

INTRODUCTION TO UNIX

CHAP # 1 Background and some Basic Commands

Introduction

This chapter introduces you to the UNIX operating system. We first look at what is an operating system and then proceed to discuss the different features of UNIX that have made it a popular operating system.

Objectives

- What is an operating system (OS)?
- Features of UNIX OS
- A Brief History of UNIX OS, POSIX and Single Unix Specification (SUS)

1. What is an operating system (OS)?

An operating system (OS) is a resource manager. It takes the form of a set of software routines that allow users and application programs to access system resources (e.g. the CPU, memory, disks, modems, printers, network cards etc.) in a **safe, efficient** and **abstract** way.

For example, an OS ensures **safe** access to a printer by allowing only one application program to send data directly to the printer at any one time. An OS encourages **efficient** use of the CPU by suspending programs that are waiting for I/O operations to complete to make way for programs that can use the CPU more productively. An OS also provides convenient **abstractions** (such as files rather than disk locations) which isolate application programmers and users from the details of the underlying hardware.

User Applications	
Application Programs	System Utilities
System Call Interface	
Operating System Kernel	
Processor/Hardware	

UNIX Operating system allows complex tasks to be performed with a few keystrokes. It doesn't tell or warn the user about the consequences of the command.

Kernighan and Pike (The UNIX Programming Environment) lamented long ago that "as the UNIX system has spread, the fraction of its users who are skilled in its application has decreased." However, the capabilities of UNIX are limited only by your imagination.

2. Features of UNIX OS

Several features of UNIX have made it popular. Some of them are:

Portable

UNIX can be installed on many hardware platforms. Its widespread use can be traced to the decision to develop it using the C language.

Multuser

The UNIX design allows multiple users to concurrently share hardware and software

Multitasking

UNIX allows a user to run more than one program at a time. In fact more than one program can be running in the background while a user is working foreground.

Networking

While UNIX was developed to be an interactive, multiuser, multitasking system, networking is also incorporated into the heart of the operating system. Access to another system uses a standard communications protocol known as Transmission Control Protocol/Internet Protocol (TCP/IP).

Organized File System

UNIX has a very organized file and directory system that allows users to organize and maintain files.

Device Independence

UNIX treats input/output devices like ordinary files. The source or destination for file input and output is easily controlled through a UNIX design feature called redirection.

Utilities

UNIX provides a rich library of utilities that can be use to increase user productivity.

3. A Brief History of UNIX

In the late 1960s, researchers from General Electric, MIT and Bell Labs launched a joint project to develop an ambitious multi-user, multi-tasking OS for mainframe computers known as MULTICS (Multiplexed Information and Computing System). MULTICS failed, but it did inspire Ken Thompson, who was a researcher at Bell Labs, to have a go at writing a simpler operating system himself. He wrote a simpler version of MULTICS on a PDP7 in assembler and called his attempt UNICS (Uniplexed Information and Computing System). Because memory and CPU power were at a premium in those days, UNICS (eventually shortened to UNIX) used short commands to minimize the space needed to store them and the time needed to decode them - hence the tradition of short UNIX commands we use today, e.g. `ls`, `cp`, `rm`, `mv` etc.

Ken Thompson then teamed up with Dennis Ritchie, the author of the first C compiler in 1973. They rewrote the UNIX kernel in C - this was a big step forwards in terms of the system's portability - and released the Fifth Edition of UNIX to universities in 1974. The Seventh Edition, released in 1978, marked a split in UNIX development into two main branches: SYSV (System 5) and BSD (Berkeley Software Distribution). BSD arose from the University of California at Berkeley where Ken Thompson spent a sabbatical year. Its development was continued by students at Berkeley and other research institutions. SYSV was developed by AT&T and other commercial companies. UNIX flavors based on SYSV have traditionally been more conservative, but better supported than BSD- based flavors.

Until recently, UNIX standards were nearly as numerous as its variants. In early days, AT&T published a document called System V Interface Definition (SVID). X/OPEN (now The Open Group), a consortium of vendors and users, had one too, in the X/Open Portability Guide (XPG). In the US, yet another set of standards, named Portable Operating System Interface for Computer Environments (POSIX), were developed at the behest of the Institution of Electrical and Electronics Engineers (IEEE).

In 1998, X/OPEN and IEEE undertook an ambitious program of unifying the two standards. In 2001, this joint initiative resulted in a single specification called the Single UNIX Specification, Version 3 (SUSV3), that is also known as IEEE 1003.1:2001 (POSIX.1). In 2002, the International Organization for Standardization (ISO) approved SUSV3 and IEEE 1003.1:2001.

Some of the commercial UNIX based on system V are:

- IBM's AIX
- Hewlett-Packard's HPUX
- SCO's Open Server Release 5
- Silicon Graphics' IRIS
- DEC's Digital UNIX
- Sun Microsystems' Solaris 2

Some of the commercial UNIX based on BSD are:

- SunOS 4.1.X (now Solaris)
- DEC's Ultrix
- BSD/OS, 4.4BSD

Some Free UNIX are:

- Linux, written by Linus Torvalds at University of Helsinki in Finland.
- FreeBSD and NetBSD, a derivative of 4.4BSD

Conclusion

In this chapter we defined an operating system. We also looked at history of UNIX and features of UNIX that make it a popular operating system. We also discussed the convergence of different flavors of UNIX into Single Unix Specification (SUS) and Portable Operating System Interface for Computing Environments (POSIX).

The UNIX Architecture and Command Usage


Introduction

In order to understand the subsequent chapters, we first need to understand the architecture of UNIX and the concept of division of labor between two agencies viz., the shell and the kernel. This chapter introduces the architecture of UNIX. Next we discuss the rich collection of UNIX command set, with a specific discussion of command structure and usage of UNIX commands. We also look at the man command, used for obtaining online help on any UNIX command. Sometimes the keyboard sequences don't work, in which case, you need to know what to do to fix them. Final topic of this chapter is troubleshooting some terminal problems.

Objectives

- The UNIX Architecture
- Locating Commands
- Internal and External Commands
- Command Structure and usage
- Flexibility of Command Usage
- The man Pages, apropos and whatis
- Troubleshooting the terminal problems

1. The UNIX Architecture



UNIX architecture comprises of two major components viz., the shell and the kernel. The kernel interacts with the machine's hardware and the shell with the user.

The kernel is the core of the operating system. It is a collection of routines written in C. It is loaded into memory when the system is booted and communicates directly with the hardware. User programs that need to access the hardware use the services of the kernel via use of system calls and the kernel performs the job on behalf of the user. Kernel is also responsible for managing system's memory, schedules processes, decides their priorities.

The shell performs the role of command interpreter. Even though there's only one kernel running on the system, there could be several shells in action, one for each user who's logged in. The shell is responsible for interpreting the meaning of metacharacters if any, found on the command line before dispatching the command to the kernel for execution.

The File and Proces

A file is an array of bytes that stores information. It is also related to another file in the sense that both belong to a single hierarchical directory structure.

A process is the second abstraction UNIX provides. It can be treated as a time image of an executable file. Like files, processes also belong to a hierarchical structure. We will be discussing the processes in detail in a subsequent chapter.

2. Locating Files

All UNIX commands are single words like `ls`, `cd`, `cat`, etc. These names are in lowercase. These commands are essentially *files* containing programs, mainly written in C. Files are stored in directories, and so are the binaries associated with these commands. You can find the location of an executable program using `type` command:

```
$ type ls
ls is /bin/ls
```

This means that when you execute `ls` command, the shell locates this file in `/bin` directory and makes arrangements to execute it.

The Path

The sequence of directories that the shell searches to look for a command is specified in its own `PATH` variable. These directories are colon separated. When you issue a command, the shell searches this list in the sequence specified to locate and execute it.

3. Internal and External Commands

Some commands are implemented as part of the shell itself rather than separate executable files. Such commands that are built-in are called internal commands. If a command exists both as an internal command of the shell as well as an external one (in `/bin` or `/usr/bin`), the shell will accord top priority to its own internal command with the same name. Some built-in commands are `echo`, `pwd`, etc.

4. Command Structure

UNIX commands take the following general form:

```
verb [options] [arguments]
```

where verb is the command name that can take a set of optional options and one or more optional arguments.

Commands, options and arguments have to be separated by spaces or tabs to enable the shell to interpret them as words. A contiguous string of spaces and tabs together is called a whitespace. The shell compresses multiple occurrences of whitespace into a single whitespace.

Options

An option is preceded by a minus sign (-) to distinguish it from filenames.

Example: `$ ls -l`

There must not be any whitespaces between - and l. Options are also arguments, but given a special name because they are predetermined. Options can be normally compined with only one - sign. i.e., instead of using

```
$ ls -l -a -t
```

we can as well use,

```
$ ls -lat
```

Because UNIX was developed by people who had their own ideas as to what options should look like, there will be variations in the options. Some commands use + as an option prefix instead of -.

Filename Arguments

Many UNIX commands use a filename as argument so that the command can take input from the file. If a command uses a filename as argument, it will usually be the last argument, after all options.

Example: `cp file1 file2 file3 dest_dir`
`rm file1 file2 file3`

The command with its options and argumens is known as the command line, which is considered as complete after *[Enter]* key is pressed, so that the entire line is fed to the shell as its input for interpretation and execution.

Exceptions

Some commands in UNIX like `pwd` do not take any options and arguments. Some commands like `who` may or may not be specified with arguments. The `ls` command can run without arguments (`ls`), with only options (`ls -l`), with only filenames (`ls f1 f2`), or using a combination of both (`ls -l f1 f2`). Some commands compulsorily take options (`cut`). Some commands like `grep`, `sed` can take an expression as an argument, or a set of instructions as argument.

5. Flexibility of Command Usage

UNIX provides flexibility in using the commands. The following discussion looks at how permissive the shell can be to the command usage.

Combining Commands

Instead of executing commands on separate lines, where each command is processed and executed before the next could be entered, UNIX allows you to specify more than one command in the single command line. Each command has to be separated from the other by a ; (semicolon).

```
wc sample.txt ; ls -l sample.txt
```

You can even group several commands together so that their combined output is redirected to a file.

```
(wc sample.txt ; ls -l sample.txt) > newfile
```

When a command line contains a semicolon, the shell understands that the command on each side of it needs to be processed separately. Here ; is known as a metacharacter.

Note: When a command overflows into the next line or needs to be split into multiple lines, just press enter, so that the secondary prompt (normally >) is displayed and you can enter the remaining part of the command on the next line.

Entering a Command before previous command has finished

You need not have to wait for the previous command to finish before you can enter the next command. Subsequent commands entered at the keyboard are stored in a buffer (a temporary storage in memory) that is maintained by the kernel for all keyboard input. The next command will be passed on to the shell for interpretation after the previous command has completed its execution.

6. man: Browsing The Manual Pages Online

UNIX commands are rather cryptic. When you don't remember what options are supported by a command or what its syntax is, you can always view man (short for manual) pages to get online help. The man command displays online documentation of a specified command.

A pager is a program that displays one screenful information and pauses for the user to view the contents. The user can make use of internal commands of the pager to scroll up and scroll down the information. The two popular pagers are more and less. more is the Berkeley's pager, which is a superior alternative to original pg command. less is the standard pager used on Linux systems. less is modeled after a popular editor called vi and is more powerful than more as it provides vi-like navigational and search facilities. We can use pagers with commands like ls | more. The man command is configured to work with a pager.

7. Understanding The man Documentation

The man documentation is organized in eight (08) sections. Later enhancements have added subsections like 1C, 1M, 3N etc.) References to other sections are reflected as SEE ALSO section of a man page.

A man page is divided into a number of compulsory and optional sections. Every command doesn't need all sections, but the first three (NAME, SYNOPSIS and DESCRIPTION) are generally seen in all man pages. NAME presents a one-line introduction of the command. SYNOPSIS shows the syntax used by the command and DESCRIPTION provides a detailed description.

The SYNOPSIS follows certain conventions and rules:

- If a command argument is enclosed in rectangular brackets, then it is optional; otherwise, the argument is required.
- The ellipsis (a set of three dots) implies that there can be more instances of the preceding word.
- The | means that only one of the options shows on either side of the pipe can be used.

All the options used by the command are listed in OPTIONS section. There is a separate section named EXIT STATUS which lists possible error conditions and their numeric representation.

Note: You can use man command to view its own documentation (\$ man man). You can also set the pager to use with man (\$ PAGER=less ; export PAGER). To understand which pager is being used by man, use \$ echo \$PAGER.

The following table shows the organization of man documentation.

Section	Subject (SVR4)	Subject (Linux)
1	User programs	User programs
2	Kernel's system calls	Kernel's system calls
3	Library functions	Library functions
4	Administrative file formats	Special files (in /dev)
5	Miscellaneous	Administrative file formats
6	Games	Games
7	Special files (in /dev)	Macro packages and conventions
8	Administration commands	Administration commands

8. Further Help with man -k, apropos and whatis

man -k: Searches a summary database and prints one-line description of the command.

