

Chapter 1, An Introduction to Linux

John M. Morrison

September 11, 2014

Contents

1	Introduction	2
2	In the Beginning ...	3
3	The Anatomy of a UNIX Command	4
4	Managing Directories	5
4.1	Processes and Directories	8
5	A Field Trip	8
6	Making and Listing Regular Files	10
7	Renaming and Deleting Files	13
7.1	Everything is a Computer Program	14
8	Editing Files with vi	18
8.1	A Note for Ubuntu Users	18
8.2	Launching vi	18
8.3	vi Modes	19
8.4	Cut and Paste	22
8.5	Cutting and Pasting with External Files	22
8.6	Searching	23

9	Visual Mode	24
10	Copy Paste from a GUI	25
10.1	Permissions	27
11	The Octal Representation for Permissions	28
12	The Man	29
13	Redirection of Standard Output and Standard Input	32
14	Using Pipes	35
14.1	The <code>sort</code> filter	35
14.2	The Filters <code>head</code> , <code>tail</code> , and <code>uniq</code>	36
14.3	The <code>grep</code> Command	36
14.4	And Now to Pipes	37

1 Introduction

Appendix A give details on getting set up to run in a Linux environment. You will need to consult it to get started.

You are probably used to running a computer with a graphical user interface (GUI). These interfaces feature windows, icons, buttons, toolbars, and other graphical features you navigate and use with the aid of your keyboard and mouse. You likely run Windoze, MacOSX, or you may be running a Linux GUI on your machine. The GUI allows you to communicate with the operating system, which is the master program of your computer that manages such things as running applications and maintaining your file system.

In this book, we will study the Linux operating system in its command-line guise. You will control the computer by entering commands into a text window called a *terminal window*. Typically they look something like this.



The name “terminal” name harkens back to the days when you had an actual appliance on your desk consisting of a (heavy) CRT screen and keyboard that was connected to a computer elsewhere. Your terminal window is just a software version of this old appliance. You will enter commands into this window to get the remote machine you are communicating with to perform various tasks.

The text string `morrison@odonata` appearing in the window is called the **prompt**. Its presence indicates that the computer is waiting for you type in a command.

2 In the Beginning . . .

As we progress, everything will seem unfamiliar, but actually relate very directly to some very familiar things you have seen working with a computer having a GUI. We assume you have basic proficiency using some kind of computer such as a Mac or a Windoze box, and we will relate the things you do in Linux to those you do in your usual operating system.

Log in to your UNIX account. If you are working in a Linux GUI, or a Mac, just open a terminal session. The first thing you will see after any password challenge will resemble this

```
[yourUserName}@hostName yourUserName] $
```

or this

```
[yourUserName]@hostName ~]$
```

On a Mac, it will resemble this

```
John-Morrisons-MacBook-Pro:~ morrison$
```

The presence of the prompt means that Linux is waiting for you to enter a command. The token `yourUserName` will show your login name. Your prompt may have an appearance different from the ones shown here; this depends on how your system administrator sets up your host or on the type of Linux distribution you are using. The appearance of your prompt does not affect the underlying performance of UNIX. In fact, the properties of your session are highly configurable, and you can customize your prompt to have any appearance you wish.

To keep things simple and uniform throughout the book, we will use `$` to represent the UNIX prompt. You will interact with the operating system by entering commands, instead of using a mouse to push buttons or click in windows.

When you see this terminal, a program called a *shell* is running. The shell takes the commands you type and passes them on to the operating system (kernel) for action. You will type a command, then hit the ENTER key; this causes the command to be shipped to the OS by the shell. The shell then conveys the operating system's reply to your terminal screen. Think of the shell as a telephone through which you communicate with the operating system. This analogy is only fitting since UNIX was originally developed at AT&T Bell Labs.

We will begin by learning how to interact with the file system. This is what give you access to your programs and data, so it is very fundamental.

3 The Anatomy of a UNIX Command

Every UNIX command has three parts: a name, options, and zero or more arguments. They all have the following appearance

```
commandName -option(s) argument(s)
```

Notice how the options are preceded by a `-`. Certain “long-form” options are preceded by a `--`.

A command always has a name. Grammatically you should think of a command as an imperative sentence. For example, the command `passwd` means, “Change my password!” You can type this command at the prompt, hit enter, and follow the instructions to change your password any time you wish.

Arguments are supplementary information that is sometimes required and sometimes optional, depending on the command. Grammatically, arguments

are nouns: they are things that the command acts upon.

Options are always, well, ... optional. Options modify the basic action of the command and they behave grammatically as adverbs. We will assume you have used a computer with a GUI before. All familiar features of a graphics-based machine are present in Linux, you will just invoke them with a text command instead of a mouse click. We will go through some examples so you get familiar with all the parts of a Linux command.

Two very basic Linux commands are `whoami` and `hostname`. Here is a typical response. These commands give, respectively, your user name and the name of the host you log in to.

Now we run them. We show the results here; your computer will show your login name and your host name. Here is what they look like on a server.

```
$ whoami
morrison
$ hostname
carbon.ncssm.edu
$
```

Here is their appearance on a PC (A Mac in this instance).

```
$ whoami
morrison
$ hostname
John-Morrisons-MacBook-Pro.local
$
```

We will next turn to the organization of the file system.

4 Managing Directories

We will do a top-down exploration of the file system. In this spirit, we will first learn how to manage directories; this is the UNIX name for folders. You will want to know how to create and manage folders, and how to navigate through them.

Directories in UNIX work just like folders on a graphical operating system. You have been in a directory all along without knowing it. Whenever you start a UNIX session, you begin in your *home directory*. Every user on a UNIX system owns a home directory. This is where you will keep all of your stuff.

To see your home directory, type `pwd` at the UNIX prompt. This command means, “Print working directory!” You will see something like this.

```
$ pwd
/home/faculty/morrison
```

This directory is your *home directory*. Whenever you start a new session, you will begin here. This is the directory where all the stuff that belongs to you is kept.

In this example, `morrison` is a directory inside of `faculty`, which is inside a directory `home`, which is inside the *root directory* `/`. Your home directory will likely be slightly different. It is very common for UNIX systems to keep all user directories inside of a directory named `home`. Often, several different types of users are organized into sub-directories of `home`. You will later see that all directories live inside of the root directory, `/`. Enter the `pwd` command on your machine and compare the result to what was shown here. Become familiar with your home directory's appearance so you can follow what goes on in the rest of this chapter.

If you are using Linux on your PC, your home directory will likely look like this.

```
/home/morrison
```

This directory structure is exactly the same as your hierarchy of folders and files on a Mac or a Windows box. You already know that folders can contain files and other folders. This is also true in a UNIX environment. We will show UNIX commands that do all the things you already know how to do in a graphical environment.

