OPERATING SYSTEM WIN7 WINSERVER 2008 LINUX

DIT Part 1ST

BY SARDAR AZEEM
(PICT COMPUTER CENTER LINK ROAD ABBOTTABAD)

SARDARAZEEM

Pre pared by: Sardar Azeem (MBA (BGF) Computer HW And Network Engineer: Pict Computer Center Link Road Abbottabad.
Email: Azeem_aag2000@yahoo.com/Website:www.pictcomputeratd.com

What is an Operating System and its role in the Computing environment?

An **operating system** (**OS**) is a set of programs that manage computer hardware resources and provide common services for application software. The operating system is the most important type of system software in a computer system. A user cannot run an application program on the computer without an operating system, unless the application program is self booting.

For hardware functions such as input and output and memory allocation, the operating system acts as an intermediary between application programs and the computer hardware, although the application code is usually executed directly by the hardware and will frequently call the OS or be interrupted by it. Operating systems are found on almost any device that contains a computer—from cellular phones and video game consoles to supercomputers and web servers.

Examples of popular modern operating systems include Android, iOS, Linux, Mac OS X, all of which have their roots in Unix, and Microsoft Windows.

Roles of operating system

• Operating systems are designed to provide uniform abstraction across multiple applications: fair sharing of resources

• General purpose OS like Solaris in wizard.cse.nd.edu Mail, web, samba server, telnet, emacs …Memory fs, afs, ufs …Fibre channel devices, floppy disks …

• What about applications/services such as video games, data base servers, mail servers
OS gets in the way of these applications in the name of fairness (MSDOS is the ideal OS!!)
• Create multiple virtual machines that each user can control all to themselves
IBM 360/370 …
• Monolithic kernel: Linux One kernel provides all services.
New paradigms are harder to implement
May not be optimal for any one application
• Microkernel: Mach
Microkernel provides minimal service
Application servers provide OS functionality
• Nanokernel: OS is implemented as application level libraries
• Shared memory multiprocessor The Multics hardware architecture supports multiple CPUs sharing the same physical memory. All processors are equivalent.
• Multi-language supportIn addition to PL/I, Multics supports BCPL, BASIC, APL, FORTRAN,LISP, C, COBOL, ALGOL 68 and Pascal. Routines in these languages can call each other.
• Relational database Multics provided the first commercial relational database product, the Multics Relational Data Store (MRDS), in 1978.
• Security
Multics was designed to be secure from the beginning. In the 1980s, the system was awarded the B2 security rating by the US government NCSC, the first (and for years only) system to get a B2 rating.

SARDAR AZEEM(PICT)

Components of An Operating system

The components of an operating system all exist in order to make the different parts of a computer work together. All software—from financial databases to film editors—needs to go through the operating system in order to use any of the hardware, whether it be as simple as a mouse or keyboard or complex as an Internet connection.

Some major components of operating system are.

1. Kernal:

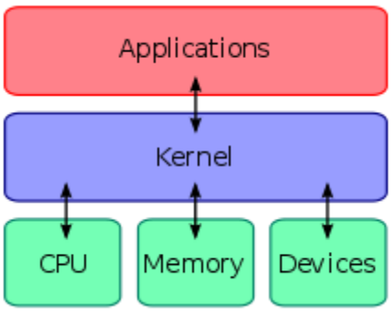

A kernel connects the application software to the hardware of a computer.With the aid of the firmware and device drivers, the kernel provides the most basic level of control over all of the computer's hardware devices. It manages memory access for programs in the RAM, it determines which programs get access to which hardware resources, it sets up or resets the CPU's operating states for optimal operation at all times, and it organizes the data for long-term non-volatile storage with file systems on such media as disks, tapes, flash memory, etc.

2. Program Execution:

The operating system provides an interface between an application program and the computer hardware, so that an application program can interact with the hardware only by obeying rules and procedures programmed into the operating system. The operating system is also a set of services which simplify development and execution of application programs. Executing an application program involves the creation of a process by the operating system kernel which assigns memory space and other resources, establishes a priority for the process in multi-tasking systems, loads program binary code into memory, and initiates execution of the application program which then interacts with the user and with hardware devices.

3.Interrupts:

Interrupts are central to operating systems, as they provide an efficient way for the operating system to interact with and react to its environment. The alternative — having the operating system "watch" the various sources of input for events (polling) that require action — can be found in older systems with very small stacks (50 or 60 bytes) but are unusual in modern systems with large stacks. Interrupt-based programming is directly supported by most modern CPUs. Interrupts provide a computer with a way of automatically saving local register contexts, and running specific code in response to events. Even very basic computers support hardware interrupts, and allow the programmer to specify code which may be run when that event takes place.

When an interrupt is received, the computer's hardware automatically suspends whatever program is currently running, saves its status, and runs computer code previously associated with the interrupt; this is analogous to placing a bookmark in a book in response to a phone call. In modern operating systems, interrupts are handled by the operating system's kernel. Interrupts may come from either the computer's hardware or from the running program.

SARDAR AZEEM (PICT)

When a hardware device triggers an interrupt, the operating system's kernel decides how to deal with this event, generally by running some processing code. The amount of code being run depends on the priority of the interrupt (for example: a person usually responds to a smoke detector alarm before answering the phone). The processing of hardware interrupts is a task that is usually delegated to software called device driver, which may be either part of the operating system's kernel, part of another program, or both. Device drivers may then relay information to a running program by various means.

A program may also trigger an interrupt to the operating system. If a program wishes to access hardware for example, it may interrupt the operating system's kernel, which causes control to be passed back to the kernel. The kernel will then process the request. If a program wishes additional resources (or wishes to shed resources) such as memory, it will trigger an interrupt to get the kernel's attention.

4. Modes

Privilege rings for the x86 available in protected mode. Operating systems determine which processes run in each mode.

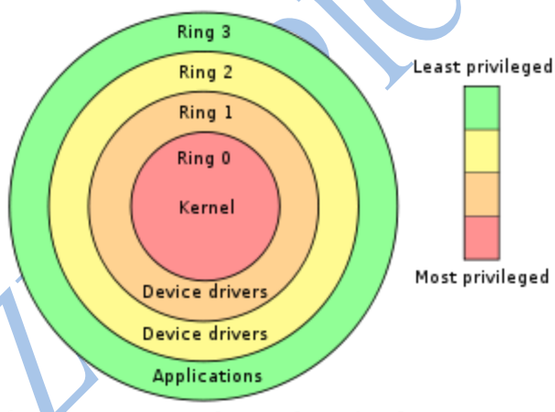

Modern CPUs support multiple modes of operation. CPUs with this capability use at least two modes: protected mode and supervisor mode. The supervisor mode is used by the operating system's kernel for low level tasks that need unrestricted access to hardware, such as controlling how memory is written and erased, and communication with devices like graphics cards. Protected mode, in contrast, is used for almost everything else. Applications operate within protected mode, and can only use hardware by communicating with the kernel, which controls everything in supervisor mode. CPUs might have other modes similar to protected mode as well, such as the virtual modes in order to emulate older processor types, such as 16-bit processors on a 32-bit one, or 32-bit processors on a 64-bit one.

When a computer first starts up, it is automatically running in supervisor mode. The first few programs to run on the computer, being the BIOS or EFI, bootloader, and the operating system have unlimited access to hardware - and this is required because, by definition, initializing a protected environment can only be done outside of one. However, when the operating system passes control to another program, it can place the CPU into protected mode.

In protected mode, programs may have access to a more limited set of the CPU's instructions. A user program may leave protected mode only by triggering an interrupt, causing control to be passed back to the kernel. In this way the operating system can maintain exclusive control over things like access to hardware and memory.

The term "protected mode resource" generally refers to one or more CPU registers, which contain information that the running program isn't allowed to alter. Attempts to alter these resources generally causes a switch to supervisor mode, where the operating system can deal with the illegal operation the program was attempting (for example, by killing the program).

5. Memory Management:

Among other things, a multiprogramming operating system kernel must be responsible for managing all system memory which is currently in use by programs. This ensures that

SARDAR AZEEM PICT

a program does not interfere with memory already in use by another program. Since programs time share, each program must have independent access to memory.

6. Virtual Memory

Many operating systems can "trick" programs into using memory scattered around the hard disk and RAM as if it is one continuous chunk of memory, called virtual memory.

The use of virtual memory addressing (such as paging or segmentation) means that the kernel can choose what memory each program may use at any given time, allowing the operating system to use the same memory locations for multiple tasks.

If a program tries to access memory that isn't in its current range of accessible memory, but nonetheless has been allocated to it, the kernel will be interrupted in the same way as it would if the program were to exceed its allocated memory. (See section on memory management.) Under UNIX this kind of interrupt is referred to as a page fault.

When the kernel detects a page fault it will generally adjust the virtual memory range of the program which triggered it, granting it access to the memory requested. This gives the kernel discretionary power over where a particular application's memory is stored, or even whether or not it has actually been allocated yet.

In modern operating systems, memory which is accessed less frequently can be temporarily stored on disk or other media to make that space available for use by other programs. This is called swapping, as an area of memory can be used by multiple programs, and what that memory area contains can be swapped or exchanged on demand.

"Virtual memory" provides the programmer or the user with the perception that there is a much larger amount of RAM in the computer than is really there.

7. Multitasking:

Multitasking refers to the running of multiple independent computer programs on the same computer; giving the appearance that it is performing the tasks at the same time. Since most computers can do at most one or two things at one time, this is generally done via time-sharing, which means that each program uses a share of the computer's time to execute.

An operating system kernel contains a piece of software called a scheduler which determines how much time each program will spend executing, and in which order execution control should be passed to programs. Control is passed to a process by the kernel, which allows the program access to the CPU and memory. Later, control is returned to the kernel through some mechanism, so that another program may be allowed to use the CPU. This so-called passing of control between the kernel and applications is called a context switch.

An early model which governed the allocation of time to programs was called cooperative multitasking. In this model, when control is passed to a program by the kernel, it may execute for as long as it wants before explicitly returning control to the kernel. This means that a malicious or malfunctioning program may not only prevent any other programs from using the CPU, but it can hang the entire system if it enters an infinite loop.

Modern operating systems extend the concepts of application preemption to device drivers and kernel code, so that the operating system has preemptive control over internal run-times as well.

The philosophy governing preemptive multitasking is that of ensuring that all programs are given regular time on the CPU. This implies that all programs must be limited in how

SARDAR AZEEM (PICT)

much time they are allowed to spend on the CPU without being interrupted. To accomplish this, modern operating system kernels make use of a timed interrupt. A protected mode timer is set by the kernel which triggers a return to supervisor mode after the specified time has elapsed. (See above sections on Interrupts and Dual Mode Operation.)

On many single user operating systems cooperative multitasking is perfectly adequate, as home computers generally run a small number of well tested programs. The AmigaOS is an exception, having pre-emptive multitasking from its very first version. Windows NT was the first version of Microsoft Windows which enforced preemptive multitasking, but it didn't reach the home user market until Windows XP (since Windows NT was targeted at professionals).

9. Disk access and file systems:

Filesystems allow users and programs to organize and sort files on a computer, often through the use of directories (or "folders")Access to data stored on disks is a central feature of all operating systems. Computers store data on disks using files, which are structured in specific ways in order to allow for faster access, higher reliability, and to make better use out of the drive's available space. The specific way in which files are stored on a disk is called a file system, and enables files to have names and attributes. It also allows them to be stored in a hierarchy of directories or folders arranged in a directory tree.

Early operating systems generally supported a single type of disk drive and only one kind of file system. Early file systems were limited in their capacity, speed, and in the kinds of file names and directory structures they could use. These limitations often reflected limitations in the operating systems they were designed for, making it very difficult for an operating system to support more than one file system.

Various differences between file systems make supporting all file systems difficult. Allowed characters in file names, case sensitivity, and the presence of various kinds of file attributes makes the implementation of a single interface for every file system a daunting task. Operating systems tend to recommend using (and so support natively) file systems specifically designed for them; for example, NTFS in Windows and ext3 and ReiserFS in GNU/Linux. However, in practice, third party drives are usually available to give support for the most widely used file systems in most general-purpose operating systems (for example, NTFS is available in GNU/Linux through NTFS-3g, and ext2/3 and ReiserFS are available in Windows through FS-driver and rfstool).

Support for file systems is highly varied among modern operating systems, although there are several common file systems which almost all operating systems include support and drivers for. Operating systems vary on file system support and on the disk formats they may be installed on. Under Windows, each file system is usually limited in application to certain media; for example, CDs must use ISO 9660 or UDF, and as of Windows Vista, NTFS is the only file system which the operating system can be installed on. It is possible to install GNU/Linux onto many types of file systems. Unlike other operating systems, GNU/Linux and UNIX allow any file system to be used regardless of the media it is stored in, whether it is a hard drive, a disc (CD,DVD...), a USB flash drive, or even contained within a file located on another file system.

10. Device Drivers:

A device driver is a specific type of computer software developed to allow interaction with hardware devices. Typically this constitutes an interface for communicating with the device, through the specific computer bus or communications subsystem that the hardware is connected to, providing commands to and/or receiving data from the device, and on the other end, the requisite interfaces to the operating system and software applications. It is a specialized hardware-dependent computer program which is also operating system specific that enables another program, typically an operating system or applications software package or computer program running under the operating system

SARDAR AZEEM (PICT)

kernel, to interact transparently with a hardware device, and usually provides the requisite interrupt handling necessary for any necessary asynchronous time-dependent hardware interfacing needs.

11.Networking

Currently most operating systems support a variety of networking protocols, hardware, and applications for using them. This means that computers running dissimilar operating systems can participate in a common network for sharing resources such as computing, files, printers, and scanners using either wired or wireless connections. Networks can essentially allow a computer's operating system to access the resources of a remote computer to support the same functions as it could if those resources were connected directly to the local computer. This includes everything from simple communication, to using networked file systems or even sharing another computer's graphics or sound hardware. Some network services allow the resources of a computer to be accessed transparently, such as SSH which allows networked users direct access to a computer's command line interface.

12.Security:

A computer being secure depends on a number of technologies working properly. A modern operating system provides access to a number of resources, which are available to software running on the system, and to external devices like networks via the kernel.

The operating system must be capable of distinguishing between requests which should be allowed to be processed, and others which should not be processed. While some systems may simply distinguish between "privileged" and "non-privileged", systems commonly have a form of requester *identity*, such as a user name. To establish identity there may be a process of *authentication*. Often a username must be quoted, and each username may have a password. Other methods of authentication, such as magnetic cards or biometric data, might be used instead. In some cases, especially connections from the network, resources may be accessed with no authentication at all (such as reading files over a network share). Also covered by the concept of requester **identity** is *authorization*; the particular services and resources accessible by the requester once logged into a system are tied to either the requester's user account or to the variously configured groups of users to which the requester belongs.

Internal security is especially relevant for multi-user systems; it allows each user of the system to have private files that the other users cannot tamper with or read. Internal security is also vital if auditing is to be of any use, since a program can potentially bypass the operating system, inclusive of bypassing auditing.

13.User Interface:

Every computer that is to be operated by an individual requires a user interface. The user interface is not actually a part of the operating system—it generally runs in a separate program usually referred to as a shell, but is essential if human interaction is to be supported. The user interface requests services from the operating system that will acquire data from input hardware devices, such as a keyboard, mouse or credit card reader, and requests operating system services to display prompts, status messages and such on output hardware devices, such as a video monitor or printer. The two most common forms of a user interface have historically been the command-line interface, where computer commands are typed out line-by-line, and the graphical user interface, where a visual environment (most commonly a WIMP) is present.

Graphical user interfaces

A screenshot of the KDE graphical user interface. Programs take the form of images on the screen, and the files, folders (directories), and applications take the form of icons and

SARDAR AZEEM (PICT)

symbols. A mouse is used to navigate the computer.Most of the modern computer systems support graphical user interfaces (GUI), and often include them. In some computer systems, such as the original implementation of Mac OS, the GUI is integrated into the kernel.

CDI(commanad driven interface)

Provides command base access to resources of a computer.

Like Ms dos in Ms windows.

MDI(menu driven interface)

Provides interactive menus to access computer resources.

Tasks of an operating system

1.Multitasking:

Running more thenone tasks at a time
In computing, **multitasking** is a method where multiple tasks, also known as processes, share common processing resources such as a CPU. In the case of a computer with a single CPU, only one task is said to be *running* at any point in time, meaning that the CPU is actively executing instructions for that task. Multitasking solves the problem by scheduling which task may be the one running at any given time, and when another waiting task gets a turn

2. Multiprocessing:

Multiprocessing is the use of two or more central processing units (CPUs) within a single computer system. The term also refers to the ability of a system to support more than one processor and/or the ability to allocate tasks between them. There are many variations on this basic theme, and the definition of multiprocessing can vary with context, mostly as a function of how CPUs are defined (multiple cores on one die, multiple dies in one package, multiple packages in one system unit, etc.).

Symmetric multiprocessing

In computing, **symmetric multiprocessing** (SMP) involves a multiprocessor computer hardware architecture where two or more identical processors are connected to a single shared main memory and are controlled by a single OS instance.

Asymmetric multiprocessing

Asymmetric multiprocessing, or AMP, was a software stopgap for handling multiple CPUs before symmetric multiprocessing, or SMP, was available.

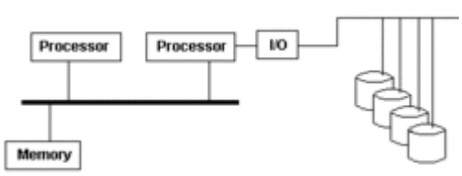

Asymmetric multiprocessing

Multiprocessing is the use of more than one CPU in a computer system. The CPU is the arithmetic and logic engine that executes user applications; an I/O interface such as a GPU, even if it is implemented using an embedded processor, does not constitute a CPU because it does not run the user's application program. With multiple CPUs, more than one set of program instructions can be executed at the same time. All of the CPUs have

SARDAR AZEEM(PICT)

the same user-mode instruction set, so a running job can be rescheduled from one CPU to another.

3. Multithreading:

Multithreading as a widespread programming and execution model allows multiple threads to exist within the context of a single process.

Preemptive multithreading is generally considered the superior approach, as it allows the operating system to determine when a context switch should occur. The disadvantage to preemptive multithreading is that the system may make a context switch at an inappropriate time, causing lock convoy, priority inversion or other negative effects which may be avoided by cooperative multithreading.

Cooperative multithreading, on the other hand, relies on the threads themselves to relinquish control once they are at a stopping point. This can create problems if a thread is waiting for a resource to become available.

4. Multi Programming:

In *multiprogramming* systems, the running task keeps running until it performs an operation that requires waiting for an external event (e.g. reading from a tape) or until the computer's scheduler forcibly swaps the running task out of the CPU. Multiprogramming systems are designed to maximize CPU usage.

5. Time Sharing:

In *time-sharing* systems, the running task is required to relinquish the CPU, either voluntarily or by an external event such as a hardware interrupt. Time sharing systems are designed to allow several programs to execute apparently simultaneously.

Types of Operating system

Real-time Operating system

A real-time operating system is a multitasking operating system that aims at executing real-time applications. Real-time operating systems often use specialized scheduling algorithms so that they can achieve a deterministic nature of behavior. The main objective of real-time operating systems is their quick and predictable response to events. They have an event-driven or time-sharing design and often aspects of both. An event-driven system switches between tasks based on their priorities or external events while time-sharing operating systems switch tasks based on clock interrupts.

Multi-user vs. Single-user

A multi-user operating system allows multiple users to access a computer system concurrently. Time-sharing system can be classified as multi-user systems as they enable a multiple user access to a computer through the sharing of time. Single-user operating systems, as opposed to a multi-user operating system, are usable by a single user at a time. Being able to have multiple accounts on a Windows operating system does not make it a multi-user system. Rather, only the network administrator is the real user. But for a Unix-like operating system, it is possible for two users to login at a time and this capability of the OS makes it a multi-user operating system.

Single user operating system allows a single user to access the computer at a time. These computers have a single processor and execute a single program. The resources such as CPU and I/O devices are constantly available to the user in a single user operating system for operating the system. As a result, the CPU sites idlefor most of the time and is not utilized to its maximum.

Multi user operating system allows various users to access the different resources of a computer simultaneously. The access is provided using a network that consists of various presonal computers attached to a mainframe computer. These computers send and recieve information to multi user mainframe computer system.

Multi-tasking vs. Single-tasking

SARDAR AZEEM(PICT)

When only a single program is allowed to run at a time, the system is grouped under a single-tasking system. However, when the operating system allows the execution of multiple tasks at one time, it is classified as a multi-tasking operating system. Multi-tasking can be of two types: pre-emptive or co-operative. In pre-emptive multitasking, the operating system slices the CPU time and dedicates one slot to each of the programs. Unix-like operating systems such as Solaris and Linux support pre-emptive multitasking, as does AmigaOS. Cooperative multitasking is achieved by relying on each process to give time to the other processes in a defined manner. MS Windows prior to Windows 2000 and Mac OS prior to OS X used to support cooperative multitasking.

Distributed

A distributed operating system manages a group of independent computers and makes them appear to be a single computer. The development of networked computers that could be linked and communicate with each other gave rise to distributed computing. Distributed computations are carried out on more than one machine. When computers in a group work in cooperation, they make a distributed system.

Embedded

Embedded operating systems are designed to be used in embedded computer systems. They are designed to operate on small machines like PDAs with less autonomy. They are able to operate with a limited number of resources. They are very compact and extremely efficient by design. Windows CE and Minix 3 are some examples of embedded operating systems.

File Systems

A **file system** (or **filesystem**) is a means to organize data expected to be retained after a program terminates by providing procedures to store, retrieve and update data, as well as manage the available space on the device(s) which contain it. A file system organizes data in an efficient manner and is tuned to the specific characteristics of the device. There is usually a tight coupling between the operating system and the file system. Some filesystems provide mechanisms to control access to the data and metadata. Ensuring reliability is a major responsibility of a filesystem. Some filesystems provide a means for multiple programs to update data in the same file at nearly the same time.

Without a filesystem programs would not be able to access data by file name or directory and would need to be able to directly access data regions on a storage device.

Windows makes use of the FAT and NTFS file systems.

Windows uses a *drive letter* abstraction at the user level to distinguish one disk or partition from another. For example, the path **C:\WINDOWS** represents a directory WINDOWS on the partition represented by the letter C.

Network drives may also be mapped to drive letters.

FAT

The family of FAT file systems is supported by almost all operating systems for personal computers, including all versions of Windows and MS-DOS/PC DOS and DR-DOS. (PC DOS is an OEM version of MS-DOS, MS-DOS was originally based on SCP's 86-DOS. DR-DOS was based on Digital Research's Concurrent DOS.) The FAT file systems are therefore well-suited as an universal exchange format between computers and devices of most any type and age.

The FAT file system traces its roots back to an (incompatible) 8-bit FAT precursor in the short-lived M-DOS project and Standalone disk BASIC.

Over the years, the filesystem has been expanded from FAT12 to FAT16 and FAT32. Various features have been added to the file system including sub-directories, codepage support, extended attributes, and long filenames. Third-parties such as Digital Research

SARDAR AZEEM (PICT)

have incorporated optional support for deletion tracking, and volume/directory/file-based multi-user security schemes to support file and directory passwords and permissions such as read/write/delete/execute access rights. Most of these extensions are not supported by Windows.

The FAT12 and FAT16 file systems had a limit on the number of entries in the root directory of the file system and had restrictions on the maximum size of FAT-formatted disks or partitions.

FAT32 addresses the limitations in FAT12 and FAT16, except for the file size limit of close to 4 GB, but it remains limited compared to NTFS.

FAT12, FAT16 and FAT32 also have a limit of 8 characters for the file name, and 3 characters for the extension (such as .exe). This is commonly referred to as the 8.3 filename limit. VFAT, an optional extension to FAT12, FAT16 and FAT32, introduced in Windows 95 and Windows NT 3.5, allowed long file names (LFN) to be stored in the FAT filesystem in a backwards compatible fashion.

NTFS

NTFS, introduced with the Windows NT operating system, allowed ACL-based permission control. Other features also supported by NTFS include hard links, multiple file streams, attribute indexing, quota tracking, sparse files, encryption, compression, and reparse points (directories working as mount-points for other file systems, symlinks, junctions, remote storage links), though not all these features are well-documented.

exFAT

exFAT is a proprietary and patent-protected file system with certain advantages over NTFS with regards to file system overhead.exFAT is not backwards compatible with FAT file systems such as FAT12, FAT16 or FAT32. The file system is supported with newer Windows systems, such as Windows 2003, Windows Vista, Windows 2008, Windows 7 and more recently, support has been added for Windows XP. Support in other operating systems is sparse since Microsoft has not published the specifications of the file system and implementing support for exFAT requires a license.

Win 7 System Requirements

Windows 7 system requirements

If you want to run Windows 7 on your PC, here's what it takes:

- 1 gigahertz (GHz) or faster 32-bit (x86) or 64-bit (x64) processor
- 1 gigabyte (GB) RAM (32-bit) or 2 GB RAM (64-bit)
- 16 GB available hard disk space (32-bit) or 20 GB (64-bit)
- DirectX 9 graphics device with WDDM 1.0 or higher driver

Additional requirements to use certain features:

- Internet access (fees may apply)
- Depending on resolution, video playback may require additional memory and advanced graphics hardware
- Some games and programs might require a graphics card compatible with DirectX 10 or higher for optimal performance
- For some Windows Media Center functionality a TV tuner and additional hardware may be required
- Windows Touch and Tablet PCs require specific hardware
- HomeGroup requires a network and PCs running Windows 7

SARDAR AZEEM(PICT)

- DVD/CD authoring requires a compatible optical drive
- BitLocker requires Trusted Platform Module (TPM) 1.2
- BitLocker To Go requires a USB flash drive
- Windows XP Mode requires an additional 1 GB of RAM and an additional 15 GB of available hard disk space.
- Music and sound require audio output

Installing Windows 7

Setting up your Computer & BIOS changes (If required)

In most cases if you do this you'll automatically boot in to the Windows 7 Installation DVD. But in some cases if the Boot device order is changed in the BIOS it may boot in to your older OS, instead of our DVD. In that case you'll need to change the BIOS settings to get it done.

Booting Up and First Installation Steps

Steps from here are pretty straight forward. Read the descriptions in each pages before clicking the **Next** button to avoid any disasters. IF there is options to choose in these steps you may find them with each images.