Example:

```
$ man -k awk
awk awk(1)      -pattern scanning and processing language
nawk nawk(1)    -pattern scanning and processing language
```

apropos: lists the commands and files associated with a keyword.

Example:

```
$ apropos FTP
ftp ftp(1)      -file transfer program
ftpd in.ftpd(1m) -file transfer protocol server
```

```
ftpusers ftpusers(4) -file listing users to be disallowed
ftp login privileges
```

whatis: lists one-liners for a command.

Example:

```
$ whatis cp
cp cp(1) -copy files
```

9. When Things Go Wrong

Terminals and keyboards have no uniform behavioral pattern. Terminal settings directly impact the keyboard operation. If you observe a different behavior from that expected, when you press certain keystrokes, it means that the terminal settings are different. In such cases, you should know which keys to press to get the required behavior. The following table lists keyboard commands to try when things go wrong.

Keystroke or command	Function
<i>[Ctrl-h]</i>	Erases text
<i>[Ctrl-c]</i> or <i>Delete</i>	Interrupts a command
<i>[Ctrl-d]</i>	Terminates login session or a program that expects its input from keyboard
<i>[Ctrl-s]</i>	Stops scrolling of screen output and locks keyboard
<i>[Ctrl-q]</i>	Resumes scrolling of screen output and unlocks keyboard
<i>[Ctrl-u]</i>	Kills command line without executing it
<i>[Ctrl-^]</i>	Kills running program but creates a core file containing the memory image of the program
<i>[Ctrl-z]</i>	Suspends process and returns shell prompt; use fg to resume job
<i>[Ctrl-j]</i>	Alternative to <i>[Enter]</i>
<i>[Ctrl-m]</i>	Alternative to <i>[Enter]</i>
<i>stty sane</i>	Restores terminal to normal status

Conclusion

In this chapter, we looked at the architecture of UNIX and the division of labor between two agencies viz., the shell and the kernel. We also looked at the structure and usage of UNIX commands. The man documentation will be the most valuable source of documentation for UNIX commands. Also, when the keyboard sequences won't sometimes work as expected because of different terminal settings. We listed the possible remedial keyboard sequences when that happens.

CHAP 2 # The File System

Introduction

In this chapter we will look at the file system of UNIX. We also look at types of files their significance. We then look at two ways of specifying a file viz., with absolute pathnames and relative pathnames. A discussion on commands used with directory files viz., cd, pwd, mkdir, rmdir and ls will be made. Finally we look at some of the important directories contained under UNIX file system.

Objectives

- Types of files
- UNIX Filenames
- Directories and Files
- Absolute and Relative Pathnames
- pwd – print working directory
- cd – change directory
- mkdir – make a directory
- rmdir – remove directory
- The PATH environmental variable
- ls – list directory contents
- The UNIX File System

1. Types of files

A simple description of the UNIX system is this:

“On a UNIX system, everything is a file; if something is not a file, it is a process.”

A UNIX system makes no difference between a file and a directory, since a directory is just a file containing names of other files. Programs, services, texts, images, and so forth, are all files. Input and output devices, and generally all devices, are considered to be files, according to the system.

Most files are just files, called *regular* files; they contain normal data, for example text files, executable files or programs, input for or output from a program and so on.

While it is reasonably safe to suppose that everything you encounter on a UNIX system is a file, there are some exceptions.

Directories: files that are lists of other files.

Special files or Device Files: All devices and peripherals are represented by files. To read or write a device, you have to perform these operations on its associated file. Most special files are in /dev.

Links: a system to make a file or directory visible in multiple parts of the system's file tree.

(Domain) sockets: a special file type, similar to TCP/IP sockets, providing inter-process networking protected by the file system's access control.

Named pipes: act more or less like sockets and form a way for processes to communicate with each other, without using network socket semantics.

Ordinary (Regular) File

This is the most common file type. An ordinary file can be either a text file or a binary file. A text file contains only printable characters and you can view and edit them. All C and Java program sources, shell scripts are text files. Every line of a text file is terminated with the *newline* character.

A binary file, on the other hand, contains both printable and nonprintable characters that cover the entire ASCII range. The object code and executables that you produce by compiling C programs are binary files. Sound and video files are also binary files.

Directory File

A directory contains no data, but keeps details of the files and subdirectories that it contains. A directory file contains one entry for every file and subdirectory that it houses. Each entry has two components namely, the filename and a unique identification number of the file or directory (called the *inode number*).

When you create or remove a file, the kernel automatically updates its corresponding directory by adding or removing the entry (filename and inode number) associated with the file.

Device File

All the operations on the devices are performed by reading or writing the file representing the device. It is advantageous to treat devices as files as some of the commands used to access an ordinary file can be used with device files as well.

Device filenames are found in a single directory structure, /dev. A device file is not really a stream of characters. It is the attributes of the file that entirely govern the operation of the device. The kernel identifies a device from its attributes and uses them to operate the device.

2. Filenames in UNIX

On a UNIX system, a filename can consist of up to 255 characters. Files may or may not have extensions and can consist of practically any ASCII character except the / and the Null character. You are permitted to use control characters or other nonprintable characters in a filename. However, you should avoid using these characters while naming a file. It is recommended that only the following characters be used in filenames:

Alphabets and numerals.

The period (.), hyphen (-) and underscore (_).

UNIX imposes no restrictions on the extension. In all cases, it is the application that imposes that restriction. Eg. A C Compiler expects C program filenames to end with .c, Oracle requires SQL scripts to have .sql extension.

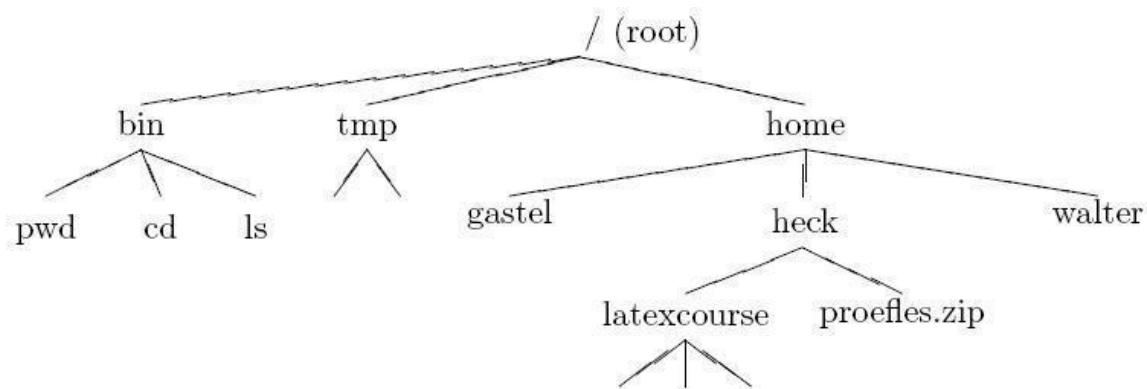
A file can have as many dots embedded in its name. A filename can also begin with or end with a dot.

UNIX is case sensitive; cap01, Chap01 and CHAP01 are three different filenames that can coexist in the same directory.

3. Directories and Files

A file is a set of data that has a name. The information can be an ordinary text, a user-written computer program, results of a computation, a picture, and so on. The file name may consist of ordinary characters, digits and special tokens like the underscore, except the forward slash (/). It is permitted to use special tokens like the ampersand (&) or spaces in a filename.

Unix organizes files in a tree-like hierarchical structure, with the *root directory*, indicated by a forward slash (/), at the top of the tree. See the Figure below, in which part of the hierarchy of files and directories on the computer is shown.



4. Absolute and relative paths

A path, which is the way you need to follow in the tree structure to reach a given file, can be described as starting from the trunk of the tree (the / or root directory). In that case, the path starts with a slash and is called an absolute path, since there can be no mistake: only one file on the system can comply.

Paths that don't start with a slash are always relative to the current directory. In relative paths we also use the . and .. indications for the current and the parent directory.

The HOME variable

When you log onto the system, UNIX automatically places you in a directory called the *home directory*. The shell variable HOME indicates the home directory of the user.

E.g.,
\$ echo \$HOME
/home/kumar

What you see above is an absolute pathname, which is a sequence of directory names starting from root (/). The subsequent slashes are used to separate the directories.

5. pwd - print working directory

At any time you can determine where you are in the file system hierarchy with the *pwd*, print working directory, command,

E.g.,:
\$ pwd

/home/frank/src

6. cd - change directory

You can change to a new directory with the **cd**, change directory, command. **cd** will accept both absolute and relative path names.

Syntax

cd [directory]

Examples

cd changes to user's home directory

cd / changes directory to the system's root

cd .. goes up one directory level

cd ../.. goes up two directory levels

cd /full/path/name/from/root changes directory to absolute path named
(note the leading slash)

cd path/from/current/location changes directory to path relative to current
location (no leading slash)

7. mkdir - make a directory

You extend your home hierarchy by making sub-directories underneath it. This is done with the **mkdir**, make directory, command. Again, you specify either the full or relative path of the directory.

Examples

mkdir patch Creates a directory *patch* under current directory

mkdir patch dbs doc Creates three directories under current directory

mkdir pis pis/progs pis/data Creates a directory tree with *pis* as a directory under the current directory and *progs* and *data* as subdirectories under *pis*

Note the order of specifying arguments in example 3. The parent directory should be specified first, followed by the subdirectories to be created under it.

The system may refuse to create a directory due to the following reasons:

1. The directory already exists.
2. There may be an ordinary file by the same name in the current directory.
3. The permissions set for the current directory don't permit the creation of files and directories by the user.

8. rmdir - remove directory

A directory needs to be empty before you can remove it. If it's not, you need to remove the files first. Also, you can't remove a directory if it is your present working directory; you must first change out of that directory. You cannot remove a subdirectory unless you are placed in a directory which is hierarchically *above* the one you have chosen to remove.

E.g.

`rmdir patch` Directory must be empty
`rmdir pis pis/progs pis/data` Shows error as *pis* is not empty. However `rmdir` silently deletes the lower level subdirectories *progs* and *data*.

9. The PATH environment variable

Environmental variables are used to provide information to the programs you use. We have already seen one such variable called HOME.

A command runs in UNIX by executing a disk file. When you specify a command like *date*, the system will locate the associated file from a list of directories specified in the PATH variable and then executes it. The PATH variable normally includes the current directory also.

Whenever you enter any UNIX command, you are actually specifying the name of an executable file located somewhere on the system. The system goes through the following steps in order to determine which program to execute:

1. Built in commands (such as `cd` and `history`) are executed within the shell.
2. If an absolute path name (such as `/bin/ls`) or a relative path name (such as `./myprog`), the system executes the program from the specified directory.
3. Otherwise the PATH variable is used.

10. ls - list directory contents

The command to list your directories and files is *ls*. With options it can provide information about the size, type of file, permissions, dates of file creation, change and access.

Syntax

`ls [options] [argument]`

Common Options

When no argument is used, the listing will be of the current directory. There are many very useful options for the `ls` command. A listing of many of them follows. When using the command, string the desired options together preceded by "-".

- a** Lists all files, including those beginning with a dot (`.`).
- d** Lists only names of directories, not the files in the directory
- F** Indicates type of entry with a trailing symbol: executables with `*`, directories with `/` and symbolic links with `@`
- R** Recursive list
- u** Sorts filenames by last access time
- t** Sorts filenames by last modification time
- i** Displays inode number
- l** Long listing: lists the mode, link information, owner, size, last modification (time). If the file is a symbolic link, an arrow (`-->`) precedes the pathname of the linked-to file.

The **mode field** is given by the **-l** option and consists of 10 characters. The first character is one of the following:

CHARACTER	IF ENTRY IS A
d	directory

- plain file
b block-type special file
c character-type special file
l symbolic link
s socket

The next 9 characters are in 3 sets of 3 characters each. They indicate the **file access permissions**: the first 3 characters refer to the permissions for the **user**, the next three for the users in the Unix **group** assigned to the file, and the last 3 to the permissions for **other** users on the system.

Designations are as follows:

r read permission
w write permission
x execute permission
 - no permission

Examples

1. To list the files in a directory:

```
$ ls
```

2. To list all files in a directory, including the hidden (dot) files:

```
$ ls -a
```

3. To get a long listing:

```
$ ls -al
total 24
drwxr-sr-x 5 workshop acs 512 Jun 7 11:12 .
drwxr-xr-x 6 root sys 512 May 29 09:59 ..
-rwxr-xr-x 1 workshop acs 532 May 20 15:31 .cshrc
-rw - - - - - 1 workshop acs 525 May 20 21:29 .emacs
-rw - - - - - 1 workshop acs 622 May 24 12:13 .history
-rwxr-xr-x 1 workshop acs 238 May 14 09:44 .login
-rw-r--r-- 1 workshop acs 273 May 22 23:53 .plan
-rwxr-xr-x 1 workshop acs 413 May 14 09:36 .profile
-rw - - - - - 1 workshop acs 49 May 20 20:23 .rhosts
drwx - - - - - 3 workshop acs 512 May 24 11:18 demofiles
drwx - - - - - 2 workshop acs 512 May 21 10:48 frank
drwx - - - - - 3 workshop acs 512 May 24 10:59 linda
```

11. The UNIX File System

The root directory has many subdirectories. The following table describes some of the subdirectories contained under root.