To make a new directory in Mac or Windows, you right click in the open folder and choose a menu for making a new folder. In UNIX, the `mkdir` command makes a one or more new directories. It requires at least one argument, the name(s) of the director(ies) you are creating. Let us make a directory by typing

```
$ mkdir Projects
```

makes a directory called `Projects`; this directory is now empty. We can always get rid of an empty directory or directories by typing a command like there

```
$ rmdir garbageDirectory(ies)
```

where `garbageDirector(ies)` stands for the directory or directories you wish removed.

The `rmdir` command will not remove a directory unless it is empty. There is a way to snip off directories with their contents, but we will avoid it for now because it is very dangerous. For now, you can delete the contents of a directory, then remove the directory. Be warned as you proceed: *When you remove files or directories in LINUX, they are gone for good! There is no "undelete."*

To get into our new directory `Projects`, enter this command:

```
$ cd Projects
```

and type `ls`. You will see no files. This is because the directory `Projects` is empty, and `ls` by default only shows you the files in the directory you are currently occupying. The command `cd` means, “Change directory!” Having done this now type

```
$ pwd
```

You will see a directory path now ending in `Projects`.

There is a command called `touch` which will create an empty file(s) with a name(s) you specify. Create files named `moo` and `baa` with `touch` as follows

```
$ touch moo baa
```

Then enter `ls` at the command line. This command means “list stuff.” You will see just the files you created.

As we said before, The command `ls` displays only files in the directory you are currently occupying. This directory is called your *current working directory*, or `cwd` for short. Every terminal session has a working directory. When you first log in, your working directory is always your home directory. You will see that every user on a UNIX system has a home directory that containing all that user’s data.

```
/home/yourUserName/Projects
```

This directory is the `Projects` directory you just created.

If you type `cd` without arguments, you will go straight back to your home directory. This should make you will feel like Dorothy going back to Kansas. Now if we use `pwd` again we see our home directory printed out.

You can also see your home directory anywhere you are by typing

```
$ echo $HOME
```

The fearsome-looking object `$HOME` is just a symbol that points to your home directory. There are various items like this present in your system. They are called *environment variables*. Another environment variable is `$PWD`; this is just your current working directory.

4.1 Processes and Directories

We know that when we log in, we are starting a program called a *shell*. The shell is a process, or running program. Every process has a `cwd` (current working directory). When you type `pwd` into your shell, you are asking the OS to tell you your shells current working directory. If you log in to a UNIX server in several terminal windows, each runs in a separate shell, so each has its own working directory.

Observe that, much of the time, your shell is idle. When you finish typing a command and hit the enter key, that command launches a program, that program runs, and any output is directed to your terminal window.

The command `cd` is a computer program. What it does is it changes the `cwd` of the shell that calls it. Now you know what it means to be “in” a directory: it means the `cwd` of your shell is that directory.

Programming Exercises

1. Enter

```
$ cd $HOME/Projects
```

and see what happens.
2. Make these directories inside of `Projects` `labors`, `feats` and `chores`
3. Type `cd labors` at the command line then `pwd`.
4. Type `cd ..` at the command line then `pwd`. What happened?
5. Type `cd ..` at the command line again, then `pwd`. What happened?
6. What do you think `..` is?
7. Type `cd .` at the comand line then `pwd`. What happened?
8. Type `ls .` at the comand line then `pwd`. What happened?
9. What do you think `.` is?

5 A Field Trip

To get to our first destination, type `cd /`. The directory `/` is the “root” directory. If you think of the directory structure as an upside-down (Australian) tree (root at top), the directory `/` is at the top. Type `pwd` and see where you are. Type `ls`; you should see that the directory `home` listed with several other directories. Here is what the directory structure looks like on a PC running Red Hat Fedora Core 9. Yours may have a slightly different appearance.


```

$ cd /
$ ls
bin    dev    initrd.img    lib64    mnt    root    selinux    tmp    vmlinuz
boot  etc    initrd.img.old  lost+found  opt    run    srv        usr    vmlinuz.old
cdrom  home  lib            media     proc   sbin   sys        var
$

```

Now type if we type `cd home` then `ls`, you will see one or more directories. On the machine being used here, you would see

```

$ cd home
$ ls
guest lost+found morrison

```

This machine has two users, `morrison` and `tt guest`. Since it is a personal computer, it does not have many users. You may be working on a server in which there could be dozens, or even hundreds of other users who are organized into various directories.

See if you can follow this all the way down to another users home directory. You may be able to list the files there, or even read them, depending on that users permissions. From this modest demonstration, you see that you can step down through the directory structure using `cd`. Now we will learn how to step up.

Try typing `cd ..`; the special symbol `..` represents the directory above your `cwd`. Now you can climb up and down the directory structure! The `..` symbol works like the up-arrow in a file chooser dialog box in Mac or Windoze. You saw this when you did the last group of exercises.

Practice this; go back to your home directory. Make a new directory called `mudpies`. Put some files in it. Make new directories in `mudpies`, got down inside these and make more directories and files. Practice using `cd` to navigate the tree you create. When you are done, get rid of the whole mess; remember you have to go to the bottom, empty out the files using `rm` and then use `rmdir` to get rid of the empty directories.

If you type `ls` in a directory, notice how any directories inside it are in differently colored type than regular files. This color is often blue. You can use the `-F` option in `ls` to print directory names with a slash (`/`) after them. Try this; it was an important option back in the days of monochrome monitors. If you use the `-l` option in `ls`, you will see that in the *permissions column*, the column begins with a `d` for any directory. Here is a possible sample

```

-rw-rw-r--    1 morrison morrison          0 Jun  9 14:54 bar
-rw-rw-r--    1 morrison morrison          0 Jun  9 14:54 foo
drwxrwxr-x    2 morrison morrison       4096 Jun  9 14:54 junk

```

You can see there that `bar` and `foo` are empty files. Notice the `d` at the beginning of the line in `junk`; this tells you `junk` is a directory.

6 Making and Listing Regular Files

In Chapter 0, we learned that the operating system is responsible for maintaining the file system. The file system maintained by a UNIX system consists of a hierarchy of files. Two types of files will be of interest to us: *directories* (folders) and *regular files*, i.e. files that are not directories. Regular files may hold data or programs. They may consist of text or be binary files that are not human-readable.

You are used to working with regular files and directories in Windows or MacOSX. Things in UNIX work the same way, but we will use commands to manage files instead of mouse clicking or dragging.