Screen 1

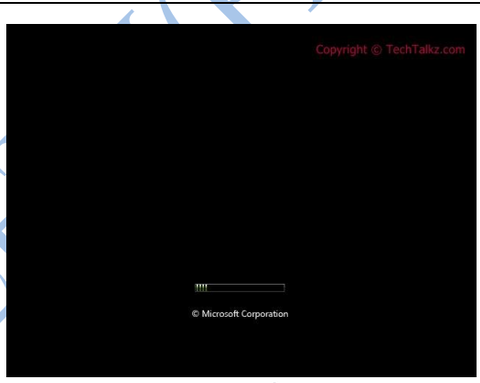

screen 2

screen 3

screen 4

You may choose your Language options by selecting the dropdowns 'Language to install', 'Time and Currency format' and 'Keyboard or input method' here. I decided to leave everything to 'US' but it would be better for selecting the correct settings here for non-english users.

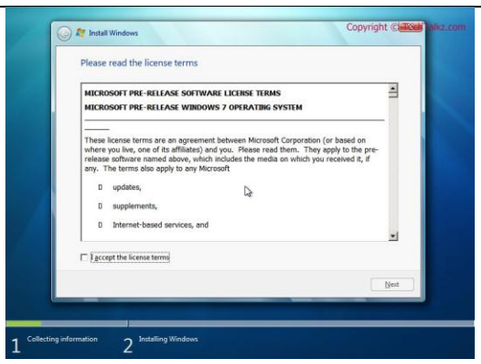

This screen is where the installation wizard begins. The install now options will leads to the advanced install options. For repairing a corrupt installation the 'Repair your computer' button located at the bottom-left can be used. For fresh installs just click the **Install now** button.

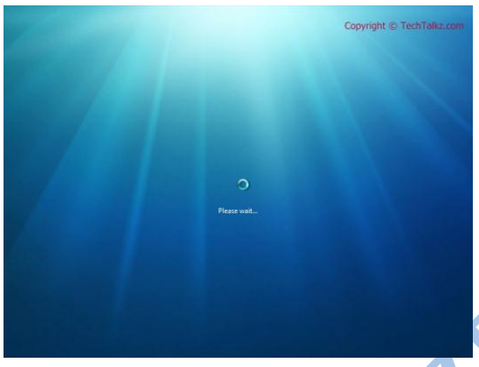

Tick the 'I accept the license terms' checkbox and click 'Next' to proceed.

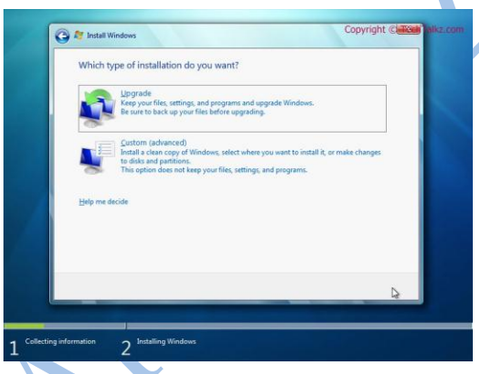

Which type of installation do you want?
This screen provides two options, **Upgrade** and **Custom (advanced)**. The upgrade option is for those who wish to upgrade an existing installation of older version of Windows to Windows 7. It is confirmed that Windows Vista can be upgraded to Windows 7 with out any issues, but Windows XP is still a problem. We will opt for the second option here, the **Custom** install.

SARDAR AZEEM(PICT)

Where do you want to install Windows?

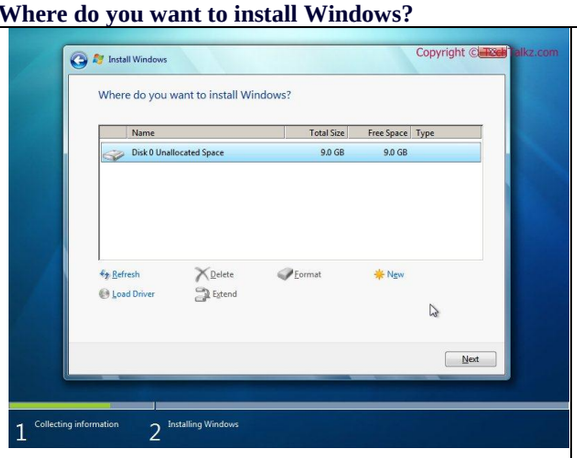

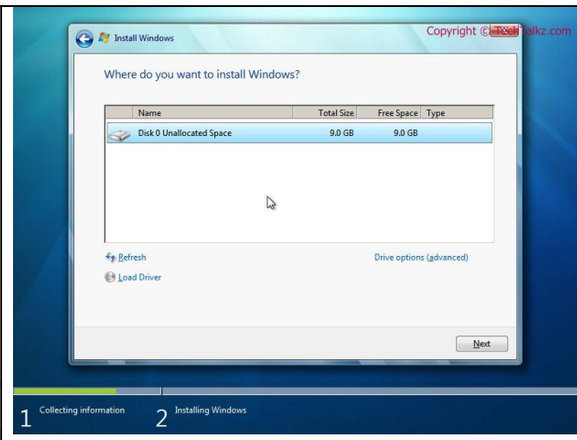

Clicking on the 'Custom' button brings the install location selection screen. In the test machine we have an un-partitioned empty disk. But in the case of a normal installation all your hard drive partitions (e.g. C:\ , D:\ etc.) will be listed here. Choose the drive as you like (a 15 GB size is recommended). Make sure the drive don't have any important data or the Windows 7 installation will wipe-out the contents of that partition. You can backup the data to another partition (e.g. for installing in D:\ drive move important files from there to say E:\ drive or to an external usb drive) for safe keeping.

Partitioning /Creating or Modifying Partitions

Click the **Drive options (advanced)** for advanced partition management options like 'Delete', 'Format', 'New' and 'Extend'. To create a new partition click the **New** button. But if you have list of partitions in the previous screen, choose the one where you are planning to install windows 7 and click **Format**. Then click 'Next' to proceed.

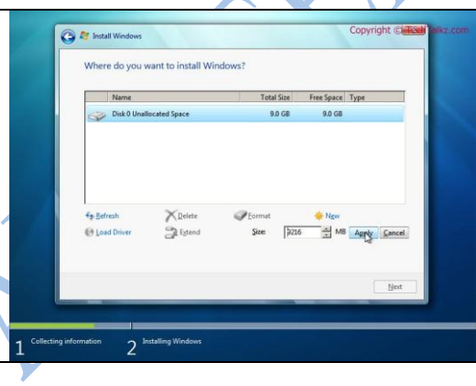

The 'New' option provides a text box to enter the size of the partition you wish to create in MBs. A 15 GB (15*1024 = 15360 MB) is recommended. In this case I opted for the full size of my virtual drive, i.e. 9216 MB. Click the **Apply** button to continue. You may be greeted with a message **To ensure that all Windows features work correctly, windows might create additional partitions for system files.** . This is a new feature in Windows 7 to have a small boot partition for system files. Click **Ok**to continue.

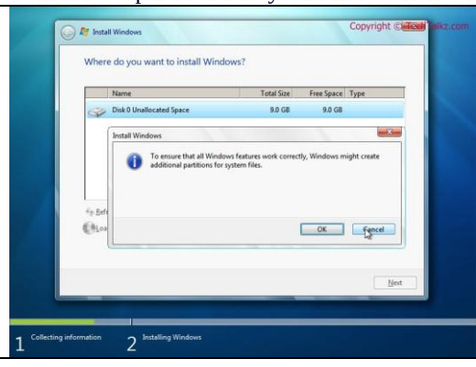

So here we are, A new primary partition of 8.8 GB is created along with a **System** type partition of 200 MB. Select the partition you just created and click **Next** to continue.

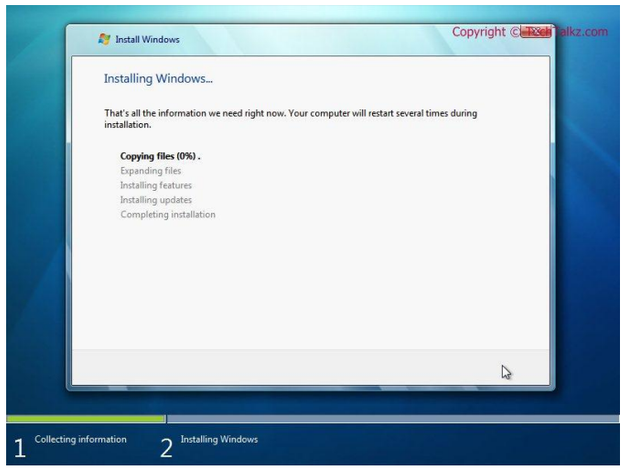

All right, we are all set. The Windows 7 installation begins. time to grab a cup of coffee for and wait for 15 - 30 minutes for the all new OS!
We discussed the initial steps of Windows 7 installation in the first part of this tutorial. In this session the instillation proceeds though the series of screen shots which involves a reboot. No user interaction is required for this Part.

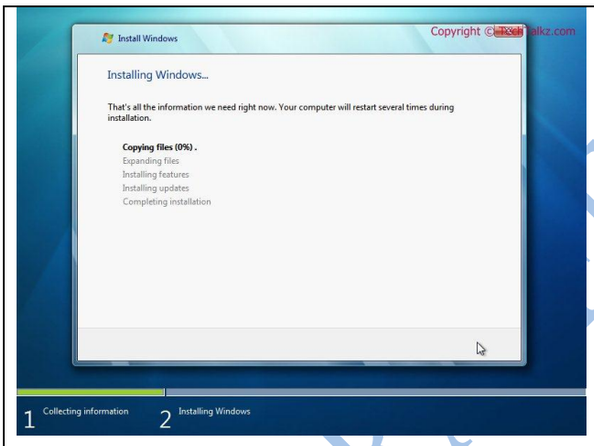

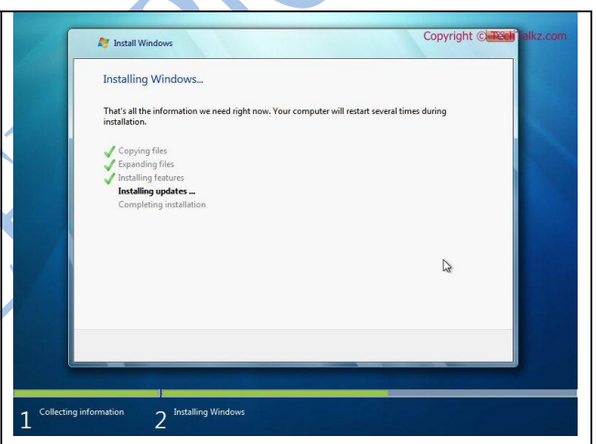

This part of Windows 7 installation guide covers the Windows Product code (Serial key) input, Computer name setup, user account and password setup, time zone selection, Windows update configuration, location selection, Home group setup etc.

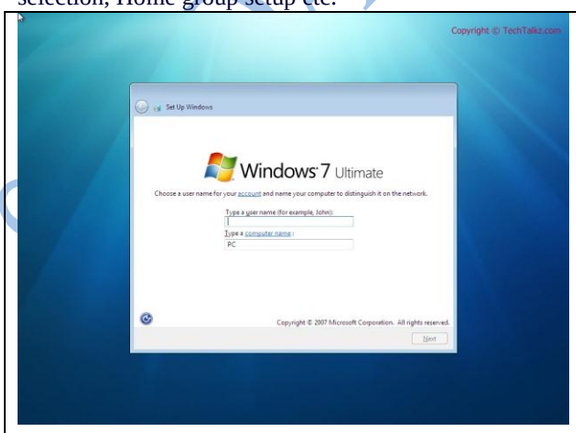

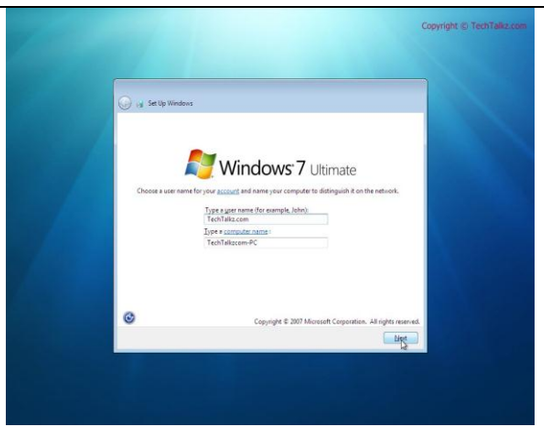

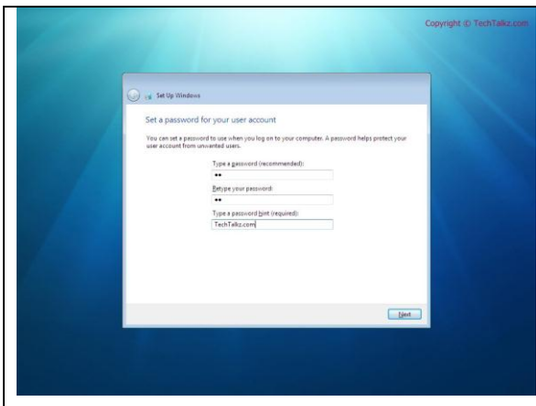
Copyright © TechTalkz.com
Set Up Windows
Set a password for your user account
You can set a password to use when you log on to your computer. A password helps protect your user account from unwanted users.
Type a password (recommended):
Retype your password:
Type a password hint (required):
TechTalkz.com
Next

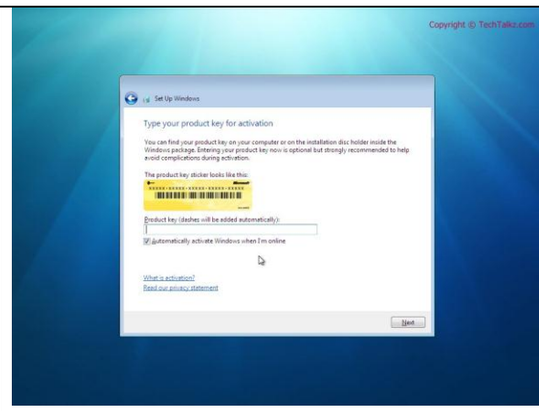
Copyright © TechTalkz.com
Set Up Windows
Type your product key for activation
You can find your product key on your computer or on the installation disc holder inside the Windows package. Entering your product key now is optional but strongly recommended to help avoid complications during activation.
The product key sticker looks like this:
Product key (dashes will be added automatically):
Automatically activate Windows when I'm online
What is activation?
Read our privacy statement
Next

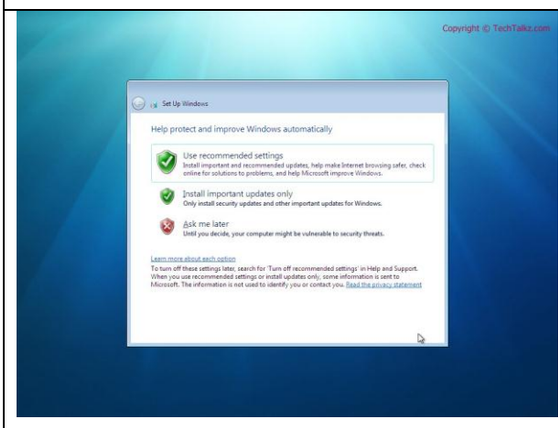
Copyright © TechTalkz.com
Set Up Windows
Help protect and improve Windows automatically
Use recommended settings
Install important updates only
Only install security updates and other important updates for Windows.
Ask me later
Until you decide, your computer might be vulnerable to security threats.
Learn more about each option
To turn off these settings later, search for 'Turn off recommended settings' in Help and Support. When you use recommended settings or install updates only, some information is sent to Microsoft. The information is not used to identify you or contact you. Read the privacy statement

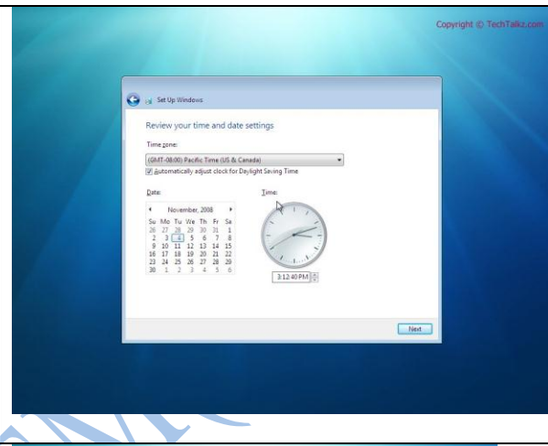
Copyright © TechTalkz.com
Set Up Windows
Review your time and date settings
Time zone:
(GMT-08:00) Pacific Time (US & Canada)
Automatically adjust clock for Daylight Saving Time
Date:
Time:
November, 2008
3:12:40 PM
Next


Copyright © TechTalkz.com
Set Up Windows
Select your computer's current location
This computer is connected to a network. Windows will automatically apply the correct network settings based on the network's location.
Home network
Choose this for a home or similar location. You can participate in a homegroup, and you can see other network computers.
Work network
Choose this for a workplace or similar location. You can't participate in a homegroup, but you can see other network computers.
Public network
Choose this for airports, coffee shops, other public places, or connections to the Internet without a router. You can't participate in a homegroup or access network computers.
If you aren't sure, select Public network.

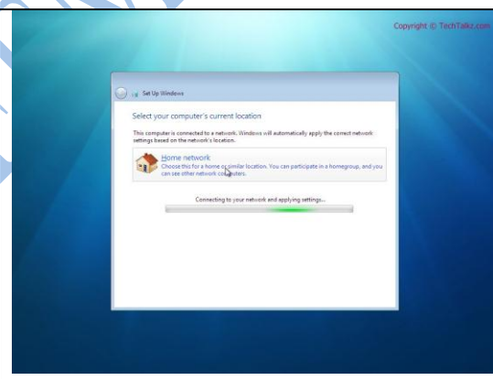
Copyright © TechTalkz.com
Set Up Windows
Select your computer's current location
This computer is connected to a network. Windows will automatically apply the correct network settings based on the network's location.
Home network
Choose this for a home or similar location. You can participate in a homegroup, and you can see other network computers.
Connecting to your network and applying settings...

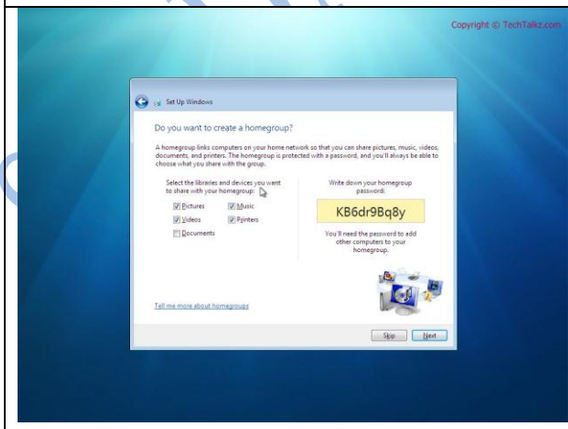
Copyright © TechTalkz.com
Set Up Windows
Do you want to create a homegroup?
A homegroup links computers on your home network so that you can share pictures, music, videos, documents, and printers. The homegroup is protected with a password, and you'll always be able to choose what you share with the group.
Select the libraries and devices you want to share with your homegroup:
Pictures
Music
Videos
Printers
Documents
Write down your homegroup password:
KB6dr9Bq8y
You'll need the password to add other computers to your homegroup.
Tell me more about homegroups
Skip
Next

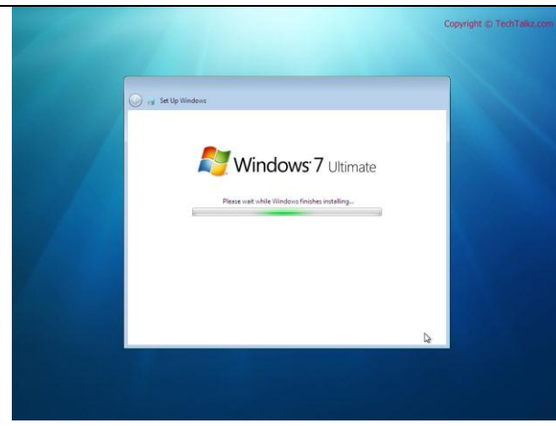
Copyright © TechTalkz.com
Set Up Windows
Windows 7 Ultimate
Please wait while Windows finishes installing...

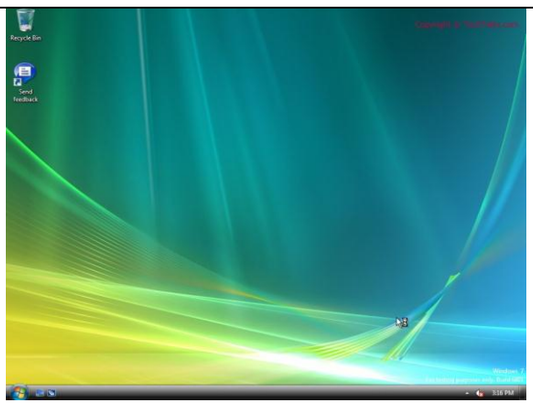

Configure disks, Partitions, volumes, and device drivers

Basic Disks vs Dynamic Disks:

Basic disks and dynamic disks are two types of hard disk configurations in Windows. Most personal computers are configured as basic disks, which are the simplest to manage. Advanced users and IT professionals can make use of dynamic disks, which use multiple hard disks within a computer to manage data, usually for increased performance or reliability.

A basic disk uses primary partitions, extended partitions, and logical drives to organize data. A formatted partition is also called a volume (the terms volume and partition are often used interchangeably). In this version of Windows, basic disks can have either four primary partitions or three primary and one extended partition. The extended partition can contain multiple logical drives (up to 128 logical drives are supported). The partitions on a basic disk cannot share or split data with other partitions. Each partition on a basic disk is a separate entity on the disk.

Dynamic disks can contain a large number of dynamic volumes (approximately 2000) that function like the primary partitions used on basic disks. In Windows 7, you can combine separate dynamic hard disks into a single dynamic volume (called spanning), split data among several hard disks (called striping) for increased performance (no fault tolerance), or duplicate the data on one disk to another (called mirroring).

Let's take a look at these configurations in a little more detail.

- **Spanned volumes** - A spanned volume is a formatted partition which data is stored on more than one hard disk, yet appears as one volume. Spanned volumes are a non-RAID drive architecture. If you extend a simple volume to another dynamic disk, it automatically becomes a spanned volume. You can extend a simple volume to make it a spanned volume, and you can also further extend a spanned volume to add disk storage capacity to the volume. After a volume is spanned, you cannot stripe or mirror it.
- **Striped volumes** - Striped volumes are dynamic volumes that contain disk space from two to thirty-two hard disks. Data that is written to a striped volume is divided by the operating system into chunks of 64KB. The operating system stores each chunk on a separate disk. Since, in a striped volume, a large amount of data is divided into identical portions, it is faster to read or write the data from a striped volume than from a spanned volume. Striped volumes are not fault

SARDAR AZEEM (PICT)

tolerant and are also referred to as RAID-0. A striped volume's capacity is limited to the space available on the disk with the smallest amount of available space.

- **Mirrored volumes** - A mirrored volume, also known as RAID-1, is a fault-tolerant volume that duplicates data on two different physical disks. If one of the disks fails, the data is not lost as an exact copy remains on the surviving disk. Mirroring costs disk space - for example, if you mirror two 100 GB disks, you are left with just 100GB of space rather than 200GB. While great for redundancy, mirroring isn't ideal for performance as all data has to be written twice.

While Windows 7 only supports RAID levels 0 and 1 via software, a 3rd party solution can provide additional RAID services such as RAID-5. For more information about RAID, read Hardware and Software RAID.

Note: Dynamic disks are supported only on computers that use the Integrated Drive Electronics (IDE), Small Computer System Interface (SCSI), Fibre Channel, or Serial Storage Architecture (SSA). Laptops and other mobile devices, removable disks, and disks connected via Universal Serial Bus (USB) or FireWire (IEEE 1394) interfaces are not supported for dynamic storage. Dynamic disks are also not supported on hard drives with a sector size less than 512 bytes and the disk must have at least 1mb of free space for the dynamic disk database.

Disk Management Tool:

Disk Management is a useful built-in Windows 7 partition manager that makes hard disk partitioning quick and simple. Windows 7 Disk Management includes:

- A built-in partition manager
- A graphical user interface (GUI)
- Ability to create new disk partitions within Windows 7
- Ability to shrink existing disk partitions

With Disk Management, you can initialize disks, create volumes, format volumes with file systems FAT, exFAT, FAT32 or NTFS. You can also extend a disk, reduce a disk, check if a disk is healthy or unhealthy, create partitions, delete partitions, or change a drive letter. Disk Management enables you to perform most disk-related tasks without restarting the system, and most changes take effect immediately. To access Disk management, perform these steps:

1. Click *Start*, right-click *Computer*, and click *Manage*.
2. In the left pane click *Disk Management*.

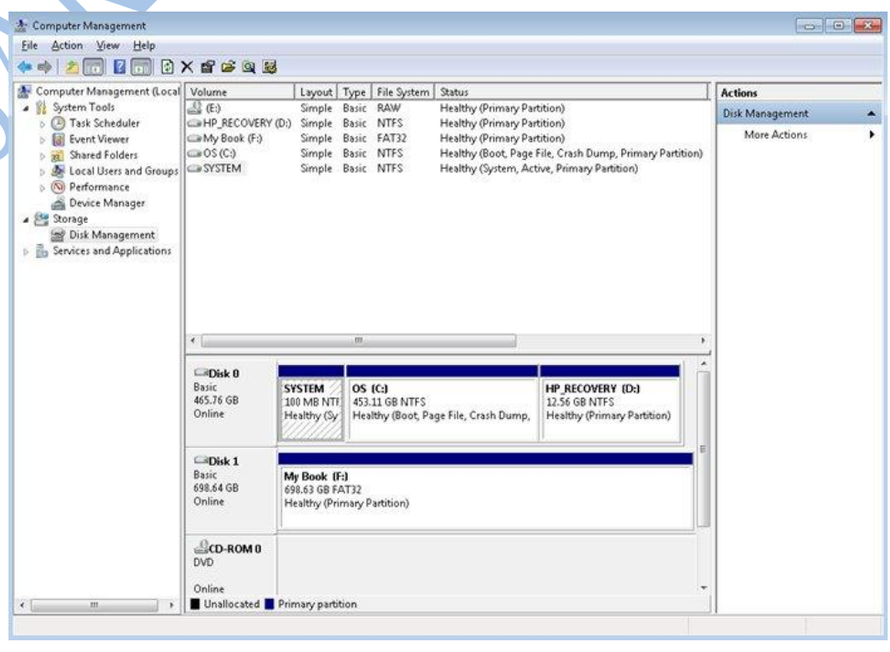

SARDAR AZEEM(PICT)

The Disk Management console shown above lists each volume in alphabetical order. Each hard disk is then broken down into Type, File System, Status, Capacity, and Free Space. In the second horizontal column, each logical drive is labeled by its letter and given a color. Right-clicking on each drive opens a menu where users can extend volumes, shrink volumes, or create new logical drive. You need Administrator or Backup Operator credentials to perform most Disk Management tasks.

