't for example insert a Dell blade server into an HP chassis. The reason why blade server design is proprietary is because unlike a rack server, which is fully self-contained, blade servers lack some of the components that they need to function. For example, blade servers do not have power supplies. The blade server chassis is designed to accept various modular components, including the blade servers themselves.

For example, a chassis might contain a power supply unit, a cooling unit, and a blade server. The idea is that the power supply, input/output interfaces, cooling and networking are all independent, and the blades contain a processor and storage, which can be easily upgraded. The actual chassis design varies from one vendor to the next, but most blade server chassis are designed to accommodate multiple power supplies, multiple blade servers, and a variety of other components (such as network adapters, storage modules, and cooling modules). With the exception of the cooling components, individual blade servers are mapped to the individual modules or components. An example for blade servers is the Dell PowerEdge M series.

Figure 5.5. Dell PowerEdge M Blade Server

Server Components

Let's briefly discuss the components of a server.

Server Case Designs

After our discussion on server platforms, let's move our focus to some important components of a server case. These components, though, not crucial parts of a server as memory or processor is, they play significant role in affecting the proper functionality of a server.

Cooling Features

Computers must be kept cool in order to function properly. It's a much undesired situation when a system's processor heats. Servers use devices such as heat sinks and CPU fans to keep components cool, but the case design also plays a vital role in regulating the server's overall temperature. Blade servers however, are the exception. The blade server chassis is responsible for cooling the individual blade servers. For rack mount and tower servers however, the case is specifically designed to keep the server cool.

The actual cooling features that are found in a server's case vary from one server to the next. At a minimum, a server typically has at least a couple of case fans. The fans are arranged in a way so that cool air is pulled into the server through vents. This cool air passes over the individual system components and is then exhausted from the system's case. Servers are also typically equipped with temperature sensors. Sometimes these sensors are integrated into the system board or into individual components, but they may also be built into the case. If the temperature inside the case begins to rise then the fans will spin faster in an effort to keep the server operating at a safe temperature. The system BIOS usually also has a safety mechanism that can shut down the server before damage occurs if the temperature gets too high.

Filtering

Sometimes fans alone are not enough to keep a server cool. In industrial areas a server's fans can suck in dust, dirt, or other contaminants. These particles can block air vents and may also prevent fans from operating efficiently. In order to prevent this type of contamination, some (but certainly not all) server cases are equipped with built in air filters. Air filters do a good job of keeping stray particles of contaminants out of a server case, but they require routine maintenance. Air filters are usually washable and must be cleaned on a regular basis. A dirty air filter can actually restrict air flow just as badly as if the server's vents were completely obstructed.

Status Indicators

PCs and servers use LED lights as a various forms of status indicators though the server indicators outweigh that of PCs. Some servers have locate light indicator which is used to identify the machine we are managing within a datacenter. This indicator is used in large datacenters with several servers as it might be difficult to spot which physical machine we are connecting to. In addition to the status indicator LEDs, there are a few server cases that have integrated thermometers that visually display the temperature inside the case.

Ports

In most datacenters, servers are only accessed remotely. Even so, servers typically include keyboard, video, and mouse ports. These ports can be used during the initial setup or in those rare situations in which it is necessary to perform diagnostics locally. In the case of blade servers, individual blade servers do not include keyboard, video, and mouse ports. These ports are provided through the blade server chassis. USB ports are also included in servers.

Drive Slots

Most servers contain at least one or two drive bays or inlets. The biggest determining factor in the number of drive bays that a server can accommodate is its form factor. For example, a 1U server might only have one or two drive bays, while some 4U servers offer dozens of drive bays.

Swing Out Bezels

Another set of features that are often found on rack mount servers are swing out bezels. Swing out bezels allow us to slide the entire server forward and open the server's case without having to remove the server from the rack..

Server Processors

In general, processor selection will be based on the server role and server operating system that will be installed. But most of today's server come as a Multi-core processor which is a single chip containing more than one microprocessor core. This arrangement multiplies performance with the number of cores. Operating system must be designed to support and recognize multiple cores or processors.