As we have already seen us now use the UNIX command `touch` to create new files. This command creates an empty file for each argument given it. At your UNIX prompt, enter

```
$ touch stuff
```

This creates the empty file named `stuff` in your account.

Now let us analyze the anatomy of this command. The name of the command is `touch`; its purpose is to create an empty file. Since you do not see a `-` sign, there are no options being used. The argument is `stuff`. This is the name of the file you created. Create a few more empty files. Enter these commands

```
$ touch foo
$ touch bar
```

You may create several files at once by making a space-separated list as we show here.

```
$ touch aardvark buffalo cougar dingo elephant
```

Now you have eight new files in your account. Next we will see how to list the files. Enter this command at the UNIX prompt

```
$ ls
```

The command `ls` lists your files. Notice we had neither options nor arguments. If you created the files using `touch` as instructed, they should appear on your screen like this

```
$ aardvark bar buffalo cougar dingo elephant foo stuff
```

The command `ls` has several options. One option is the `l` option; it list the files in long format. To invoke it, type

```
$ ls -l
```

You will see a listing like this

```
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 aardvark
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 bar
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 buffalo
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 cougar
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 dingo
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:50 elephant
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:49 foo
-rw-rw-r-- 1 morrison morrison      0 Jun  9 10:49 stuff
```

The first column reflects the *permissions* for the files. The sequence

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

indicates that you and your group have read/write permission and that others (“the world”) has read permission. We will discuss permissions in more detail when we discuss the management of directories.

You can see the name here is listed in two columns; on this machine `morrison` is in his own group. On another system, you may live in group with several other people; if so you will see the name of that group in one of these columns. The zero indicates the size of the file; each file we created is empty. Then there is a date, a time and the file name. This is the long format for file listing; it is seen by using the `-l` option in the `ls` command.

Another option is the `-a` option. This lists all files, including “hidden” files. Hidden files have a dot (`.`) preceding their name. To see them, enter

```
$ ls -a
```

at the command line. One thing you are guaranteed to see, no matter where you are are the directories `..` (parent) and `.` (current). If you are in you home directory, You will see the files you saw using `ls` and several hidden files with mysterious names like `.bash_profile`. Do not delete these; they provide the configuration for your account and do things like record preferences for applications you have used or installed. You can also list all files including hidden files by entering

```
$ ls --all
```

You can use more than one option at once. For example, entering

```
$ ls -al
```

or

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

```
$ ls --all -l
```

shows all of your files and hidden files in long format. Try this now on your machine.

Note to Mac Users Mac users should precede verbose commands with a single `-`. So on a Mac, you type

```
$ ls -all -l
```

and not

```
$ ls --all -l
```

Otherwise, your Mac will respond with a cryptic error message.

Next we will show how to display a file to the screen. UNIX commands that process files are called *filters*. Let us peek inside your `.bash_profile` file. Enter the command

```
$ cat .bash_profile
```

The command name is `cat`, short for catalog (the file to the screen). The `cat` command is a filter that does not filtering at all; it simply dumps the an entire file to the screen all at once. We are using no options, but the file name is an argument to `cat`. If a file is long and you want to see it one screenful at a time, use the filter `more`. The command `more` takes a file name as an argument and shows it on the screen a screenful at a time. You can hit the space bar to see the next screenful or use the down-arrow or enter key to scroll down one line at a time. To exit `more` at any time, type a `q` and `more` quits. You can use several arguments in `cat` or `more` and the indicated files will be displayed *in seriatum*.

7 Renaming and Deleting Files

Three commands every beginner should know are: `cp`, `rm` and `mv`. These are, respectively, copy, remove and move(rename). Here are their usages

```
cp oldFile newFile
rm garbageFile(s)
mv oldFile newFile
```

Warning! Pay heed before you proceed! To *clobber* a file means to unlink it from your file system. When you clobber a file it is lost and there is virtually no chance you will recover its contents. There is no undelete facility as you might find on other computing systems you have used.

If you remove a file it is clobbered, and there is no way to get it back. If you copy or rename onto an existing file, that file is clobbered, and it is gone forever. Always check to see if the file name you are copying or moving to is unoccupied! When in doubt, do an `ls` to look before you leap. All three of these commands have an option `-i`, which warns you before clobbering a file. Using this is a smart precaution.

The first command copies *oldFile* to *newFile*. If *newFile* does not exist, it creates *newFile*; otherwise it will overwrite any existing *newFile*.

```
Try this at your UNIX prompt: cp .bash_profile quack
```

Notice that the command `cp` has two arguments: the source file and the recipient file. If you executed the last command successfully, you made a copy of your `.bash_profile` file to a file called `quack`.

Next, let's get rid of all the animals in the zoo we had created before. The command `rm` will accept one or more arguments and remove the named files. We can accomplish this in one blow with

```
$ rm aardvark buffalo cougar dingo elephant
```

Now enter

```
$ ls -l
```

You will see that `quack`'s size is nonzero because it has a copy of the contents of your `.bash_profile` file in it. The file shown here has size 191. The size is the number of bytes contained in the file; yours may be larger or smaller. You will also see that the menagerie has been sent packing.

```
-rw-rw-r-- 1 morrison morrison    0 Jun  9 10:50 bar
```

```
-rw-rw-r-- 1 morrison morrison    0 Jun  9 10:49 foo
-rw-r--r-- 1 morrison morrison  191 Jun  9 11:25 quack
-rw-rw-r-- 1 morrison morrison    0 Jun  9 10:49 stuff
```

Let us now remove the file `stuff`. We are going to use the `-i` option. Enter this at the UNIX prompt.

```
$ rm -i stuff
```

The system will then ask you if you are sure you want to remove the file. Tell it yes by typing the letter `y`. Be reminded that the `-i` option is also available with `cp` and `mv`. You should you use it to avoid costly mistakes.

Finally, we shall use `mv` This “moves” a file to have a new name. Let’s change the name of `quack` to `honk` and back again. To change `quack` to `honk`, proceed as follows.

```
$ mv quack honk
```

Once you do this, list the files in long format. Then change it back.

Now you know how to copy, move, and create files. You can show them to the screen and you can list all the files you have. So far, we can create files two ways, we can create an empty file with `touch` or copy an existing file to a new file with `cp`.

7.1 Everything is a Computer Program

Now let us take a little look under the hood. When you log in, a program called a *shell* is launched. You can think of the shell as a telephone; entering a command at the prompt is how you speak into the phone. The shell accepts commands you enter at the prompt and sends them to the *kernel*, or operating system, which runs the program. This can cause output to be put to the screen, as in `ls`, or happen without comment, as in `rm`.