Configuring Disks:

Follow the steps below to create a new partition in Windows 7:

1. In the Disk Management Console, right-click the unallocated space and select *New Simple Volume*, and click *Next*.
2. Select the size for the new volume or partition in MB.
3. Assign the drive letter to the new partition.
4. Format the partition with the appropriate file system and select the check box *Perform a Quick Format*. To enable compression, select the check box *Enable File and Folder compression*.
5. Click *Finish*.

Perform these steps to shrink an existing partition in order to create a new partition:

1. In the Disk Management Console, right-click on the partition which you want to resize. The system displays the capacity of the drive and the option to enter an amount you'd like to "shrink" your partition by. Click *Shrink*.
2. You can now see the unallocated space on your hard drive in the capacity you specified, situated just after your now resized original partition.
3. Right-click the unallocated volume, select *New Simple Volume*, assign a drive letter, and quick format the volume using the NTFS file system and default allocation unit size.

Extending a partition:

1. In the Disk Management Console, right-click the partition that you want to extend and select *Extend Volume*.
2. Click *Next*. The system displays the capacity of the drive and the option to enter an amount you'd like to extend your partition by. Click *Next*.
3. Click *Finish*.

Deleting a partition:

1. In the Disk Management Console, right-click the partition that you want to delete and select *Delete Volume*.
2. Click *Yes* to continue the deletion process.
3. Click *Yes* to delete the partition.

Changing the drive letter:

1. In the Disk Management Console, right-click on a partition and select *Change drive letters and paths*.
2. The current drive letter will display. The *Add* button typically allows the partition to be placed inside an existing NTFS folder.
3. Click *Change* to assign a new drive letter.

To convert from an MBR partition to a GPT partition, or vice versa, follow these steps:

1. Back up or move the data on the basic MBR disk you want to convert.
2. Open Computer Management (Local).
3. In the console tree, click *Computer Management* (Local), click *Storage*, and then click *Disk Management*.

SARDAR AZEEM (PICT)

4. The disk must not contain any partitions or volumes. If these exist, right-click any volumes on the disk and then click *Delete Partition* or Delete Volume.
5. Right-click the MBR disk that you want to change into a GPT disk, and then click *Convert to GPT Disk*.

Converting a basic disk to a dynamic disk:

1. In the Disk Management Console, simply right-click the disk you want to convert and click *Convert To Dynamic Disk*. If you want to convert from a dynamic disk to a basic disk, you must first delete all volumes, hence all data, on the disk.

Windows 7 will disallow any of the changes listed above if the partition is currently used as a system, boot, or page file drive.

In addition to the Disk Management Console, Windows also includes a command line utility called DiskPart that can be used to configure disks. For more information, read A Description of the Diskpart Command-Line Utility.

Disk Maintenance:

Windows provides a number of tools that can help keep your disks healthy and optimized. Let's start with Disk Cleanup. If you want to reduce the number of unnecessary files on your hard disk to free up disk space and help your computer run faster, use Disk Cleanup. It removes temporary files, empties the Recycle Bin, and removes a variety of system files and other items that you no longer need. Follow these steps:

1. Open *Computer*, right click on the drive you wish to clean and select *Properties*.

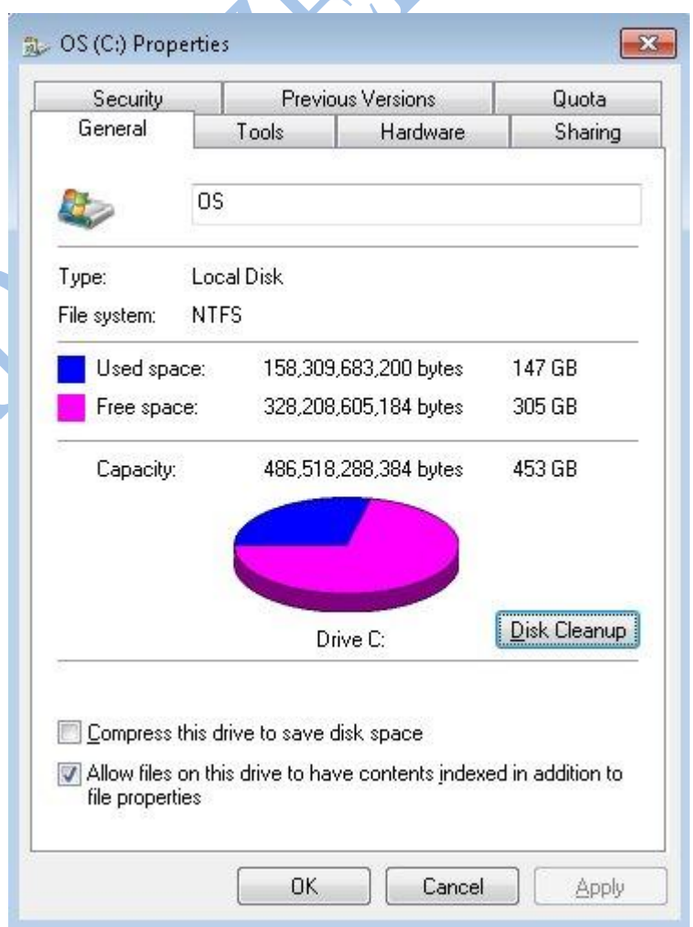

2. Click the *Disk Cleanup* button on the General tab and Windows will calculate how much space it can free up.

3. Select the items from the list that you wish to delete and click *OK*.

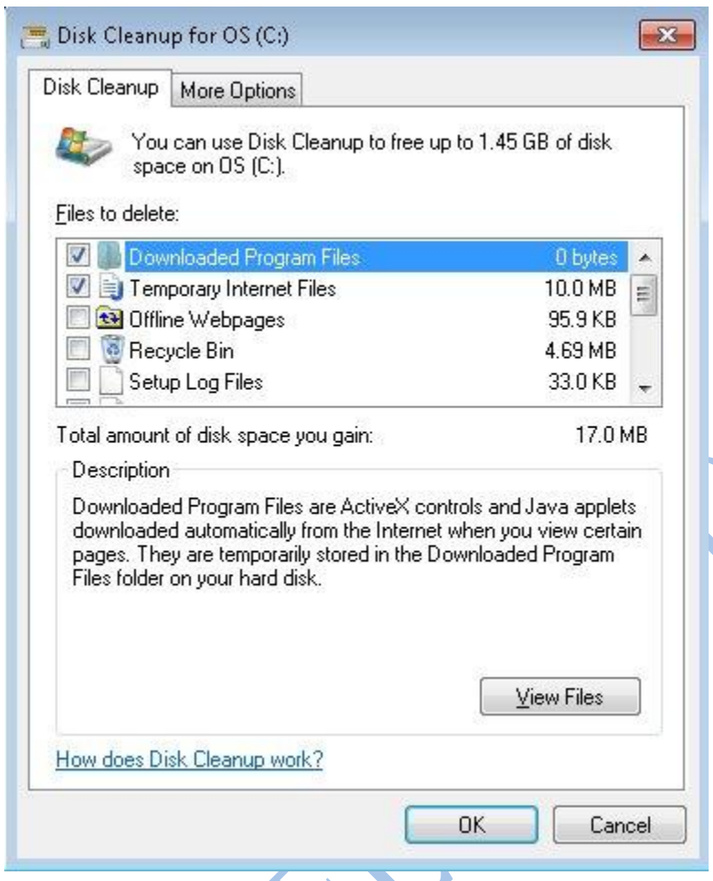

4. When asked "Are you sure you want to permanently delete these files, click *Delete Files*.

As you use your hard drive, it can develop bad sectors. Bad sectors slow down hard disk performance and sometimes make data writing (such as file saving) difficult or even impossible. The Error Checking utility scans the hard drive for bad sectors and scans for file system errors to see whether certain files or folders are misplaced. To scan a disk, follow these steps:

1. Open *Computer*, right click on the drive you wish to check and select *Properties*.
2. Click the *Tools* tab and then click the *Check Now* button.

SARDAR AZEEM (PICT)

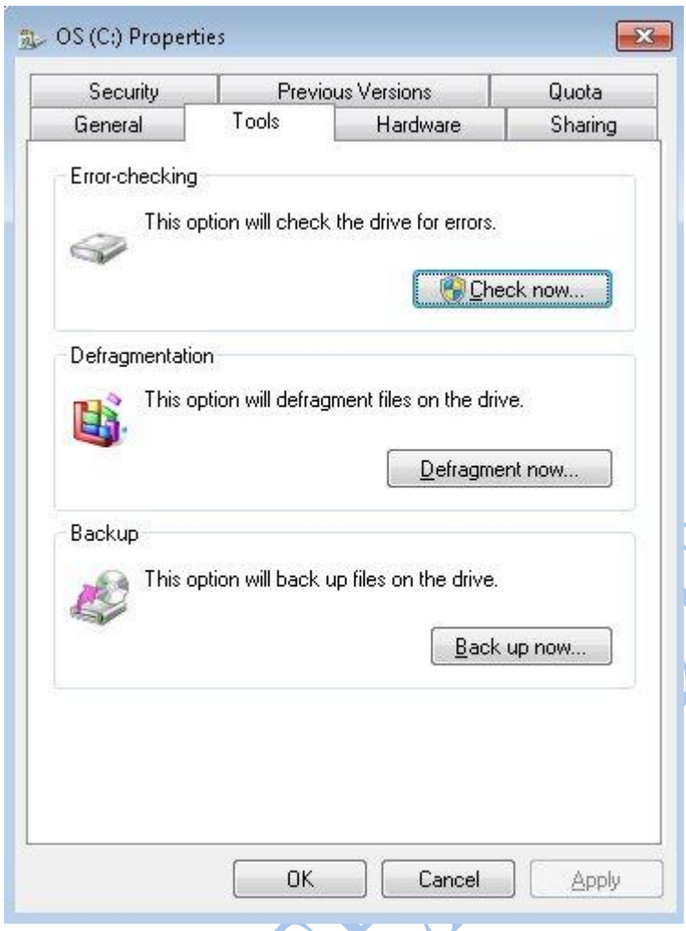

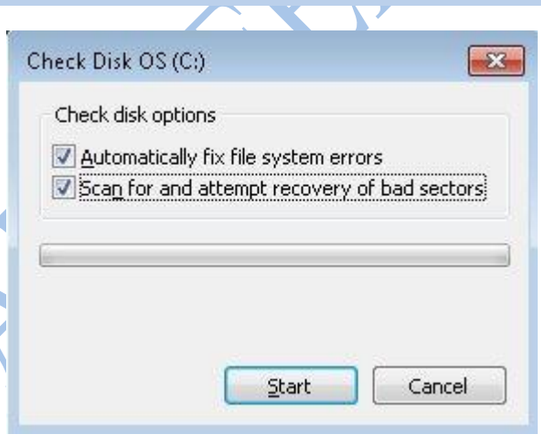

3. Click *Start* to begin checking the drive.

Fragmentation makes your hard disk do extra work that can slow down your computer. Removable storage devices such as USB flash drives can also become fragmented. Disk Defragmenter rearranges fragmented data so your disks and drives can work more efficiently. Disk Defragmenter runs on a schedule, but you can also analyze and defragment your disks and drives manually. To do this, follow these steps:

1. Open *Computer*, right click on the drive you wish to check and select *Properties*.

2. Click the *Tools* tab and then click the *Defragment Now* button.

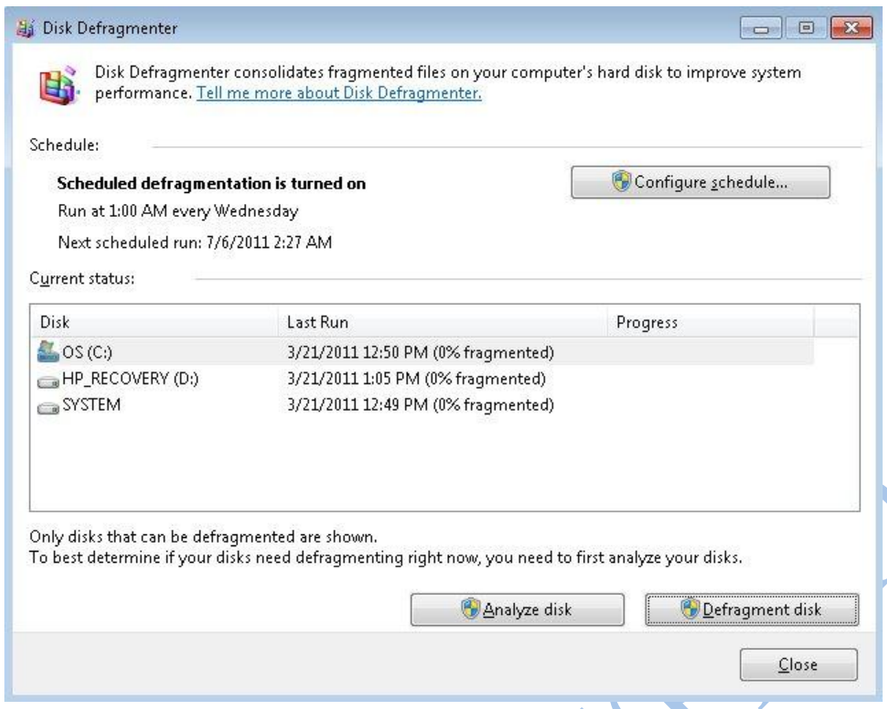

3. On the screen shown above, click on *Analyze disk* button to check the drive for fragmentation.
4. If the drive needs to be defragmented, click the *Defragment disk* button. This process can take a long time.

In most cases, manual defragmentation won't need to be done because, by default, this process is scheduled to run every week. By clicking on the *Configure schedule* button on the screen shown above, you can change the interval and select which drives you want automatically defragmented.

Removable Storage Access Policies:

Removable media can also pose a security threat as it can be lost or stolen, and some administrators may need to lock down client computers' ability to read, write, or execute files on such media. Local and group policy provide a method to prevent or limit users' abilities to interact with removable media. On a stand-alone client computer, you can do this through Local Group Policy Editor. In an enterprise, you would edit domain Group Policy at a domain controller and apply it to all clients in the domain.

To modify these settings in local policy, follow these steps:

1. Click *Start* then type *group policy* into the search box.
2. Click *Edit Group Policy* to open the Local Group Policy Editor.
3. Browse to Computer Configuration/Administrative Templates/System/Removable Storage Access.

SARDARAZEEM(PICT)

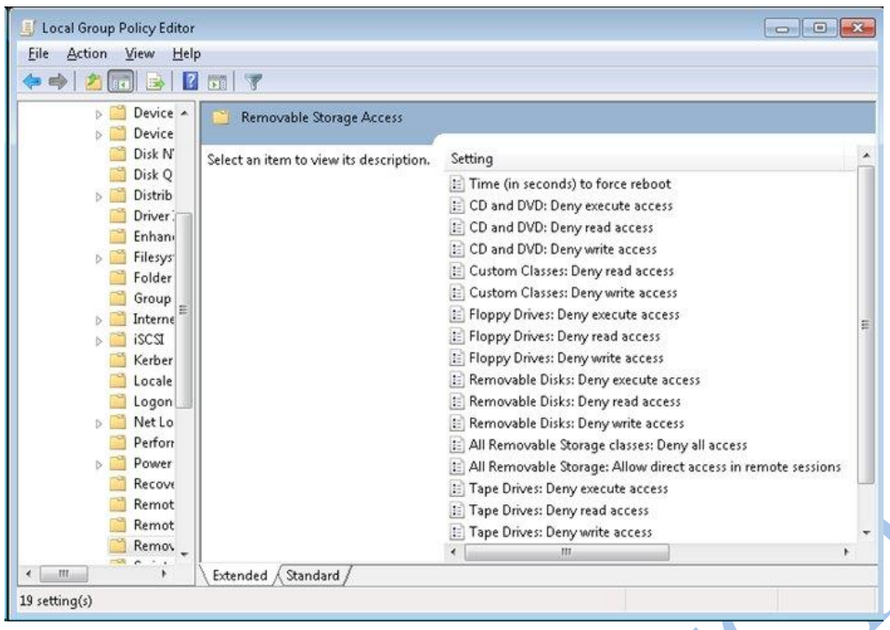

4. Here you can double click on a policy to edit it. These are pretty self-explanatory so we won't go into them here.

Configure file access and printers on a Windows 7 client computer

Device Manager:

Device Manager provides you the facility to graphically view the hardware that is installed on your computer. The Device Manager shows you the devices that are integrated in and connected to your computer, and their drivers. You can use Device Manager to manage devices only on a local computer. Using Device Manager you can:

- Determine whether the hardware on your computer is working properly.
- Change hardware configuration settings.
- Identify the device drivers that are loaded for each device, and obtain information about each device driver.
- Change advanced settings and properties for devices. Install updated device drivers.
- Enable, disable, and uninstall devices.
- Roll back to the previous version of a driver.
- View the devices based on their type, by their connection to the computer, or by the resources they use.
- Show or hide hidden devices that are not critical to view but might be necessary for advanced troubleshooting.

Accessing Device Manager:

There are a few ways to open the device manager and you should be familiar with them. For the purposes of this guide, we are going to open it as follows:

1. Click *Start*.
2. Right click on *Computer* and select *Properties*.
3. In the left menu, click on *Device Manager*

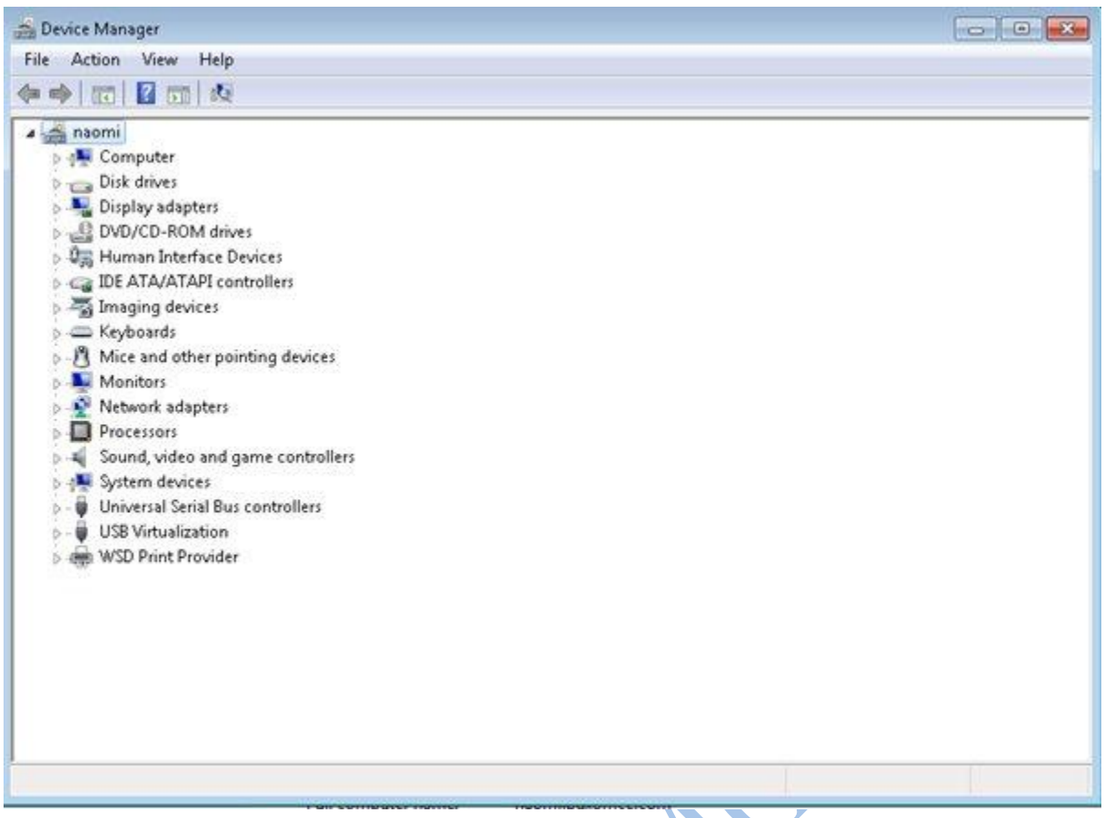

If you see a yellow exclamation point next to a device, it means that there is a conflict or problem with the driver as shown below:

Right clicking on the device will allow you to update, disable, or uninstall a problem driver.

Read Device Manager for more information.

Signed Drivers:

A signed driver is a device driver that includes a digital signature. A digital signature is an electronic security mark that can indicate the publisher of the software, as well as whether someone has changed the original contents of the driver package. If a driver has been signed by a publisher that has verified its identity with a certification authority, you can be confident that the driver actually comes from that publisher and hasn't been altered.

Steps for Staging a Device Driver Package in the Driver Store

Devices and Printers Control Panel:

Another place for device management is in the Devices and Printers Control Panel. Here you can add a device or printer, view and modify drivers, eject devices, and other tasks.

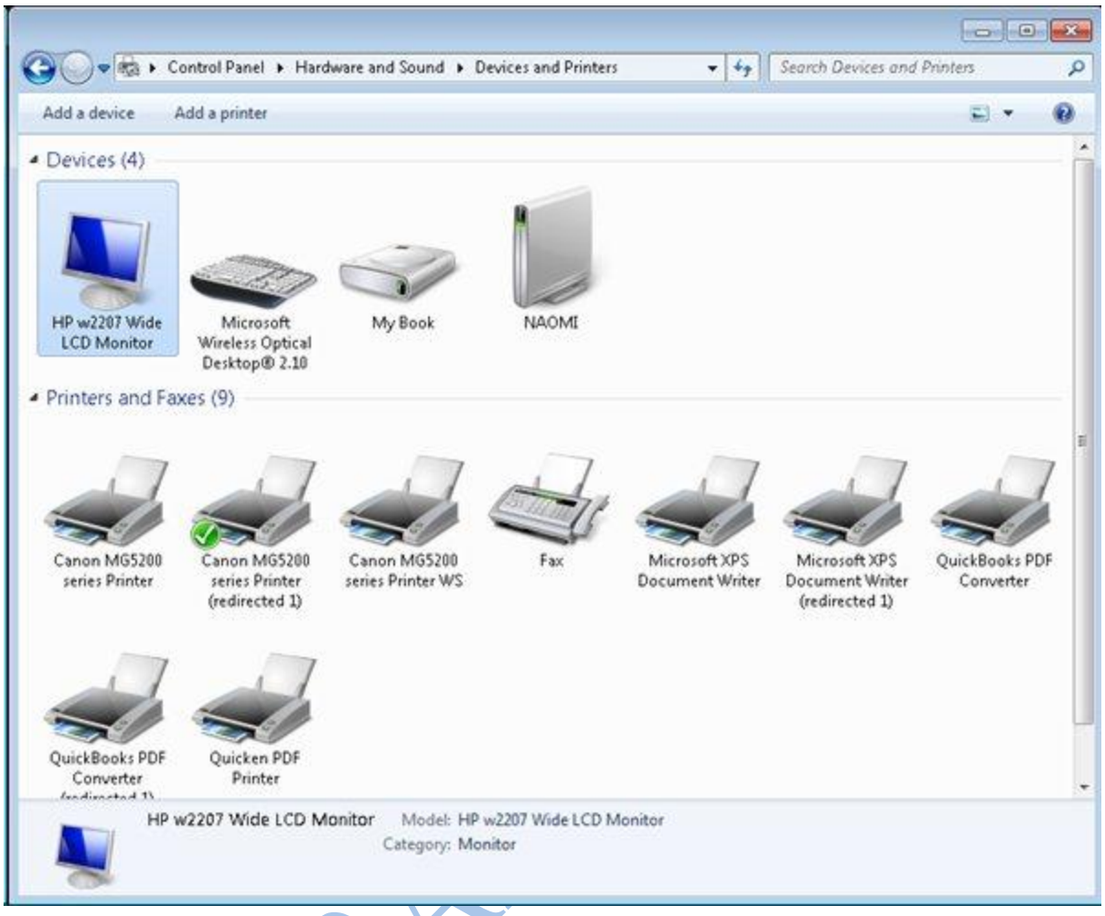

The top menu changes depending on which device is selected.

Configure network connectivity on a windows 7 client computer

Windows 7 networking starts with the Network and Sharing Center which provides a centralized location where you can view, create, and modify local area network (LAN), wireless local area network (WLAN), virtual private network (VPN), dial-up, and Broadband connections on your client and server computers. In addition, you can configure connections to the local computer and sharing options that specify the content that is available to other computers and devices on the network; and you can use Network and Sharing Center tools like Network Map and Network Location to view and specify additional settings about networks and network profiles. It can also be used to troubleshoot network connectivity issues. The Network and Sharing center can be accessed via the control panel.

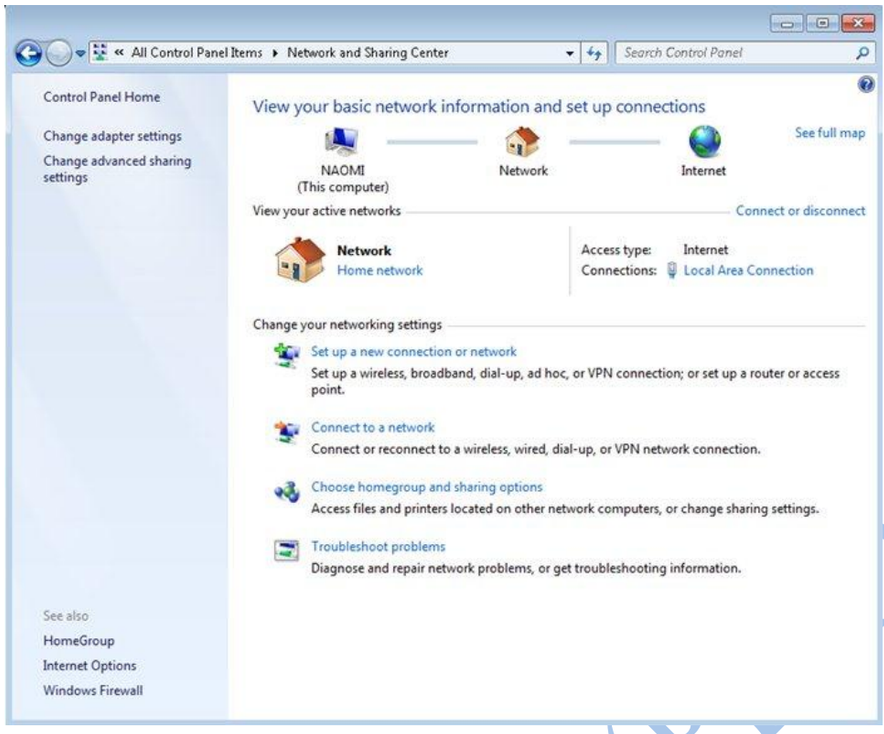

Adding a Network Connection:

If your computer has a network adapter that is connected to a local area network, you do not need to manually create a LAN connection, because Windows automatically creates and configures the connection when you start your computer. Other types of connections such as VPN and dial-up can be configured as follows:

1. Open the *Network and Sharing Center* control panel.
2. In Change your network settings, click *Set up a new connection or network*. The Set up a new connection or network wizard opens.