Network Components

Network interface controller (NIC) is a hardware device that connects a server's interface to a computer network and allows a network-capable device to access that network. A NIC is both a physical and data link layer device on the OSI Reference Model. A NIC is typically available in 10/100/1000 Mbps. Modern servers typically are configured with at least two 10/100/1000 Mbps adapters. It is also possible to team together multiple server NICs to perform as one, thus allowing for twice the throughput as a single NIC.

Server Memory

When it comes to server memory, two main concepts need to be understood NUMA and parity

NUMA, Non Uniform Memory Access, is a memory included in a server to overcome the waiting time CPUs encounter while trying to access the main memory. As modern servers come with more than one CPU all of which need to interact with the memory, processor has to wait in line to get access to memory and this wastes CPU's time hugely. One of the ways in which server manufacturers have attempted to reduce the amount of "get and wait" that a CPU experiences is by introducing NUMA memory. NUMA memory is separated into individual compartments which are known as NUMA nodes. That way, a separate NUMA node can be allocated to each processor. In most cases there is a one to one relationship between physical CPU sockets and NUMA nodes, regardless of how many cores exist within a physical CPU. Dedicating a NUMA node to each of the server's CPUs helps to improve performance in two ways. First, the use of separate NUMA nodes helps to prevent multiple CPUs from competing for memory access and are able to access memory simultaneously because each CPU is accessing a separate NUMA node. The other way in which NUMA improves performance is by limiting the amount of memory that is available to a single CPU. If a CPU is using a dedicated NUMA node exclusively then it is dealing with a smaller amount of memory than what the system is equipped with as a whole making memory management more efficient though it also has several other problems.

Parity

Internal VS. External Storage

A server's physical size often determines the amount of physical storage that the server can accommodate. Many organizations choose to use 1U servers connected to external storage because it is often less expensive to connect several 1U servers to an external storage pool than it is to purchase several, fully loaded 4U servers. External storage is often used as an alternative to internal storage for performance reasons.

Server Hard Drive Types

Most of the servers that are being sold today can either accommodate SATA hard drives, like most PCs, or SAS hard drives. The SATA hard drives that are sold for server use are actually very similar to desktop hard drives in that they use a similar controller, and are 3.5 inches in size. But the cost of server hard drive is way expensive than that of desktop hard drives. It is possible to use, thus, desktop hard drive to a server. But the duty cycle of the desktop hard drive is less than the server one as server hard drives are specifically designed for much heavier use.

Solid State Drives

Solid State Drives (SSDs) use semiconductor memory such as NAND flash memory instead of the magnetic platters found in traditional hard drives. The memory used in SSDs is very similar to the type of memory found in USB flash drives and in SD cards. SSD drives offer a number of advantages over traditional hard drives (HDDs). Most of these advantages are a direct result of the fact that unlike a HDD, an SSD does not contain any moving parts. One advantage is their less power consumption and less heat dissemination. Another advantage of SSD is that its silent operation. The third and main advantage to SSD is its better performance and reliability. The reason why SSDs offer such high performance is because the drive's speed is not limited by mechanical components. The operating system will never have to wait for a SSD drive to spin up and there is no waiting for a mechanical drive head to move across the platter to the required location. The reliability comes from the no moving parts and there aren't any mechanical mechanisms that can wear out. However, in spite of the many advantages to using solid state drives, there are at least two major disadvantages that must be taken into account – price and capacity.

Hybrid Drives

Most of the SSDs on the market use a SATA interface. This allows SSDs to be used in any computer that has a compatible SATA controller. However, the SATA interface was originally designed for use with HDDs and it is the SATA interface that is usually the performance bottleneck when SSDs are used. Which has led to the creation of hybrid SSDs. A hybrid SSD does not use a SATA interface. Instead, the drive is mounted directly on a PCIe card. This allows for much greater performance since the PCIe bus is faster than the SATA bus. The reason why hybrid SSDs are called hybrid drives is because the PCI card makes use of both SSD and HDD storage. The basic idea is that data which is frequently accessed is stored on the SSD portion of the drive so that it can be retrieved quickly. Data that is not used as often is written to a standard HDD that is mounted directly onto the PCIe card. This design not only offers high performance, but it also allows for greater storage capacity than what is currently possible with a SSD drive.