Directory	Content
/bin	Common programs, shared by the system, the system administrator and the users.
/dev	Contains references to all the CPU peripheral hardware, which are represented as files with special properties.
/etc	Most important system configuration files are in /etc, this directory contains data similar to those in the Control Panel in Windows
/home	Home directories of the common users.
/lib	Library files, includes files for all kinds of programs needed by the system and the users.

/sbin	Programs for use by the system and the system administrator.
/tmp	Temporary space for use by the system, cleaned upon reboot, so don't use this for saving any work!
/usr	Programs, libraries, documentation etc. for all user-related programs.
/var	Storage for all variable files and temporary files created by users, such as log files, the mail queue, the print spooler area, space for temporary storage of files downloaded from the Internet, or to keep an image of a CD before burning it.

Conclusion

In this chapter we looked at the UNIX file system and different types of files UNIX understands. We also discussed different commands that are specific to directory files viz., pwd, mkdir, cd, rmdir and ls. These commands have no relevance to ordinary or device files. We also saw filename conventions in UNIX. Difference between the absolute and relative pathnames was highlighted next. Finally we described some of the important subdirectories contained under root (/).

CHAP # 3 Basic File Attributes

The UNIX file system allows the user to access other files not belonging to them and without infringing on security. A file has a number of attributes (properties) that are stored in the inode. In this chapter, we discuss,

- `ls -l` to display file attributes (properties)
- Listing of a specific directory
- Ownership and group ownership
- Different file permissions

Listing File Attributes

`ls` command is used to obtain a list of all filenames in the current directory. The output in UNIX lingo is often referred to as the listing. Sometimes we combine this option with other options for displaying other attributes, or ordering the list in a different sequence. `ls` look up the file's inode to fetch its attributes. It lists seven attributes of all files in the current directory and they are:

- File type and Permissions
- Links
- Ownership
- Group ownership
- File size
- Last Modification date and time
- File name

The file type and its permissions are associated with each file. Links indicate the number of file names maintained by the system. This does not mean that there are so many copies of the file. File is created by the owner. Every user is attached to a group owner. File size in bytes is displayed. Last modification time is the next field. If you change only the permissions or ownership of the file, the modification time remains unchanged. In the last field, it displays the file name.

For example,

```
$ ls -l
total 72
-rw-r--r--    1 kumar  metal 19514 may 10 13:45  chap01
-rw-r--r--    1 kumar  metal  4174 may 10 15:01  chap02
-rw-rw-rw-    1 kumar  metal   84  feb 12 12:30  dept.lst
-rw-r--r--    1 kumar  metal  9156 mar 12 1999  genie.sh
drwxr-xr-x    2 kumar  metal   512 may 9 10:31  helpdir
drwxr-xr-x    2 kumar  metal   512 may 9 09:57  progs
```

Listing Directory Attributes

ls -d will not list all subdirectories in the current directory
For example,

```
ls -ld helpdir progs
drwxr-xr-x 2 kumar metal    512 may 9 10:31 helpdir
drwxr-xr-x 2 kumar metal    512 may 9 09:57 progs
```

Directories are easily identified in the listing by the first character of the first column, which here shows a d. The significance of the attributes of a directory differs a good deal from an ordinary file. To see the attributes of a directory rather than the files contained in it, use ls -ld with the directory name. Note that simply using ls -d will not list all subdirectories in the current directory. Strange though it may seem, ls has no option to list only directories.

File Ownership

When you create a file, you become its owner. Every owner is attached to a group owner. Several users may belong to a single group, but the privileges of the group are set by the owner of the file and not by the group members. When the system administrator creates a user account, he has to assign these parameters to the user:

The user-id (UID) – both its name and numeric representation

The group-id (GID) – both its name and numeric representation

File Permissions

UNIX follows a three-tiered file protection system that determines a file's access rights. It is displayed in the following format:

Filetype owner (rwx) groupowner (rwx) others (rwx)

For Example:

```
-rwxr-xr-- 1 kumar metal 20500 may 10 19:21 chap02
```

r w x

r - x

r - -

owner/user

group owner

others

The first group has all three permissions. The file is readable, writable and executable by the owner of the file. The second group has a hyphen in the middle slot, which indicates the absence of write permission by the group owner of the file. The third group has the write and executes bits absent. This set of permissions is applicable to others.

You can set different permissions for the three categories of users – owner, group and others. It's important that you understand them because a little learning here can be a dangerous thing. Faulty file permission is a sure recipe for disaster

Changing File Permissions

A file or a directory is created with a default set of permissions, which can be determined by umask. Let us assume that the file permission for the created file is -rw-r--r--. Using **chmod** command, we can change the file permissions and allow the owner to execute his file. The command can be used in two ways:

- In a relative manner by specifying the changes to the current permissions
- In an absolute manner by specifying the final permissions

Relative Permissions

chmod only changes the permissions specified in the command line and leaves the other permissions unchanged. Its syntax is:

```
chmod category operation permission filename(s)
```

chmod takes an expression as its argument which contains:

- user category (user, group, others)
- operation to be performed (assign or remove a permission)
- type of permission (read, write, execute)

Category	operation	permission
u - user	+ assign	r - read
g - group	- remove	w - write
o - others	= absolute	x - execute
a - all (ugo)		

Let us discuss some examples:

Initially,

```
-rw-r--r-- 1 kumar metal 1906 sep 23:38 xstart
```

```
chmod u+x xstart
```

```
-rwxr--r-- 1 kumar metal 1906 sep 23:38 xstart
```

The command assigns (+) execute (x) permission to the user (u), other permissions remain unchanged.

```
chmod ugo+x xstart or
chmod a+x xstart or
chmod +x xstart
```

```
-rwxr-xr-x 1 kumar metal 1906 sep 23:38 xstart
```

chmod accepts multiple file names in command line

```
chmod u+x note note1 note3
```

Let initially,

```
-rwxr-xr-x 1 kumar metal 1906 sep 23:38 xstart
```

```
chmod go-r xstart
```

Then, it becomes

```
-rwx--x--x 1 kumar metal 1906 sep 23:38 xstart
```

Absolute Permissions

Here, we need not to know the current file permissions. We can set all nine permissions explicitly. A string of three octal digits is used as an expression. The permission can be represented by one octal digit for each category. For each category, we add octal digits. If we represent the permissions of each category by one octal digit, this is how the permission can be represented:

- Read permission – 4 (octal 100)
- Write permission – 2 (octal 010)
- Execute permission – 1 (octal 001)

Octal	Permissions	Significance
0	- - -	no permissions
1	- - x	execute only
2	- w -	write only
3	- w x	write and execute
4	r - -	read only
5	r - x	read and execute
6	r w -	read and write
7	r w x	read, write and execute

We have three categories and three permissions for each category, so three octal digits can describe a file's permissions completely. The most significant digit represents user and the least one represents others. chmod can use this three-digit string as the expression.

Using relative permission, we have,

```
chmod a+rw xstart
```

Using absolute permission, we have,

```
chmod 666 xstart
```

```
chmod 644 xstart
```

```
chmod 761 xstart
```

will assign all permissions to the owner, read and write permissions for the group and only execute permission to the others.

777 signify all permissions for all categories, but still we can prevent a file from being deleted. 000 signifies absence of all permissions for all categories, but still we can delete a file. It is the directory permissions that determine whether a file can be deleted or not. Only owner can change the file permissions. User can not change other user's file's permissions. But the system administrator can do anything.

The Security Implications

Let the default permission for the file xstart is

```
-rw-r--r--
```

```
chmod u-rw, go-r xstart
```

or

```
chmod 000 xstart
```

```
-----
```

This is simply useless but still the user can delete this file
On the other hand,

```
chmod a+rwx xstart
```

```
chmod 777 xstart
```

```
-rwxrwxrwx
```

The UNIX system by default, never allows this situation as you can never have a secure system. Hence, directory permissions also play a very vital role here

We can use chmod Recursively.

```
chmod -R a+x shell_scripts
```

This makes all the files and subdirectories found in the shell_scripts directory, executable by all users. When you know the shell meta characters well, you will appreciate that the * doesn't match filenames beginning with a dot. The dot is generally a safer but note that both commands change the permissions of directories also.

Directory Permissions

It is possible that a file cannot be accessed even though it has read permission, and can be removed even when it is write protected. The default permissions of a directory are,

```
rwxr-xr-x (755)
```

A directory must never be writable by group and others

Example:

```
mkdir c_progs
```

```
ls -ld c_progs
```

```
drwxr-xr-x  2 kumar metal 512 may 9 09:57 c_progs
```

If a directory has write permission for group and others also, be assured that every user can remove every file in the directory. As a rule, you must not make directories universally writable unless you have definite reasons to do so.

Changing File Ownership

Usually, on BSD and AT&T systems, there are two commands meant to change the ownership of a file or directory. Let kumar be the owner and metal be the group owner. If sharma copies a file of kumar, then sharma will become its owner and he can manipulate the attributes

chown changing file owner and **chgrp** changing group owner

On BSD, only system administrator can use chown

On other systems, only the owner can change both

chown

Changing ownership requires superuser permission, so use **su** command

```
ls -l note
```

```
-rwxr----x  1 kumar metal 347 may 10 20:30 note
```

chown sharma note; ls -l note

```
-rwxr---x    1 sharma metal 347 may 10 20:30 note
```

Once ownership of the file has been given away to sharma, the user file permissions that previously applied to Kumar now apply to sharma. Thus, Kumar can no longer edit *note* since there is no write privilege for group and others. He can not get back the ownership either. But he can copy the file to his own directory, in which case he becomes the owner of the copy.

chgrp

This command changes the file's group owner. No superuser permission is required.

```
ls -l dept.lst
```

```
-rw-r--r--    1 kumar metal 139 jun 8 16:43 dept.lst
```

```
chgrp dba dept.lst; ls -l dept.lst
```

```
-rw-r--r--    1 kumar dba 139 jun 8 16:43 dept.lst
```

MORE FILE ATTRIBUTES

Apart from permissions and ownership, a UNIX file has several other attributes, and in this chapter, we look at most of the remaining ones. A file also has properties related to its time stamps and links. It is important to know how these attributes are interpreted when applied to a directory or a device.

This chapter also introduces the concepts of file system. It also looks at the inode, the lookup table that contained almost all file attributes. Though a detailed treatment of the file systems is taken up later, knowledge of its basics is essential to our understanding of the significance of some of the file attributes. Basic file attributes has helped us to know about - ls -l to display file attributes (properties), listing of a specific directory, ownership and group ownership and different file permissions. ls -l provides attributes like - permissions, links, owner, group owner, size, date and the file name.

File Systems and inodes

The hard disk is split into distinct partitions, with a separate file system in each partition. Every file system has a directory structure headed by root.

n partitions = n file systems = n separate root directories

(index node), accessed by the inode number. The inode contains the following attributes of a file:

- File type
- File permissions
- Number of links
- The UID of the owner
- The GID of the group owner
- File size in bytes
- Date and time of last modification
- Date and time of last access
- Date and time of last change of the inode
- An array of pointers that keep track of all disk blocks used by the file

Please note that, neither the name of the file nor the inode number is stored in the inode. To know inode number of a file:

```
ls -il tulec05
```

```
9059 -rw-r--r-- 1 kumar metal 51813 Jan 31 11:15 tulec05
```

Where, 9059 is the inode number and no other file can have the same inode number in the same file system.