Programs that are running in UNIX are called *processes*. Every process has an owner and an integer associated with it called a process ID (PID). The user who spawns a process will generally be its owner. You are the owner of all processes you spawn. Many, such as `ls`, last such a short time you never notice them; they terminate in a fraction of a second after you enter them. When you log into your host, you actually are launching a program; this is your shell. When the shell terminates, your terminal session will be gone. At the command line, enter `ps` and you will see something like this.

```
$ ps
PID TTY
```

```
10355 pts/1
10356 pts/1
$
TIME CMD
00:00:00 bash
00:00:00 ps
```

The `ps` command shows all processes currently running spawned by your shell. On this machine, the shells (bash) process ID is 10355. By entering `ps aux` at the command line, you can see all processes running on your UNIX server, along with their process IDs and an abundance of other information. Try this at several different times. If you are using a server, you will see processes spawned by other users. You will also see other processes being run by the system to support your machine's operation.

An example of a program that does not finish its work immediately is the program `bc`. We show a sample `bc` session here; this application is a simple arbitrary-precision calculator.

```
$ bc
bc 1.06.94
Copyright 1991-1994, 1997, 1998, 2000, 2004, 2006
Free Software Foundation, Inc.
This is free software with ABSOLUTELY NO WARRANTY.
For details type warranty.
3+4
7
4*5
20
2^20
1048576
2^100
1267650600228229401496703205376
quit
```

When you type `bc` at the command prompt, the shell runs the `bc` program. This program continues to run until you stop it by typing `quit`. To see `bc`'s process ID, start `bc` and then type `Control-Z`. This interrupts the `bc` process, puts it in the background, and returns you to your shell. Then enter `ps` at the command prompt to see the process ID for your `bc` session.

```
$ bc
bc 1.06.94
Copyright 1991-1994, 1997, 1998, 2000, 2004, 2006 Free Software Foundation, Inc.
This is free software with ABSOLUTELY NO WARRANTY.
```

For details type warranty.

```
[1]+  
$ ps  
PID  
14110  
14253  
14254  
$  
Stopped  
TTY  
pts/4  
pts/4  
pts/4  
bc  
TIME  
00:00:00  
00:00:00  
00:00:00  
CMD  
bash  
bc  
ps
```

Try typing `exit` to log out; you will see something like this.

```
$ exit  
exit  
There are stopped jobs.  
$
```

Now type `jobs` at the command prompt. You will see this.

```
$ jobs  
[1]+ Stopped  
$  
bc
```

You can end the job `bc` labeled `[1]` by doing the following

```
$ kill %1  
$ jobs  
[1]+ Terminated  
$ jobs  
$  
bc
```


If several jobs are stopped, each will be listed with a number. You can end any you wish to by entering a `kill` command for each job. When you type `jobs` at the command line the first time, it will tell you what jobs it has suspended. After that, you will see a (possibly empty, like here) list of jobs still in the background. Do not dismiss a shell with running jobs; end them to preserve system resources.

You can bring your stopped job into the foreground by entering `fg` at the command prompt.

Exercises

1. Start up a session of `bc` and put it into the background using `control-Z`. Do this for several sessions. Type in some calculations into some of the sessions and see if they reappear when you bring the `bc` session containing that calculation into the foreground.
2. The `bc` calculator has variables which allow you to store numbers under a name. These play the role of the symbols described in Chapter 0, but they are limited to storing numbers. Here we show some variables being created and some expressions being evaluated.

```
morrison@ghent:~$ bc
bc 1.06.94
Copyright 1991-1994, 1997, 1998, 2000, 2004, 2006 Free Software Foundation, Inc.
This is free software with ABSOLUTELY NO WARRANTY.
For details type warranty.
cow = 5
pig = 2
horse = 7
horse + cow
12
horse/pig
3
pig/horse
0
cow^horse
78125
```

Replicate this session. Put it into the background and bring it into the foreground. Were your variables saved? Notice that this calculator does integer arithmetic. The `=` sign you see is actually assignment, which was discussed in Chapter 0.

3. Look at one of the algorithms for converting a binary number into a decimal number described in Chapter 0. Can you step through the process using `bc` and make it work?

8 Editing Files with vi

We can create files with `touch` and use `cp` to copy them. How do we edit text files and place information in them? This is the role of the UNIX text editor, `vi`. The O'Reilly book [?] on it comes highly recommended if you want to become a power user (you do). A second text editor, `emacs` is also available. It is powerful and extensible. Like `vi` it is a serious tool requiring serious learning, and like `vi` there is an O'Reilly book on it, too. You may use `emacs` instead of `vi` if you wish. Both of these are just tools for creating and editing text files, and both do a great job. You may create or modify any text file with either program. Ubuntu users can also use `gedit` or `gvim`, which have some nice advantages.

8.1 A Note for Ubuntu Users

Ubuntu by default installs the package `vi-tiny`. We want `vi` with all bells and whistles. To get this, make sure you are connected to the Internet, then type the following command in an terminal window.

```
$ sudo apt-get install vim
```

You will be asked to enter your password, then it will install the full `vi` package. The `sudo` command tells Ubuntu you are behaving as a system administrator, so you must enter your password to proceed. It will ask you to confirm you wish to install, and then it will download the package from the repositories, install, and configure it for you. Ubuntu has lots of programs and packages that are freely available, and you use `sudo apt-get install` to obtain them.

8.2 Launching vi

To create a new file or open an existing file, type

```
$ vi someFileName
```

at the UNIX command line. If the file `someFileName` exists, it will be opened; otherwise, it will be created. Now let us open the file `bar` we created with `touch`. You will see this:

```
~  
~  
~  
~  
~  
~
```


keystroke	Action
i	insert characters before the cursor
I	insert characters at the beginning of the line
a	append characters after the cursor
A	append characters at the end of the line
o	open a new line below the cursor
O	open a new line above the cursor

There is **one** way to get out of insert mode. You do this by hitting the escape (ESC) key. Let's now try this out. Go into your file `bar` and hit the `i` key to enter text. Type some text. Then hit ESC. To save your current effort type this anywhere:

```
:w
```

This will write the file; a message will appear at the bottom of the window indicating this has happened. Do not panic; that message is **not** a part of the file that is saved. To quit, type

```
:q
```

this will quit you out of the file. You can type

```
:wq
```

to write and quit. The command `:qw` does not work for obvious reasons. You have just done a simple `vi` session. You can reopen the file `bar` by typing

```
$ vi bar
```

at the UNIX command line; the contents you last saved will be re-displayed. You should take a few minutes to try all of the ways of getting into insert mode. Change the file and save it. Quit and display it to the screen with `cat` and `more`. At first, `vi` will seem clunky and awkward. However, as you ascend the learning curve, you will see that `vi` is blazingly fast and very efficient.