Serial Attached SCSI, SAS

SAS is generally the preferred storage medium for servers today. SAS drives perform well, are relatively inexpensive, have high capacities, and are almost universally accepted. The fact that SAS makes use of a serial bus means that SAS not only supports more devices than legacy SCSI, but also achieves much higher throughput. One of the big limitations of the parallel SCSI bus was that all of the devices on the bus shared the bus's bandwidth. In a SAS implementation, each device has a dedicated connection to the SCSI Initiator (usually the controller card). As such, SAS devices are not impacted by the bandwidth that is consumed by other devices. In fact, SAS devices typically support a transfer speed of either 3 Gbps or 6 Gbps. Another advantage to the serial bus is that because only a single device is connected to each SAS port, there is no need for termination.

Server Management

Being able to effectively manage servers is critical in any data center. Management capabilities are built into the server hardware by vendors but we also need server management software since servers are very complex and management needs to be performed on many different levels. The server hardware level monitoring features are used to monitor the hardware's well-being while the software management tools are used to monitor the software level health. Example of hardware management tools are: HP's HP Systems Insight Manager and Dell's OpenManage.

When we talk about hardware based management, it is not a must for the system be functional, like the case in software management tools. In fact, some forms of hardware management do not even require that the server be turned on. As long as the server is connected to a power source it can be managed at the hardware level. Hardware management requires management software that is able to interact with the server at the hardware level.

There are two main standards for hardware based management – IPMI and BMC. Most servers support one management standard or the other, but some servers support both standards.

Servers that support IPMI include a dedicated management port. This port is simply a low bandwidth network port. In order to manage such a server, an administrator connects to the server through this port from another machine that is running IPMI compliant management software. Once the connection has been established, the management software is able to assess the server's health by examining various universal status codes. IPMI management software works with any IPMI compliant server regardless of the server's manufacturer.

One thing we need to note regarding hardware server management is that, although hardware level management is important, its capabilities are somewhat limited. They mainly detect low level problems and do not offer advanced management capabilities.

Conclusion

A server can be either hardware or software which is used to manage network resources of a system. Server differs from a desktop pc in its performance capacity and redundancy of resources it utilize. Server platform to be used in an organization must be chosen based on assessment of the business' current needs and future expansion. There are three common server platforms each with its own pros and cons. Some of the components of a server are server case, processor, memory, storage and network components.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. What is a server farm?
2. What are datacenters?
3. What advantages are there when rack server is used?
4. What is the advantage of using server hardware management tools?
5. Discuss the main differences between magnetic drives and SDDs?

Activity 3: Automating Routine Tasks Using Scripts

Introduction

Cron is a tool for configuring scheduled tasks on Unix systems, used to schedule commands or scripts to run periodically and at fixed intervals. Tasks range from backing up the user's home folders every day at midnight, to logging CPU information every hour. In this activity, we will discuss how to use cron from the command line and how to read its configuration file.

What is Cron?

Cron is a scheduling utility that allows you to assign tasks to run at preconfigured times. A basic tool, cron can be utilized to automate almost anything on your system that must happen at regular intervals. Equally adept at managing task that must be performed hourly or daily and large routines that must be done once or twice a year, cron is an essential tool for a system administrator.

How Cron Works

Cron is started at boot and runs in the background as a daemon. This means that it runs without user interaction and waits for certain events to happen to decide when to execute. In the case of cron, these events are certain moments in time. Cron runs in the background and checks its configuration file once every minute to see if an event is scheduled to run that minute.