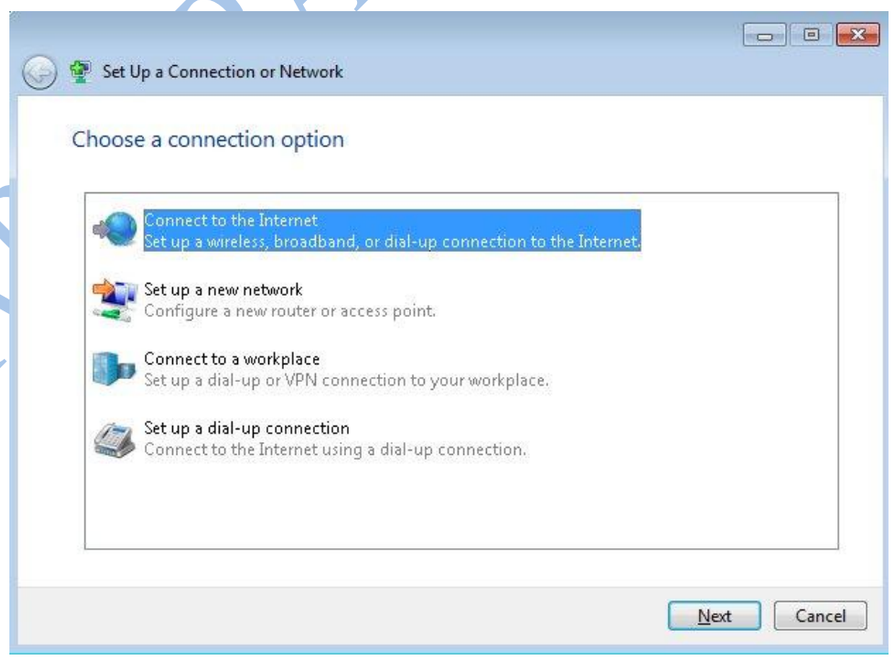

3. Select the type of network or connection you wish to establish and complete the rest of the wizard.

Network Locations:

The first time that you connect to a network, you must choose a network location. This automatically sets the appropriate firewall and security settings for the type of network that you connect to. If you connect to networks in different locations (for example, a network at your home, at a local coffee shop, or at work), choosing a network location can help ensure that your computer is always set to an appropriate security level.

There are four network locations:

- **Home Network** - For home networks or when you know and trust the people and devices on the network. Computers on a home network can belong to a homegroup. Network discovery is turned on for home networks, which allows you to see other computers and devices on the network and allows other network users to see your computer. For more information, see What is network discovery?
- **Work Network** - For small office or other workplace networks. Network discovery, which allows you to see other computers and devices on a network and allows other network users to see your computer, is on by default, but you can't create or join a homegroup.
- **Public Network** - For networks in public places (such as internet cafes or airports). This location is designed to keep your computer from being visible to other computers around you and to help protect your computer from any malicious software on the Internet. HomeGroup is not available on public networks, and network discovery is turned off. You should also choose this option if you're connected directly to the Internet without using a router, or if you have a mobile broadband connection.
- **Domain Network** - Used for domain networks such as those at enterprise workplaces. This type of network location is controlled by your network administrator and can't be selected or changed.

About IPv4:

Once your network connection(s) are established, it is time to configure IP settings. Every IP address can be broken down into 2 parts, the Network ID and the Host ID. All hosts on the same network must have the same netid. Each of these hosts must have a hostid that is unique in relation to the netid. IPv4 addresses are divided into 4 octets with each having a maximum value of 255. We view IPv4 addresses in decimal notation such as 124.35.62.181, but it is actually utilized as binary data.

IP addresses are divided into 3 classes as shown below:

Class	Range
A	1-126
B	<128-191
C	192-223

NOTE: 127.x.x.x is reserved for loopback testing on the local system and is not used on live systems. The following address ranges are reserved for private networks:

10.0.0.0 - 10.254.254.254
172.16.0.0 - 172.31.254.254
192.168.0.0 - 192.168.254.254

About IPv6:

IPv4 has nearly run out of available IP addresses due to the large influx of internet users and expanding networks. As a result, the powers that be had to create a new addressing scheme to deal with this situation and developed IPv6. This new addressing scheme utilizes a 128 bit address (instead of 32) and utilizes a hex numbering method in order to avoid long addresses such as 132.64.34.26.64.156.143.57.1.3.7.44.122.111.201.5. The hex address format will appear in the form of 3FFE:B00:800:2::C for example. The IPv6

equivalent of IPv4's loopback address is 0:0:0:0:0:0:0:1. This can be abbreviated as ::1.

Windows 7 supports both IPv4 and IPv6 through a dual-IP-layer architecture and both are installed and enabled by default. This architecture enables you to tunnel IPv6 traffic across an IPv4 network in addition to tunneling IPv4 traffic across an IPv6 network.

Configuring IP Settings:

You can configure network settings on your Windows 7 computer either statically or dynamically. Static network settings include assigning the IP address, and other related information like gateway, DNS etc manually which enable it to become a part of a network. Dynamic settings make use of Dynamic Host Configuration Protocol (DHCP) to assign IP address and other networking information to your system automatically from a pre-set pool of addresses.

Follow these steps to configure TCP/IP settings manually on your computer:

1. Open Network Connections by clicking the *Start* button, then click *Control Panel* and *Network and Internet* inside the control panel.
2. Click *Network and Sharing Center* on your computer.
3. Click *Local Area Connections* and then click *Properties* to configure network addresses and other information.

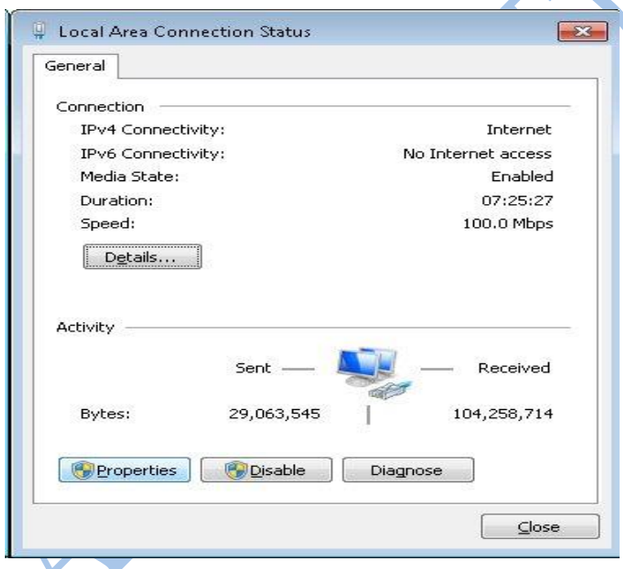

4. Click the *Networking* tab and then, click either *Internet Protocol Version 4 (TCP/IPv4)* or *Internet Protocol Version 6 (TCP/IPv6)* and then click *Properties*.

SARDAR AZEEM

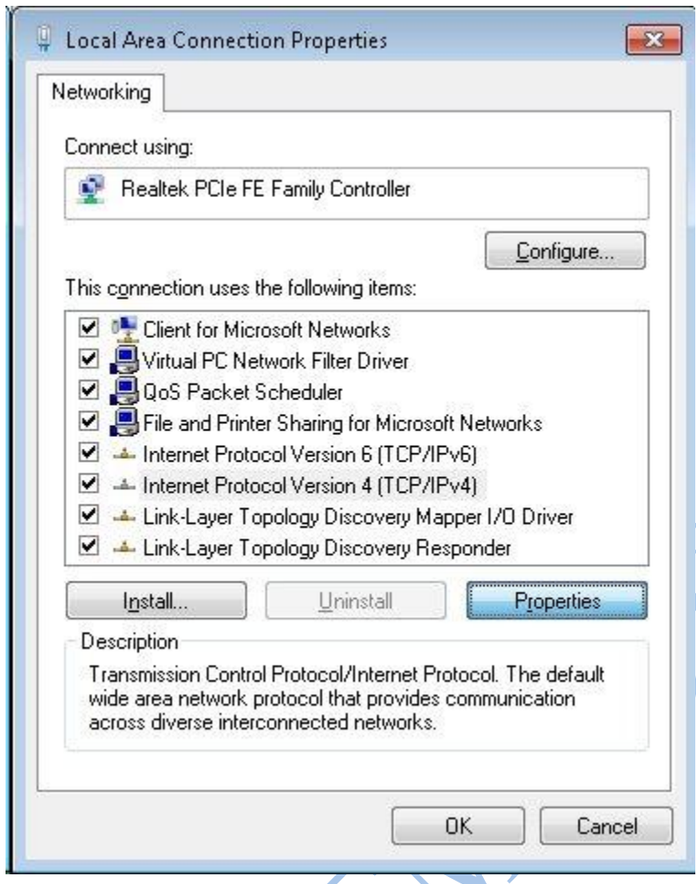

5. To specify IPv4 IP address settings:
   - To configure IP address automatically, click *Obtain an IP address automatically*, and then click *OK*. This option will only work if there is a DHCP server on your network with available addresses to lease.
   - To specify an IP address manually, click *Use the following IP address*, and then, in the IP address, Subnet mask, and Default gateway boxes, type the IP address settings.
6. To specify DNS server address settings:
   - To get the DNS server address automatically, click *Obtain DNS server address automatically*, then click the *Advanced* button. Select the *WINS* tab. Under *NetBIOS setting*, select the *Default* and then click *OK*.
   - To specify a DNS server address manually, click *Use the following DNS server addresses* radio button, and then, for the Preferred DNS server and Alternate DNS server, type the addresses of the primary and secondary DNS servers.
7. Click *OK*. This will make the appropriate changes in the TCP/IP configuration of tour computer.

Automatic Private IP Addressing (APIPA):

A Windows 7 computer that is configured to use DHCP can automatically assign itself an Internet Protocol (IP) address if a DHCP server is not available. For example, this could occur on a network without a DHCP server, or on a network if a DHCP server is temporarily down for maintenance.

The Internet Assigned Numbers Authority (IANA) has reserved 169.254.0.0-169.254.255.255 for Automatic Private IP Addressing. As a result, APIPA provides an address that is guaranteed not to conflict with routable addresses.

After the network adapter has been assigned an IP address, the computer can use TCP/IP

to communicate with any other computer that is connected to the same LAN and that is also configured for APIPA or has the IP address manually set to the 169.254.x.y (where x.y is the client's unique identifier) address range with a subnet mask of 255.255.0.0. Note that the computer cannot communicate with computers on other subnets, or with computers that do not use automatic private IP addressing. This also means that a computer with an APIPA address cannot connect to the internet, only other computers with APIPA addresses. Automatic private IP addressing is enabled by default.

If a DHCP enabled computer is using an IP address in the APIPA range, it often indicates that the computer is unable to contact the DHCP server.

Link-Local Multicast Name Resolution:

LLMNR is a Microsoft designed protocol that can be used on private networks where there is no DNS server, as a mechanism for providing name resolution like DNS does. It is one of many protocols that do similar things for zero-configuration networks - they basically allow private networks to function as IP networks without requiring hosts to be configured with addresses.

Resolving Connectivity Issues:

Windows offers a number of tools and utilities for troubleshooting connectivity and other network problems. A good place to start is by clicking on the *Troubleshoot problems* option in the Network and Sharing Center. This opens the Windows Network Diagnostics tool. If Windows 7 detects the problem, it may be able to automatically fix it, or possibly offer a solution.

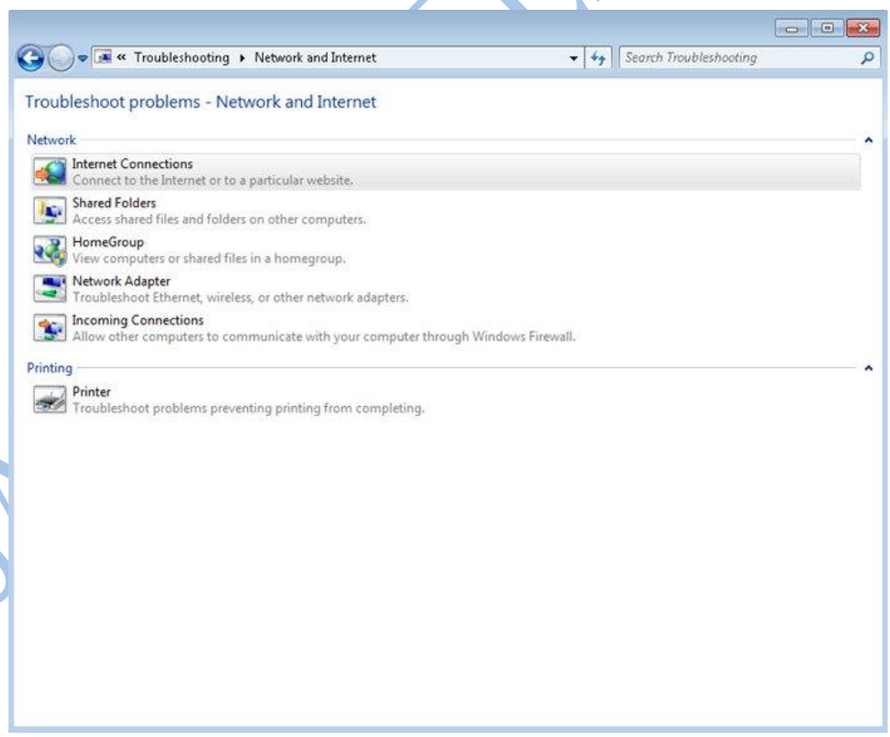

If this wizard is unable to fix the problem or offer a solution, there are a number of other tools listed below that can help.

- **IPCONFIG** - This command is used to view network settings from a Windows computer command line. Below are the ipconfig switches that can be used at a command prompt:
  - *ipconfig /all* will display all of your IP settings.
  - *ipconfig /renew* forces the DHCP server, if available to renew a lease.
  - *ipconfig /renew6* renews n IPv6 DHCP lease.
  - *ipconfig /flushdns* purges the DNS resolver cache.

  - *ipconfig /registerdns* refreshes all DHCP leases and reregisters DNS names.
  - *ipconfig /release* forces the release of a lease.
- **PING (Packet InterNet Groper)** - PING is a command-line utility used to verify connections between networked devices. PING uses ICMP echo requests that behave similarly to SONAR pings. The standard format for the command is *ping [IP address or hostname]*. If successful, the ping command will return replies from the remote host with the time it took to receive the reply. If unsuccessful, you will likely recieve and error message. This is one of the most important tools for determining network connectivity between hosts.

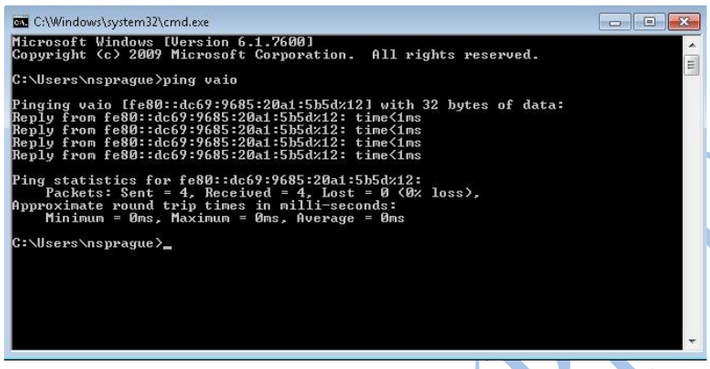

- **TRACERT** - A command-line troubleshooting tool that enables you to view the route to a specified host. This will show how many hops the packets have to travel and how long it takes. Basic usage looks like: *tracert [IP address or hostname]*.

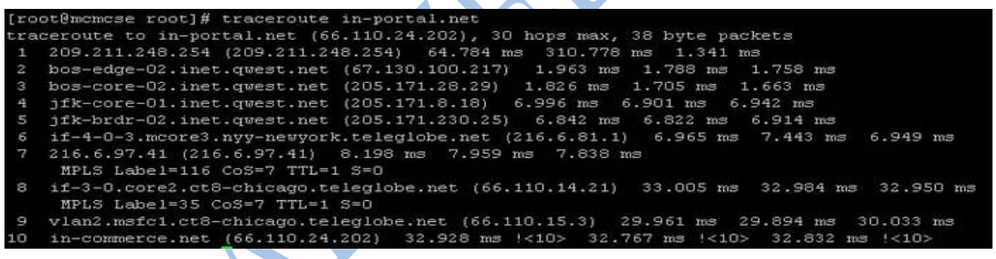

- **PATHPING** - This tool is very similar to tracert, however, pathping provides more detailed statistics on individual hops.
- **ARP (Address Resolution Protocol)** - A host PC must have the MAC and IP addresses of a remote host in order to send data to that remote host, and it's ARP that allows the local host to request the remost host to send the local host its MAC address through an ARP Request. The *ARP -a [IP Address]* command will show you the MAC address associated with a computer or device's IP address.

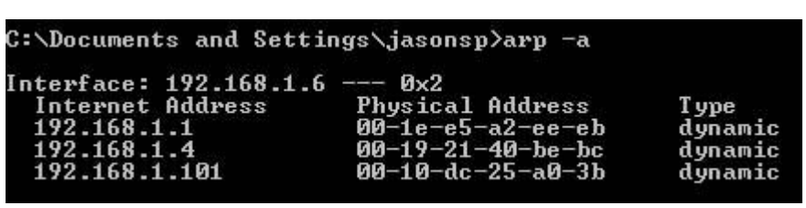

- **NSLOOKUP** - This is a command that queries a DNS server for machine name and address information. To use nslookup, type *nslookup [IP address or computer name or domain name]*. NSLOOKUP will return the name, all known IP addresses and all known aliases (which are just alternate names) for the identified machine. NSLOOKUP is a useful tool for troubleshooting DNS problems.

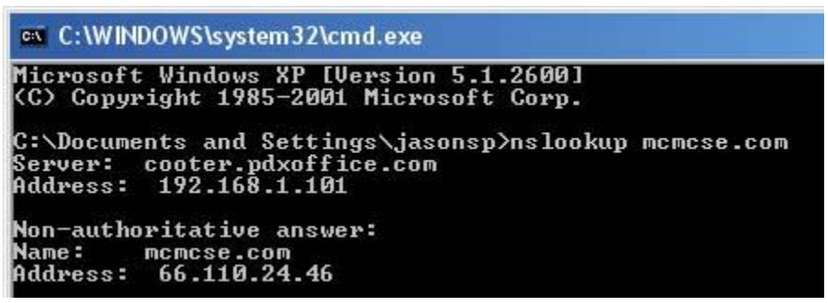

Configure wireless network connectivity on a windows 7 client computer.(Bluetooth, Wi Fi)

To configure a PEAP-TLS wireless profile for computers running Windows 7 and Windows Vista

1. Open the New Wireless Network (IEEE 802.11) Policies Properties dialog box.

2. On the **General** tab, in **Policy Name**, type a new name for your policy, or leave the default.

3. In **Description**, type a description of your policy.

4. Select **Use Windows to configure wireless network settings for clients** to specify that WLAN AutoConfig is used to configure wireless network adapter settings.

5. On the **General** tab, do one of the following:

   - To add and configure a new profile, click **Add**, and then select **Infrastructure**.
   - To edit an existing profile, select the profile you want to modify, and then click **Edit**.

6. On the **Connection** tab, in **Profile Name**, if you are adding a new profile, type a name for the profile. If you are editing a profile that is already added, use the existing profile name, or modify the name as needed.

7. In **Network Name(s) (SSID)**, type the service set identifier (SSID) for your wireless APs, and then click **Add**.

   If your deployment uses multiple SSIDs and each wireless AP uses the same wireless security settings, repeat this step to add the SSID for each wireless AP to which you want this profile to apply.

   If your deployment uses multiple SSIDs and the security settings for each SSID do not match, configure a separate profile for each group of SSIDs that use the same security settings. For example, if you have one group of wireless APs configured to use WPA2-Enterprise and AES, and another group of wireless APs to use WPA-Enterprise and TKIP, configure a profile for each group of wireless APs.

SARDAR AZEEM(PICT)

8. To specify that wireless clients automatically connect to wireless APs for which the SSID is specified in **Network Name(s) (SSID)**, select **Connect automatically when this network is in range**.

9. To specify that wireless clients connect to networks in order of preference, select **Connect to a more preferred network if available**.

10. If you deployed wireless access points that are configured to suppress the broadcast beacon, select **Connect even if the network is not broadcasting**.

> **Security Note**
>
> Enabling this option can create a security risk because wireless clients will probe for and attempt connections to any wireless network. By default, this setting is not enabled.

11. Click the **Security** tab. In **Select the security methods for this network**, in **Authentication**, select **WPA2-Enterprise** if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **WPA-Enterprise**.

> **Note**
>
> Selecting WPA2 exposes settings for Fast Roaming that are not displayed if WPA is selected. The default settings for Fast Roaming are sufficient for most wireless deployments.

12. In **Encryption**, select **AES**, if it is supported by your wireless AP and wireless client network adapters. Otherwise, select **TKIP**.

> **Note**
>
> The settings for both **Authentication** and **Encryption** must match the settings configured on your wireless AP.

13. In **Select a network authentication method**, select **Microsoft: Protected EAP (PEAP)**.

14. In **Authentication mode**, select from the following, depending on your needs: **User or Computer authentication**, **Computer authentication**, **User authentication**, **Guest authentication**. By default, **User or Computer authentication** is selected.

15. In **Max Authentication Failures**, specify the maximum number of failed attempts allowed before the user is notified that authentication has failed. By default, this value is set to "1."

16. To specify that user credentials are held in cache, select **Cache user information for subsequent connections to this network**.

17. Click **Advanced**, and then configure the following:

a. To configure advanced 802.1X settings, in **IEEE 802.1X**, select **Enforce advanced 802.1X settings**, and then configure the following settings, depending on your needs: **Max Eapol-Start Msgs**, **Held Period**, **Start Period**, and **Auth Period**.

When the advanced 802.1X settings are enforced, the default values are sufficient for most wireless deployments.

b. To enable Single Sign On, select **Enable Single Sign On for this network**.

c. To specify when Single Sign On occurs, select either **Perform immediately before User Logon** or **Perform immediately after User Logon**, depending on your needs.

The remaining default values in **Single Sign On** are sufficient for typical wireless deployments.

d. To specify the maximum amount of time, in seconds, in which 802.1X authentication must complete and authorize network access, in **Max delay for connectivity (seconds)**, enter a value, depending on your needs.

e. To allow dialogs during Single Sing On, select **Allow additional dialogs to be displayed during Single Sign On**.

f. To specify that wireless computers are placed on one virtual local area network (VLAN) at startup, and then transitioned to a different network after the user logs on to the computer, select **This network uses different VLAN for authentication with machine and user credentials**.

g. To enable Fast Roaming, in **Fast Roaming**, select **Enable Pairwise Master Key (PMK) Caching**. The default values for **PMK Time to Live (minutes)** and **Number of entries in PMK Cache** are typically sufficient for Fast Roaming.

h. Select **This network uses pre-authentication**, if your wireless AP is configured for pre-authentication. The default value of 3 is typically sufficient for **Maximum Pre-authentication attempts**.

i. To specify that cryptography adheres to the FIPS 140-2 certified mode, select **Perform cryptography in FIPS 140-2 certified mode**.

18. Click **OK** to save your settings and return to the **Security** tab.

19. Click **Properties**. The **Protected EAP Properties** dialog box opens.

20. In **Protected EAP Properties**, verify that **Validate server certificate** is selected.

21. In **Trusted Root Certification Authorities**, select the trusted root certification authority (CA) that issued the server certificate to your server running Network Policy Server (NPS).

SARDAR AZEEM(PICT)

Note

This setting limits the trusted root CAs that clients trust to the selected CAs. If no trusted root CAs are selected, then clients trust all root CAs listed in their trusted root certification authority store.

22. To specify which Remote Authentication Dial-In User Service (RADIUS) servers your wired access clients must use for authentication and authorization, in **Connect to these servers**, type then name of each RADIUS server, exactly as it appears in the subject field of the server certificate. Use semicolons to specify multiple RADIUS server names.

23. For improved security and a better user experience, select **Do not prompt user to authorize new servers or trusted certification authorities**.

24. In **Select Authentication Method**, select **Smart Card or other certificate**.

25. To enable PEAP Fast Reconnect, select **Enable Fast Reconnect**.

26. To specify that Network Access Protection (NAP) performs system health checks on clients to ensure they meet health requirements, before connections to the network are permitted, select **Enforce Network Access Protection**.

27. To require cryptobinding Type-Length-Value (TLV), select **Disconnect if server does not present cryptobinding TLV**.

28. To configure your clients so that they will not send their identity in plaintext before the client has authenticated the RADIUS server, select **Enable Identity Privacy**, and then in **Anonymous Identity**, type a name or value, or leave the field empty.

    For example, if **Enable Identity Privacy** is enabled and you use “guest” as the anonymous identity value, the identity response for a user with identity alice@realm is guest@realm. If you select **Enable Identity Privacy** but do not provide an anonymous identity value, the identity response is @realm.

29. Click **Configure**. In the **Smart Card or other Certificate Properties** dialog box, in **When connecting**, select either **Use my smart card** or select both **Use a certificate on this computer** and **Use simple certificate selection (Recommended)**.

30. To require that access clients validate the NPS server certificate, select **Validate server certificate**.

31. To specify which RADIUS servers your wired access clients must use for authentication and authorization, in **Connect to these servers**, type then name of each RADIUS server, exactly as it appears in the subject field of the server’s certificate. Use semicolons to specify multiple RADIUUS server names.