A Reassuring Note If you are in command mode and hit ESC, your computer will just beep at you. This is its way of letting you know you were already in command mode. Nothing additional happens. If you are unsure what mode you are in, hit ESC and you will be back in command mode.

The figure below will help you to remember the structure of `vi`. When you first start editing a file, you enter in in command mode. Typing `i`, `a`, `o`, `I`, `A` or `O` all put you into insert mode. You can also see in the diagram how to get out of insert mode by typing ESC.

```

-----
|          command mode          |
| -----i, a, o, I, A, O to get in----- |
| |   insert mode:              |
| |   insert characters          |
| |   paste in text with GUI    |
| -----ESC to get out----- |
|          search, replace      |
|          copy with yy, paste with p |
|          delete lines with dd  |
|          all "colon commands" (commands that start with :) |
-----

```

Let's go back in our file now and learn some more useful commands. We will look at command mode commands now.

Sometimes, line numbers will be helpful; these are especially useful when you program. To see them, you get into command mode and type the *colon command* `:set number`. Do this and watch them appear. Now type `:set nonu` or `:set nonumber` and watch them disappear. Line numbers are not a part of the file; however, they are a helpful convenience.

Here are some useful command mode mobility features. Experiment with them in a file.

Command	Action
<code>:lineNumber</code>	go to indicated line number
<code>^</code>	go to the beginning of the current line
<code>\$</code>	go to the end of the current line
<code>G</code>	go to the end of the file
<code>gg</code>	go to the beginning of the file

The colon commands in this table will allow you to alter your editing environment. The last two are useful editing tricks that are sometimes quite convenient. Open a file and try them all.

Command	Action
<code>:set number</code>	display line numbers
<code>:set nonu</code>	get rid of line numbers
<code>:set autoindent</code>	This causes <code>vi</code> to autoindent.
<code>:set noautoindent</code>	This causes <code>vi</code> to turn off autoindent.
<code>r</code> (then a character)	replace character under cursor
<code>~</code>	change case upper \rightarrow ;lower or lower \rightarrow upper

8.4 Cut and Paste

The vi editor has a space of memory called the *unstable buffer*, which we nickname Mabel. Mabel provides a temporary cache for holding things while we are editing and she is very helpful for doing quick copy-paste jobs.

This buffer is unstable because it loses its contents every time new text is placed in it. Do not use it to store things for a long time; instead write those things to files and retrieve them later. You will learn several ways to do this.

We show here a table with some cut, copy, and paste commands you will find helpful.

dd	Yank line to Mabel
dd	Delete line starting at the cursor; this cuts to Mabel
dw	Delete word; this cuts to Mabel
cw	Delete word, then enter insert mode(change word) The changed word is cut to Mabel.
p	Paste Mabel's contents at the cursor.
D	Cut line at cursor; this cuts the stricken text to Mabel
C	Cut line at cursor and enter insert mode; this cuts the stricken text to Mabel

All of these commands can be preceded by a number, and they will happen that number of times. For example typing `10yy` in command mode will yank ten lines, starting at the cursor, to Mabel. Since so many of these commands place new text in Mabel, you should know that if you copy or cut to Mabel and intend to use the text, paste it right away. You should open a file and experiment with these.

8.5 Cutting and Pasting with External Files

You can select a range of line numbers before each of these commands, or select in visual mode and use these commands.

<code>:w fileName</code>	Write to the file <code>fileName</code>
<code>:w! fileName</code>	Write selection to existing file <code>fileName</code> , and clobber it.
<code>:w >> fileName</code>	Append selection to file <code>fileName</code> .
<code>:r fileName</code>	Read in file <code>fileName</code> starting at the cursor

For example

```
:20,25 w foo.txt
```

will write lines 20-25 to the file `foo.txt`. If you want to write the entire file, omit the line numbers and that will happen. If you want to write from line 20 to the end of the file, the usage is as follows.

```
:20,$ w foo.txt
```

Note the use of `$` to mean “end of file.” When you learn about visual mode (just ahead), you can use these command to act on things you select in visual mode as well.

Housekeeping Tip If you use this facility, adopt a naming convention for these files you create on a short-term basis. When you are done editing, get rid of them or they become a choking kudzu and a source of confusion in your file system. Use names such as `buf`, `buf1`, etc as a signal to yourself that these files quickly outlive their usefulness and can be chucked.

8.6 Searching

Finally we shall look at search capabilities. These all belong to command mode. Enter

```
/someString
```

in command mode and `vi` will seek out the first instance of that string in the file or tell you it is not found. Type an `n` to find the next instance. You can enter

```
?someString
```

to search for `someString` backwards from the cursor. Type `n` to find the previous instance. Your machine may be configured to highlight every instance of the string you searched for. If you find this feature annoying, you can deactivate it with

```
:set nohlsearch
```

To do a global search and replace enter

```
:1,$ s/oldString/newString/g
```

to find every instance of `oldString` and convert it to `newString`. If you leave the `g` off the end, you will only change the first instance of `oldString` to `newString` on each line. (Think of `g` for global). You can do this on a range of line numbers by typing

```
<beginLine>,<endLine>; s/oldString/newString/g
```

One last thing bears mentioning. Do not use the scrollbar and the mouse to move around the screen. Use the arrow keys or go to a line number. If you use the scrollbar, you will see other parts of your terminal session, but if you attempt to edit, you will “snap” back into your file. You can use `$` for `<endLine>` to have the substitution carried out from the beginning line to the end of the file.

9 Visual Mode

The third mode of the `vi` editor, visual mode is actually three modes in one: line mode, character mode, and block mode. To enter line mode from command mode, hit `V`; to enter character mode hit `v`, and to enter block mode, hit `Control-v`. You can exit any of these by hitting the `ESC` key; this places you back in command mode. Visual mode has one purpose: it allows you to select text using keyboard commands; you may then perform various operations on these selections. First, let us see the selection mechanism at work.

Go into a file and position your cursor in the middle of a line. Hit `v` to enter visual character mode. Now use the arrow keys; notice how the selected text changes in response to arrow key movement. Try entering `gg` and `G` and see what happens. Hit `ESC` to finish. Now enter visual mode and use the `/` search facility to search up something on the page. What happens? Search backward and try that too.

Now enter visual line mode by hitting `V`; now try the keystrokes we just indicated and see how the selection behaves. This mode only selects whole lines.

Finally if you enter `Control-V` and you enter visual block mode, you can select a rectangular block of text from the screen by using the keyboard.