If an event is scheduled, cron executes whatever predetermined command has been given to it and then goes back into the background for another minute. If no event was scheduled, it waits 60 seconds and checks again.

Because of this minute-by-minute scheduling, it is extremely flexible and configurable. Upon installing your distribution, cron is already configured to run a variety of tasks.

How to Read a Crontab

Cron decides which commands to run at what time by reading a series of files, each known as a "crontab". We can see the system-wide crontab by looking at "/etc/crontab":

```
less /etc/crontab
SHELL=/bin/sh
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin
# m h dom mon dow user  command
17 *    * * *   root    cd / && run-parts --report /etc/cron.hourly
```

```
25 6    * * *   root    test -x /usr/sbin/anacron || ( cd / && run-parts --report
    /etc/cron.daily )

47 6    * * 7   root    test -x /usr/sbin/anacron || ( cd / && run-parts --report
    /etc/cron.weekly )

52 6    1 * *   root    test -x /usr/sbin/anacron || ( cd / && run-parts --report
    /etc/cron.monthly )
```

This is the system crontab and should not be edited in most cases. Most of the time it is preferable to use your own crontab. The system file could be replaced in an update and your changes would be lost.

The file has a few interesting parts that we need to understand.

The first two lines specify the shell that will execute the commands listed and the path to check for the programs.

The rest of the file specifies the actual commands and scheduling. The lines in this list each represent a record, or row, in a table. The "tab" in "crontab" stands for table. Each table cell is represented by a column separated by spaces or tabs.

The commented line above the table gives a hint as to what each of the columns represent:

```
# m h dom mon dow user  command
```

Scheduling Hours and Minutes with Cron

The first column is the minute (0-59) of the hour that the command should run. The second column is the hour of the day, from 0-23, that it should run. An asterisk (*) means "every possible value", and is used as a wildcard.

By combining these first two columns, we can get a time value for the command. For instance, the second line in the table has 25 in the minutes column and 6 in the hours column:

```
25 6    * * *   root    test -x /usr/sbin/anacron ll (cd / && run-parts --report /etc/cron.daily)
```

This means the second line should be run at 6:25 in the morning.

Similarly, the first line means that the command should be run every hour, at 17 minutes passed the hour:

```
17 *    * * *   root    cd / && run-parts --report /etc/cron.hourly
```

So it will be run at 1:17am, 2:17am, 3:17am, etc.

Scheduling Days with Cron

The third, fourth, and fifth column determine which days the command should be run. The third column specifies a day of the month, 1-31 (be careful when scheduling for late in the month, as not all months have the same number of days).

The fourth column specifies which months, from 1-12, a command should be run, and the fifth column is reserved to specify which day of the week a command should be run, with 0 and 7 both meaning Sunday. This last one allows you to schedule by week instead of by month.

If both the day of the week and the day of the month columns have values that are not wildcards, then the command will execute if either of the columns match.

Days of the week and months can also be specified with the first three letters of their name. We can also use ranges with hyphens (-) and select multiple values with commas (,).

We can also specify an interval by following a value with / and a number. For instance, to execute the command every other hour, we could place "*/2" in the hours column.

If we look at the crontab, you will see that the third record is run every Sunday at 6:47am:

```
47 6    * * 7   root    test -x /usr/sbin/anacron || ( cd / && run-parts --report /etc/cron.weekly )
```

The fourth record is run on the first of the month at 6:52am:

```
52 6    1 * *   root    test -x /usr/sbin/anacron || ( cd / && run-parts --report /etc/cron.monthly )
```

Using Time Shortcuts to Schedule

You can replace the first five columns of each record with a named shortcut if you have simple requirements. The syntax for these is "@" followed by the named interval.

For instance, we can schedule something to be executed every week by specifying "@weekly" instead of creating the five column configuration. Other choices are "@yearly", "@monthly", "@daily", and "@hourly".

There is also a special shortcut called "@reboot" which runs as soon as cron is started. This usually only happens when the system starts, which is why it is called "reboot" instead of "cron-restart" or something similar.