Hard Links

The link count is displayed in the second column of the listing. This count is normally 1, but the following files have two links,

```
-rwxr-xr-- 2 kumar metal 163 Jul 13 21:36 backup.sh  
-rwxr-xr-- 2 kumar metal 163 Jul 13 21:36 restore.sh
```

All attributes seem to be identical, but the files could still be copies. It's the link count that seems to suggest that the files are linked to each other. But this can only be confirmed by using the `-li` option to `ls`.

```
ls -li backup.sh restore.sh
```

```
478274 -rwxr-xr-- 2 kumar metal163 jul 13 21:36 backup.sh  
478274 -rwxr-xr-- 2 kumar metal163 jul 13 21:36 restore.sh
```

ln: Creating Hard Links

A file is linked with the `ln` command which takes two filenames as arguments (`cp` command). The command can create both a hard link and a soft link and has syntax similar to the one used by `cp`. The following command links `emp.lst` with `employee`:

```
ln emp.lst employee
```

The `-i` option to `ls` shows that they have the same inode number, meaning that they are actually one and the same file:

```
ls -li emp.lst employee
```

```
29518 -rwxr-xr-x 2 kumar metal 915 may 4 09:58 emp.lst
29518 -rwxr-xr-x 2 kumar metal 915 may 4 09:58 employee
```

The link count, which is normally one for unlinked files, is shown to be two. You can increase the number of links by adding the third file name `emp.dat` as:

```
ln employee emp.dat ; ls -l emp*
```

```
29518 -rwxr-xr-x 3 kumar metal 915 may 4 09:58 emp.dat
29518 -rwxr-xr-x 3 kumar metal 915 may 4 09:58 emp.lst
29518 -rwxr-xr-x 3 kumar metal 915 may 4 09:58 employee
```

You can link multiple files, but then the destination filename must be a directory. A file is considered to be completely removed from the file system when its link count drops to zero. `ln` returns an error when the destination file exists. Use the `-f` option to force the removal of the existing link before creation of the new one

Where to use Hard Links

```
ln data/ foo.txt input_files
```

It creates link in directory *input_files*. With this link available, your existing programs will continue to find `foo.txt` in the *input_files* directory. It is more convenient to do this that modifies all programs to point to the new path. Links provide some protection against accidental deletion, especially when they exist in different directories. Because of links, we don't need to maintain two programs as two separate disk files if there is very little difference between them. A file's name is available to a C program and to a shell script. A single file with two links can have its program logic make it behave in two different ways depending on the name by which it is called.

We can't have two linked filenames in two file systems and we can't link a directory even within the same file system. This can be solved by using symbolic links (soft links).

Symbolic Links

Unlike the hard linked, a symbolic link doesn't have the file's contents, but simply provides the pathname of the file that actually has the contents.

```
ln -s note note.sym
```

```
ls -li note note.sym
```

```
9948 -rw-r--r-- 1 kumar group 80 feb 16 14:52 note
9952 lrwxrwxrwx 1 kumar group 4 feb16 15:07note.sym ->note
```

Where, l indicate symbolic link file category. -> indicates note.sym contains the pathname for the filename note. Size of symbolic link is only 4 bytes; it is the length of the pathname of note.

It's important that this time we indeed have two files, and they are not identical. Removing note.sym won't affect us much because we can easily recreate the link. But if we remove note, we would lose the file containing the data. In that case, note.sym would point to a nonexistent file and become a dangling symbolic link.

Symbolic links can also be used with relative pathnames. Unlike hard links, they can also span multiple file systems and also link directories. If you have to link all filenames in a directory to another directory, it makes sense to simply link the directories. Like other files, a symbolic link has a separate directory entry with its own inode number. This means that rm can remove a symbolic link even if its points to a directory.

A symbolic link has an inode number separate from the file that it points to. In most cases, the pathname is stored in the symbolic link and occupies space on disk.

However, Linux uses a fast symbolic link which stores the pathname in the inode itself provided it doesn't exceed 60 characters.

The Directory

A directory has its own permissions, owners and links. The significance of the file attributes change a great deal when applied to a directory. For example, the size of a directory is in no way related to the size of files that exists in the directory, but rather to the number of files housed by it. The higher the number of files, the larger the directory size. Permission acquires a different meaning when the term is applied to a directory.

```
ls -l -d progs
```

```
drwxr-xr-x 2 kumar metal 320 may 9 09:57 progs
```

The default permissions are different from those of ordinary files. The user has all permissions, and group and others have read and execute permissions only. The permissions of a directory also impact the security of its files. To understand how that can happen, we must know what permissions for a directory really mean.

Read permission

Read permission for a directory means that the list of filenames stored in that directory is accessible. Since ls reads the directory to display filenames, if a directory's read permission is removed, ls wont work. Consider removing the read permission first from the directory *progs*,

```
ls -ld progs
```

```
drwxr-xr-x 2 kumar metal 128 jun 18 22:41 progs
```

```
chmod -r progs ; ls progs
```

```
progs: permission denied
```

Write permission

We can't write to a directory file. Only the kernel can do that. If that were possible, any user could destroy the integrity of the file system. Write permission for a directory implies that you are permitted to create or remove files in it. To try that out, restore the read permission and remove the write permission from the directory before you try to copy a file to it.

```
chmod 555 progs ; ls -ld progs
```

```
dr-xr-xr-x 2 kumar metal 128 jun 18 22:41 progs
```

```
cp emp.lst progs
```

```
cp: cannot create progs/emp.lst: permission denied
```

- The write permission for a directory determines whether we can create or remove files in it because these actions modify the directory
- Whether we can modify a file depends on whether the file itself has write permission. Changing a file doesn't modify its directory entry

Execute permission

If a single directory in the pathname doesn't have execute permission, then it can't be searched for the name of the next directory. That's why the execute privilege of a directory is often referred to as the search permission. A directory has to be searched for the next directory, so the cd command won't work if the search permission for the directory is turned off.

```
chmod 666 progs ; ls -ld progs
```

```
drw-rw-rw- 2 kumar metal 128 jun 18 22:41 progs
```

```
cd progs
```

```
permission denied to search and execute it
```

umask: DEFAULT FILE AND DIRECTORY PERMISSIONS

When we create files and directories, the permissions assigned to them depend on the system's default setting. The UNIX system has the following default permissions for

all files and directories.

rw-rw-rw- (octal 666) for regular files

rwxrwxrwx (octal 777) for directories

The default is transformed by subtracting the user mask from it to remove one or more permissions. We can evaluate the current value of the mask by using *umask* without arguments,

```
$ umask
022
```

This becomes 644 (666-022) for ordinary files and 755 (777-022) for directories *umask 00*. This indicates, we are not subtracting anything and the default permissions will remain unchanged. Note that, changing system wide default permission settings is possible using *chmod* but not by *umask*

MODIFICATION AND ACCESS TIMES

- A UNIX file has three time stamps associated with it. Among them, two are:
- Time of last file modification ls -l
 - Time of last access ls -lu
 - Time of lastinode modification ls -lc

The access time is displayed when *ls -l* is combined with the *-u* option. Knowledge of file's modification and access times is extremely important for the system administrator. Many of the tools used by them look at these time stamps to decide whether a particular file will participate in a backup or not.

TOUCH COMMAND – changing the time stamps

To set the modification and access times to predefined values, we have,

```
touch options expression filename(s)
```

```
touch emp.lst (without options and expression)
```

Then, both times are set to the current time and creates the file, if it doesn't exist.

touch command (without options but with expression) can be used. The expression consists of MMDDhhmm (month, day, hour and minute).

```
touch 03161430 emp.lst ; ls -l emp.lst
```

```
-rw-r--r-- 1 kumar metal 870 mar 16 14:30 emp.lst
```

```
ls -lu emp.lst
```

```
-rw-r--r-- 1 kumar metal 870 mar 16 14:30 emp.lst
```

It is possible to change the two times individually. The `-m` and `-a` options change the modification and access times, respectively:

touch command (with options and expression)

`-m` for changing modification time

`-a` for changing access time

```
touch -m 02281030 emp.lst ; ls -l emp.lst
```

```
-rw-r--r-- 1 kumar metal 870 feb 28 10:30 emp.lst
```

```
touch -a 01261650 emp.lst ; ls -lu emp.lst
```

```
-rw-r--r-- 1 kumar metal 870 jan 26 16:50 emp.lst
```

find : locating files

It recursively examines a directory tree to look for files matching some criteria, and then takes some action on the selected files. It has a difficult command line, and if you have ever wondered why UNIX is hated by many, then you should look up the cryptic *find* documentation. However, *find* is easily tamed if you break up its arguments into three components:

```
find path_list selecton_criteria action
```

where,

- Recursively examines all files specified in `path_list`
- It then matches each file for one or more selection-criteria
- It takes some action on those selected files

The `path_list` comprises one or more subdirectories separated by white space. There can also be a host of `selection_criteria` that you use to match a file, and multiple actions to dispose of the file. This makes the command difficult to use initially, but it is a program that every user must master since it lets him make file selection under practically any condition.

CHAP 4# The VI Editor

To write and edit some programs and scripts, we require editors. UNIX provides vi editor for BSD system – created by Bill Joy. Bram Moolenaar improved vi editor and called it as vim (vi improved) on Linux OS.

vi Basics

To add some text to a file, we invoke,

```
vi <filename>
```

In all probability, the file doesn't exist, and vi presents you a full screen with the filename shown at the bottom with the qualifier. The cursor is positioned at the top and all remaining lines of the screen show a ~. They are non-existent lines. The last line is reserved for commands that you can enter to act on text. This line is also used by the system to display messages. This is the command mode. This is the mode where you can pass commands to act on text, using most of the keys of the keyboard. This is the default mode of the editor where every key pressed is interpreted as a command to run on text. You will have to be in this mode to copy and delete text

For, text editing, vi uses 24 out of 25 lines that are normally available in the terminal. To enter text, you must switch to the input mode. First press the key i, and you are in this mode ready to input text. Subsequent key depressions will then show up on the screen as text input.

After text entry is complete, the cursor is positioned on the last character of the last line. This is known as current line and the character where the cursor is stationed is the current cursor position. This mode is used to handle files and perform substitution. After the command is run, you are back to the default command mode. If a word has been misspelled, use ctrl-w to erase the entire word.

Now press esc key to revert to command mode. Press it again and you will hear a beep. A beep in vi indicates that a key has been pressed unnecessarily. Actually, the text entered has not been saved on disk but exists in some temporary storage called a buffer. To save the entered text, you must switch to the execute mode (the last line mode). Invoke the execute mode from the command mode by entering a: which shows up in the last line.

The Repeat Factor

vi provides repeat factor in command and input mode commands. Command mode command k moves the cursor one line up. 10k moves cursor 10 lines up. To undo whenever you make a mistake, press

Esc u

To clear the screen in command mode, press

ctrl-l

Don't use (caps lock) - vi commands are case-sensitive

Avoid using the PC navigation keys

Input Mode – Entering and Replacing Text

It is possible to display the mode in which is user is in by typing,

:set showmode

Messages like INSERT MODE, REPLACE MODE, CHANGE MODE, etc will appear in the last line.

Pressing 'i' changes the mode from command to input mode. To append text to the right of the cursor position, we use *a*, *text*. I and A behave same as i and a, but at line extremes I inserts text at the beginning of line. A appends text at end of line. o opens a new line below the current line

- *r*<letter> replacing a single character
- *s*<text/word> replacing text with s
- *R*<text/word> replacing text with R
- Press esc key to switch to command mode after you have keyed in text

Some of the input mode commands are:

COMMAND	FUNCTION
i	inserts text
a	appends text
I	inserts at beginning of line
A	appends text at end of line
o	opens line below
O	opens line above
r	replaces a single character
s	replaces with a text
S	replaces entire line

Saving Text and Quitting – The ex Mode

When you edit a file using vi, the original file is not distributed as such, but only a copy of it that is placed in a buffer. From time to time, you should save your work by writing the buffer contents to disk to keep the disk file current. When we talk of saving a file, we actually mean saving this buffer. You may also need to quit vi after or without saving the buffer. Some of the save and exit commands of the ex mode is:

Command	Action
:W	saves file and remains in editing mode
:x	saves and quits editing mode
:wq	saves and quits editing mode
:w <filename>	save as
:w! <filename>	save as, but overwrites existing file
:q	quits editing mode
:q!	quits editing mode by rejecting changes made
:sh	escapes to UNIX shell
:recover	recovers file from a crash

Navigation

A command mode command doesn't show up on screen but simply performs a function. To move the cursor in four directions,

k	moves cursor up
j	moves cursor down
h	moves cursor left
l	moves cursor right

Word Navigation

Moving by one character is not always enough. You will often need to move faster along a line. vi understands a word as a navigation unit which can be defined in two ways, depending on the key pressed. If your cursor is a number of words away from your desired position, you can use the word-navigation commands to go there directly. There are three basic commands:

b	moves back to beginning of word
e	moves forward to end of word
w	moves forward to beginning word

Example,

5b takes the cursor 5 words back

3w takes the cursor 3 words forward

Moving to Line Extremes

Moving to the beginning or end of a line is a common requirement. To move to the first character of a line

0 or |

30| moves cursor to column 30

\$ moves to the end of the current line
The use of these commands along with b, e, and w is allowed

Scrolling

Faster movement can be achieved by scrolling text in the window using the control keys. The two commands for scrolling a page at a time are

ctrl-f	scrolls forward
ctrl-b	scrolls backward

10ctrl-fscroll 10 pages and navigate faster

ctrl-d	scrolls half page forward
ctrl-u	scrolls half page backward

The repeat factor can also be used here.