Now lets see what you can do with these selections. First let us look at character and line mode, as block mode behaves a little differently. You can delete the selected text by hitting `d`. You can yank it into Mabel by hitting `y`. Upon typing either command, you will be put back into command mode. Once any text is yanked into Mabel, you can paste it with `p` as you would any other text yanked there. If you hit `c`, the selection will be deleted and you will be in insert mode so you can change the text.

In block mode, things are a little different. If you hit `d`, the selected block will be deleted, and the lines containing it shortened. The stricken text is cut to Mabel. If you hit `y`, the block will be yanked just as in any other visual mode, and its line structure will be preserved. If you hit `c`, and enter text, the same change will be made on all line selected provided you do not hit the `ENTER` key. If you do, the change will only be carried out on the first line. You can insert text rather than change by hitting `I`, entering your text, and then hitting `ESC`. If the text you enter has no newline in it, the same text will be added to

each line; if it has a newline, only the first line is changed.

If you hit `r` then any character in any visual mode, all selected characters are changed to that character.

Here is a very common use for character or line visual mode. Suppose you are editing a document and the lines end in very jagged fashion. This sort of thing will commonly happen when maintaining a web or if you are editing a `LaTeX` document such as this one, where the page that is subjected to repeated edits. Use visual mode to select the affected paragraphs and hit `gq` (think Gentleman's Quarterly) and your paragraphs will be tidied up.

You can also do search-and-replace using visual mode to select the text to be acted upon. Simply select the text in visual mode. Then hit

```
: s/outText/inText/g
```

to perform the substitution in the selected text. For example if you select text in visual mode and change every `w` to a `v`, you will see this.

```
: '<,'>s/w/v/g
```

The `<,'>` is a quirky way of indicating you are doing a visual-mode search-replace operation.

10 Copy Paste from a GUI

You can copy and paste with the mouse in a window or between windows. The way you do it varies by OS so we will quickly discuss each. If you are pasting into a file you are editing with `vi`, it is a smart idea to use the colon command `:set paste`. This will prevent the “mad spraying” of text. For certain types of files, this turns off automatic indentation or formatting. You can use `:set nopaste` to turn off the paste mode.

Windoze If you are copying *from* a Windoze application into a terminal window, select the text you want to copy and use control-C in the usual way. This places the text in your Windoze system clipboard. Now go into your terminal window and get into insert mode where you want to paste. Right-click to paste the contents of your system clipboard into the terminal window. Many of you will say, “Why did the beginning of the text I copied get cut off or why didn't it appear at all?” This will occur if you are not in insert mode when you paste. It is important to be in insert mode before pasting to avoid unpleasant surprises. If this happens, hit `ESC` then `u` in command mode. The `u` command undoes the last `vi` command. Then you can take a fresh run at it.

If you are copying from a terminal window, select the text you wish to copy; PuTTY will place the text in your system clipboard. Then go into the window in which you wish to paste it. If the window is another terminal, get into insert mode and right-click on the mouse. If it's a Windoze app, use control-V as you usually do.

Mac Use apple-c to copy and apple-v to paste to or from a terminal window, just as you would with any other mac app. Mac gets this right.

Linux If you use a Linux box, use control-shift-C for copying in terminal windows and control-shift-V for pasting to terminal windows.

A Reprise: A Warning About autoindent and paste Before pasting with the mouse make sure you have autoindent turned off. Otherwise, your text will “go mad and spray everywhere,” especially if you are copying a large block of text with indents in it. You can turn autoindent on with `:set autoindent` and off with `:set noautoindent`. This feature can be convenient when editing certain types of files. You can use the command `:set paste` to turn off all smart indentation; when finished use `:set nopaste` to set things back to their original state.

A Warning about Line Numbers If you copy-paste to a GUI, line numbers will get copied. To prevent this from happening, use the colon command `:set nonu` before copying.

Experiment with these new techniques in some files. Deliberately make mistakes and see what happens. Then when you are editing files, you will know what to expect and how to recover.

There are a lot of excellent tutorials on vi on the web; avail yourself of these to learn more. Remember the most important thing: you never stop learning vi! Here are some useful vi resources on the web.

- The site [1] for is complete, organized and well-written. You can download the whole shebang in a PDF. Read this in little bits and try a few new tricks at a time.
- The site [?], vi for Smarties will introduce you to vi with a bit of churlish sneery attitude. It's pretty cool. And it's sneery like the author of this august volume.
- The link <ftp://ftp.vim.org/pub/vim/doc/book/vimbook-OPL.pdf> will download The Vim book for you. It is a very comprehensive guide, and it has excellent coverage on the visual mode.

10.1 Permissions

Now we will see how you can use permissions to control the visibility of your files on the system. You are the owner of your home directory and all directories and files it contains. This is your “subtree” of the systems directories belonging to you. You may grant, revoke or configure permissions for all the files and directories you own as you wish. UNIX was designed with the fundamental idea that your data are your property, and you can control what others see of them.

There are three layers of permission: you, your group, and others. You is letter **u**, your group is letter **g** and others is letter **o**. There are three types of permission for each of these: read, write and execute. Read means that level can read the file, write means that level can execute the file, and execute means that level can execute the file. In the example above, the file **bar** has the *permission string*

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

which means the following.

- You can read or write to the file. You cannot execute.
- Your group can read or write but not execute.
- The world can read this file but neither write nor execute.

For the user to execute this file, use the **chmod** command as follows

```
$ chmod u+x bar
```

The u(ser’s, that’s you) permission changed to allow the user to execute the file.

If you do not want the world to see this file you could enter

```
$ chmod o-r bar
```

and revoke permission for the world to see the file **bar**. Since you are the owner of the file, you have this right. In general you can do

```
$ chmod (u or g or o)(+ or -)(r or w or x) fileName
```

to manage permissions. You can omit the **u**, **g** or **o** and the permission will be added or deleted for all three categories. In the next subsection, we discuss the octal representation of the permissions string. This will allow you to change all three levels of permissions at once quickly and easily.

11 The Octal Representation for Permissions

There is also a numerical representation for permissions. This representation is a three-digit octal (base 8) number. Each permission has a number value as follows.

- The permission `r` has value 4.
- The permission `w` has value 2.
- The permission `x` has value 1.
- Lack of any permission has value 0.

We show how to translate a string in this example.

```
-rw-r--r--  
 6 4 4
```

The only way to get a sum of 6 from 1,2 and 4 is $4 + 2$. therefore 6 is read-write permission. The string translates into three digits 0-7; this file has 644 permissions. It is a simple exercise to look at all the digits 0-7 and see what permissions they convey.