Keep in mind that these shortcuts do not provide fine-grained control over when they are run. They are also all configured to run at the first possible moment of the matching time.

For example, "@monthly" will run at midnight of the first of the month. This can lead to many commands scheduled to run at one time if they all fall on the same time. You are unable to stagger these events like you can with the conventional scheduling syntax.

Specifying Commands and Users with Cron

The next columns involve the actual execution of the commands scheduled.

The sixth column, which is only present in the system crontab that we are looking at, names the user that the command should be executed as.

The final column specifies the actual command that should be executed. The command can contain a percent sign (%), which means that everything beyond the first percent sign is passed to the command as standard input.

Every record needs to end with a new-line character. This is not a problem for most entries, but be sure that you have a blank line after the final entry, or else the command will not run properly.

Using run-parts and Cron Directories

If you look at the commands specified in the system crontab, you will see a mention of "anacron", which we will discuss later, and "run-parts".

The run-parts command is a simple command that runs every executable located within a specified directory. It is used extensively with cron because it allows you to run multiple scripts at a specified time by placing them in a single location.

This has the advantage of allowing the crontab to be kept clean and simple, and allowing you to add additional scripts by simply placing them in or linking them to the appropriate directory instead of adjusting the crontab.

By default, most distributions set up folders for each interval, where they place the scripts or links to the scripts that they would like to run at that interval.

For instance, Ubuntu has folders named "cron.daily", "cron.hourly", "cron.monthly", and "cron.weekly". Inside of these folders are the appropriate scripts.

Using User Specific Cron tabs

Now that you understand the syntax of cron, you can use it to create scheduled tasks for your own user. We can do this with the "crontab" command.Because the commands in your crontab will run with your user privileges, the "user" column does not exist in user-specific crontabs.

To see your current crontab, type:

```
crontab -l
```

You will probably not have one unless you've specifically created it by hand. If you do have a crontab, it is best to backup the current copy before editing so that any changes you make can be reverted.

To store your backup in a file called "cron.bak" in your home directory, run:

```
crontab -l > ~/cron.back
```

To edit your crontab, type:

```
crontab -e
```

```
no crontab for demouser - using an empty one
```

Select an editor. To change later, run 'select-editor'.

```
1. /bin/nano  <---- easiest
2. /usr/bin/vim.basic
3. /usr/bin/vim.tiny
```

You might be given a selection prompt similar to the one above your first time using this command. Select the editor you prefer to continue.

You will be dropped into a commented file that you can edit to create your own rules.

As a nonsensical example, if we wanted to echo the date into a file every 15 minutes every Wednesday, we could place this line into the file:

```
*/15 * * * 3 echo "$(date)" >> /home/demouser/file
```

We can then save the file and now, when we run "crontab -l", we should see the rule we just created:

```
crontab -l
*/15 * * * 3 echo "$(date)" >> /home/demouser/file
```

If you need to edit the crontab of a specific user, you can also add the "-u username" option. You will only be able to do this as root or with an account with administrative privileges.

For instance, if you would like to add something to the "root" crontab, you could issue:

```
sudo crontab -u root -e
```

Using Anacron with Cron

One of cron's biggest weaknesses is that it assumes that your server or computer is always on. If your machine is off and you have a task scheduled during that time, the task will never run.

This is a serious problem with systems that cannot be guaranteed to be on at any given time. Due to this scenario, a tool called "anacron" was developed. Anacron stands for anachronistic, and it is used compensate for this problem with cron.

Anacron uses parameters that are not as detailed as cron's options. The smallest increment that anacron understands is days. This means that anacron should be used to complement cron, not to replace it.

Anacron's advantage is that it uses time-stamped files to find out when the last time its commands were executed. This means, if a task is scheduled to be run daily and the computer was turned off during that time, when anacron is run, it can see that the task was last run more than 24 hours ago and execute the task correctly.