Absolute Movement

The editor displays the total number of lines in the last line

Ctrl-g	to know the current line number
40G	goes to line number 40
1G	goes to line number 1
G	goes to end of file

Editing Text

The editing facilities in vi are very elaborate and invoke the use of operators. They use operators, such as,

d	delete
y	yank (copy)

Deleting Text

x	deletes a single character
dd	delete entire line
yy	copy entire line
6dd	deletes the current line and five lines below

Moving Text

Moving text (p) puts the text at the new location.

p and P place text on right and left only when you delete parts of lines. But the same keys get associated with “below” and “above” when you delete complete lines

Copying Text

Copying text (y and p) is achieved as,

yy copies current line
10yy copies current line & 9 lines below

Joining Lines

J to join the current line and the line following it
4J joins following 3 lines with current line

Undoing Last Editing Instructions

In command mode, to undo the last change made, we use u
To discard all changes made to the current line, we use U

vim (LINUX) lets you undo and redo multiple editing instructions. u behaves differently here; repeated use of this key progressively undoes your previous actions. You could even have the original file in front of you. Further 10u reverses your last 10 editing actions. The function of U remains the same.

You may overshoot the desired mark when you keep u pressed, in which case use ctrl-r to redo your undone actions. Further, undoing with 10u can be completely reversed with 10ctrl-r. The undoing limit is set by the execute mode command: set undolevels=n, where n is set to 1000 by default.

Repeating the Last Command

The . (dot) command is used for repeating the last instruction in both editing and command mode commands

For example:

2dd deletes 2 lines from current line and to repeat this operation, type. **(dot)**

Searching for a Pattern

/ search forward
? search backward

/printf
The search begins forward to position the cursor on the first instance of the word

?pattern

Searches backward for the most previous instance of the pattern

Repeating the Last Pattern Search

n repeats search in same direction of original search

n doesn't necessarily repeat a search in the forward direction. The direction depends on the search command used. If you used ? printf to search in the reverse direction in the first place, then n also follows the same direction. In that case, N will repeat the search in the forward direction, and not n.

Search and repeat commands

Command

Function

/pat searches forward for pattern pat

?pat searches backward for pattern pat

n repeats search in same direction along which previous search was made

N repeats search in direction opposite to that along which previous search was made

Substitution – search and replace

We can perform search and replace in execute mode using :s. Its syntax is,

:address/source_pattern/target_pattern/flags

:1,\$s/director/member/g can also use % instead of 1,\$

:1,50s/unsigned//g deletes unsigned everywhere in lines 1 to 50

:3,10s/director/member/g substitute lines 3 through 10

:.s/director/member/g only the current line

:\$s/director/member/g only the last line

Interactive substitution: sometimes you may like to selectively replace a string. In that case, add the c parameter as the flag at the end:

:1,\$s/director/member/gc

Each line is selected in turn, followed by a sequence of carets in the next line, just below the pattern that requires substitution. The cursor is positioned at the end of this caret sequence, waiting for your response.

The ex mode is also used for substitution. Both search and replace operations also use regular expressions for matching multiple patterns.

The features of vi editor that have been highlighted so far are good enough for a beginner who should not proceed any further before mastering most of them. There are many more functions that make vi a very powerful editor. Can you copy three words or even the entire file using simple keystrokes? Can you copy or move multiple sections of text from one file to another in a single file switch? How do you compile your C and Java programs without leaving the editor? vi can do all this.

CHAP 5 # The Shell

Introduction

In this chapter we will look at one of the major component of UNIX architecture – The Shell. Shell acts as both a command interpreter as well as a programming facility. We will look at the interpretive nature of the shell in this chapter.

Objectives

- The Shell and its interpretive cycle
- Pattern Matching – The wild-cards
- Escaping and Quoting
- Redirection – The three standard files
- Filters – Using both standard input and standard output
- /dev/null and /dev/tty – The two special files
- Pipes
- tee – Creating a tee
- Command Substitution
- Shell Variables

1. The shell and its interpretive cycle

The shell sits between you and the operating system, acting as a command interpreter. It reads your terminal input and translates the commands into actions taken by the system. The shell is analogous to *command.com* in DOS. When you log into the system you are given a default shell. When the shell starts up it reads its startup files and may set environment variables, command search paths, and command aliases, and executes any commands specified in these files. The original shell was the Bourne shell, *sh*. Every Unix platform will either have the Bourne shell, or a Bourne compatible shell available.

Numerous other shells are available. Some of the more well known of these may be on your Unix system: the Korn shell, *ksh*, by David Korn, C shell, *csh*, by Bill Joy and the Bourne Again SHell, *bash*, from the Free Software Foundations GNU project, both based on *sh*, the T-C shell, *tcsh*, and the extended C shell, *cshe*, both based on *csh*.

Even though the shell appears not to be doing anything meaningful when there is no activity at the terminal, it swings into action the moment you key in something.

The following activities are typically performed by the shell in its interpretive cycle:

- The shell issues the prompt and waits for you to enter a command.
- After a command is entered, the shell scans the command line for metacharacters and expands abbreviations (like the * in *rm **) to recreate a simplified command line.
- It then passes on the command line to the kernel for execution.
- The shell waits for the command to complete and normally can't do any work while the command is running.

- After the command execution is complete, the prompt reappears and the shell returns to its waiting role to start the next cycle. You are free to enter another command.

2. Pattern Matching – The Wild-Cards

A pattern is framed using ordinary characters and a metacharacter (like *) using well-defined rules. The pattern can then be used as an argument to the command, and the shell will expand it suitably before the command is executed.

The metacharacters that are used to construct the generalized pattern for matching filenames belong to a category called wild-cards. The following table lists them:

Wild-Card	Matches
*	Any number of characters including none
?	A single character
[ijk]	A single character – either an i, j or k
[x-z]	A single character that is within the ASCII range of characters x and x
[!ijk]	A single character that is not an i,j or k (Not in C shell)
[!x-z]	A single character that is not within the ASCII range of the characters x and x (Not in C Shell)
{pat1,pat2...}	Pat1, pat2, etc. (Not in Bourne shell)

Examples:

To list all files that begin with *chap*, use

```
$ ls chap*
```

To list all files whose filenames are six character long and start with chap, use

```
$ ls chap??
```

Note: Both * and ? operate with some restrictions. for example, the * doesn't match all files beginning with a . (dot) or the / of a pathname. If you wish to list all hidden filenames in your directory having at least three characters after the dot, the dot must be matched explicitly.

```
$ ls .???*
```

However, if the filename contains a dot anywhere but at the beginning, it need not be matched explicitly.

Similarly, these characters don't match the / in a pathname. So, you cannot use

```
$ cd /usr?local      to change to /usr/local.
```

The character class

You can frame more restrictive patterns with the character class. The character class comprises a set of characters enclosed by the rectangular brackets, [and], but it matches a single character in the class. The pattern [abd] is character class, and it matches a single character – an a,b or d.

Examples:

```
$ls chap0[124]      Matches chap01, chap02, chap04 and lists if found.
```

```
$ ls chap[x-z]      Matches chapx, chapy, chapz and lists if found.
```

You can negate a character class to reverse a matching criteria. For example,

- To match all filenames with a single-character extension but not the .c or .o files, use `*.[!co]`
- To match all filenames that don't begin with an alphabetic character, use `[!a-zA-Z]*`

Matching totally dissimilar patterns

This feature is not available in the Bourne shell. To copy all the C and Java source programs from another directory, we can delimit the patterns with a comma and then put curly braces around them.

```
$ cp $HOME/prog_sources/*.{c,java} .
```

The Bourne shell requires two separate invocations of `cp` to do this job.

```
$ cp /home/srm/{project,html,scripts}/* .
```

The above command copies all files from three directories (project, html and scripts) to the current directory.

3. Escaping and Quoting

Escaping is providing a `\` (backslash) before the wild-card to remove (escape) its special meaning.

For instance, if we have a file whose filename is `chap*` (Remember a file in UNIX can be names with virtually any character except the `/` and null), to remove the file, it is dangerous to give command as `rm chap*`, as it will remove all files beginning with `chap`. Hence to suppress the special meaning of `*`, use the command `rm chap*`

To list the contents of the file `chap0[1-3]`, use

```
$ cat chap0\[1-3\]
```

A filename can contain a whitespace character also. Hence to remove a file named `My Document.doc`, which has a space embedded, a similar reasoning should be followed:

```
$ rm My\ Document.doc
```

Quoting is enclosing the wild-card, or even the entire pattern, within quotes. Anything within these quotes (barring a few exceptions) are left alone by the shell and not interpreted. When a command argument is enclosed in quotes, the meanings of all enclosed special characters are turned off.

Examples:

```
$ rm 'chap*'           Removes file chap*
$ rm "My Document.doc" Removes file My Document.doc
```

4. Redirection : The three standard files

The shell associates three files with the terminal – two for display and one for the keyboard. These files are streams of characters which many commands see as input and output. When a user logs in, the shell makes available three files representing three streams. Each stream is associated with a default device:

Standard input: The file (stream) representing input, connected to the keyboard.

Standard output: The file (stream) representing output, connected to the display.

Standard error: The file (stream) representing error messages that emanate from the command or shell, connected to the display.

The standard input can represent three input sources:

- The keyboard, the default source.
- A file using redirection with the < symbol.
- Another program using a pipeline.

The standard output can represent three possible destinations:

- The terminal, the default destination.
- A file using the redirection symbols > and >>.
- As input to another program using a pipeline.

A file is opened by referring to its pathname, but subsequent read and write operations identify the file by a unique number called a file descriptor. The kernel maintains a table of file descriptors for every process running in the system. The first three slots are generally allocated to the three standard streams as,

- 0 – Standard input
- 1 – Standard output
- 2 – Standard error

These descriptors are implicitly prefixed to the redirection symbols.

Examples:

Assuming file2 doesn't exist, the following command redirects the standard output to file *myOutput* and the standard error to file *myError*.

```
$ ls -l file1 file2 1>myOutput 2>myError
```

To redirect both standard output and standard error to a single file use:

```
$ ls -l file1 file2 1>| myOutput 2>| myError OR  
$ ls -l file1 file2 1> myOutput 2>& 1
```

5. Filters: Using both standard input and standard output

UNIX commands can be grouped into four categories viz.,

1. Directory-oriented commands like `mkdir`, `rmdir` and `cd`, and basic file handling commands like `cp`, `mv` and `rm` use neither standard input nor standard output.
2. Commands like `ls`, `pwd`, `who` etc. don't read standard input but they write to standard output.
3. Commands like `lp` that read standard input but don't write to standard output.
4. Commands like `cat`, `wc`, `cmp` etc. that use both standard input and standard output.

Commands in the fourth category are called filters. Note that filters can also read directly from files whose names are provided as arguments.

Example: To perform arithmetic calculations that are specified as expressions in input file `calc.txt` and redirect the output to a file `result.txt`, use

```
$ bc < calc.txt > result.txt
```

6. /dev/null and /dev/tty : Two special files

/dev/null: If you would like to execute a command but don't like to see its contents on the screen, you may wish to redirect the output to a file called /dev/null. It is a special file that can accept any stream without growing in size. It's size is always zero.

/dev/tty: This file indicates one's terminal. In a shell script, if you wish to redirect the output of some select statements explicitly to the terminal. In such cases you can redirect these explicitly to /dev/tty inside the script.

7. Pipes

With piping, the output of a command can be used as input (piped) to a subsequent command.

```
$ command1 | command2
```

Output from command1 is piped into input for command2.

This is equivalent to, but more efficient than:

```
$ command1 > temp
```

```
$ command2 < temp
```

```
$ rm temp
```

Examples

```
$ ls -al | more
```

```
$ who | sort | lpr
```

When a command needs to be ignorant of its source

If we wish to find total size of all C programs contained in the working directory, we can use the command,

```
$ wc -c *.c
```

However, it also shows the usage for each file(size of each file). We are not interested in individual statistics, but a single figure representing the total size. To be able to do that, we must make wc ignorant of its input source. We can do that by feeding the concatenated output stream of all the .c files to wc -c as its input:

```
$ cat *.c | wc -c
```

8. Creating a tee

tee is an external command that handles a character stream by duplicating its input. It saves one copy in a file and writes the other to standard output. It is also a filter and hence can be placed anywhere in a pipeline.

Example: The following command sequence uses tee to display the output of who and saves this output in a file as well.

```
$ who | tee users.lst
```

9. Command substitution

The shell enables the connecting of two commands in yet another way. While a pipe enables a command to obtain its standard input from the standard output of another command, the shell enables one or more command arguments to be obtained from the standard output of another command. This feature is called command substitution.

Example:

```
$ echo Current date and time is `date`
```

Observe the use of backquotes around date in the above command. Here the output of the command execution of date is taken as argument of echo. The shell executes the enclosed command and replaces the enclosed command line with the output of the command.

Similarly the following command displays the total number of files in the working directory.

```
$ echo "There are `ls | wc -l` files in the current directory"
```

Observe the use of double quotes around the argument of echo. If you use single quotes, the backquote is not interpreted by the shell if enclosed in single quotes.