32. In **Trusted Root Certification Authorities**, select the CA that issued certificates to your NPS servers.

SARDAR AZEEM(PICT)

33. To specify that clients use an alternate name for the access attempt, select **Use a different user name for the connection**.

34. To prevent users from being prompted to trust a server certificate if that certificate is incorrectly configured, is not already trusted, or both, select **Do not prompt user to authorize new servers or trusted certification authorities**. (Recommended)

35. Click **OK** to close the **Smart card or other Certificate Properties** dialog box, and then click **OK** again to close the **Protected EAP (PEAP) Properties** dialog box, returning you to **New Wireless Network Policy Properties**.

Secure Windows 7 client desktop computers.

The Windows operating system helps protect files, applications, and other resources from unauthorized use through a process of matching user accounts and group membership against the rights, privileges, and permissions associated with those accounts and group memberships. The topics in this section will show you how to assign or set privileges and permissions. In addition, understanding privileges and permissions, why they are necessary, and how they function can help you manage shared resources effectively. Understanding these processes can also help you avoid unnecessary risks and troubleshoot any access control problems you might encounter.

Access control is the process of authorizing users, groups, and computers to access objects on the network or computer.

To understand and manage access control, you need to understand the relationship between:

- Objects (files, printers, and other resources)
- Access tokens
- Access control lists (ACLs) and access control entries (ACEs)
- Subjects (users or applications)
- The operating system
- Permissions
- User rights and privileges

Before a subject can gain access to an object, the subject must identify itself to the security subsystem for the operating system. This identity is contained within an access token that is re-created every time a subject logs on. Before allowing the subject to access an object, the operating system checks to determine whether the access token for the subject is authorized to access the object and complete the desired task. It does this by comparing information in the access token with access control entries (ACEs) for the object.

ACEs can allow or deny a number of different behaviors, depending on the type of object. For example, options on a file object can include Read, Write, and Execute. On a printer, the ACEs that are available include Print, Manage printers, and Manage documents.

SARDARAZEEM(PICT)

Individual ACEs for an object are combined in an access control list (ACL). The security subsystem checks the object's ACL for ACEs that apply to the user and the groups that the user belongs to. It steps through each ACE until it finds one that either allows or denies access to the user or one of the user's groups, or until there are no more ACEs to check. If it comes to the end of the ACL and the desired access is still not explicitly allowed or denied, the security subsystem denies access to the object.

Permissions

Permissions define the type of access granted to a user or group for an object or object property. For example, the Finance group can be granted Read and Write permissions for a file named Payroll.dat.

Using the access control user interface, you can set NTFS permissions for objects such as files, Active Directory objects, registry objects, or system objects such as processes. Permissions can be granted to any user, group, or computer. It is a good practice to assign permissions to groups because it improves system performance when verifying access to an object.

For any object, you can grant permissions to:

- Groups, users, and other objects with security identifiers in the domain.
- Groups and users in that domain and any trusted domains.
- Local groups and users on the computer where the object resides.

The permissions attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. Some permissions, however, are common to most types of objects. These common permissions are:

- Read
- Modify
- Change owner
- Delete

When you set permissions, you specify the level of access for groups and users. For example, you can let one user read the contents of a file, let another user make changes to the file, and prevent all other users from accessing the file. You can set similar permissions on printers so that certain users can configure the printer and other users can only print.

When you need to change the permissions on a file, you can run Windows Explorer, right-click the file name, and click **Properties**. On the **Security** tab, you can change permissions on the file. For more information, see Managing Permissions.

Ownership of objects

An owner is assigned to an object when that object is created. By default, the owner is the creator of the object. No matter what permissions are set on an object, the owner of the object can always change the permissions on an object. For more information, see Managing Object Ownership.

Inheritance of permissions

Inheritance allows administrators to easily assign and manage permissions. This feature automatically causes objects within a container to inherit all the inheritable permissions of that container. For example, the files within a folder, when created, inherit the permissions of the folder. Only permissions marked to be inherited will be inherited.

SARDARAZEEM(PICT)

User rights and privileges

User rights grant specific privileges and logon rights to users and groups in your computing environment. Administrators can assign specific rights to group accounts or to individual user accounts. These rights authorize users to perform specific actions, such as logging on to a system interactively or backing up files and directories.

User rights are different from permissions because user rights apply to user accounts, and permissions are attached to objects. Although user rights can apply to individual user accounts, user rights are best administered on a group account basis. There is no support in the access control user interface to grant user rights; however, user rights assignment can be administered through the Local Security Policy snap-in under **Local Policies\User Rights Assignment**. For more information, see User Rights and Privileges.

Object auditing

With administrator's rights, you can audit users' successful or failed access to objects. You can select which object access to audit by using the access control user interface, but first you must enable the audit policy by selecting **Audit object access** under **Local Policy\Audit Policy\Local Policies** in the Local Security Policy snap-in. You can then view these security-related events in the Security log in Event Viewer.

Optimize and maintain the performance and reliability of a windows 7 client computer

Windows Performance Monitor

You can use Windows Performance Monitor to examine how programs you run affect your computer's performance, both in real time and by collecting log data for later analysis.

Windows Performance Monitor uses performance counters, event trace data, and configuration information, which can be combined into Data Collector Sets.

Performance counters are measurements of system state or activity. They can be included in the operating system or can be part of individual applications. Windows Performance Monitor requests the current value of performance counters at specified time intervals.

Event trace data is collected from trace providers, which are components of the operating system or of individual applications that report actions or events. Output from multiple trace providers can be combined into a **trace session**.

Configuration information is collected from key values in the Windows registry. Windows Performance Monitor can record the value of a registry key at a specified time or interval as part of a log file.

- Overview of Windows Performance Monitor
- Using Performance Monitor
- Creating Data Collector Sets
- Scheduling and Managing Data in Windows Performance Monitor
- User Interface: Windows Performance Monitor

Using Performance Monitor

Performance Monitor is a simple yet powerful visualization tool for viewing performance data, both in real time and from log files. With it, you can examine performance data in a graph, histogram, or report.

Membership in the local **Performance Log Users** group, or equivalent, is the minimum required to complete this procedure.

To start Performance Monitor

1. Click **Start**, click in the **Start Search** box, type **perfmon**, and press ENTER.

2. In the navigation tree, expand **Monitoring Tools**, and then click **Performance**

Monitor.

You can also use Performance Monitor to view real-time performance data on a remote computer.

Membership in the target computer's **Performance Log Users** group, or equivalent, is the minimum required to complete this procedure.

To connect to a remote computer with Performance Monitor

1. Start Performance Monitor.

2. In the navigation tree, right-click **Reliability and Performance**, and then click **Connect to another computer**.

3. In the **Select Computer** dialog box, type the name of the computer you want to monitor, or click **Browse** to select it from a list.

4. Click **OK**.

Creating Data Collector Sets

A Data Collector Set is the building block of performance monitoring and reporting in Windows Performance Monitor. It organizes multiple data collection points into a single component that can be used to review or log performance. A Data Collector Set can be created and then recorded individually, grouped with other Data Collector Set and incorporated into logs, viewed in Performance Monitor, configured to generate alerts when thresholds are reached, or used by other non-Microsoft applications. It can be associated with rules of scheduling for data collection at specific times. Windows Management Interface (WMI) tasks can be configured to run upon the completion of Data Collector Set collection.

Data Collector Sets can contain the following types of data collectors:

- Performance counters
- Event trace data
- System configuration information (registry key values)

Scheduling and Managing Data in Windows Performance Monitor

You schedule data collection and manage the storage of log data on a Data Collector Set basis in Windows Performance Monitor. Reports can be stored after log data has been deleted, which gives you access to performance statistics without storing individual counter values.

- Schedule Data Collection in Windows Performance Monitor

SARDAR AZEEM(PICT)

- Manage Data in Windows Performance Monitor
- View Reports in Windows Performance Monitor

Remote access setting for a windows 7 client computer.

Introduction to Remote Desktop:

The Remote Desktop (RDP) utility allows you to connect to a computer on a network and access all of your programs, files, and network resources as if you were sitting in front of that computer. Remote desktop is often used by security professionals to administer servers, and fix problems on client computers without having to be in front of them. In fact, they could be in another country.

You cannot use Remote Desktop Connection to connect to computers running Windows 7 Starter, Windows 7 Home Basic, and Windows 7 Home Premium. In other words, only Windows 7 Professional, Ultimate, and Enterprise editions allow a computer to connect to them via RDP. All versions of Windows 7 have the Remote Desktop client software that allows them to make outgoing connections.

To use Remote Desktop and Remote Assistance, you have to use TCP port 3389. Therefore, it needs to be opened using the Windows Firewall and any other firewalls between your computer and the remote host. Additional requirements include:

- You must have permission to connect to the remote computer.
- The remote computer must be turned on or have Wake on LAN enabled.
- Both computers must be connected to a network.
- The remote computer must be configured to accept incoming connections (see next section). By default this is turned off.

Enabling Remote Desktop Connections:

Follow these steps to enable the remote desktop connection in Windows 7:

1. Click *Start*, then right-click the Computer and select *properties*.
2. Click the *Remote settings* option in the window.
3. Enable *Allow connections from computers running any version of Remote Desktop* in the System Properties dialogue box.
4. Click *Apply* and the remote desktop connections feature will be enabled on your Windows 7 computer.

Note: When you enable Remote Desktop, Windows Firewall automatically updates rules to allow Remote Desktop connections to be made to the computer. If you reset Windows Firewall to its default settings, the firewall will no longer allow connections. Simply disable and then re-enable Remote Desktop to correct this problem.

Establishing a Remote Desktop Connection:

1. Click *Start* and then click *All Programs* and then click *Accessories*.

SARDAR AZEEM(PICT)

2. Click the *Remote Desktop Connection* option. The Remote Desktop dialogue box is displayed; specify the IP address or hostname of the remote machine to which you want to connect.
3. Click *Connect* and if the computer is running and remote connections are enabled on it, a connection will be made.
4. Enter the the username and password for that computer.

If a user other than yourself is logged into the remote machine, they will be presented with an alert that someone is trying to establish a remote desktop connection with the computer. They can choose to accept the connection or not.

Configuring Remote Desktop:

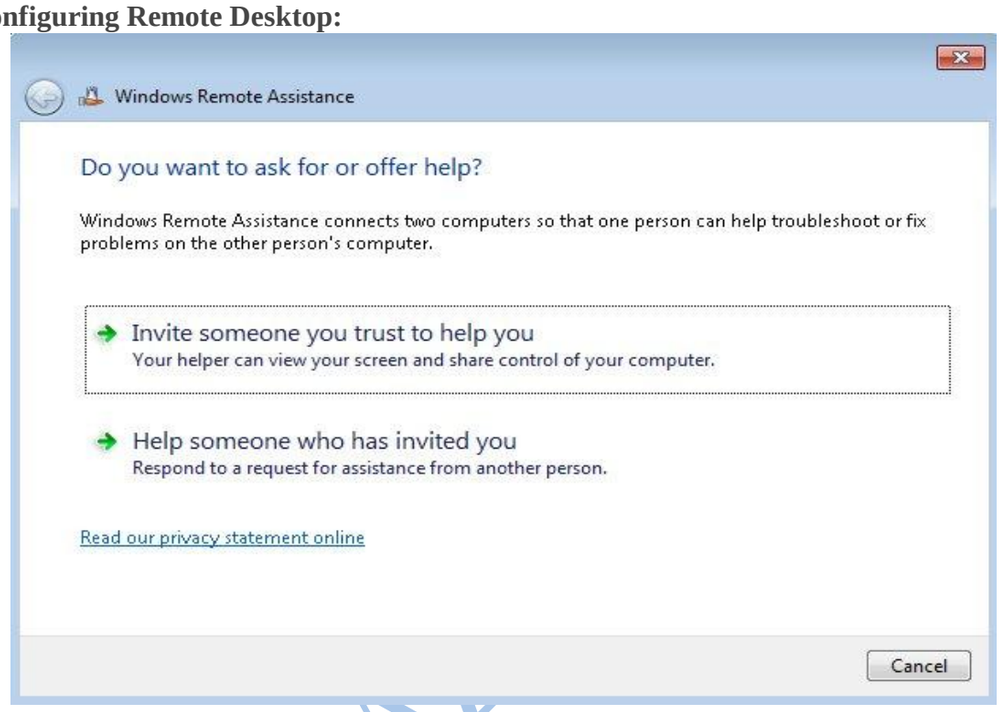

To configure remote access, follow these steps:

1. In the Control Panel, first click *System And Security*, and then click *System*.
2. Click *Remote Settings* in the left pane and the System Properties dialog box to the Remote tab opens.
3. If you want to disable the Remote Desktop, select *Don't Allow Connections To This Computer*, then click *OK* and skip the remaining steps.
4. To enable Remote Desktop, choose either of the two options:
   - Select *Allow Connections From Computers Running Any Version Of Remote Desktop*. This allows connections from any version of Windows.
   - Select *Allow Connections Only From Computers Running Remote Desktop With Network Level Authentication*. This allows connections only from Windows 7 or later computers and computers with secure network authentication.
5. Click *Select Users* to open the Remote Desktop Users dialog box. To grant Remote Desktop access to any user, click *Add*. This opens the *Select Users* dialog box. In the Select Users dialog box, click *Locations* to select the computer or domain in which the users are located with whom you want to work. Type the name of a user and enter the object names to the selected fields, and then click *Check Names*. If matches are found, then you can select the account you want to use and then click *OK*. If no matches are found, update the name you entered and search again.
6. To revoke remote access permissions for any user account, select the account and then click *Remove*.
7. Click *OK* when you have finished.

Remote Assistance:

Since Windows XP, Remote Assistance has provided a handy way to get or give a helping hand from a distance. The main difference between Remote Desktop and Remote Assistance is that with Remote Assistance, the remote user must request a connection, and when connected, both the local user and the remote user can see what is happening on the screen at the same time. Windows Remote Assistance is enabled by default on computers running Windows 7.

In previous versions of Windows, the primary way to initiate a Remote Assistance connection was by creating an "invitation" file with info on how to find and connect to your system, and sending it the person you are requesting help from via e-mail. You can still use invitation files in Windows 7, and if your helper is running Vista or XP, you'll have to. However, if both parties have Windows 7, a new feature called Easy Connect can simplify the connection process by eliminating e-mail as a middleman.

To request remote assistance in Windows 7, search for *assistance* from the Start menu, then click *Windows Remote Assistance*.

After you click *Invite someone you trust to help you*, you'll see *Easy Connect* along with the two e-mail-based invitation options. Choose *Easy Connect*, and then you should see a Windows Remote Assistance window displaying the 12-character password needed for access to your computer. This automatic password generation is another new Windows 7 feature, and it occurs whether you use Easy Connect or invitations. It forces you to use a strong password to increase security over that in Vista or XP.

If Easy Connect is grayed out, one of the following reasons may be the cause.

- Both computers aren't running Windows 7. In order to use Easy Connect with Remote Assistance, both of the computers must be running Windows 7.
- Access to the Internet is limited. If access to the Internet is limited on either computer, Easy Connect is disabled. Internet access might be limited if you're on a corporate network.
- Your router doesn't support Easy Connect. Easy Connect uses the Peer Name Resolution Protocol (PNRP) to transfer the Remote Assistance invitation over the Internet. One possible issue is that your router doesn't support UPnP, or doesn't have it enabled. You may also want to try enabling port 3540 (UDP) on your router. You can check your router by using the Internet Connectivity Evaluation Tool on the Microsoft website. If you're running Windows Server, you need to install the Peer Name Resolution Protocol.

Once connected and with the remote user's permission, you can "remote control" their computer like you would with Remote Desktop, only the user will be able to see what you are doing.

After you've successfully established a Remote Assistance session with someone via Easy Connect, connecting to that person in the future will be even easier. The next time you run Remote Assistance you'll see a list of people you've previously connected to. Select a name and the Windows Remote Assistance window will launch, and when your helper connects to you, you'll be connected without having to see or enter a password because the one from your last session is cached. This subsequent connect feature only works when the helper is using the same computer they were on initially.

Windows Remote Management Service:

The Windows Remote Management service allows you to execute commands on a remote computer, either from the command prompt using WinRS or from Windows PowerShell. Before you can use WinRS or Windows PowerShell for remote management tasks, it is necessary to configure the target computer using the WinRM command. To configure the target computer, run the command *WinRM quickconfig* from a command prompt.

SARDAR AZEEM (PICT)

You can use Windows Remote Management service (WinRS) to execute command-line utilities or scripts on a remote computer. To use WinRS, open a command prompt and prefix the command that you want to run on the remote computer with the *WinRS –r:[RemoteComputerName]* command. For example, to execute the Ipconfig command on a computer named *Naomi*, issue the command:

WinRS –r:Naomi ipconfig

If the computer is on the local network, you can use its NetBIOS name. If the computer is on a remote network, you may need to specify its fully qualified domain name (FQDN). It is also possible to specify credentials to be used on the remote computer, for example, to run the command net accounts, which displays information about a computer's password policy on a computer named Naomi.7-seconds.pdxoffice using the *NaomiS* user account, issue the following command:

WinRS –r:http://Naomi.7-seconds.pdxoffice –u:NaomiS net accounts

If you do not specify a password using the –p:password option, you are prompted to enter a password after you execute the command. You can configure WinRS options through Group Policy in the *Computer Configuration\Administrative Templates\Windows Components\Windows Remote Shell* node.

PowerShell:

Windows PowerShell utilities give you the ability to remotely configure and administer a Windows 7 machine. Windows PowerShell is a command-line scripting utility that allows you to remotely execute commands on a Windows 7 machine. Windows PowerShell is a command line utility that was specifically designed for system administrators to allow for remote administration. One of the advantages of Windows PowerShell is that it introduced the concept of a cmdlet. A cmdlet is a command that is built into Windows PowerShell. There are more than 100 built-in cmdlets, and you can build your own cmdlets and allow others to use them as well.

Another advantage of Windows PowerShell is that it allows you to gain access to a file system on a computer. Windows PowerShell also allows you to access the Registry, digital certificate stores, and other data stores.

The following features are new with PowerShell in Windows 7:

- **New cmdlets** - Windows PowerShell includes over 100 new cmdlets, like Get-Hotfix, Send-MailMessage, Get-ComputerRestorePoint, New-WebServiceProxy, Debug-Process, Add-Computer, Rename-Computer, Reset-ComputerMachinePassword, and Get-Random.
- **Remote management** - You can run commands on one computer or more computers with a single command. You can establish an interactive session with a single computer, and computers can receive remote commands from multiple computers.
- **PowerShell Integrated Scripting Environment (ISE)** - Windows PowerShell ISE is a graphical user interface for Windows PowerShell with which you can run commands, and write, edit, run, test, and debug scripts in the same window. It offers eight independent execution environments and includes a inbuilt debugger, multiline editing, selective execution, syntax colors, line and column numbers, and context-sensitive Help.
- **Background jobs** - With Windows PowerShell background jobs, you can run commands asynchronously in the background and can continue to work in your session. You can run background jobs on a local or remote computer, and can store the results locally or remotely.
- **Debugger** - The Windows PowerShell debugger can help you to debug functions and scripts. You can step through code , set and remove breakpoints, check the values of variables, and display a call-stack trace.

SARDAR AZEEM(PICT)

- **Modules** - Windows PowerShell modules allow you to organize your Windows PowerShell scripts and functions into independent, self-contained units. You can package your cmdlets, scripts, functions, and other files into modules that can be distributed to other users. Modules are easier to install and use as compared to Windows PowerShell snap-ins. Modules can include any type of file, like audio files, images, Help files, and icons. Modules run in a separate session so as to avoid name conflicts.
- **Transactions** - Windows PowerShell now also supports transactions, through which you can manage a set of commands as a logical unit. A transaction can be committed, or it can be completely undone to undo the changes and the affected data is not changed by the transaction.
- **Events** Windows PowerShell includes a new event infrastructure with which you can create events, subscribe to system and application events, then you can listen, forward, and act on the events synchronously and asynchronously.
- **The Advanced functions** - Advanced functions are similar to cmdlets, but they are written in the Windows PowerShell scripting language instead of in C#.
- **Script internationalization** - Scripts and functions display messages and Help text to users in various languages.

Using Windows PowerShell:

1. Click *Start*, then click *All Programs* and then click *Accessories*.
2. Click *Windows PowerShell* and then again click *Windows PowerShell* to access it. Windows PowerShell Integrated Scripting Environment (ISE) is a new host application which allows you to run commands and write, test, and debug scripts in a friendly, syntax-colored. It can be accesses by clicking *Windows PowerShell ISE*.
3. When the Windows PowerShell utility starts, type *Help* and press *Enter*. This will show you the Windows PowerShell syntax and some of the commands included with Windows PowerShell. You can type *Help ** at the Windows command prompt. This will show you all of the cmdlet commands that you can use.

Following are few Windows PowerShellcmdlets:

Clear-History	Deletes entries from the command history
Invoke-command	Runs commands on local or remote computers
Start-job	Starts a Windows PowerShell background job
Stop-job	Stops a Windows PowerShell background job
Remove-job	Deletes a Windows PowerShell background job
Import-Module	Adds modules to the current session
Receive-job	Gets the results of a Windows PowerShell background job
Format-table	Shows the results in a table format
Out-file	Sends the job results to a file
Get-Date	Gets the date and time
Set-Date	Sets the system time and date on a computer
Get-event	Gets an event in the event queue
New-event	Creates a new event
Trace-command	Configures and starts a trace of a command on a machine.

PAZEEMPICT

Installing Windows Server 2008

How to Install Windows Server 2008 Step by Step

Installing Windows Server 2008 is pretty straightforward and is very much like installing Windows Vista, but I thought I'd list the necessary steps here for additional information. For those of you who have never installed Vista before, the entire installation process is different than it used to be in previous Microsoft operating systems, and notably much easier to perform.

To use Windows Server 2008 you need to meet the following hardware requirements:

Component	**Requirement**
Processor	• Minimum: 1GHz (x86 processor) or 1.4GHz (x64 processor) • Recommended: 2GHz or faster Note: An Intel Itanium 2 processor is required for Windows Server 2008 for Itanium-based Systems
Memory	• Minimum: 512MB RAM • Recommended: 2GB RAM or greater • Maximum (32-bit systems): 4GB (Standard) or 64GB (Enterprise and Datacenter) • Maximum (64-bit systems): 32GB (Standard) or 2TB (Enterprise, Datacenter and Itanium-based Systems)
Available Disk Space	• Minimum: 10GB • Recommended: 40GB or greater Note: Computers with more than 16GB of RAM will require more disk space for paging, hibernation, and dump files
Drive	DVD-ROM drive
Display and Peripherals	• Super VGA (800 x 600) or higher-resolution monitor • Keyboard • Microsoft Mouse or compatible pointing device

Upgrade notes:

I will not discuss the upgrade process in this article, but for your general knowledge, the upgrade paths available for Windows Server 2008 shown in the table below:

If you are currently running:	**You can upgrade to:**
Windows Server 2003 Standard Edition (R2, Service Pack 1 or Service Pack 2)	Full Installation of Windows Server 2008 Standard Edition
	Full Installation of Windows Server 2008 Enterprise Edition
Windows Server 2003 Enterprise Edition (R2, Service Pack 1 or Service Pack 2)	Full Installation of Windows Server 2008 Enterprise Edition
Windows Server 2003 Datacenter Edition (R2, Service Pack 1 or Service Pack 2)	Full Installation of Windows Server 2008 Datacenter Edition

Follow this procedure to install Windows Server 2008:

1. Insert the appropriate **Windows Server 2008 installation media** into your DVD drive. If you don't have an installation DVD for Windows Server 2008, you can download one for free from Microsoft's Windows 2008 Server Trial website.

2. **Reboot** the computer.

SARDARAZEEM(PICT)

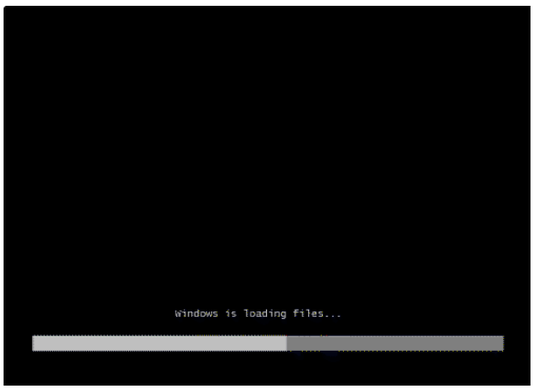

3. When prompted for an **installation language** and other regional options, make your selection and press **Next**.

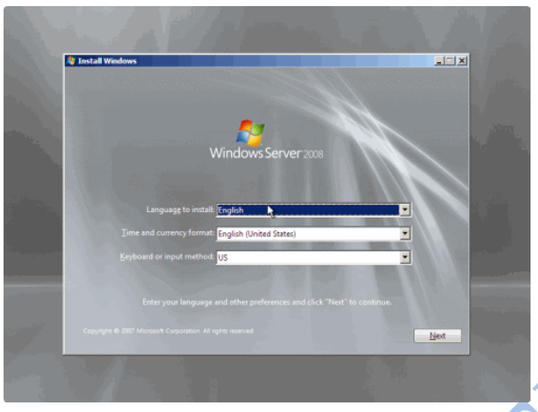

4. Next, press **Install Now** to begin the installation process.

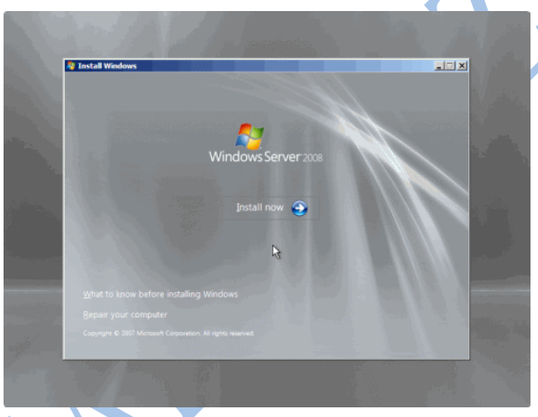

5. Product activation is now also identical with that found in Windows Vista. Enter your **Product ID** in the next window, and if you want to automatically activate Windows the moment the installation finishes, click **Next**.