The anacron utility has a scheduling table just like cron does. It is appropriately named "anacrontab" and is located in the "/etc" directory as well. Let's see how it is formatted:

```
less /etc/anacrontab

# /etc/anacrontab: configuration file for anacron

# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh

PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin

# These replace cron's entries

1       5       cron.daily      nice run-parts --report /etc/cron.daily

7       10      cron.weekly     nice run-parts --report /etc/cron.weekly

@monthly        15      cron.monthly nice run-parts --report /etc/cron.monthly
```

We can see that it follows a similar format to the "crontab" files, but there are fewer columns and some noticeable differences.

The first column specifies how often the command should be run. It is given as an interval in days. A value of "1" will run every day, while a value of "3" will run every three days.

The second column is the delay to use before executing the commands. Anacron is not a daemon. It is run explicitly at one time. This field allows you to stagger execution so that not every task is running at the same time.

For example, the first line runs every day, five minutes after anacron is called:

```
1       5       cron.daily      nice run-parts --report /etc/cron.daily
```

The following line is run weekly (every 7 days), ten minutes after anacron is called:

```
7       10      cron.weekly     nice run-parts --report /etc/cron.weekly
```

The third column contains the name that the job will be known as in the anacron's messages and log files. The fourth field is the actual command that is run.

You can see that anacron is set to run some of the same scripts that are run by cron. Distributions handle this clash differently, by creating a preference for either cron or anacron and making the other program not execute the rule.

For instance, on Ubuntu, the "/etc/crontab" tests if anacron is available on the system, and only executes the scripts in the cron.* directories with cron if anacron is not found.

Other distributions have cron update the anacron's time-stamps every time it runs the contents of these directories, so that anacron does not execute when it is called.

Conclusion

Both cron and anacron are useful tools for when you need to automate processes. Understanding how to leverage their strengths and work around their weaknesses will allow you to utilize them easily and effectively.

Although the configuration syntax may be confusing at first, these tools will save you time in the long run and usually do not have to be adjusted often once you have a good working schedule.

Assessment: Essay Type Questions

Instruction

Write short answers for the following questions

1. What is a server farm?
2. What are datacenters?
3. What advantages are there when rack server is used?

UNIT SUMMARY

In this unit we discussed about how installation and removal of applications are done in both Windows and Linux, server hardware that is available and its types, and automation of routine tasks using Linux cron command.

Unit Assessment

Check your understanding!

Assessment: Essay Type Questions

Instructions

Write your answer for the following questions

1. List server components.
2. Why do we require servers to be reliable and durable?
3. What factors do affect the choice of server platform?
4. How do you schedule hours with cron ?

Grading Scheme

This unit's assessments worth a total of 20% from which

a. Activity Assessments: 5%

a. Unit Assessment I:3 %

a. Unit Assessment II (Lab Project): 12%

Answers

Rubrics for assessing activities

Assessment Criteria	Doesn't attempt both assessment or the answers are all wrong	Attempt both assessments and provide partial answers or partially correct answers	Attempt all assessments with full and correct answer
Grading Scales	Fail	Got half mark for each assessment	Score full mark for each assessment

Unit Readings and Other Resources

- http://www.psychocats.net/ubuntu/installingsoftware, Installing Software in Ubuntu, February 24, 2016
- http://windows.microsoft.com/en-us/windows/install-program#1TC=windows-7, Install a Program, February 24, 2016
- http://www.unixmen.com/removing-applications-using-terminal-ubuntu/, Removing Applications Using Terminal in Ubuntu, Enock Seth Nyamador ,February 24, 2016
- https://www.digitalocean.com/community/tutorials/how-to-schedule-routine-tasks-with-cron-and-anacron-on-a-vps, How To Schedule Routine Tasks With Cron and Anacron on a VPS,Justin Ellingwood, February 24,2016
- https://www.raspberrypi.org/documentation/linux/usage/cron.md, Scheduling Tasks with Cron, February 24, 2016
- Brian Hilton, Michael Welch, Server Primer : Understanding the current state of the industry, Golisano Institute for Sustainability

MODULE SUMMARY

Operating system administration module covers topics which lets you manage and administer operating systems on your machine. Topics including users and group administration, file system and disk management, backup, recovery and server virtualization, securing operating system and monitoring, special topics like applications on operating systems, server hardware knowledge, and automating routine tasks has been covered.