10. Shell variables

Environmental variables are used to provide information to the programs you use. You can have both global environment and local shell variables. Global environment variables are set by your login shell and new programs and shells inherit the environment of their parent shell. Local shell variables are used only by that shell and are not passed on to other processes. A child process cannot pass a variable back to its parent process.

To declare a local shell variable we use the form *variable=value* (no spaces around =) and its evaluation requires the \$ as a prefix to the variable.

Example:

```
$ count=5
$ echo $count
5
```

A variable can be removed with **unset** and protected from reassignment by **readonly**. Both are shell internal commands.

Note: In C shell, we use **set** statement to set variables. Here, there either has to be whitespace on both sides of the = or none at all.

```
$ set count=5
$ set size = 10
```

Uses of local shell variables

1. Setting pathnames: If a pathname is used several times in a script, we can assign it to a variable and use it as an argument to any command.
2. Using command substitution: We can assign the result of execution of a command to a variable. The command to be executed must be enclosed in backquotes.
3. Concatenating variables and strings: Two variables can be concatenated to form a new variable.

Example: \$ base=foo ; ext=.c

```
$ file=$base$ext  
$ echo $file // prints foo.c
```

CHAP 6 # The Process

Introduction

A process is an OS abstraction that enables us to look at files and programs as their time image. This chapter discusses processes, the mechanism of creating a process, different states of a process and also the ps command with its different options. A discussion on creating and controlling background jobs will be made next. We also look at three commands viz., at, batch and cron for scheduling jobs. This chapter also looks at nice command for specifying job priority, signals and time command for getting execution time usage statistics of a command.

Objectives

- Process Basics
- ps: Process Status
- Mechanism of Process Creation
- Internal and External Commands
- Process States and Zombies
- Background Jobs
- nice: Assigning execution priority
- Processes and Signals
- job Control
- at and batch: Execute Later
- cron command: Running Jobs Periodically
- time: Timing Usage Statistics at process runtime

1. Process Basics

UNIX is a multiuser and multitasking operating system. *Multiuser* means that several people can use the computer system simultaneously (unlike a single-user operating system, such as MS-DOS). *Multitasking* means that UNIX, like Windows NT, can work on several tasks concurrently; it can begin work on one task and take up another before the first task is finished.

When you execute a program on your UNIX system, the system creates a special environment for that program. This environment contains everything needed for the system to run the program as if no other program were running on the system. Stated in other words, a process is created. A process is a program in execution. A process is said to be born when the program starts execution and remains alive as long as the program is active. After execution is complete, the process is said to die.

The kernel is responsible for the management of the processes. It determines the time and priorities that are allocated to processes so that more than one process can share the CPU resources.

Just as files have attributes, so have processes. These attributes are maintained by the kernel in a data structure known as process table. Two important attributes of a process are:

1. The Process-Id (PID): Each process is uniquely identified by a unique integer called the PID, that is allocated by the kernel when the process is born. The PID can be used to control a process.
2. The Parent PID (PPID): The PID of the parent is available as a process attribute.

There are three types of processes viz.,

1. Interactive: Initiated by a shell and running in the foreground or background
2. batch: Typically a series of processes scheduled for execution at a specified point in time
3. daemon: Typically initiated at boot time to perform operating system functions on demand, such as LPD, NFS, and DNS

The Shell Process

As soon as you log in, a process is set up by the kernel. This process represents the login shell, which can be either sh(Bourne Shell), ksh(korn Shell), bash(Bourne Again Shell) or csh(C Shell).

Parents and Children

When you enter an external command at the prompt, the shell acts as the parent process, which in turn starts the process representing the command entered. Since every parent has a parent, the ultimate ancestry of any process can be traced back to the first process (PID 0) that is set up when the system is booted. It is analogous to the root directory of the file system. A process can have only one parent. However, a process can spawn multiple child processes.

Wait or not Wait?

A parent process can have two approaches for its child:

- It may wait for the child to die so that it can spawn the next process. The death of the child is intimated to the parent by the kernel. Shell is an example of a parent that waits for the child to terminate. However, the shell can be told not to wait for the child to terminate.
- It may not wait for the child to terminate and may continue to spawn other processes. init process is an example of such a parent process.

2. ps: Process Status

Because processes are so important to getting things done, UNIX has several commands that enable you to examine processes and modify their state. The most frequently used command is ps, which prints out the process status for processes running on your system. Each system has a slightly different version of the ps command, but there are two main variants, the System V version (POSIX) and the Berkeley version. The following table shows the options available with ps command.

POSIX	BSD	Significance
-f	f	Full listing showing PPID of each process
-e or -A	aux	All processes (user and system) processes

-u <i>user</i>	U <i>user</i>	Processes of user <i>user</i> only
-a		Processes of all users excluding processes not associated with terminal
-l	l	Long listing showing memory related information
-t <i>term</i>	t <i>term</i>	Processes running on the terminal <i>term</i>

Examples

```
$ ps
PID  TTY    TIME    CMD
4245 pts/7  00:00:00 bash
5314 pts/7  00:00:00 ps
```

The output shows the header specifying the PID, the terminal (TTY), the cumulative processor time (TIME) that has been consumed since the process was started, and the process name (CMD).

```
$ ps -f
UID      PID    PPID  C  STIME     TTY    TIME  COMMAND
root    14931  136   0  08:37:48 ttys0  0:00  rlogind
sartin  14932  14931 0  08:37:50 ttys0  0:00  -sh
sartin  15339  14932 7  16:32:29 ttys0  0:00  ps -f
```

The header includes the following information:

- UID** – Login name of the user
- PID** – Process ID
- PPID** – Parent process ID
- C** – An index of recent processor utilization, used by kernel for scheduling
- STIME** – Starting time of the process in hours, minutes and seconds
- TTY** – Terminal ID number
- TIME** – Cumulative CPU time consumed by the process
- CMD** – The name of the command being executed

System processes (-e or -A)

Apart from the processes a user generates, a number of system processes keep running all the time. Most of them are not associated with any controlling terminal.

They are spawned during system startup and some of them start when the system goes into multiuser mode. These processes are known as daemons because they are called without a specific request from a user. To list them use,

```
$ ps -e
PID     TTY      TIME    CMD
0        ?        0:34    sched
1        ?        41:55   init
23274   Console  0:03    sh
272     ?        2:47    cron
7015    term/12  20:04   vi
```

3. Mechanism of Process Creation

There are three distinct phases in the creation of a process and uses three important system calls viz., *fork*, *exec*, and *wait*. The three phases are discussed below:

- **Fork:** A process in UNIX is created with the *fork* system call, which creates a copy of the process that invokes it. The process image is identical to that of the calling process, except for a few parameters like the PID. The child gets a new PID.
- **Exec:** The forked child overwrites its own image with the code and data of the new program. This mechanism is called *exec*, and the child process is said to *exec* a new program, using one of the family of *exec* system calls. The PID and PPID of the *exec*'d process remain unchanged.
- **Wait:** The parent then executes the *wait* system call to *wait* for the child to complete. It picks up the exit status of the child and continues with its other functions. Note that a parent need not decide to wait for the child to terminate.

To get a better idea of this, let us explain with an example. When you enter *ls* to look at the contents of your current working directory, UNIX does a series of things to create an environment for *ls* and the run it:


- The shell has UNIX perform a *fork*. This creates a new process that the shell will use to run the *ls* program.
- The shell has UNIX perform an *exec* of the *ls* program. This replaces the shell program and data with the program and data for *ls* and then starts running that new program.
- The *ls* program is loaded into the new process context, replacing the text and data of the shell.
- The *ls* program performs its task, listing the contents of the current directory. In the meanwhile, the shell executes *wait* system call for *ls* to complete.

When a process is forked, the child has a different PID and PPID from its parent. However, it inherits most of the attributes of the parent. The important attributes that are inherited are:

- **User name of the real and effective user (RUID and EUID):** the owner of the process. The real owner is the user issuing the command, the effective user is the one determining access to system resources. RUID and EUID are usually the same, and the process has the same access rights the issuing user would have.
- **Real and effective group owner (RGID and EGID):** The real group owner of a process is the primary group of the user who started the process. The effective group owner is usually the same, except when SGID access mode has been applied to a file.
- The current directory from where the process was run.
- The file descriptors of all files opened by the parent process.
- Environment variables like HOME, PATH.

The inheritance here means that the child has its own copy of these parameters and thus can alter the environment it has inherited. But the modified environment is not available to the parent process.

How the Shell is created?



- When the system moves to multiuser mode, **init** forks and execs a **getty** for every active communication port.
- Each one of these **getty**'s prints the login prompt on the respective terminal and then goes off to sleep.
- When a user tries to log in, **getty** wakes up and fork-exec the **login** program to verify login name and password entered.
- On successful login, **login** for-exec the process representing the login shell.
- **init** goes off to sleep, waiting for the children to terminate. The processes **getty** and **login** overlay themselves.
- When the user logs out, it is intimated to **init**, which then wakes up and spawns another **getty** for that line to monitor the next login.

4. Internal and External Commands

From the process viewpoint, the shell recognizes three types of commands:

1. External commands: Commonly used commands like **cat**, **ls** etc. The shell creates a process for each of these commands while remaining their parent.
2. Shell scripts: The shell executes these scripts by spawning another shell, which then executes the commands listed in the script. The child shell becomes the parent of the commands that feature in the shell.
3. Internal commands: When an internal command is entered, it is directly executed by the shell. Similarly, variable assignment like `x=5`, doesn't generate a process either.

Note: Because the child process inherits the current working directory from its parent as one of the environmental parameters, it is necessary for the `cd` command not to spawn a child to achieve a change of directory. If this is allowed, after the child dies, control would revert to the parent and the original directory would be restored. Hence, `cd` is implemented as an internal command.

5. Process States and Zombies

At any instance of time, a process is in a particular state. A process after creation is in the *runnable* state. Once it starts running, it is in the *running* state. When a process requests for a resource (like disk I/O), it may have to wait. The process is said to be in *waiting* or *sleeping* state. A process can also be *suspended* by pressing a key (usually *Ctrl-z*).

When a process terminates, the kernel performs clean-up, assigns any children of the exiting process to be adopted by **init**, and sends the death of a child signal to the parent process, and converts the process into the zombie state.

A process in zombie state is not alive; it does not use any resources nor does any work. But it is not allowed to die until the exit is acknowledged by the parent process.

It is possible for the parent itself to die before the child dies. In such case, the child becomes an **orphan** and the kernel makes **init** the parent of the orphan. When this adopted child dies, **init** waits for its death.

6. Running Jobs in Background

The basic idea of a background job is simple. It's a program that can run without prompts or other manual interaction and can run in parallel with other active processes.

Interactive processes are initialized and controlled through a terminal session. In other words, there has to be someone connected to the system to start these processes; they are not started automatically as part of the system functions. These processes can run in the foreground, occupying the terminal that started the program, and you can't start other applications as long as this process is running in the foreground.

There are two ways of starting a job in the background – with the shell's **&** operator and the **nohup** command.

&: No Logging out

Ordinarily, when the shell runs a command for you, it waits until the command is completed. During this time, you cannot communicate with the shell. You can run a command that takes a long time to finish as a background job, so that you can be doing something else. To do this, use the **&** symbol at the end of the command line to direct the shell to execute the command in the background.

```
$ sort -o emp.dat emp.dat &  
[1] 1413           The job's PID
```

Note:

1. Observe that the shell acknowledges the background command with two numbers. First number [1] is the *job ID* of this command. The other number 1413 is the PID.
2. When you specify a command line in a pipeline to run in the background, all the commands are run in the background, not just the last command.
3. The shell remains the parent of the background process.

nohup: Log out Safely

A background job executed using **&** operator ceases to run when a user logs out. This is because, when you logout, the shell is killed and hence its children are also killed. The UNIX system provides **nohup** statement which when prefixed to a command, permits execution of the process even after the user has logged out. You must use the **&** with it as well.

The syntax for the **nohup** command is as follows:

```
nohup command-string [input-file] output-file &
```

If you try to run a command with **nohup** and haven't redirected the standard error, UNIX automatically places any error messages in a file named **nohup.out** in the directory from which the command was run.

In the following command, the sorted file and any error messages are placed in the file `nohup.out`.

```
$ nohup sort sales.dat &  
1252  
Sending output to nohup.out
```

Note that the shell has returned the PID (1252) of the process.

When the user logs out, the child turns into an orphan. The kernel handles such situations by reassigning the PPID of the orphan to the system's init process (PID 1) - the parent of all shells. When the user logs out, init takes over the parentage of any process run with `nohup`. In this way, you can kill a parent (the shell) without killing its child.

Additional Points

When you run a command in the background, the shell disconnects the standard input from the keyboard, but does not disconnect its standard output from the screen. So, output from the command, whenever it occurs, shows up on screen. It can be confusing if you are entering another command or using another program. Hence, make sure that both standard output and standard error are redirected suitably.

```
$ find . -name "*.log" -print > log_file 2> err.dat &  
OR $ find . -name "*.log" -print > log_file 2> /dev/null &
```

Important:

1. You should relegate time-consuming or low-priority jobs to the background.
2. If you log out while a background job is running, it will be terminated.