We show some more examples of `chmod` at work. Look at how the permissions change in response to the `chmod` commands. Suppose we are a directory containing one file named `empty`, which has permission string `-rw-r--r--`, or 644. We begin by revoking the read permission from others.

```
$ chmod o-r empty
```

We now list the files in the directory

```
$ ls -l  
total 0  
-rwxr----- 1 morrison morrison 0 2008-08-26 10:52 empty  
$ ls  
empty
```

We can now restore the original permissions all at once by using the octal number representation for our permissions.

```
chmod 644 empty  
$ ls -l  
total 0  
-rw-r--r-- 1 morrison morrison 0 2008-08-26 10:52 empty
```

Notice what happens when we try to use a 9 for a permission string.

```
$ chmod 955 empty
chmod: invalid mode: 955
Try chmod --help for more information.
```

Try typing the `chmod --help` command at your prompt and it will show you some useful information about the `chmod` command. Almost all UNIX commands have this help feature.

Directories must have executable permissions, or they cannot be entered, and their contents are invisible. Here we use the `-a` option on `ls`. Notice that the current working directory and the directory above it have execute permissions at all levels. Try revoking execute permissions from one of your directories and attempt to enter it with `cd`; you will get a **Permission Denied** nastygram from the operating system.

```
$ ls -al
total 20
drwxr-xr-x 2 morrison faculty 4096 2008-10-17 11:51 .
drwx--x--x 9 morrison faculty 4096 2008-10-16 08:39 ..
-rw-r--r-- 1 morrison faculty 0 2008-10-17 11:51 empty
$
```

Here we shall do this so you can bear witness

```
$ mkdir fake
$ chmod u-x fake
$ cd fake
bash: cd: fake: Permission denied
$
```

Assigning 600 permissions to a file is a way to prevent anyone but yourself from seeing or modifying that file. It is a quick and useful way of hiding things from public view. Later, when you create a web page, you can use this command to hide files in your website that you do not want to be visible.

12 The Man

The command `man` is your friend. Type `man` then your favorite UNIX command to have its inner secrets exposed! For example, at the UNIX prompt, enter

```
$ man cat
```

This brings up the man(ual) page for the command `cat`. A complete list of options is furnished. Notice that some of these have the `--`, or long form.

CAT(1) User Commands CAT(1)

NAME

`cat` - concatenate files and print on the standard output

SYNOPSIS

`cat [OPTION] [FILE]...`

DESCRIPTION

Concatenate FILE(s), or standard input, to standard output.

`-A, --show-all`

equivalent to `-vET`

`-b, --number-nonblank`

number nonblank output lines

`-e` equivalent to `-vE`

`-E, --show-ends`

display `$` at end of each line

`-n, --number`

number all output lines

`-s, --squeeze-blank`

never more than one single blank line

`-t` equivalent to `-vT`

`-T, --show-tabs`

display TAB characters as `^I`

`-u` (ignored)

`-v, --show-nonprinting`

use `^` and `M-` notation, except for LFD and TAB

`--help` display this help and exit

`--version`

output version information and exit

With no FILE, or when FILE is -, read standard input.

EXAMPLES

```
cat f g
    Output f's contents, then standard input,
    then g's contents.
```

```
cat    Copy standard input to standard output.
```

AUTHOR

Written by Torbjorn Granlund and Richard M. Stallman.

REPORTING BUGS

Report bugs to <bug-coreutils@gnu.org>.

COPYRIGHT

Copyright 2006 Free Software Foundation, Inc.
This is free software. You may redistribute
copies of it under the terms of the GNU General
Public License <<http://www.gnu.org/licenses/gpl.html>>.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

The full documentation for cat is maintained
as a Texinfo manual. If the info and cat
programs are properly installed at your site,
the command info cat should give you access
to the complete manual.

cat 5.97 August 2006 CAT(1)

You can see here that even humble cat has some options to enhance its
usefulness. Here is cat at work on a file named trap.py.

```
$ cat trap.py
def trap(a, b, n, f):
    a = float(a)
    b = float(b)
    h = (b - a)/n
    list = map(lambda x: a + h*x, range(0,n+1))
    tot = .5*(f(a) + f(b))
    tot += sum(map(f, list[1:n]))
    tot *= h
return tot
def f(x):
    return x*x
print trap(0,1,10,f)
```

```
print trap(1,2,100,f)
```

Using the `-n` option causes the output to have line numbers.

```
cat -n trap.py
 1 def trap(a, b, n, f):
 2     a = float(a)
 3     b = float(b)
 4     h = (b - a)/n
 5     list = map(lambda x: a + h*x, range(0,n+1))
 6     tot = .5*(f(a) + f(b))
 7     tot += sum(map(f, list[1:n]))
 8     tot *= h
 9     return tot
10 def f(x):
11     return x*x
12 print trap(0,1,10,f)
13 print trap(1,2,100,f)
$
```

View the manual pages on commands such as `rm`, `ls`, `chmod` and `cp` to learn more about each command. Experiment with the options you see there on some junky files you create and do not care about losing.

Exercises

1. Use the `man` command to learn about the UNIX commands `more` and `less`. You will see here, that in fact, `less` is more!
2. Use the `man` command to learn about the UNIX commands `head` and `tail`. Can you create a recipe to get the first and last lines of a file?
3. What does the `ls -R` command do?

13 Redirection of Standard Output and Standard Input

UNIX treats everything in your system as a file; this includes all devices such as printers, mice and the keyboard. Things put to the screen are generally put to one of two files, `stdout`, or *standard output* and `stderr`, or standard error. It is easy to redirect standard output to a another file.

The keyboard, by default, is represented by the file `stdin`, or *standard input*. It is also possible to redirect standard input and take standard input from another file.

Sometimes a UNIX command or a program puts a large quantity of text to the screen; redirection allows you to capture the results into a file. You can open this file with `vi`, search it, or edit it. The examples here are based on the files `animalNoises.txt`

```
miao
bleat
moo
```

and `physics.txt`

```
snape
benetton
stephan
```

Create each of these files in a directory so you can follow the examples here. First we show how `cat` puts files to `stdout`.

```
$ cat animalNoises.txt physics.txt
miao
bleat
moo
snape
benetton
stephan
```

Now let us capture this critical information into the file `stuff.txt` by redirecting `stdout`. We then use `cat` to display the resulting file to `stdout`.

```
$ cat animalNoises.txt physics.txt > stuff.txt
$ cat stuff.txt
miao
bleat
moo
snape
benetton
stephan
```

The `cat` command has a second guise. It accepts a file name as an argument, but it will also accept standard input; this is no surprise since `stdin` is treated as a file. At the UNIX command line enter

```
$ cat
```

The `cat` program is now running and it awaits word from `stdin`. Enter some text and then hit the enter key; `cat` echoes back the text you type in. To finish, hit control-d (end-of-file).

```
$ cat
me too
me too
ditto
ditto
$
```

The control-d puts no response to the screen. You can also put a file to the screen with

```
$ cat < someFile
```

Here, the file `someFile` becomes `stdin` for the `cat` command. This phenomenon is shown in the man page for `cat`. Under the description of the command it says, “Concatenate FILE(s), or standard input, to standard output.”