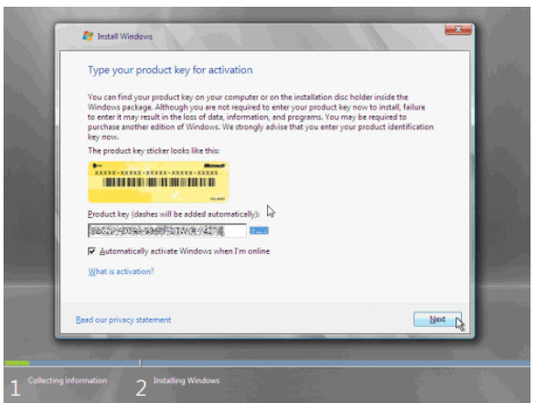

SARDAR AZEEM(PICT)

If you do not have the Product ID available right now, you can leave the box empty, and click Next. You will need to provide the Product ID later, after the server installation is over. Press No.

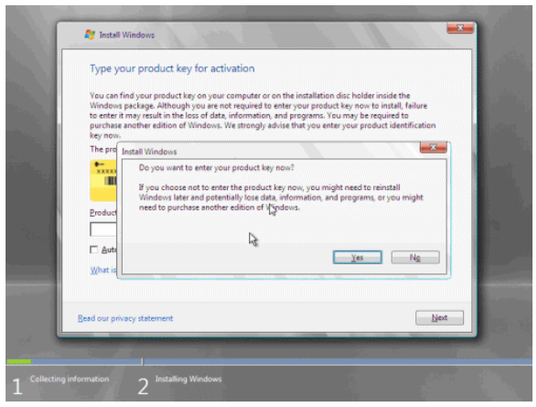

6. Because you did not provide the correct ID, the installation process cannot determine what kind of Windows Server 2008 license you own, and therefore you will be prompted to **select your correct version** in the next screen, assuming you are telling the truth and will provide the correct ID to prove your selection later on.

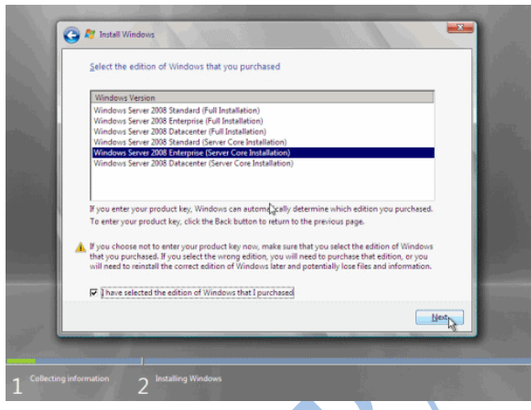

7. If you did provide the right Product ID, select the **Full version** of the right Windows version you're prompted, and click **Next**.

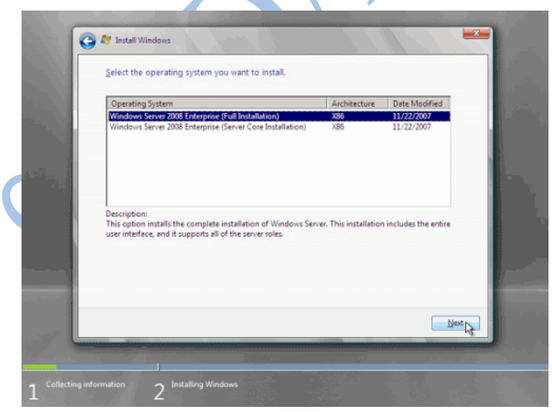

8. Read and accept the license terms by clicking to select the **checkbox** and pressing **Next**.

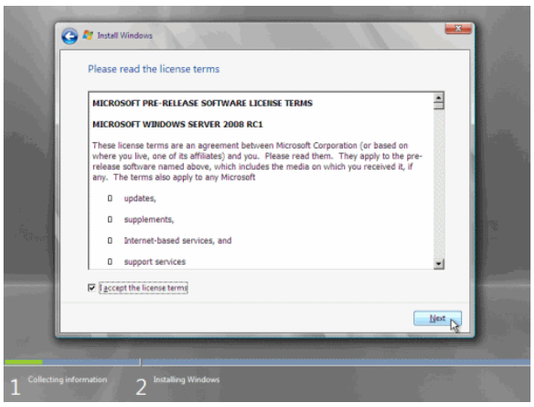

9. In the "**Which type of installation do you want?**" window, click the only available option – **Custom (Advanced).**

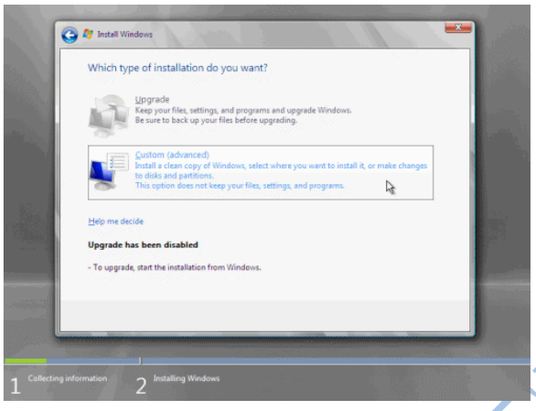

10. In the "**Where do you want to install Windows?"**, if you're installing the server on a regular IDE hard disk, click to select the **first disk**, usually **Disk 0,** and click **Next**.

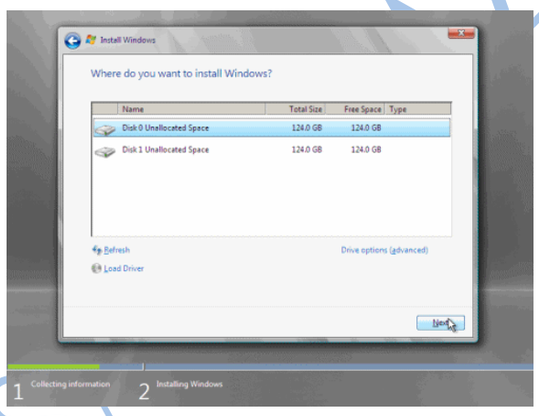

If you're installing on a hard disk that's connected to a SCSI controller, click Load Driver and insert the media provided by the controller's manufacturer.

If you're installing in a Virtual Machine environment, make sure you read the "Installing the Virtual SCSI Controller Driver for Virtual Server 2005 on Windows Server 2008"

If you must, you can also click Drive Options and manually create a partition on the destination hard disk.

11. The installation now begins, and you can go and have lunch. Copying the setup files from the DVD to the hard drive only takes about one minute. However, extracting and uncompressing the files takes a good deal longer. After 20 minutes, the operating system is installed. The exact time it takes to install server core depends upon your hardware specifications. Faster disks will perform much faster installs… Windows Server 2008 takes up approximately 10 GB of hard drive space.

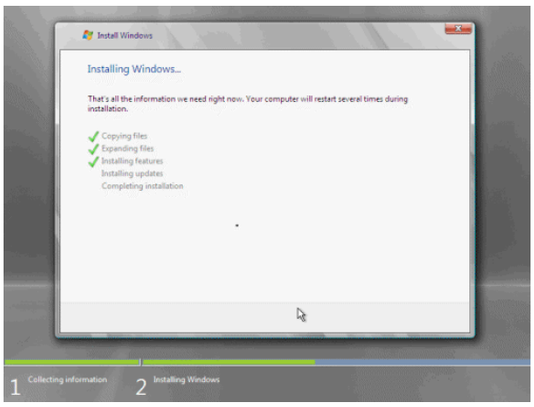

The installation process will reboot your computer, so, if in step #10 you inserted a floppy disk (either real or virtual), make sure you remove it before going to lunch, as you'll find the server hanged without the ability to boot (you can bypass this by configuring the server to boot from a CD/DVD and then from the hard disk in the booting order on the server's BIOS)

12. Then the server reboots you'll be prompted with the new Windows Server 2008 type of login screen. Press **CTRL+ALT+DEL** to log in.

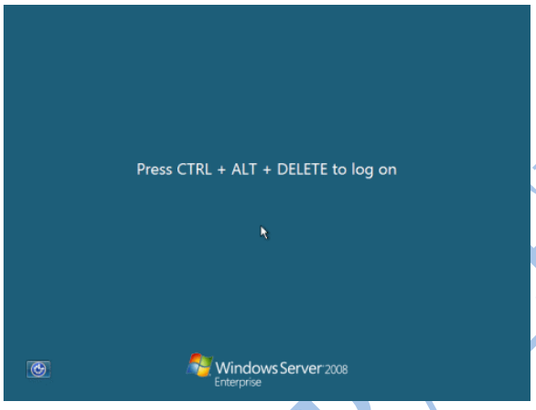

13. Click on **Other User.**

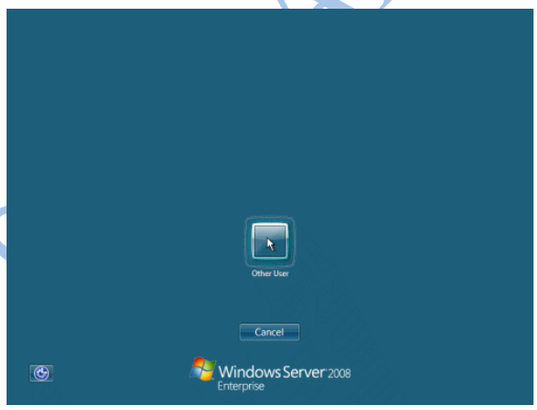

14. The default **Administrator** is **blank**, so just type **Administrator** and press **Enter**.

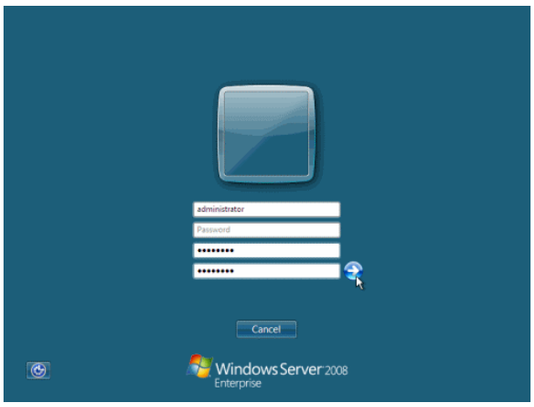

15. You will be prompted to change the user's password. You have no choice but to press **Ok**.

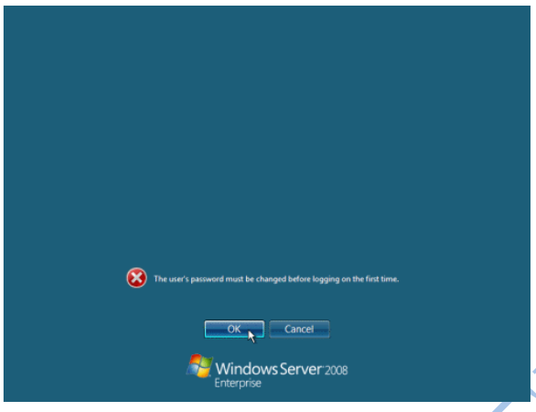

16. In the password changing dialog box, leave the d**efault password blank** (duh, read step #15…), and enter a new, complex, at-least-7-characters-long new password twice. A password like "topsecret" is not valid (it's not complex), but one like "T0pSecreT!" sure is. Make sure you remember it.

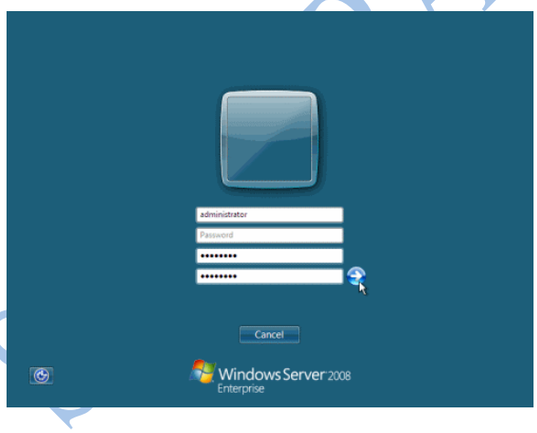

17. Someone thought it would be cool to nag you once more, so now you'll be prompted to accept the fact that the password had been changed. Press **Ok**.

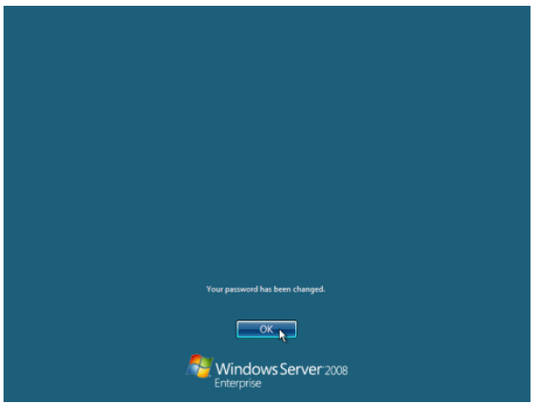

18. Finally, the desktop appears and that's it, you're logged on and can begin working. You will be greeted by an assistant for the **initial server configuration,** and after performing some initial configuration tasks, you will be able to start working.

Introduction to Administrative Tasks in Windows Server 2008 Environment

(1) Prepare your Domain for the Windows Server 2008 R2 Domain Controller

Before installing the first Windows Server 2008 R2 domain controller (DC) into an existing Windows 2000, Windows Server 2003 or Windows Server 2008 domain, you must prepare the AD forest and domain. You do so by running a tool called **ADPREP**.

ADPREP extends the Active Directory schema and updates permissions as necessary to prepare a forest and domain for a domain controller that runs the Windows Server 2008 R2 operating system.

Note: You may remember that ADPREP was used on previous operating systems such as Windows Server 2003, Windows Server 2003 R2 and Windows Server 2008. This article focuses on Windows Server 2008 R2.

What does ADPREP do? ADPREP has parameters that perform a variety of operations that help prepare an existing Active Directory environment for a domain controller that runs Windows Server 2008 R2. Not all versions of ADPREP perform the same operations, but generally the different types of operations that ADPREP can perform include the following:

- Updating the Active Directory schema
- Updating security descriptors
- Modifying access control lists (ACLs) on Active Directory objects and on files in the SYSVOL shared folder
- Creating new objects, as needed
- Creating new containers, as needed

To prepare the forest and domain for the installation of the first Windows Server 2008 R2 domain controller please perform these tasks:

Lamer note: The following tasks are required ONLY before adding the first Windows Server 2008 R2 **domain controller**. If you plan on simply joining a Windows Server 2008 R2 Server to the domain and configuring as a regular member server, none of the following tasks are required.

Another lamer note: Please make sure you read the system requirements for Windows Server 2008 R2. For example, you cannot join a Windows Server 2008 R2 server to a

SARDARAZEEM(PICT)

Windows NT 4.0 domain, not can it participate as a domain controller in a mixed domain. If any domain controllers in the forest are running Windows 2000 Server, they must be running Service Pack 4 (SP4).

First, you should review and understand the schema updates and other changes that ADPREP makes as part of the schema management process in Active Directory Domain Services (AD DS). You should test the ADPREP schema updates in a lab environment to ensure that they will not conflict with any applications that run in your environment.

You must make a system state backup for your domain controllers, including the schema master and at least one other domain controller from each domain in the forest (you do have backups, don't you?).

Also, make sure that you can log on to the schema master with an account that has sufficient credentials to run **adprep /forestprep**. You must be a member of the Schema Admins group, the Enterprise Admins group, and the Domain Admins group of the domain that hosts the schema master, which is, by default, the forest root domain.

Next, insert the Windows Server 2008 R2 DVD media into your DVD drive. Note that if you do not have the media handy, you may use the evaluation version that is available to download from Microsoft's website. You can also use an MSDN or Technet ISO image, if you have a subscription to one of them.

Browse to the ***X:\support\adprep*** folder, where *X:* is the drive letter of your DVD drive. Find a file called ***adprep.exe*** or ***adprep32.exe***.

Note: Unlike in Windows Server 2008 where you had to use either the 32-bit or 64-bit installation media to get the right version of ADPREP, Windows Server 2008 R2 ADPREP is available in a 32-bit version **and** a 64-bit version. The 64-bit version runs by default. If you need to run ADPREP on a 32-bit computer, run the 32-bit version (*adprep32.exe*).

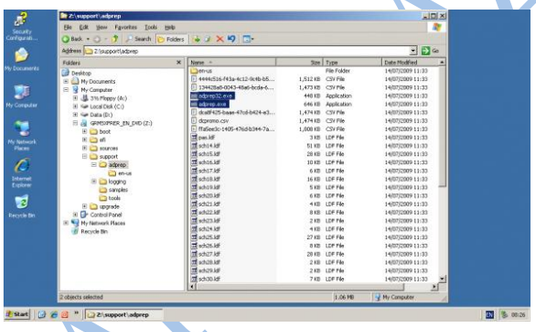

To perform this procedure, you must use an account that has membership in all of the following groups:

- Enterprise Admins
- Schema Admins
- Domain Admins for the domain that contains the schema master

Open a Command Prompt window by typing CMD and pressing ENTER in the Run menu.

Drag the ***adprep.exe*** file from the Windows Explorer window to the Command Prompt window. Naturally, if you want, you can always manually type the path of the file in the Command Prompt window if that makes you feel better...

Note: You must run *adprep.exe* from an elevated command prompt. To open an elevated command prompt, click Start, right-click Command Prompt, and then click Run as administrator.

Note: If your existing DCs are Windows Server 2008, dragging and dropping into a Command Prompt window will not work, as that feature was intentionally disabled in windows Server 2008 and Windows Vista.

In the Command Prompt window, type the following command:

adprep /forestprep

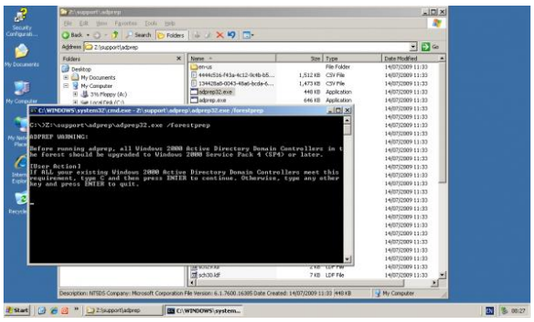

You will be prompted to type the letter "c" and then press ENTER. After doing so, process will begin.

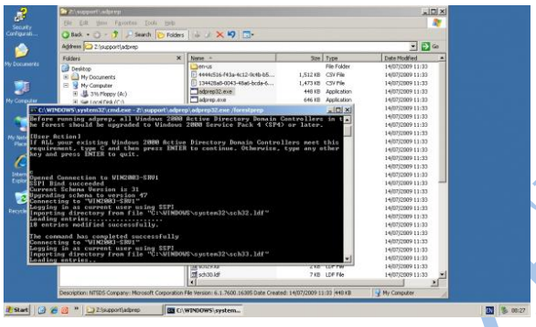

ADPREP will take several minutes to complete. During that time, several LDF files will be imported into the AD Schema, and messages will be displayed in the Command Prompt window. File sch47.ldf seems to be the largest one.

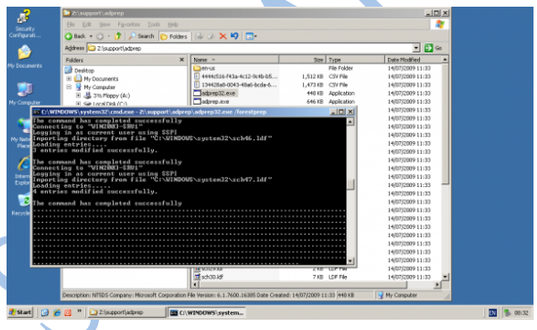

When completed, you will receive a success message.

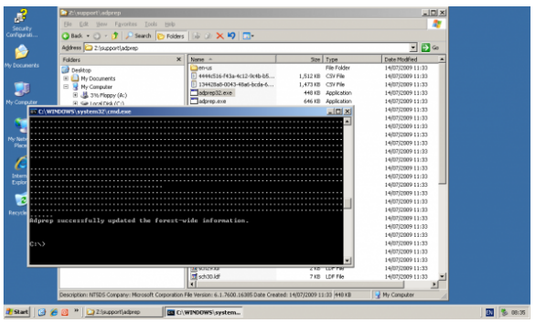

Note: As mentioned above, ADPREP should only be run on an existing DC. When trying to run it from a non-DC, you will get this error:

Adprep cannot run on this platform because it is not an Active Directory Domain Controller.
[Status/Consequence]
Adprep stopped without making any changes.
[User Action]
Run Adprep on a Active Directory Domain Controller.

Allow the operation to complete, and then allow the changes to replicate throughout the forest before you prepare any domains for a domain controller that runs Windows Server 2008 R2.

In the Command Prompt window, type the following command:

adprep /domainprep

Process will take less than a second.

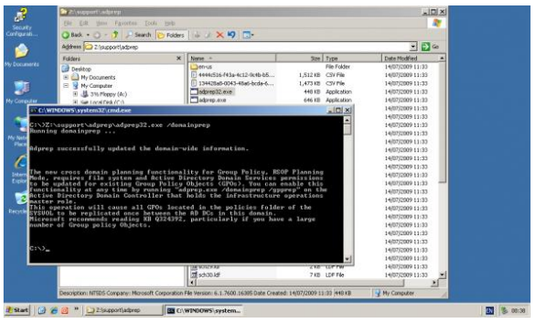

ADPREP must only be run in a Windows 2000 Native Mode or higher. If you attempt to run in Mixed Mode you will get this error:

Adprep detected that the domain is not in native mode
[Status/Consequence]
Adprep has stopped without making changes.
[User Action]
Configure the domain to run in native mode and re-run domainprep

Allow the operation to complete, and then allow the changes to replicate throughout the forest before you prepare any domains for a domain controller that runs Windows Server 2008 R2.

If you're running a Windows 2008 Active Directory domain, that's it, no additional tasks are needed.

If you're running a Windows 2000 Active Directory domain, you must also the following command:

adprep /domainprep /gpprep

Allow the operation to complete, and then allow the changes to replicate throughout the forest before you prepare any domains for a domain controller that runs Windows Server 2008 R2.

If you're running a Windows 2003 Active Directory domain, that's it, no additional tasks are needed. However, if you're planing to run Read Only Domain controllers (RODCs), you must also

type the following command:

adprep /rodcprep

SARDARAZEEM(PICT)

If you already ran this command for Windows Server 2008, you do not need to run it again for Windows Server 2008 R2.

Process will complete in less than a second.

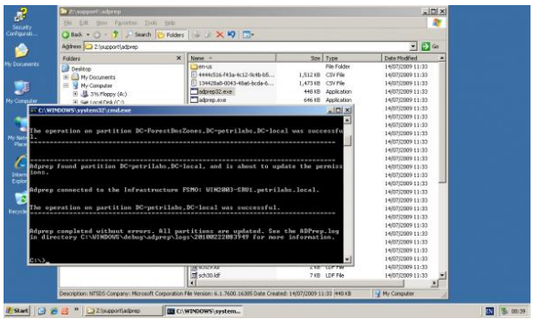

Allow the operation to complete, and then allow the changes to replicate throughout the forest before you prepare any domains for a domain controller that runs Windows Server 2008 R2.

To verify that *adprep /forestprep* completed successfully please perform these steps:

1. Log on to an administrative workstation that has ADSIEdit installed. ADSIEdit is installed by default on domain controllers that run Windows Server 2008 or Windows Server 2008 R2. On Windows Server 2003 you must install the Resource Kit Tools.

2. Click Start, click Run, type ADSIEdit.msc, and then click OK.

3. Click Action, and then click Connect to.

4. Click Select a well known Naming Context, select Configuration in the list of available naming contexts, and then click OK.

5. Double-click Configuration, and then double-click *CN=Configuration,DC=forest_root_domain* where *forest_root_domain* is the distinguished name of your forest root domain.

6. Double-click *CN=ForestUpdates*.

7. Right-click *CN=ActiveDirectoryUpdate*, and then click Properties.

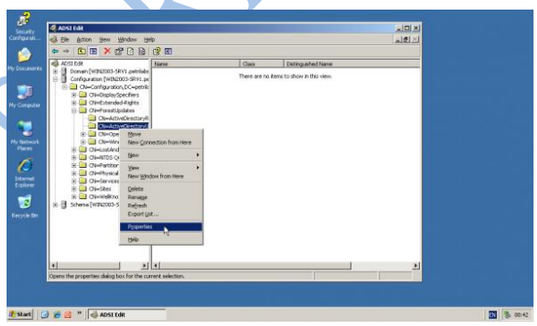

8. If you ran *adprep /forestprep* for Windows Server 2008 R2, confirm that the *Revision* attribute value is **5**, and then click OK.

SARDAR AZEEM(PICT)

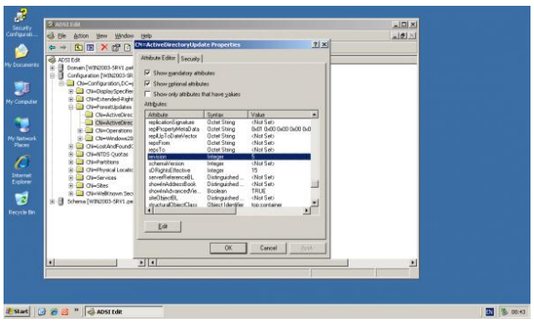

9. Click ADSI Edit, click Action, and then click Connect to.

10. Click Select a Well known naming context, select Schema in the list of available naming contexts, and then click OK.

11. Double-click Schema.

12. Right-click *CN=Schema,CN=Configuration,DC=forest_root_domain*, and then click Properties.

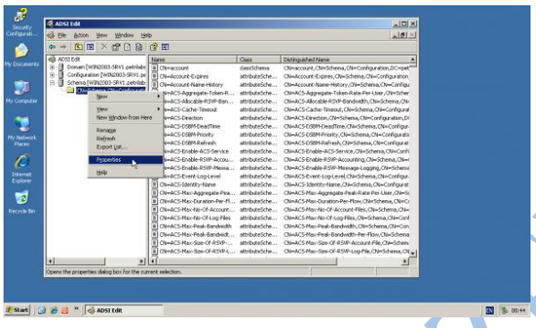

13. If you ran *adprep /forestprep* for Windows Server 2008 R2, confirm that the *objectVersion* attribute value is set to **47**, and then click OK.

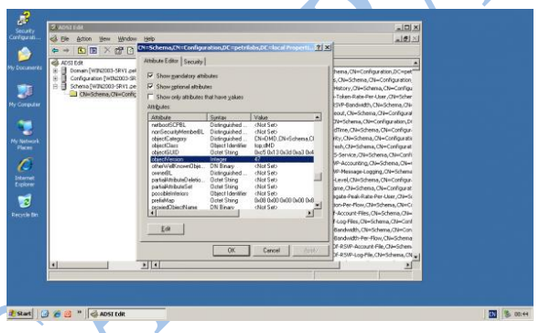

SARDAR AZEEM(PICT)

Install Windows Server 2008 Server Roles with Server Manager

The following server roles are available in Windows Server 2008.