7. nice: Job Execution with Low Priority

Processes in UNIX are sequentially assigned resources for execution. The kernel assigns the CPU to a process for a time slice; when the time elapses, the process is placed in a queue. How the execution is scheduled depends on the priority assigned to the process.

The *nice* command is used to control background process dispatch priority.

The idea behind *nice* is that background jobs should demand less attention from the system than interactive processes.

Background jobs execute without a terminal attached and are usually run in the background for two reasons:

1. the job is expected to take a relatively long time to finish, and
2. the job's results are not needed immediately.

Interactive processes, however, are usually shells where the speed of execution is critical because it directly affects the system's apparent response time. It would therefore be nice for everyone (others as well as you) to let interactive processes have priority over background work.

nice values are system dependent and typically range from 1 to 19.

A high *nice* value implies a lower priority. A program with a high nice number is friendly to other programs, other users and the system; it is not an important job. The lower the nice number, the more important a job is and the more resources it will take without sharing them.

Example:

```
$ nice wc -l hugefile.txt  
OR $ nice wc -l hugefile.txt &
```

The default nice value is set to 10.

We can specify the nice value explicitly with `-n number` option where *number* is an offset to the default. If the `-n number` argument is present, the priority is incremented by that amount up to a limit of 20.

Example: `$ nice -n 5 wc -l hugefile.txt &`

8. Killing Processes with Signals

When you execute a command, one thing to keep in mind is that commands do not run in a vacuum. Many things can happen during a command execution that are not under the control of the command. The user of the command may press the interrupt key or send a kill command to the process, or the controlling terminal may become disconnected from the system. In UNIX, any of these events can cause a **signal** to be sent to the process. The default action when a process receives a signal is to terminate.

When a process ends normally, the program returns its *exit status* to the parent. This exit status is a number returned by the program providing the results of the program's execution. Sometimes, you want or need to terminate a process.

The following are some reasons for stopping a process:

- It's using too much CPU time.
- It's running too long without producing the expected output.
- It's producing too much output to the screen or to a disk file.
- It appears to have locked a terminal or some other session.
- It's using the wrong files for input or output because of an operator or programming error.
- It's no longer useful.

If the process to be stopped is a background process, use the kill command to get out of these situations. To stop a command that isn't in the background, press `<ctrl-c>`.

To use kill, use either of these forms:

```
kill PID(s)    OR    kill -s NUMBER PID(s)
```

To kill a process whose PID is 123 use,
`$ kill 123`

To kill several processes whose PIDs are 123, 342, and 73 use,
`$ kill 123 342 73`

Issuing the kill command sends a signal to a process. The default signal is SIGTERM signal (15). UNIX programs can send or receive more than 20 signals, each of which is represented by a number. (Use kill -l to list all signal names and numbers)

If the process ignores the signal SIGTERM, you can kill it with SIGKILL signal (9) as,
\$ kill -9 123 OR \$ kill -s KILL 123

The system variable \$! stores the PID of the last background job. You can kill the last background job without knowing its PID by specifying \$ kill \$!

Note: You can kill only those processes that you own; You can't kill processes of other users. To kill all background jobs, enter kill 0.

9. Job Control

A job is a name given to a group of processes that is typically created by piping a series of commands using pipeline character. You can use job control facilities to manipulate jobs. You can use job control facilities to,

1. Relegate a job to the background (bg)
2. Bring it back to the foreground (fg)
3. List the active jobs (jobs)
4. Suspend a foreground job ([Ctrl-z])
5. Kill a job (kill)

The following examples demonstrate the different job control facilities.

Assume a process is taking a long time. You can suspend it by pressing [Ctrl-z].

```
[1] + Suspended                   wc -l hugefile.txt
```

A suspended job is not terminated. You can now relegate it to background by,
\$ bg

You can start more jobs in the background any time:

```
$ sort employee.dat > sortedlist.dat &
```

```
[2]    530
```

```
$ grep 'director' emp.dat &
```

```
[3]    540
```

You can see a listing of these jobs using jobs command,

```
$ jobs
```

```
[3] +  Running                   grep 'director' emp.dat &
```

```
[2] -  Running                   sort employee.dat > sortedlist.dat &
```

```
[1]    Suspended    wc -l hugefile.txt
```

You can bring a job to foreground using fg %jobno OR fg %jobname as,

```
$ fg %2            OR    $ fg %sort
```

10. at And batch: Execute Later

UNIX provides facilities to schedule a job to run at a specified time of day. If the system load varies greatly throughout the day, it makes sense to schedule less important jobs at a time when the system load is low. The at and batch commands make such job scheduling possible.

at: One-Time Execution

To schedule one or more commands for a specified time, use the `at` command. With this command, you can specify a time, a date, or both.

For example,

```
$ at 14:23 Friday
at> lp /usr/sales/reports/*
at> echo "Files printed, Boss!" | mail -s"Job done" boss
[Ctrl-d]
commands will be executed using /usr/bin/bash
job 1041198880.a at Fri Oct 12 14:23:00 2007
```

The above job prints all files in the directory `/usr/sales/reports` and sends a user named `boss` some mail announcing that the print job was done.

All `at` jobs go into a queue known as `at queue`. `at` shows the job number, the date and time of scheduled execution. This job number is derived from the number of seconds elapsed since the Epoch. A user should remember this job number to control the job.

```
$ at 1 pm today
at> echo  ``^G^GLunch with Director at 1 PM^G^G" >
/dev/term/43
```

The above job will display the following message on your screen (`/dev/term/43`) at 1:00 PM, along with two beeps (^G^G).

Lunch with Director at 1 PM

To see which jobs you scheduled with `at`, enter `at -l`. Working with the preceding examples, you may see the following results:

```
job 756603300.a at Tue Sep 11 01:00:00 2007
job 756604200.a at Fri Sep 14 14:23:00 2007
```

The following forms show some of the keywords and operations permissible with `at` command:

<code>at hh:mm</code>	Schedules job at the hour (<i>hh</i>) and minute (<i>mm</i>) specified, using a 24-hour clock
<code>at hh:mm month day year</code>	Schedules job at the hour (<i>hh</i>), minute (<i>mm</i>), month, day, and year specified
<code>at -l</code>	Lists scheduled jobs
<code>at now +count time-units</code>	Schedules the job right now plus <i>count</i> number of <i>timeunits</i> ; time units can be minutes, hours, days, or weeks
<code>at -r job_id</code>	Cancels the job with the job number matching <i>job_id</i>

batch: Execute in Batch Queue

The `batch` command lets the operating system decide an appropriate time to run a process. When you schedule a job with `batch`, UNIX starts and works on the process whenever the system load isn't too great.

To sort a collection of files, print the results, and notify the user named boss that the job is done, enter the following commands:

```
$ batch
sort /usr/sales/reports/* | lp
echo "Files printed, Boss!" | mailx -s"Job done" boss
```

The system returns the following response:

```
job 7789001234.b at Fri Sep 7 11:43:09 2007
```

The date and time listed are the date and time you pressed <Ctrl-d> to complete the batch command. When the job is complete, check your mail; anything that the commands normally display is mailed to you. Note that any job scheduled with batch command goes into a special at queue.

11. cron: Running jobs periodically

cron program is a daemon which is responsible for running repetitive tasks on a regular schedule. It is a perfect tool for running system administration tasks such as backup and system logfile maintenance. It can also be useful for ordinary users to schedule regular tasks including calendar reminders and report generation.

Both *at* and *batch* schedule commands on a one-time basis. To schedule commands or processes on a regular basis, you use the cron (short for *chronograph*) program. You specify the times and dates you want to run a command in crontab files. Times can be specified in terms of minutes, hours, days of the month, months of the year, or days of the week.

cron is listed in a shell script as one of the commands to run during a system boot-up sequence. Individual users don't have permission to run cron directly.

If there's nothing to do, cron "goes to sleep" and becomes inactive; it "wakes up" every minute, however, to see if there are commands to run.

cron looks for instructions to be performed in a control file in
/var/spool/cron/crontabs

After executing them, it goes back to sleep, only to wake up the next minute.

To a create a crontab file,

First use an editor to create a crontab file say cron.txt

Next use crontab command to place the file in the directory containing crontab files. crontab will create a file with filename same as user name and places it in /var/spool/cron/crontabs directory.

Alternately you can use crontab with -e option.

You can see the contents of your crontab file with crontab -l and remove them with crontab -r.

The cron system is managed by the cron daemon. It gets information about which programs and when they should run from the system's and users' crontab entries. The

crontab files are stored in the file `/var/spool/cron/crontabs/<user>` where `<user>` is the login-id of the user. Only the root user has access to the system crontabs, while each user should only have access to his own crontabs.

A typical entry in crontab file

A typical entry in the crontab file of a user will have the following format.

minute hour day-of-month month-of-year day-of-week command

where, Time-Field Options are as follows:

Field	Range
<i>minute</i>	00 through 59 Number of minutes after the hour
<i>hour</i>	00 through 23 (midnight is 00)
<i>day-of-month</i>	01 through 31
<i>month-of-year</i>	01 through 12
<i>day-of-week</i>	01 through 07 (Monday is 01, Sunday is 07)

The first five fields are time option fields. You must specify all five of these fields. Use an asterisk (*) in a field if you want to ignore that field.

Examples:

```
00-10 17 * 3.6.9.12 5 find / -newer .last_time -print >backuplist
```

In the above entry, the find command will be executed every minute in the first 10 minutes after 5 p.m. every Friday of the months March, June, September and December of every year.

```
30 07 * * 01 sort /usr/www/sales/weekly |mail -s"Weekly Sales" srm
```

In the above entry, the sort command will be executed with `/usr/www/sales/weekly` as argument and the output is mailed to a user named `srm` at 7:30 a.m. each Monday.

12. time: Timing Processes

The time command executes the specified command and displays the time usage on the terminal.

Example: You can find out the time taken to perform a sorting operation by preceding the sort command with time.

```
$ time sort employee.dat > sortedlist.dat
real    0m29.811s
user    0m1.370s
sys     0m9.990s
```

where,

the *real* time is the clock elapsed from the invocation of the command until its termination.

the *user* time shows the time spent by the program in executing itself.

the *sys* time indicates the time used by the kernel in doing work on behalf of a user process.

The sum of user time and sys time actually represents the CPU time. This could be significantly less than the real time on a heavily loaded system.

Conclusion

In this chapter, we saw an important abstraction of the UNIX operating system viz., processes. We also saw the mechanism of process creation, the attributes inherited by the child from the parent process as well as the shell's behavior when it encounters internal commands, external commands and shell scripts. This chapter also discussed background jobs, creation and controlling jobs as well as controlling processes using signals. We finally described three commands viz., at, batch and cron for process scheduling, with a discussion of time command for obtaining time usage statistics of process execution.

CHAP # 8 SIMPLE FILTERS

Filters are the commands which accept data from standard input manipulate it and write the results to standard output. Filters are the central tools of the UNIX tool kit, and each filter performs a simple function. Some commands use delimiter, pipe (|) or colon (:). Many filters work well with delimited fields, and some simply won't work without them. The piping mechanism allows the standard output of one filter serve as standard input of another. The filters can read data from standard input when used without a filename as argument, and from the file otherwise

The Simple Database

Several UNIX commands are provided for text editing and shell programming. (emp.lst) - each line of this file has six fields separated by five delimiters. The details of an employee are stored in one single line. This text file designed in fixed format and containing a personnel database. There are 15 lines, where each field is separated by the delimiter |.

```
$ cat emp.lst
```

```
2233 | a.k.shukla | g.m | sales | 12/12/52 | 6000
9876 | jai sharma | director | production | 12/03/50 | 7000
5678 | sumit chakrobarty | d.g.m. | marketing | 19/04/43 | 6000
2365 | barun sengupta | director | personnel | 11/05/47 | 7800
5423 | n.k.gupta | chairman | admin | 30/08/56 | 5400
1006 | chanchal singhvi | director | sales | 03/09/38 | 6700
6213 | karuna ganguly | g.m. | accounts | 05/06/62 | 6300
1265 | s.n. dasgupta | manager | sales | 12/09/63 | 5600
4290 | jayant choudhury | executive | production | 07/09/50 | 6000
2476 | anil aggarwal | manager | sales | 01/05/59 | 5000
6521 | lalit chowdury | directir | marketing | 26/09/45 | 8200
3212 | shyam saksena | d.g.m. | accounts | 12/12/55 | 6000
3564 | sudhir agarwal | executive | personnel | 06/07/47 | 7500
2345 | j. b. sexena | g.m. | marketing | 12/03/45 | 8000
0110 | v.k.agrawal | g.m.| marketing | 31/12/40 | 9000
```

pr : paginating files

We know that,

```
cat dept.lst
```

```
01|accounts|6213
02|progs|5423
03|marketing|6521
```