Let us now come back to `stdout`. Next create a new file named `sheck.txt` with these contents.

```
roach
stag beetle
tachnid wasp
```

Were we to invoke the command

```
$ cat animalNoises.txt physics.txt > sheck.txt
```

we would clobber the file `sheck.txt` and lose its valuable contents. This may be our intent; if so very well. If we want to add new information to the end of the file we use the `>>` *append operator* to append it to the end of the receiving file. If we do this

```
$ cat animalNoises.txt physics.txt >> sheck.txt
```

we get the following result if we use the original file `sheck.txt`.

```
$ cat sheck.txt
roach
stag beetle
tachnid wasp
```

```
miao
bleat
moo
snape
benetton
stephan
```

The `>>` redirection operator will automatically create a file for you if the file to which you are redirecting does not already exist.

14 Using Pipes

It is very common to want to use `stdout` from one command to be `stdin` for another command. This will grant us the ability to chain the actions of the existing filters we have `cat`, `more` and `less` with some new filters to do a wide variety of tasks. To achieve this, we use a device called a *pipe*. Pipes allow you to chain the action of various UNIX commands. We shall add to our palette of UNIX commands to give ourselves a larger and more interesting collection of examples. These commands are extremely useful for manipulating files of data.

14.1 The sort filter

Bring up the `man` page for the command `sort`. This command accepts a file (or `stdin`) and it sorts the lines in the file.

This begs the question: how does it sort? It sorts alphabetically in a case-insensitive manner, and it “alphabetizes” non-alphabetical characters by ASCII value. The `sort` command several four helpful options.

<code>-b</code>	<code>--ignore-leading-blanks</code>	ignores leading blanks
<code>-d</code>	<code>--dictionary-order</code>	pays heed to alphanumeric characters and blanks and ignores other characters
<code>-f</code>	<code>--ignore-cases</code>	ignores case
<code>-r</code>	<code>--reverse</code>	reverses comparisons

Here we put the command to work with `stdin`; use a control-d on its own line to get the prettiest format. Here we put the items `moose`, `jaguar`, `cat` and `katydid` each on its own line into `stdin`. Without comment, a sorted list is produced.

```
$ sort -f
```

```
moose
jaguar
cat
katydid      (now hit control-d)
cat
jaguar
katydid
moose
$
```

You should try various lists with different options on the `sort` command to see how it works for yourself. You can also run `sort` on a file and send a sorted copy of the file to `stdout`. Of course, you can redirect this result into a file using `>` or `>>`.

14.2 The Filters `head`, `tail`, and `uniq`

The commands `head` and `tail` put the top or bottom of a file to `stdout`; the default amount is 10 lines. To show the first 5 lines of the file `foo.txt`, enter the following at the UNIX command line.

```
$ head -5 foo.txt
```

You can do exactly the same thing with `tail` with an entirely predictable result. The command `uniq` weeds out consecutive duplicate lines in a file, leaving only the first copy in place. These three commands have many useful options; explore them in the man pages.

14.3 The `grep` Command

This command is incredibly powerful; here we will just scratch the surface of its protean powers. You can search and filter files using `grep`; it can be used to search for needles in haystacks. In its most basic form `grep` will inspect a file line-by-line and put all lines to `stdout` containing a specified string. Here is a sample session.

```
$ grep gry /usr/share/dict/words
angry
hungry
$
```

The file `/usr/share/dict/words` is a dictionary file containing a list of words, one word to a line in (mostly) lower-case characters. Here we are searching

the dictionary for all lines containing the character sequence gry; the result is the two words angry and hungry. There is an option `-i` to ignore the case of alphabetical characters.

14.4 And Now to Pipes

Pipes allow you to feed `stdout` from one command into `stdin` to another without creating any temporary files yourself. Pipes can be used along with redirection of `stdin` and `stdout` to accomplish a huge array of text-processing chores.

Now let us do a practical example. Suppose we want to print the first 5 lines alphabetically in a file named `sampleFile.txt`. We know that `sort` will sort the file asciicographically; we will use the `-i` option to ignore case. The command `head -5` will print the first five lines of a file passed it or the first five lines of `stdin`. So, what we want to do is sort the file ignoring case, and pass the result to `head -5` to print out the top five lines. You join two processes with a pipe; it is represented by the symbol `|`, which is found by hitting the shift key and the key just above the enter key on a standard keyboard. Our command would be

```
$ sort -i sampleFile.txt | head -5
```

The pipe performs two tasks. It redirects the output of `sort -i` into a temporary buffer and then it feeds the contents of the buffer as standard input to `head -5`. The result: the first five lines in the alphabet in the file `sampleFile.txt` are put to `stdout`.

Suppose you wanted to save the results in a file named `results.txt`. To do this, redirect `stdout` as follows

```
$ (sort -i sampleFile.txt | head -5) > results.txt
```

Note the use of defensive parentheses to make our intent explicit. We want the five lines prepared, then stored in the file `results.txt`.

Programming Exercises Here are some other commands: `wc`, `echo`. You will use the `man` pages to determine their action and to use them to solve the problems below.

1. Tell how to put the text “Cowabunga, Turtle soup!” to `stdout`.
2. Tell how to get the text “This is written in magic ink” into a text file without using a text editor of any kind.
3. The `ls` command has an option `-R`, for “list files recursively.” This lists all of the sub-directories and all of their contents within the directory being listed. Use this command along with `grep` to find a file containing a specified string in a file path.

4. Put a list of names in a file in `lastName, firstName` format. Put them in any old order and put in duplicates. Use pipes to eliminate duplicates in this file and sort the names in alphabetical order.
5. Find the word in the system dictionary occupying line 10000.
6. How do you count all of the words in the system dictionary containing the letter `x`?
7. Find all words in the system dictionary occupying lines 50000-50500.
8. Tell how, in one line, to take the result of the previous exercise, place it in reverse alphabetical order and store in in a file named `myWords.txt`.
9. Tell how to display these w

References

- [1] University of Hawai'i at Manoa College of Engineering. *Mastering the vi Editor*. <http://www.eng.hawaii.edu/Tutor/vi.html>.

References

- [1] University of Hawai'i at Manoa College of Engineering. *Mastering the vi Editor*. <http://www.eng.hawaii.edu/Tutor/vi.html>.