As a multiuser OS, Linux relies upon user accounts and groups to maintain security and keep the system usable to all its users. Over the years, manual methods, text–based tools, and GUI tool shave been developed to manage these accounts and groups. However you do it, though, it's important that you understand how Linux handles its accounts and groups—the relationship between usernames and UIDs or group names and GIDs; where passwords are stored; and how users are assigned to groups. Understanding these topics will allow you to effectively manage your user base, whether it's just one or two people or thousands. These subjects are also critically important for securing the files and processes your system uses.

The file system defines how disks are structured. Individual disks are hidden from the casual user. All access to a file is through the directory structure, and file permissions are used to control user access to the system. Different disks can be formatted with different file systems and must be mounted into the directory structure to be accessed. A disk or partition can be mounted into any directory, which is then referred to as the mount point for that disk. Disk quota mechanism should be implemented in systems where several users are available. Linux implements disk quota to limit and manage disk usage of users. This module addressed disk quota and file system management in Linux.

This module tried to address backup and recovery systems and server virtualization in brief. It talked about backup, backup methods as well as backup algorithms. It also elaborated on storage devices and media to be used for backup. Data recovery and the various options to be used as a recovery strategy had been touched. The discussion continued to concept of virtualization and in particular focused on server virtualization; what it is, it's importance, it's evolutions and finally approaches to server virtualization.

Security is such a big topic that there is always more to be said. We have included the specific tools that we use or have had recommended to us. You should search the Internet and talk to other system administrators to find the specific products that meet the needs of your situation. Test everything. Be very thorough and always watch your distribution's homepage for security advisories and updated packages.

In this module we discussed also about how installation and removal of applications are done in both Windows and Linux, server hardware that is available and its types, and automation of routine tasks using Linux cron command.

Module Assessment

Instructions: the module assessment is considered as your final exam and consists of three parts. 10 questions of type multiple choice, 5 questions of matching type, and 10 questions of type Essay. Attempt all questions accordingly.

Part I. 10 Multiple Choice questions are given below. Choose the best answer from the given choices and put the letter in the space provided

1._____ One of the following is not a service provided by an operating system

A. File management

B. Resource Allocation

C. Accounting

D. Debugging

2._____In an operating system with a command line interface, CLI,

A. Users communicate with the system through command typing

B. Several users can communicate the system simultaneously

C. Users communicate with the system through icon clicks

D. All of the above

3._____Why is it not recommended to login to your system as a root unless required

A. Because root login is for novice users

B. Because root login damages the system

C. Because root login restricts the privilege a user has

D. Because root login is meant for experts and system administrators

4._____In Linux operating system, mount command is used for

A. Mounting the operating system

B. Accessing a file

C. Installing applications

D. All of the above

5._____One is not the goal of having backups

A. Buffering

B. Transferring data between non-networked devices

C. System recovery

D. Data maintenance while upgrading

6._____One of the following is true regarding backup strategy algorithms?

A. In a Volume/Calendar backup, the backup media required is minimal

B. In Grandfather/Father/Son(GFS) Backup, better granularity of backed up data is achieved

C. In Tower of Hanoi backup, backed up data restorage is simple and efficient

D. All of the above

7 _____Virtualization offers

A. Efficient computation

B. Resource utilization

C. Extra cost

D. Vulnerability

8._____All of the following are similar except,

A. Trojan Horse

B. PGP

C. Checksum

D. Password Hashing

9._____One can't be prevented through a firewall?