04|personnel|2365
05|production|9876
06|sales|1006

pr command adds suitable headers, footers and formatted text. pr adds five lines of margin at the top and bottom. The header shows the date and time of last modification of the file along with the filename and page number.

```
pr dept.lst
```

```
May 06 10:38 1997 dept.lst page 1
```

```
01:accounts:6213  
02:progs:5423  
03:marketing:6521  
04:personnel:2365  
05:production:9876  
06:sales:1006
```

...blank lines...

pr options

The different options for pr command are:

- k prints k (integer) columns
- t to suppress the header and footer
- h to have a header of user's choice
- d double spaces input
- n will number each line and helps in debugging
- on offsets the lines by n spaces and increases left margin of page

```
pr +10 chap01
```

starts printing from page 10

```
pr -l 54 chap01
```

this option sets the page length to 54

head – displaying the beginning of the file

The command displays the top of the file. It displays the first 10 lines of the file, when used without an option.

```
head emp.lst
```

-n to specify a line count
head -n 3 emp.lst

will display the first three lines of the file.

tail – displaying the end of a file

This command displays the end of the file. It displays the last 10 lines of the file, when used without an option.

```
tail emp.lst
```

-n to specify a line count

```
tail -n 3 emp.lst
```

displays the last three lines of the file. We can also address lines from the beginning of the file instead of the end. The +count option allows to do that, where count represents the line number from where the selection should begin.

```
tail +11 emp.lst  
Will display 11th line onwards
```

Different options for tail are:

- Monitoring the file growth (-f)
- Extracting bytes rather than lines (-c)

Use tail -f when we are running a program that continuously writes to a file, and we want to see how the file is growing. We have to terminate this command with the interrupt key.

cut – slitting a file vertically

It is used for slitting the file vertically. head -n 5 emp.lst | tee shortlist will select the first five lines of emp.lst and saves it to *shortlist*. We can cut by using -c option with a list of column numbers, delimited by a comma (cutting columns).

```
cut -c 6-22,24-32 shortlist
```

```
cut -c -3,6-22,28-34,55- shortlist
```

The expression 55- indicates column number 55 to end of line. Similarly, -3 is the same as 1-3.

Most files don't contain fixed length lines, so we have to cut fields rather than columns (cutting fields).

-d for the field delimiter
-f for the field list

```
cut -d \| -f 2,3 shortlist | tee cutlist1
```

will display the second and third columns of *shortlist* and saves the output in *cutlist1*. here | is escaped to prevent it as pipeline character

- To print the remaining fields, we have

```
cut -d \| -f 1,4- shortlist > cutlist2
```

paste – pasting files

When we cut with *cut*, it can be pasted back with the *paste* command, *vertically* rather than horizontally. We can view two files side by side by pasting them. In the previous topic, cut was used to create the two files cutlist1 and cutlist2 containing two cut-out portions of the same file.

```
paste cutlist1 cutlist2
```

We can specify one or more delimiters with -d

```
paste -d “|” cutlist1 cutlist2
```

Where each field will be separated by the delimiter |. Even though paste uses at least two files for concatenating lines, the data for one file can be supplied through the standard input.

Joining lines (-s)

Let us consider that the file *address book* contains the details of three persons

```
cat addressbook
```

```
paste -s addressbook -to print in one single line
```

```
paste -s -d ”||\n” addressbook -are used in a circular manner
```

sort : ordering a file

Sorting is the ordering of data in ascending or descending sequence. The sort command orders a file and by default, the entire line is sorted

sort shortlist

This default sorting sequence can be altered by using certain options. We can also sort one or more keys (fields) or use a different ordering rule.

sort options

The important sort options are:

-tchar	uses delimiter <i>char</i> to identify fields
-k n	sorts on nth field
-k m,n	starts sort on mth field and ends sort on nth field
-k m.n	starts sort on nth column of mth field
-u	removes repeated lines
-n	sorts numerically
-r	reverses sort order
-f	folds lowercase to equivalent uppercase
-m list	merges sorted files in list
-c	checks if file is sorted
-o filename	places output in file filename sort

`-t'|' -k 2 shortlist`

sorts the second field (name)

`sort -t'|' -r -k 2 shortlist` or

`sort -t'|' -k 2r shortlist`

sort order can be reversed with this `-r` option.

`sort -t'|' -k 3,3 -k 2,2 shortlist`

sorting on secondary key is also possible as shown above.

`sort -t'|' -k 5.7,5.8 shortlist`

we can also specify a character position within a field to be the beginning of sort as shown above (sorting on columns).

`sort -n numfile`

when sort acts on numerals, strange things can happen. When we sort a file containing only numbers, we get a curious result. This can be overridden by `-n` (numeric) option.

```
cut -d "|" -f3 emp.lst | sort -u | tee desigx.lst
```

Removing repeated lines can be possible using `-u` option as shown above. If we cut out the designation filed from `emp.lst`, we can pipe it to `sort` to find out the unique designations that occur in the file.

Other sort options are:

```
sort -o sortedlist -k 3 shortlist
```

```
sort -o shortlist shortlist
```

```
sort -c shortlist
```

```
sort -t "|" -c -k 2 shortlist
```

```
sort -m foo1 foo2 foo3
```

uniq command – locate repeated and nonrepeated lines

When we concatenate or merge files, we will face the problem of duplicate entries creeping in. we saw how `sort` removes them with the `-u` option. UNIX offers a special tool to handle these lines – the `uniq` command. Consider a sorted `dept.lst` that includes repeated lines:

```
cat dept.lst
```

displays all lines with duplicates. Where as,

```
uniq dept.lst
```

simply fetches one copy of each line and writes it to the standard output. Since `uniq` requires a sorted file as input, the general procedure is to sort a file and pipe its output to `uniq`. The following pipeline also produces the same output, except that the output is saved in a file:

```
sort dept.lst | uniq -> uniuqlist
```

Different `uniq` options are :

Selecting the nonrepeated lines (`-u`)

```
cut -d "|" -f3 emp.lst | sort | uniq -u
```

Selecting the duplicate lines (`-d`)

```
cut -d "|" -f3 emp.lst | sort | uniq -d
```

Counting frequency of occurrence (-c)

```
cut -d "|" -f3 emp.lst | sort | uniq -c
```

tr command – translating characters

The tr filter manipulates the individual characters in a line. It translates characters using one or two compact expressions.

tr options expn1 expn2 standard input

It takes input only from standard input, it doesn't take a filename as argument. By default, it translates each character in expression1 to its mapped counterpart in expression2. The first character in the first expression is replaced with the first character in the second expression, and similarly for the other characters.

```
tr '/' '~-' < emp.lst | head -n 3
```

```
exp1='/' ; exp2='~'
```

```
tr "$exp1" "$exp2" < emp.lst
```

Changing case of text is possible from lower to upper for first three lines of the file.

```
head -n 3 emp.lst | tr '[a-z]' '[A-Z]'
```

Different **tr options** are:

Deleting characters (-d)

```
tr -d '/' < emp.lst | head -n 3
```

Compressing multiple consecutive characters (-s)

```
tr -s ' ' < emp.lst | head -n 3
```

Complementing values of expression (-c)

```
tr -cd '/' < emp.lst
```

Using ASCII octal values and escape sequences

```
tr '|' '\012' < emp.lst | head -n 6
```

Filters Using Regular Expression grep and sed

We often need to search a file for a pattern, either to see the lines containing (or not containing) it or to have it replaced with something else. This chapter discusses two important filters that are specially suited for these tasks – grep and sed. grep takes care of all search requirements we may have. sed goes further and can even manipulate the individual characters in a line. In fact sed can do several things, some of them quite well.

grep – searching for a pattern

It scans the file / input for a pattern and displays lines containing the pattern, the line numbers or filenames where the pattern occurs. It's a command from a special family in UNIX for handling search requirements.

```
grep options pattern filename(s)
```

```
grep "sales" emp.lst
```

will display lines containing sales from the file emp.lst. Patterns with and without quotes is possible. It's generally safe to quote the pattern. Quote is mandatory when pattern involves more than one word. It returns the prompt in case the pattern can't be located.

```
grep president emp.lst
```

When grep is used with multiple filenames, it displays the filenames along with the output.

```
grep "director" emp1.lst emp2.lst
```

Where it shows filename followed by the contents

grep options

grep is one of the most important UNIX commands, and we must know the options that POSIX requires grep to support. Linux supports all of these options.

-i	ignores case for matching
-v	doesn't display lines matching expression
-n	displays line numbers along with lines
-c	displays count of number of occurrences
-l	displays list of filenames only
-e exp	specifies expression with this option
-x	matches pattern with entire line
-f file	takes patterns from file, one per line

-E treats pattern as an extended RE
-F matches multiple fixed strings
grep -i 'agarwal' emp.lst

grep -v 'director' emp.lst > otherlist

wc -l otherlist will display 11 otherlist

grep -n 'marketing' emp.lst

grep -c 'director' emp.lst

grep -c 'director' emp*.lst

will print filenames prefixed to the line count

grep -l 'manager' *.lst

will display filenames *only*

grep -e 'Agarwal' -e 'aggarwal' -e 'agrawal' emp.lst

will print matching multiple patterns

grep -f pattern.lst emp.lst

all the above three patterns are stored in a separate file *pattern.lst*

Basic Regular Expressions (BRE) – An Introduction

It is tedious to specify each pattern separately with the -e option. grep uses an expression of a different type to match a group of similar patterns. If an expression uses meta characters, it is termed a regular expression. Some of the characters used by regular expression are also meaningful to the shell.

BRE character subset

The basic regular expression character subset uses an elaborate meta character set, overshadowing the shell's wild-cards, and can perform amazing matches.

*	Zero or more occurrences
g*	nothing or g, gg, ggg, etc.
.	A single character
.*	nothing or any number of characters
[pqr]	a single character p, q or r
[c1-c2]	a single character within the ASCII range represented by c1 and c2

The character class

grep supports basic regular expressions (BRE) by default and extended regular expressions (ERE) with the `-E` option. A regular expression allows a group of characters enclosed within a pair of `[]`, in which the match is performed for a single character in the group.

```
grep "[aA]g[ar][ar]wal" emp.lst
```

A single pattern has matched two similar strings. The pattern `[a-zA-Z0-9]` matches a single alphanumeric character. When we use range, make sure that the character on the left of the hyphen has a lower ASCII value than the one on the right. Negating a class (`^`) (caret) can be used to negate the character class. When the character class begins with this character, all characters other than the ones grouped in the class are matched.

The *

The asterisk refers to the immediately preceding character. `*` indicates zero or more occurrences of the previous character.

`g*` nothing or `g`, `gg`, `ggg`, etc.

```
grep "[aA]gg*[ar][ar]wal" emp.lst
```

Notice that we don't require to use `-e` option three times to get the same output!!!!

The dot

A dot matches a single character. The shell uses `? Character` to indicate that.

`.*` signifies any number of characters or none

```
grep "j.*saxena" emp.lst
```

Specifying Pattern Locations (^ and \$)

Most of the regular expression characters are used for matching patterns, but there are two that can match a pattern at the beginning or end of a line. Anchoring a pattern is often necessary when it can occur in more than one place in a line, and we are interested in its occurrence only at a particular location.

<code>^</code>	for matching at the beginning of a line
<code>\$</code>	for matching at the end of a line

```
grep "^2" emp.lst
```

Selects lines where emp_id starting with 2

```
grep "7...$" emp.lst
```

Selects lines where emp_salary ranges between 7000 to 7999

```
grep "[^2]" emp.lst
```

Selects lines where emp_id doesn't start with 2

When meta characters lose their meaning

It is possible that some of these special characters actually exist as part of the text. Sometimes, we need to escape these characters. For example, when looking for a pattern `g*`, we have to use `\`
To look for `[`, we use `\[` To
look for `.`, we use `\.`

Extended Regular Expression (ERE) and grep

If current version of `grep` doesn't support ERE, then use `egrep` but without the `-E` option. `-E` option treats pattern as an ERE.

`+` matches one or more occurrences of the previous character

`?` Matches zero or one occurrence of the previous character

`b+` matches `b`, `bb`, `bbb`, etc.

`b?` matches either a single instance of `b` or nothing

These characters restrict the scope of match as compared to the `*`

```
grep -E "[aA]gg?arwal" emp.lst
```

```
# ?include <stdio.h>
```

The ERE set

`ch+` matches one or more occurrences of character `ch`

`ch?` Matches zero or one occurrence of character `ch`

`exp1|exp2` matches `exp1` or `exp2`

`(x1|x2)x3` matches `x1x3` or `x2x3`

Matching multiple patterns (`|`, `(and)`)

```
grep -E 'sengupta|dasgupta' emp.lst
```

We can locate both without using `-e` option twice, or

```
grep -E '(sen|das)gupta' emp.lst
```

END
ALL THE BEST