Active Directory Certificate Services. Active Directory® Certificate Services (AD CS) provides customizable services for creating and managing public key certificates used in software security systems employing public key technologies.	File Services. File Services provides technologies for storage management, file replication, distributed namespace management, fast file searching, and streamlined client access to files.
Active Directory Domain Services. Active Directory Domain Services (AD DS) stores information about users, computers, and other devices on the network. AD DS helps administrators securely manage this information and facilitates resource sharing and collaboration between users.	Hyper-V. Hyper-V provides the services that you can use to create and manage virtual machines and their resources. Each virtual machine is a virtualized computer system that operates in an isolated execution environment. This allows you to run multiple operating systems simultaneously.
Active Directory Federation Services. Active Directory Federation Services (AD FS) provides Web single-sign-on (SSO) technologies to authenticate a user to multiple Web applications by using a single user account. AD FS accomplishes this by securely federating, or sharing, user identities and access rights, in the form of digital claims, between partner organizations.	Network Policy and Access Services. Network Policy and Access Services delivers a variety of methods to provide users with local and remote network connectivity, to connect network segments, and to allow network administrators to centrally manage network access and client health policies.
Active Directory Lightweight Directory Services. Organizations that have applications that require a directory for storing application data can use Active Directory Lightweight Directory Services (AD LDS) as the data store.	Print Services. Print Services enables the management of print servers and printers. A print server reduces administrative and management workload by centralizing printer management tasks.
Active Directory Rights Management Services. Active Directory Rights Management Services (AD RMS) is information protection technology that works with AD RMS-enabled applications to help safeguard digital information from unauthorized use.	Terminal Services. Terminal Services provides technologies that enable users to access Windows-based programs that are installed on a terminal server, or to access the Windows desktop itself from almost any computing device. Users can connect to a terminal server to run programs and to use network resources on that server.
Application Server. Application Server provides a complete solution for hosting and managing high-performance distributed business applications. Integrated services, such as the .NET Framework, Web Server Support, Message Queuing, COM+, Windows Communication Foundation, and failover Ccusters boost productivity throughout the application life cycle.	Universal Description, Discovery, and Integration Services. Universal Description, Discovery, and Integration (UDDI) Services provides UDDI capabilities for sharing information about Web services within an organization's intranet, between business partners on an extranet, or on the Internet. UDDI Services can help improve the productivity of developers and IT professionals with more reliable and manageable applications.
DHCP Server. Dynamic Host Configuration Protocol (DHCP) allows servers to assign, or lease, IP addresses to computers and other devices that are enabled as DHCP clients. Deploying DHCP servers on the network automatically provides computers and other TCP/IP-based network devices with valid IP addresses and the additional configuration parameters these devices need.	Web Server. Web Server, or Internet Information Services (IIS), enables sharing of information on the Internet, an intranet, or an extranet. It is a unified Web platform that integrates IIS 7.0, ASP.NET, and Windows Communication Foundation. IIS 7.0 also features enhanced security, simplified diagnostics, and delegated administration.
DNS Server. Domain Name System (DNS) provides a standard method for associating names with numeric Internet addresses. This makes it possible for users to refer to network computers by using easy-to-remember names instead of a long series of numbers.	Windows Deployment Services. You can use Windows Deployment Services to install and configure Windows operating systems remotely on computers by using Pre-Boot Execution Environment (PXE) boot ROMs. Administration overhead is decreased through the implementation of the WdsMgmt Microsoft Management Console (MMC) snap-in, which manages all aspects of Windows Deployment Services.
Fax Server. Fax Server sends and receives faxes, and allows you to manage fax resources such as jobs, settings, reports, and fax devices on this computer or on the network.	

Scheduling

In computer science, scheduling is the method by which threads, processes or data flows are given access to system resources (e.g. processor time, communications bandwidth). This is usually done to load balance a system effectively or achieve a target quality of service. The need for a scheduling algorithm arises from the requirement for most modern systems to perform multitasking (execute more than one process at a time) and multiplexing (transmit multiple flows simultaneously).
The scheduler is concerned mainly with:
Throughput - The total number of processes that complete their execution per time unit.
Latency, specifically:
Turnaround - total time between submission of a process and its completion.
Response time - amount of time it takes from when a request was submitted until the first response is produced.
Fairness / Waiting Time - Equal CPU time to each process (or more generally appropriate times according to each process' priority).
In practice, these goals often conflict (e.g. throughput versus latency), thus a scheduler will implement a suitable compromise. Preference is given to any one of the above mentioned concerns depending upon the user's needs and objectives.

Types of operating system schedulers

Operating systems may feature up to 3 distinct types of scheduler, a long-term scheduler (also known as an admission scheduler or high-level scheduler), a mid-term or medium-term scheduler and a short-term scheduler. The names suggest the relative frequency with which these functions are performed. The scheduler is an operating system module that selects the next jobs to be admitted into the system and the next process to run.

Long-term scheduling

The long-term, or admission scheduler, decides which jobs or processes are to be admitted to the ready queue (in the Main Memory); that is, when an attempt is made to execute a program, its admission to the set of currently executing processes is either authorized or delayed by the long-term scheduler. Thus, this scheduler dictates what processes are to run on a system, and the degree of concurrency to be supported at any one time - i.e.: whether a high or low amount of processes are to be executed concurrently, and how the split between IO intensive and CPU intensive processes is to be handled. In modern operating systems, this is used to make sure that real time processes get enough CPU time to finish their tasks. Without proper real time scheduling, modern GUI interfaces would seem sluggish. The long term queue exists in the Hard Disk or the "Virtual Memory".
Long-term scheduling is also important in large-scale systems such as batch processing systems, computer clusters, supercomputers and render farms. In these cases, special purpose job scheduler software is typically used to assist these functions, in addition to any underlying admission scheduling support in the operating system.

Medium-term scheduling

The medium-term scheduler temporarily removes processes from main memory and places them on secondary memory (such as a disk drive) or vice versa. This is commonly referred to as "swapping out" or "swapping in" (also incorrectly as "paging out" or "paging in"). The medium-term scheduler may decide to swap out a process which has not been active for some time, or a process which has a low priority, or a process which is page faulting frequently, or a process which is taking up a large amount of memory in order to free up main memory for other processes, swapping the process back in later when more memory is available, or when the process has been unblocked and is no longer waiting for a resource.
In many systems today (those that support mapping virtual address space to secondary storage other than the swap file), the medium-term scheduler may actually perform the role of the long-term scheduler, by treating binaries as "swapped out processes" upon their execution. In this way, when a segment of the binary is required it can be swapped in on demand, or "lazy loaded".

Short-term scheduling

The short-term scheduler (also known as the CPU scheduler) decides which of the ready, in-memory processes are to be executed (allocated a CPU) next following a clock interrupt, an IO interrupt, an operating system call or another form of signal. Thus the short-term scheduler makes scheduling decisions much more frequently than the long-term or mid-term schedulers - a scheduling decision will at a minimum have to be made after every time slice, and these are very short. This scheduler can be preemptive, implying that it is capable of forcibly removing processes from a CPU when it decides to allocate that CPU to another process, or non-preemptive (also known as "voluntary" or "co-operative"), in which case the scheduler is unable to "force" processes off the CPU.In most cases short-term scheduler is written in assembly because it is a critical part of the operating system.

Dispatcher

Another component involved in the CPU-scheduling function is the dispatcher. The dispatcher is the module that gives control of the CPU to the process selected by the short-term scheduler.

This function involves the following:

- Switching context
- Switching to user mode
- Jumping to the proper location in the user program to restart that program

The dispatcher should be as fast as possible, since it is invoked during every process switch. The time it takes for the dispatcher to stop one process and start another running is known as the dispatch latency.

Scheduling disciplines

Scheduling disciplines are algorithms used for distributing resources among parties which simultaneously and asynchronously request them. Scheduling disciplines are used in routers (to handle packet traffic) as well as in operating systems (to share CPU time among both threads and processes), disk drives (I/O scheduling), printers (print spooler), most embedded systems, etc.
The main purposes of scheduling algorithms are to minimize resource starvation and to ensure fairness amongst the parties utilizing the resources. Scheduling deals with the problem of deciding which of the outstanding requests is to be allocated resources. There are many different scheduling algorithms. In this section, we introduce several of them.
In packet-switched computer networks and other statistical multiplexing, the notion of a scheduling algorithm is used as an alternative to first-come first-served queuing of data packets.
The simplest best-effort scheduling algorithms are round-robin, fair queuing (a max-min fair scheduling algorithm), proportionally fair scheduling and maximum throughput. If differentiated or guaranteed quality of service is offered, as opposed to best-effort communication, weighted fair queuing may be utilized.
In advanced packet radio wireless networks such as HSDPA (High-Speed Downlink Packet Access) 3.5G cellular system, channel-dependent scheduling may be used to take advantage of channel state information. If the channel conditions are favorable, the throughput and system spectral efficiency may be increased. In even more advanced systems such as LTE, the scheduling is combined by channel-dependent packet-by-packet dynamic channel allocation, or by assigning OFDMA multi-carriers or other frequency-domain equalization components to the users that best can utilize them.

First in first out

Also known as First Come, First Served (FCFS), is the simplest scheduling algorithm, FIFO simply queues processes in the order that they arrive in the ready queue.

- Since context switches only occur upon process termination, and no reorganization of the process queue is required, scheduling overhead is minimal.
- Throughput can be low, since long processes can hog the CPU

Turnaround time, waiting time and response time can be high for the same reasons above
No prioritization occurs, thus this system has trouble meeting process deadlines.

SARDAR AZEEM (PICT)

The lack of prioritization means that as long as every process eventually completes, there is no starvation. In an environment where some processes might not complete, there can be starvation.
It is based on Queuing

Similar to Shortest Job First (SJF).

With this strategy the scheduler arranges processes with the least estimated processing time remaining to be next in the queue. This requires advanced knowledge or estimations about the time required for a process to complete.
If a shorter process arrives during another process' execution, the currently running process may be interrupted (known as preemption), dividing that process into two separate computing blocks. This creates excess overhead through additional context switching. The scheduler must also place each incoming process into a specific place in the queue, creating additional overhead.
This algorithm is designed for maximum throughput in most scenarios.
Waiting time and response time increase as the process' computational requirements increase. Since turnaround time is based on waiting time plus processing time, longer processes are significantly affected by this. Overall waiting time is smaller than FIFO, however since no process has to wait for the termination of the longest process.
No particular attention is given to deadlines, the programmer can only attempt to make processes with deadlines as short as possible.
Starvation is possible, especially in a busy system with many small processes being run.

Fixed priority pre-emptive scheduling

The O/S assigns a fixed priority rank to every process, and the scheduler arranges the processes in the ready queue in order of their priority. Lower priority processes get interrupted by incoming higher priority processes.
Overhead is not minimal, nor is it significant.
FPPS has no particular advantage in terms of throughput over FIFO scheduling.
Waiting time and response time depend on the priority of the process. Higher priority processes have smaller waiting and response times.
Deadlines can be met by giving processes with deadlines a higher priority.
Starvation of lower priority processes is possible with large amounts of high priority processes queuing for CPU time.

Round-robin scheduling

The scheduler assigns a fixed time unit per process, and cycles through them.
RR scheduling involves extensive overhead, especially with a small time unit.
Balanced throughput between FCFS and SJF, shorter jobs are completed faster than in FCFS and longer processes are completed faster than in SJF.
Poor average response time, waiting time is dependent on number of processes, and not average process length.
Because of high waiting times, deadlines are rarely met in a pure RR system.
Starvation can never occur, since no priority is given. Order of time unit allocation is based upon process arrival time, similar to FCFS.

Multilevel queue scheduling

This is used for situations in which processes are easily divided into different groups. For example, a common division is made between foreground (interactive) processes and background (batch) processes. These two types of processes have different response-time requirements and so may have different scheduling needs. It is very useful for shared memory problem

SARDAR AZEEM (PICT)

Overview

Scheduling algorithm	CPU Overhead	Throughput	Turnaround time	Response time
First In First Out	Low	Low	High	Low
Shortest Job First	Medium	High	Medium	Medium
Priority based scheduling	Medium	Low	High	High
Round-robin scheduling	High	Medium	Medium	High
Multilevel Queue scheduling	High	High	Medium	Medium

Deadlock

Both processes need both resources. P1 requires additional resource R1, P2 requires additional resource R2; neither process can continue.
This article is about the computer science concept. For other uses, see Deadlock (disambiguation).
A deadlock is a situation in which two or more competing actions are each waiting for the other to finish, and thus neither ever does.

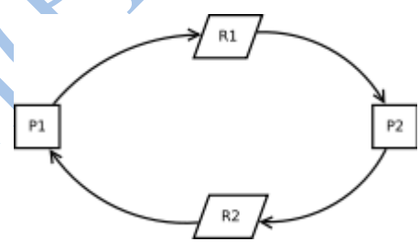

In an operating system, a deadlock is a situation which occurs when a process enters a waiting state because a resource requested by it is being held by another waiting process, which in turn is waiting for another resource. If a process is unable to change its state indefinitely because the resources requested by it are being used by other waiting process, then the system is said to be in a deadlock.
Deadlock is a common problem in multiprocessing systems, parallel computing and distributed systems, where software and hardware locks are used to handle shared resources and implement process synchronization.
In telecommunication systems, deadlocks occur mainly due to lost or corrupt signals instead of resource contention.

Examples

Deadlock situation can be compared to the classic "chicken or egg" problem. It can be also considered a paradoxical "Catch-22" situation. A real world analogical example would be an illogical statute passed by the Kansas legislature in the early 20th century, which stated:

> “ When two trains approach each other at a crossing, both shall come to a full stop and neither shall start up again until the other has gone. ”

A simple computer-based example is as follows. Suppose a computer has three CD drives and three processes. Each of the three processes holds one of the drives. If each process now requests another drive, the three processes will be in a deadlock. Each process will be waiting for the "CD drive released" event, which can be only caused by one of the other waiting processes. Thus, it results in a circular chain.

Necessary conditions

A deadlock situation can arise if and only if all of the following conditions hold simultaneously in a system:
Mutual Exclusion: At least one resource must be non-shareable. Only one process can use the resource at any given instant of time.

SARDAR AZEEM (PICT)

Hold and Wait or Resource Holding: A process is currently holding at least one resource and requesting additional resources which are being held by other processes.
No Preemption: The operating system must not de-allocate resources once they have been allocated; they must be released by the holding process voluntarily.
Circular Wait: A process must be waiting for a resource which is being held by another process, which in turn is waiting for the first process to release the resource. In general, there is a set of waiting processes, P = {P1, P2, ..., PN}, such that P1 is waiting for a resource held by P2, P2 is waiting for a resource held by P3 and so on till PN is waiting for a resource held by P1.

Deadlock handling

Most current operating systems cannot prevent a deadlock from occurring. When a deadlock occurs, different operating systems respond to them in different non-standard manners. Most approaches work by preventing one of the four Coffman conditions from occurring, especially the fourth one. Major approaches are as follows.

Ignoring deadlock

In this approach, it is assumed that a deadlock will never occur. This is also an application of the Ostrich algorithm. This approach was initially used by MINIX and UNIX. This is used when the time intervals between occurrences of deadlocks are large and the data loss incurred each time is tolerable. It is avoided in very critical systems.

Detection

Under deadlock detection, deadlocks are allowed to occur. Then the state of the system is examined to detect that a deadlock has occurred and subsequently it is corrected. An algorithm is employed that tracks resource allocation and process states, it rolls back and restarts one or more of the processes in order to remove the detected deadlock. Detecting a deadlock that has already occurred is easily possible since the resources that each process has locked and/or currently requested are known to the resource scheduler of the operating system.
Deadlock detection techniques include, but are not limited to model checking. This approach constructs a finite state-model on which it performs a progress analysis and finds all possible terminal sets in the model. These then each represent a deadlock.
After a deadlock is determined, it can be corrected by using one of the following methods:
Process Termination: One or more process involved in the deadlock may be aborted. We can choose to abort all processes involved in the deadlock. This ensures that deadlock is resolved with certainty and speed. But the expense is high as partial computations will be lost. Or, we can choose to abort one process at a time until the deadlock is resolved. This approach has high overheads because after each abortion an algorithm must detect if the system is still in deadlock. Several factors must be considered while choosing a candidate for termination, such as priority and age of the process.
Resource Preemption: Resource allocated to various processes may be successively preempted and allocated to other processes until deadlock is broken.

Prevention

Deadlock prevention works by preventing one of the four Coffman conditions from occurring.
Removing the mutual exclusion condition means that no process will have exclusive access to a resource. This proves impossible for resources that cannot be spooled. But even with spooled resources, deadlock could still occur. Algorithms that avoid mutual exclusion are called non-blocking synchronization algorithms.
The hold and wait or resource holding conditions may be removed by requiring processes to request all the resources they will need before starting up (or before embarking upon a

SARDAR AZEEM (PICT)

particular set of operations). This advance knowledge is frequently difficult to satisfy and, in any case, is an inefficient use of resources. Another way is to require processes to request resources only when it has none. Thus, first they must release all their currently held resources before requesting all the resources they will need from scratch. This too is often impractical. It is so because resource may be allocated and remain unused for long periods. Also, a process requiring a popular resource may have to wait indefinitely as such a resource may always be allocated to some process, resulting in resource starvation. (These algorithms, such as serializing tokens, are known as the all-or-none algorithms.)

The no preemption condition may also be difficult or impossible to avoid as a process has to be able to have a resource for a certain amount of time, or the processing outcome may be inconsistent or thrashing may occur. However, inability to enforce preemption may interfere with a priority algorithm. Preemption of a "locked out" resource generally implies a rollback, and is to be avoided, since it is very costly in overhead. Algorithms that allow preemption include lock-free and wait-free algorithms and optimistic concurrency control.

The final condition is the circular wait condition. Approaches that avoid circular waits include disabling interrupts during critical sections and using a hierarchy to determine a partial ordering of resources. If no obvious hierarchy exists, even the memory address of resources has been used to determine ordering and resources are requested in the increasing order of the enumeration. The Dijkstra's solution can also be used.

Avoidance

Deadlock can be avoided if certain information about processes are available to the operating system before allocation of resources, such as which resources a process will consume in its lifetime. For every resource request, the system sees if granting the request will mean that the system will enter an unsafe state, meaning a state that could result in deadlock. The system then only grants requests that will lead to safe states. In order for the system to be able to determine whether the next state will be safe or unsafe, it must know in advance at any time:

resources currently available

resources currently allocated to each process

resources that will be required and released

SARDAR AZEEM (PICT)

A memory management unit (MMU)

Sometimes called paged memory management unit (PMMU), is a computer hardware component responsible for handling accesses to memory requested by the CPU. Its functions include translation of virtual addresses to physical addresses (i.e., virtual memory management), memory protection, cache control, bus arbitration and in simpler computer architectures (especially 8-bit systems) bank switching.

How it works

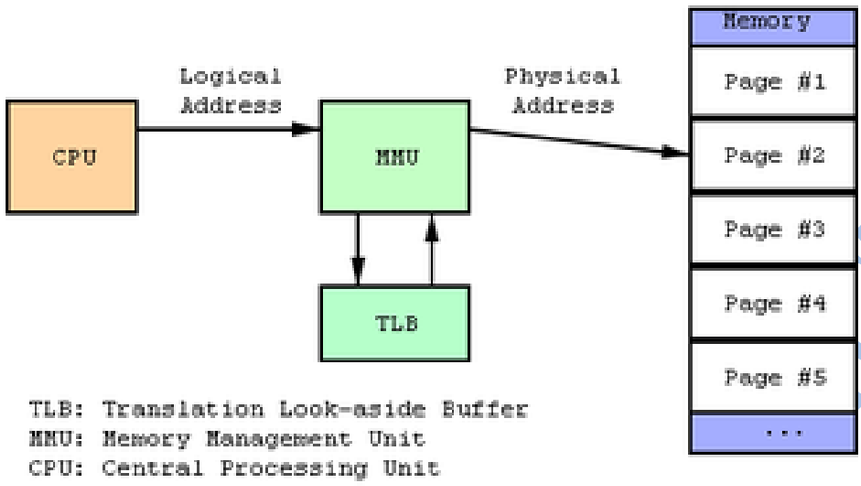

Schematic of the operation of an MMU

Modern MMUs typically divide the virtual address space (the range of addresses used by the processor) into pages, each having a size which is a power of 2, usually a few kilobytes, but they may be much larger. The bottom n bits of the address (the offset within a page) are left unchanged. The upper address bits are the (virtual) page number. The MMU normally translates virtual page numbers to physical page numbers via an associative cache called a translation lookaside buffer (TLB). When the TLB lacks a translation, a slower mechanism involving hardware-specific data structures or software assistance is used. The data found in such data structures are typically called page table entries (PTEs), and the data structure itself is typically called a page table. The physical page number is combined with the page offset to give the complete physical address.

A PTE or TLB entry may also include information about whether the page has been written to (the dirty bit), when it was last used (the accessed bit, for a least recently used page replacement algorithm), what kind of processes (user mode, supervisor mode) may read and write it, and whether it should be cached.

Sometimes, a TLB entry or PTE prohibits access to a virtual page, perhaps because no physical random access memory has been allocated to that virtual page. In this case the MMU signals a page fault to the CPU. The operating system (OS) then handles the situation, perhaps by trying to find a spare frame of RAM and set up a new PTE to map it to the requested virtual address. If no RAM is free, it may be necessary to choose an existing page (known as a victim), using some replacement algorithm, and save it to disk (this is called "paging"). With some MMUs, there can also be a shortage of PTEs or TLB entries, in which case the OS will have to free one for the new mapping.

In some cases a "page fault" may indicate a software bug. A key benefit of an MMU is memory protection: an OS can use it to protect against errant programs, by disallowing access to memory that a particular program should not have access to. Typically, an OS assigns each program its own virtual address space.

An MMU also reduces the problem of fragmentation of memory. After blocks of memory have been allocated and freed, the free memory may become fragmented (discontinuous) so that the largest contiguous block of free memory may be much smaller than the total

amount. With virtual memory, a contiguous range of virtual addresses can be mapped to several non-contiguous blocks of physical memory.

VLSI VI475 MMU "Apple HMMU"
from the Macintosh II
used with the Motorola 68020
In some early microprocessor designs, memory management was performed by a separate integrated circuit such as the VLSI VI475 or the Motorola 68851 used with the Motorola 68020 CPU in the Macintosh II or the Z8015 used with the Zilog Z80 family of processors. Later microprocessors such as the Motorola 68030 and the ZILOG Z280 placed the MMU together with the CPU on the same integrated circuit, as did the Intel 80286 and later x86 microprocessors.
While this article concentrates on modern MMUs, commonly based on pages, early systems used a similar concept for base-limit addressing, that further developed into segmentation. Those are occasionally also present on modern architectures. The x86 architecture provided segmentation rather than paging in the 80286, and provides both paging and segmentation in the 80386 and later processors (although the use of segmentation is not available in 64-bit operation).

Examples

Most modern systems divide memory into pages that are 4 KB to 64 KB in size, often with the possibility to use huge pages from 2 MB to 512 MB in size. Page translations are cached in a TLB. Some systems, mainly older RISC designs, trap into the OS when a page translation is not found in the TLB. Most systems use a hardware-based tree walker. Most systems allow the MMU to be disabled; some disable the MMU when trapping into OS code.

ACTIVE DIRECTORY (AD)

Is a directory service created by Microsoft for Windows domain networks? It is included in most Windows Server operating systems. Server computers that run Active Directory are called domain controllers.
Active Directory provides a central location for network administration and security. It authenticates and authorizes all users and computers in a Windows domain type network—assigning and enforcing security policies for all computers and installing or updating software. For example, when a user logs into a computer that is part of a Windows domain, Active Directory verifies the password and specifies whether the user is a system administrator or normal user.
Active Directory uses **Lightweight Directory Access Protocol (LDAP)** versions 2 and 3, Kerberos and DNS.

History

Active Directory was previewed in 1999, released first with Windows 2000 Server edition, and revised to extend functionality and improve administration in Windows Server 2003. Additional improvements were made in Windows Server 2003 R2, Windows Server 2008 and Windows Server 2008 R2, and with the release of the latter the domain controller role was renamed Active Directory Domain Services.

Structure of Active Directory

Objects

SARDAR AZEEM (PICT)

An Active Directory structure is a hierarchical arrangement of information about objects. The objects fall into two broad categories: resources (e.g., printers) and security principals (user or computer accounts and groups). Security principals are assigned unique **security identifiers (SIDs).**
Each object represents a single entity—whether a user, a computer, a printer, or a group—and its attributes. Certain objects can contain other objects. An object is uniquely identified by its name and has a set of attributes—the characteristics and information that the object represents— defined by a schema, which also determines the kinds of objects that can be stored in Active Directory.
The **schema object** lets administrators extend or modify the schema when necessary. However, because each schema object is integral to the definition of Active Directory objects, deactivating or changing these objects can fundamentally change or disrupt a deployment. Schema changes automatically propagate throughout the system. Once created, an object can only be deactivated—not deleted. Changing the schema usually requires planning.

Sites

A **Site** object in Active Directory represents a geographic location that hosts networks.

Forests, trees, and domains

The Active Directory framework that holds the objects can be viewed at a number of levels. The forest, tree, and domain are the logical divisions in an Active Directory network. Within a deployment, objects are grouped into domains. The objects for a single domain are stored in a single database (which can be replicated). Domains are identified by their DNS name structure, the namespace.
A tree is a collection of one or more domains and domain trees in a contiguous namespace, linked in a transitive trust hierarchy.
At the top of the structure is the forest. A forest is a collection of trees that share a common global catalog, directory schema, logical structure, and directory configuration. The forest represents the security boundary within which users, computers, groups, and other objects are accessible.