A. Attacks from inside

B. Attacks from outside

C. Virus attacks

D. Denial of service attacks

10._____Serially attached SCSI(SAS) is preferred storage medium for servers due to

A. Its performance efficiency

B. Large storage space

C. Minimal cost

D. All of the above

Part II. Matching. Match the questions under column A with possible answer from column B

"A"	"B"
1. Buffer Overflow	A. Firewall
2. RAID system	B. Rail
3. shell Security Attack	C.
4. IP Chains	D. Backup storage
5. Rack Servers	E. CLI

Part III. Essay Type Questions. Below are 10 questions that need your explanation. Write a precise answer to all the questions.

1. Discuss the two important tasks done during operating system installation.
2. What are the files Unix maintains regarding user accounts ? discuss each in brief
3. What are the differences between onsite and offsite backup methods ?
4. What is server virtualization ? what advantages does it offer to a system?
5. Discuss about performance monitoring and its importance
6. What is a proxy server? how does it differ from packet-filtering firewalls?
7. Explain in brief the three types of server platforms
8. Mention three components of a server
9. What is a cron in the Linux operating system ? what services ca nit provide to a system administrator?
10. What are the main features that distinguishes Linux Operating System from Windows Operating System?

References

1. Modern Operating System, Andrew S. Tanenbaum, Prentice-Hall, Inc., 3rd Edition, 2008
2. A Silberschatz, Peter B Galvin, G Gagne: Operating System Concepts, 7th edition
3. William Stallings, Operating Systems, 4th edition, 2002
4. Linux System Administration, Vicki Stanfield, Roderick W. Smith, Sybex., 2nd Edition, 2002
5. LPI – Linux System Administration, Revision 1.0, 2000
6. http://www.control-escape.com/linux/users-groups.html, Users and Groups, February 24, 2006
7. http://www.tldp.org/HOWTO/Shadow-Password-HOWTO-2.html, Why shadow your passwd file?, February 24, 2006.

8. http://www.w3resource.com/linux-system-administration/user-management.php, User Management, February 24, 2016

9. http://www.pcmag.com/article2/0,2817,899676,00.asp; Backup Methods and Rotation Schemes; BY MATTHEW D. SARREL; FEBRUARY 21, 2003

10. www.tandbergdata.com/us ; Tape Rotation BEST PRACTICES The Tandberg Data SMB Guide to Backup Data Protection Strategies for the Small-to-Medium Size Business p.10 Basic Backup Guide

11. [Moore] Moore, Sonia Marie (Ed.). "MS Windows NT Workstation 4.0 Resource Guide."1995. http://www.microsoft.com/technet/archive/ntwrkstn/reskit/default.mspx.

12. http://www.psychocats.net/ubuntu/installingsoftware, Installing Software in Ubuntu, February 24, 2016

13.http://windows.microsoft.com/en-us/windows/install-program#1TC=windows-7, Install a Program, February 24, 2016

14. http://www.unixmen.com/removing-applications-using-terminal-ubuntu/, Removing Applications Using Terminal in Ubuntu, Enock Seth Nyamador ,February 24, 2016

15. https://www.digitalocean.com/community/tutorials/how-to-schedule-routine-tasks-with-cron-and-anacron-on-a-vps, How To Schedule Routine Tasks With Cron and Anacron on a VPS, Justin Ellingwood, February 24, 2016

16. https://www.raspberrypi.org/documentation/linux/usage/cron.md, Scheduling Tasks with Cron, February 24, 2016

17. Brian Hilton, Michael Welch, Server Primer: Understanding the current state of the industry, Golisano Institute for Sustainability

The African Virtual University Headquarters

Cape Office Park

Ring Road Kilimani

PO Box 25405-00603

Nairobi, Kenya

Tel: +254 20 25283333

contact@avu.org

oer@avu.org

The African Virtual University Regional Office in Dakar

Université Virtuelle Africaine

Bureau Régional de l'Afrique de l'Ouest

Sicap Liberté VI Extension

Villa No.8 VDN

B.P. 50609 Dakar, Sénégal

Tel: +221 338670324

bureauregional@avu.org