Organizational units

The objects held within a domain can be grouped into Organizational Units (OUs). OUs can provide hierarchy to a domain, ease its administration, and can resemble the organization's structure in managerial or geographical terms. OUs can contain other OUs—domains are containers in this sense. Microsoft recommends using OUs rather than domains for structure and to simplify the implementation of policies and administration. The OU is the recommended level at which to apply group policies, which are Active Directory objects formally named Group Policy Objects (GPOs), although policies can also be applied to domains or sites (see below). The OU is the level at which administrative powers are commonly delegated, but delegation can be performed on individual objects or attributes as well.
Organizational Units are an abstraction for the administrator and do not function as containers; the underlying domain is the true container. It is not possible, for example, to create user accounts with an identical username (sAMAccountName) in separate OUs, such as "fred.staff-ou.domain" and "fred.student-ou.domain", where "staff-ou" and "student-ou" are the OUs. This is so because sAMAccountName, a user object attribute, must be unique within the domain. However, two users in different OUs can have the same Common Name (CN), the first component of the Distinguished Name (DN) of the user. Thus from the point of view of the DN, OUs do function as containers.
As the number of users in a domain increases, conventions such as "first initial, middle initial, last name" (Western order) or the reverse (Eastern order) fail for common family names like Li (李), Smith or Garcia. Workarounds include adding a digit to the end of the username. Alternatives include creating a separate ID system of unique employee/student

SARDAR AZEEM(PICT)

id numbers to use as account names in place of actual user's names, and allowing users to nominate their preferred word sequence within an acceptable use policy.
Because duplicate usernames cannot exist within a domain, account name generation poses a significant challenge for large organizations that cannot be easily subdivided into separate domains, such as students in a public school system or university who must be able to use any computer across the network.
In Active Directory, organizational units cannot be assigned as owners or trustees. Only groups are selectable, and members of OUs cannot be collectively assigned rights to directory objects.
In Microsoft's Active Directory, OUs do not confer access permissions, and objects placed within OUs are not automatically assigned access privileges based on their containing OU. This is a design limitation specific to Active Directory. Other competing directories such as Novell NDS are able to assign access privileges through object placement within an OU.
Active Directory requires a separate step for an administrator to assign an object in an OU as a member of a group also within that OU. Relying on OU location alone to determine access permissions is unreliable, because the object may not have been assigned to the group object for that OU.
A common workaround for an Active Directory administrator is to write a custom PowerShell or Visual Basic script to automatically create and maintain a user group for each OU in their directory. The scripts are run periodically to update the group to match the OU's account membership, but are unable to instantly update the security groups anytime the directory changes, as occurs in competing directories where security is directly implemented into the directory itself. Such groups are known as Shadow Groups. Once created, these shadow groups are selectable in place of the OU in the administrative tools.
Microsoft refers to shadow groups in the Server 2008 Reference documentation, but does not explain how to create them. There are no built-in server methods or console snap-ins for managing shadow groups.
The division of an organization's information infrastructure into a hierarchy of one or more domains and top-level OUs is a key decision. Common models are by business unit, by geographical location, by IT Service, or by object type and hybrids of these. OUs should be structured primarily to facilitate administrative delegation, and secondarily, to facilitate group policy application. Although OUs form an administrative boundary, the only true security boundary is the forest itself and an administrator of any domain in the forest must be trusted across all domains in the forest.

The Domain Name System (DNS)

Is a hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network. It associates various information with domain names assigned to each of the participating entities. A Domain Name Service resolves queries for these names into IP addresses for the purpose of locating computer services and devices worldwide. By providing a worldwide, distributed keyword-based redirection service, the Domain Name System is an essential component of the functionality of the Internet.
An often-used analogy to explain the Domain Name System is that it serves as the phone book for the Internet by translating human-friendly computer hostnames into IP addresses. For example, the domain name www.example.com translates to the addresses 192.0.43.10 (IPv4) and 2620:0:2d0:200::10 (IPv6). Unlike a phone book, however, DNS can be quickly updated and these updates distributed, allowing a service's location on the network to change without affecting the end users, who continue to use the same hostname. Users take advantage of this when they recite meaningful Uniform Resource Locators (URLs) and e-mail addresses without having to know how the computer actually locates the services.
The Domain Name System distributes the responsibility of assigning domain names and mapping those names to IP addresses by designating authoritative name servers for each domain. Authoritative name servers are assigned to be responsible for their particular domains, and in turn can assign other authoritative name servers for their sub-domains. This mechanism has made the DNS distributed and fault tolerant and has helped avoid

SARDAR AZEEM (PICT)

the need for a single central register to be continually consulted and updated. Additionally, the responsibility for maintaining and updating the master record for the domains is spread among many domain name registrars, who compete for the end-user's, domain-owner's, business. Domains can be moved from registrar to registrar at any time. The Domain Name System also specifies the technical functionality of this database service. It defines the DNS protocol, a detailed specification of the data structures and communication exchanges used in DNS, as part of the Internet Protocol Suite.

Introduction to Linux

Linux is a Unix-like computer operating system assembled under the model of free and open source software development and distribution. The defining component of Linux is the Linux kernel, an operating system kernel first released 5 October 1991 by Linus Torvalds.

Linux was originally developed as a free operating system for Intel x86-based personal computers. It has since been ported to more computer hardware platforms than any other operating system. It is a leading operating system on servers and other big iron systems such as mainframe computers and supercomputers. More than 90% of today's 500 fastest supercomputers run some variant of Linux, including the 10 fastest. Linux also runs on embedded systems (devices where the operating system is typically built into the firmware and highly tailored to the system) such as mobile phones, tablet computers, network routers, televisions and video game consoles; the Android system in wide use on mobile devices is built on the Linux kernel.

The development of Linux is one of the most prominent examples of free and open source software collaboration: the underlying source code may be used, modified, and distributed—commercially or non-commercially—by anyone under licenses such as the GNU General Public License. Typically Linux is packaged in a format known as a Linux distribution for desktop and server use. Some popular mainstream Linux distributions include Debian (and its derivatives such as Ubuntu), Fedora and openSUSE. Linux distributions include the Linux kernel, supporting utilities and libraries and usually a large amount of application software to fulfill the distribution's intended use.

A distribution oriented toward desktop use will typically include the X Window System and an accompanying desktop environment such as GNOME or KDE Plasma. Some such distributions may include a less resource intensive desktop such as LXDE or Xfce for use on older or less powerful computers. A distribution intended to run as a server may omit all graphical environments from the standard install and instead include other software such as the Apache HTTP Server and an SSH server such as OpenSSH. Because Linux is freely redistributable, anyone may create a distribution for any intended use. Applications commonly used with desktop Linux systems include the Mozilla Firefox web browser, theLibreOffice office application suite, and the GIMP image editor.

Installation of Red Hat Enterprise Linux 0.5

System Requirement:

RAM: 256

MB Processor: PIV

Disk Space: 10

GB

To delete partition:

1. Right-click on My Computer
2. Click Manage
3. Click Disk Management

4. Right-Click on D: drive (or any drive that you want to delete)
5. Click Delete Logical Drive…

   Message: All data on this volume will be lost. Do you want to continue?
6. Click Yes
7. Close

Boot from CD-ROM:

1 Insert Red Hat Enterprise Linux CD into CD-ROM (RHEL_5.2 i386
2 Click Start
3 Click Shut Down…
4 Select Restart
5 Click OK

To Install Red Hat Enterprise Linux 5:

1. Type boot: linux text and Press Enter Button (To install or upgrade in text mode) OR

   boot: Press Enter Button (To install or upgrade in graphical mode)
2. Press OK (To begin testing the CD media before installation)
3. Press Test (To test the CD currently in the drive)

   OR

   Press Skip (To skip the media test and start the
4. Red Hat Enterprise Linux 5, Click Next

   What language would you like to us during the installation process?
5. Select English (English)
6. Click Next

   Select the appropriate keyboard for the system.
7. Select U.S. English
8. Click Next

   Installation Number
9. Click Skip entering Installation Number
10. Click Ok

    If you're unable to locate the Installation Number, consult http://www.redhat.com/apps/support/in.html.
11. Click Skip

    Installation requires partitioning of your hard drive. You can either choose to use this or create your won.
12. Select Create custom layout
13. Click Next

    Partitioning hard drive
14. Click New

    Add Partition
15. Mount Point: Select "/" slash (Root)
16. File System Type: Select ext3

SARDARAZEEM(PICT)

17. Size (MB): Enter 7000 (7 GB)
18. Click OK
19. Click New (again)
20. Mount Point: Select /boot
21. File System Type: Select ext3
22. Size (MB): Enter 100
23. Click OK
24. Click New (again)
25. File System Type: Select swap
26. Size (MB): Enter 1024 (double of RAM size)
27. Click OK
28. Click New (again)
29. Mount Point: Type /data
30. File System Type: Select vfat
31. Size (MB): Enter 2110 (all free space available)
32. Click OK
33. Click Next
34. Click The GRUB boot loader will be installed on /devsda
35. Click Next 36. Check Other
37. Click Edit button
38. Enter "Microsoft Windows XP"
39. Click OK

Time Zone: Please click into the map to choose a region or:

40. Select Asia/Karachi
41. Click Next

The root account is used for administering the system. Enter a password for the root user.

42. Enter Root Password: 123456 (password must be minimum 6 character long)
43. Confirm: 123456
44. Click Next

Custom Installation

45. Click Customize now
46. Click Next
47. Select Desktop Environments
48. Select GNOME Desktop Environment (31 out of 36 optional packages are installed. To install all the 36 options)
49. Click Optional packages button

Packages in GNOME Desktop Environment. Please choose the packages which you would like to have installed.

50. Check all uncheck options one by one
51. Click Close button (now all the 36 packages are installed)
52. Check KDE (K Desktop Environment) (6 out of 7 optional packages are installed. To install all the 7 options):
53. Click Optional packages button
54. Check all uncheck options one by one
55. Click Close button

(Similarly select main categories from left panel and installed its optional

SARDARAZEEM(PICT)

packages that you want to install one by one.)

56. Click Next

To begin installation of Red hat Enterprise Linux Server:

57. Click Next

The software we have selected to install will require the following CDs:

a. Red hat Enterprise Linux Server 5 CD # 1 b. Red hat Enterprise Linux Server 5 CD # 2 c. Red hat Enterprise Linux Server 5 CD # 3 d. Red hat Enterprise Linux Server 5 CD # 4 e. Red hat Enterprise Linux Server 5 CD # 5

58. Click Continue button (Starting install process. This may take several minute.

SARDARAZEEM(PICT)

History of Linux

Unix

The UNIX operating system was conceived and implemented in **1969** at AT&T's Bell Laboratories in the United States by Ken Thompson, Dennis Ritchie, Douglas McElroy, and Joe Ossanna. It was first released in 1971 and was initially entirely written in assembly language, a common practice at the time. Later, in a key pioneering approach in 1973, UNIX was re-written in the programming language C by Dennis Ritchie (with exceptions to the kernel and I/O). The availability of an operating system written in a high-level language allowed easier portability to different computer platforms. With a legal glitch forcing AT&T to license the operating system's source code to anyone who asked, UNIX quickly grew and became widely adopted by academic institutions and businesses. In **1984**, AT&T divested itself of Bell Labs. Free of the legal glitch requiring free licensing, Bell Labs began selling UNIX as proprietary product.

GNU

The GNU Project, started in **1983** by Richard Stallman, had the goal of creating a "complete Unix-compatible software system" composed entirely of free software. Work began in 1984. Later, in 1985, Stallman started the Free Software Foundation and wrote the GNU General Public License (GNU GPL) in **1989**. By the early 1990s, many of the programs required in an operating system (such as libraries, compilers, text editors, a UNIX shell, and a windowing system) were completed, although low-level elements such as device drivers, daemons, and the kernel were stalled and incomplete. Linus Torvalds has said that if the GNU kernel had been available at the time (1991), he would not have decided to write his own.

BSD

Although not released until 1992 due to legal complications, development of 386BSD, from whichNetBSD and FreeBSD descended, predated that of Linux. Linus Torvalds has said that if 386BSD had been available at the time, he probably would not have created Linux.

MINIX

MINIX is an inexpensive minimal Unix-like operating system, designed for education in computer science, written by Andrew S. Tanenbaum. Starting with version in 2005, MINIX has become free and redesigned for "serious" use.
In 1991 while attending the University of Helsinki, Torvalds became curious about operating systems and frustrated by the licensing of MINIX, which limited it to educational use only. He began to work on his own operating system which eventually became the Linux kernel.
Torvalds began the development of the Linux kernel on MINIX, and applications written for MINIX were also used on Linux. Later Linux matured and further Linux development took place on Linux systems. GNU applications also replaced all MINIX components, because it was advantageous to use the freely available code from the GNU project with the fledgling operating system. (Code licensed under the GNU GPL can be reused in other projects as long as they also are released under the same or a compatible license.) Torvalds initiated a switch from his original license, which prohibited commercial redistribution, to the GNU GPL. Developers worked to integrate GNU components with Linux to make a fully functional and free operating system.

Ubuntu, a popular Linux distribution

Today, Linux systems are used in every domain, from embedded systems to supercomputers, and have secured a place in server installations often using the popular LAMP application stack. Use of Linux distributions in home and enterprise desktops has been growing. They have also gained popularity with various local and national governments. The federal government of Brazil is well known for its support for Linux. News of the Russian military creating its own Linux distribution has also surfaced, and has come to fruition as the G.H.ost Project. The Indian state of Kerala has gone to the extent of mandating that all state high schools run Linux on their computers. China uses Linux exclusively as the operating system for its Loongsonprocessor family to achieve technology independence. In Spain some

regions have developed their own Linux distributions, which are widely used in education and official institutions, like gnuLinEx in Extremadura and Guadalinex in Andalusia. Portugal is also using its own Linux distribution Caixa Mágica, used in the Magalhães netbook and the e-escolar government program. France and Germany have also taken steps toward the adoption of Linux.

User interface

Users operate a Linux-based system through a **command line interface (CLI),** a **graphical user interface (GUI)**, or through controls attached to the associated hardware, which is common for embedded systems. For desktop systems, the default mode is usually a graphical user interface, by which the **CLI** is available through terminal emulator windows or on a separate virtual console. Most low-level Linux components, including the GNU userland, use the CLI exclusively. The CLI is particularly suited for automation of repetitive or delayed tasks, and provides very simple inter-process communication. A graphical terminal emulator program is often used to access the CLI from a Linux desktop. A Linux system typically implements a CLI by a shell, which is also the traditional way of interacting with a UNIX system. A Linux distribution specialized for servers may use the CLI as its only interface.

On desktop systems, the most popular user interfaces are the extensive desktop environments KDE Plasma Desktop, GNOME, and Xfce; though a variety of additional user interfaces exist. Most popular user interfaces are based on the X Window System, often simply called "X". It provides network transparency and permits a graphical application running on one system to be displayed on another where a user may interact with the application.

Other GUIs may be classified as simple X window managers, such as FVWM, Enlightenment, and Window Maker, which provide aminimalist functionality with respect to the desktop environments. A window manager provides a means to control the placement and appearance of individual application windows, and interacts with the X Window System. The desktop environments include window managers as part of their standard installations (Mutter for GNOME, KWin for KDE, Xfwm for Xfce as of January 2012) although users may choose to use a different window manager if preferred.

Linux's directory structure

As you may have noticed, Linux organizes its files differently from Windows. First the directory structure may seem un logical and strange and you have no idea where all the programs, icons, config files, and others are. This tuXfile will take you to a guided tour through the Linux file system. This is by no means a complete list of all the directories on Linux, but it shows you the most interesting places in your file system.

< / >

The root directory. The starting point of your directory structure. This is where the Linux system begins. Every other file and directory on your system is under the root directory. Usually the root directory contains only subdirectories, so it's a bad idea to store single files directly under root.

Don't confuse the root directory with the root user account, root password (which obviously is the root user's password) or root user's home directory.

< /boot >

As the name suggests, this is the place where Linux keeps information that it needs when booting up. For example, this is where the Linux kernel is kept. If you list the contents of /boot, you'll see a file called **vmlinuz** - that's the kernel.

< /etc >

The configuration files for the Linux system. Most of these files are text files and can be edited by hand. Some interesting stuff in this directory:

/etc/inittab

A text file that describes what processes are started at system bootup and during normal operation. For example, here you can determine if you want the X Window System to start automatically at bootup, and configure what happens when a user presses Ctrl+Alt+Del.

/etc/fstab

this file contains descriptive information about the various file systems and their mount points, like floppies, cdroms, and so on.

/etc/passwd

A file that contains various pieces of information for each user account. This is where the users are defined.

< /bin, /usr/bin >

These two directories contain a lot of programs (binaries, hence the directory's name) for the system. The /bin directory contains the most important programs that the system needs to operate, such as the shells, ls, grep, and other essential things. /usr/bin in turn contains applications for the system's users. However, in some cases it really doesn't make much difference if you put the program in /bin or /usr/bin.

< /sbin, /usr/sbin >

Most system administration programs are stored in these directories. In many cases you must run these programs as the root user.

< /usr >

This directory contains user applications and a variety of other things for them, like their source codes, and pictures, docs, or config files they use. /usr is the largest directory on a Linux system, and some people like to have it on a separate partition. Some interesting stuff in /usr:

/usr/doc

Documentation for the user apps, in many file formats.

/usr/share

Config files and graphics for many user apps.

/usr/src

Source code files for the system's software, including the Linux kernel.

/usr/include

Header files for the C compiler. The header files define structures and constants that are needed for building most standard programs. A subdirectory under /usr/include contains headers for the C++ compiler.

/usr/X11R6

The X Window System and things for it. The subdirectories under /usr/X11R6 may contain some X binaries themselves, as well as documentation, header files, config files, icons, sounds, and other things related to the graphical programs.

< /usr/local >

This is where you install apps and other files for use on the local machine. If your machine is a part of a network, the /usr directory may physically be on another machine and can be shared by many networked Linux workstations. On this kind of a network, the /usr/local directory contains only stuff that is not supposed to be used on many machines and is intended for use at the local machine only.

Most likely your machine isn't a part of a network like this, but it doesn't mean that /usr/local is useless. If you find interesting apps that aren't officially a part of your distro, you should install them in /usr/local. For example, if the app would normally go to /usr/bin but it isn't a part of your distro, you should install it in /usr/local/bin instead. When you keep your own programs away from the programs that are included in your distro, you'll avoid confusion and keep things nice and clean.

< /lib >

The shared libraries for programs that are dynamically linked. The shared libraries are similar to DLL's on Wingbows.

< /home >

This is where users keep their personal files. Every user has their own directory under /home, and usually it's the only place where normal users are allowed to write files. You can configure a Linux system so that normal users can't even list the contents of other users' home directories. This means that if your family members have their own user

accounts on your Linux system, they won't see all the w4r3z you keep in your home directory. ;-)

< /root >

The superuser's (root's) home directory. Don't confuse this with the root directory (/) of a Linux system.

< /var >

This directory contains variable data that changes constantly when the system is running. Some interesting subdirectories:

/var/log

A directory that contains system log files. They're updated when the system runs, and checking them out can give you valuable info about the health of your system. If something in your system suddenly goes wrong, the log files may contain some info about the situation.

/var/mail

Incoming and outgoing mail is stored in this directory.

/var/spool

This directory holds files that are queued for some process, like printing.

< /tmp >

Programs can write their temporary files here.

< /dev >

The devices that are available to a Linux system. Remember that in Linux, devices are treated like files and you can read and write devices like they were files. For example, /dev/fd0 is your first floppy drive, /dev/cdrom is your CD drive, /dev/hda is the first IDE hard drive, and so on. All the devices that a Linux kernel can understand are located under /dev, and that's why it contains hundreds of entries.

< /mnt >

This directory is used for mount points. The different physical storage devices (like the hard disk drives, floppies, CD-ROM's) must be attached to some directory in the file system tree before they can be accessed. This attaching is called mounting, and the directory where the device is attached is called the mount point.

The /mnt directory contains mount points for different devices, like /mnt/floppy for the floppy drive, /mnt/cdrom for the CD-ROM, and so on. However, you're not forced to use the /mnt directory for this purpose, you can use whatever directory you wish. Actually in some distros, like Debian and SuSE, the default is to use /floppy and /cdrom as mount points instead of directories under /mnt.

< /proc >

This is a special directory. Well, actually /proc is just a virtual directory, because it doesn't exist at all! It contains some info about the kernel itself. There's a bunch of numbered entries that correspond to all processes running on the system, and there are also named entries that permit access to the current configuration of the system. Many of these entries can be viewed.

< /lost+found >

Here Linux keeps the files that it restores after a system crash or when a partition hasn't been uncounted before a system shutdown. This way you can recover files that would otherwise have been lost.

Shell

The Linux/Unix shell refers to a special program that allows you to interact with it by entering certain commands from the keyboard; the shell will execute the commands and display its output on the monitor. The environment of interaction is text-based (unlike the GUI-based interaction we have been using in the previous chapters) and since it is command-oriented this type of interface is termed Command Line interface or CLI. Before the advent of GUI-based computing environments, the CLI was the only way that one can interact and access a computer system.

Up until now, there was never a need to type commands into a shell; and with the modernization and creation of a lot of newer GUI-based tools, the shell is becoming increasingly un-required to perform many tasks. But that said, the shell is a very powerful

place, and a lot is achieved through it. A lot of the front-end GUI methods of doing things have similar ways and means to get done with using the shell. Professional Linux and UNIX users find the shell very powerful, and an introduction to at least the basic shell usage is useful.

Linux Commands

1.	# cal	- To display current month calendar.
2.	# cal 08 2010	- To display calendar of August 2010
3.	# date	- To display current system date and time
4.	# clear	- To clear the terminal screen
5.	# dir	- To list directory contents
6.	# ls	- To list directory contents with color
7.	# ls –l	- To display properties of file and folders etc
8.	# ls –a	- To display all files including hides files.
In linux if a file or folder starts with dot "." it means it is a hidden file. To hide a file or folder just rename file and place dot "." in front of it.		
9	# ls –l Desktop	- To display Desktop folder property
10	# ls –lh	- To show file and folder size in MB and GB.
11	# whatis ls	- To know more about command ls.
12	# info cal	- To know command ls in details.
13	# ls --help	- To display all options of ls command.
14.	ls –help | more	- To display page by page use pipe sign. To quit press q. to move line by line press enter button. To move page by page press space bar.
15	# man ls	- To display manual of ls command.
16	# cd ..	- CD space double dot to exit directory one by one.

- There is no extension of system files in linux
- Extensions are replaced with color
- **Black / White** - **Text files**
- **Blue** - **Folder**
- **Green** - **Executable or Commands**
- **Red** - **Backup, Compress and RPM (.exe)**
- **Light Blue** - **Shortcuts or link file**
- **Pink** - **Link**

There are 6 locations of commands

1- /bin
2- /usr/bin
3- /usr/local/bin
4- /sbin
5- /user/sbin

17	Copy File	Copy files from source to destination Syntax: cp source destination [root@localhost ~]# cp corvit /opt - Copy corvit file from root to opt folder
18	# ls /opt	- To confirm file copied
19	# pwd	- Print name of current / working directory i.e./root.

20	# cp /etc/fstab /opt Eg # cp /etc/named.caching.nameserve.conf	- Copy fstab file from etc folder to opt destination folder from root. - Copy from etc folder to root. "." means current folder in which we are present.
21	Copy Folder	To copy folder/directory from source to destination. Syntax: cp –r source destination - We use –r to copy folder # cp –r ISB /opt
22	Cut File/Folder	Move / Cut (rename) files or folder from source to destination. Syntax: mv source destination
23	Rename File or Folder	To rename file or folder we use mv command. Syntax: mv
24	Delete File	Remove files or directories Syntax: rm filename # rm fstab rm: remove regular file 'fstab'? y - Press Y to confirm deletion # rm /opt/fstab - To remove file in opt
25	Delete Folder	# rm –r ISB
26	Delete Empty Folder	Syntax: rmdir foldername # rmdir /opt/ISB - To delete empty folder ISB in opt Make Directory / Create Folder mkdir - Make directories Syntax: mkdir foldername # mkdir text
27	Delete all files and folders (del tree)	# rm –rf name - To delete all files and folders –rf means forcefully.

28	# fdisk –l	- Partition table manipulator for Linux - To check how many partitions are in use.
29	# df –h	- Report file system disk space usage.
30	cat Allows you to look, modify or combine a file.	Syntax cat filename [-n] [-b] [-u] [-s] [-v] Examples cat file1.txt file2.txt > file3.txt Reads file1.txt and file2.txt and combines those files to make file3.txt.
31	AT Schedules a command to be ran at a particular time, such as a print job late at night.	Syntax at [-c | -k | -s] [-f filename] [-q queuename] [-m] -t time [date] [-l] [-r] Examples at -m 01:35 < atjob Run the commands listed in the 'atjob' file at 1:35AM, in addition all output that is generated from job mail to the user running the task. When this command has been successfully enter you should receive a prompt similar to the below example. commands will be executed using /bin/csh job 1072250520.a at Wed Dec 24 00:22:00 2003 at -l This command will list each of the scheduled jobs as seen below. 1072250520.a Wed Dec 24 00:22:00 2003 at -r 1072250520.a Deletes the job just created. atrm 23 Deletes job 23. If you wish to create a job that is repeated you could modify the file that executes the commands with another command that recreates the job or better yet use the crontab command

32	GREP grep command to suppress normal output and display a count of matching lines for each input file.	Syntax grep -c "word" file grep -c "string" file Example grep -c 'var' /etc/passwd

33	• **zip** is a compression and file packaging utility for Linux and Unix (including FreeBSD, Solaris etc).
	• **unzip** will list, test, or extract files from a ZIP archive files.
		zip examples:
		Creates the archive data.zip and puts all the files in the current directory in it in compressed form
		$ zip data *
		No need to add .zip extension or suffix as it is added automatically by zip command.
		To zip up an entire directory (including all subdirectories), the command:
		$ zip -r data *
		unzip example:
		To use unzip to extract all files of the archive pics.zip into the current directory & subdirectories:
		$ unzip pics.zip

34	**Chmod** Changes the permission of a file. Permissions u - User who owns the file. g - Group that owns the file. o - Other. a - All. r - Read the file. w - Write or edit the file. x - Execute or run the file as a program. Numeric Permissions: CHMOD can also to attributed by using Numeric Permissions: 400 read by owner 040 read by group 004 read by anybody (other) 200 write by owner 020 write by group 002 write by anybody 100 execute by owner 010 execute by group 001 execute by anybody	Syntax chmod [OPTION]... MODE[,MODE]... FILE... chmod [OPTION]... OCTAL-MODE FILE... chmod [OPTION]... --reference=RFILE FILE... Examples The above numeric permissions can be added to set a certain permission, for example, a common HTML file on a Unix server to be only viewed over the Internet would be: chmod 644 file.htm This gives the file read/write by the owner and only read by everyone else (-rw-r--r--). Files such as scripts that need to be executed need more permissions. Below is another example of a common permission given to scripts. chmod 755 file.cgi This would be the following 400+040+004+200+100+010+001 = 755 where you are giving all the rights except the capability for anyone to write to the file.cgi file(-rwxr-xr-x). chmod 666 file.txt Finally, another common CHMOD permission is 666, as shown below, which is read and write by everyone.