

Contents

3	Linux Tutorial	3
3.1	Introduction.	3
3.2	Basic Linux concepts.	3
3.2.1	Creating an account.	3
3.2.2	Logging in.	4
3.2.3	Virtual consoles.	4
3.2.4	Shells and commands.	4
3.2.5	Logging out.	5
3.2.6	Changing your password.	5
3.2.7	Files and directories.	6
3.2.8	The directory tree.	6
3.2.9	The current working directory.	7
3.2.10	Referring to home directories.	7
3.3	First steps into Linux.	8
3.3.1	Moving around.	8
3.3.2	Looking at the contents of directories.	9
3.3.3	Creating new directories.	10
3.3.4	Copying files.	11
3.3.5	Moving files.	11
3.3.6	Deleting files and directories.	11
3.3.7	Looking at files.	11
3.3.8	Getting online help.	12
3.4	Accessing MS-DOS files.	12
3.5	Summary of basic UNIX commands.	13
3.6	Exploring the file system.	15
3.7	Types of shells.	18
3.8	Wildcards.	19
3.9	Linux plumbing.	21
3.9.1	Standard input and standard output.	21
3.9.2	Redirecting input and output.	22
3.9.3	Using pipes.	23
3.10	File permissions.	24
3.10.1	Concepts of file permissions.	24
3.10.2	Interpreting file permissions.	25
3.10.3	Permissions Dependencies.	26
3.10.4	Changing permissions.	26
3.11	Managing file links.	26
3.11.1	Hard links.	27
3.12	Job control.	28
3.12.1	Jobs and processes.	28
3.12.2	Foreground and background.	29
3.12.3	Backgrounding and killing jobs.	29
3.12.4	Stopping and restarting jobs.	31

Linux Installation and Getting Started

from: www.linuxdoc.org/LDP/gs/gs.html

Matt Welsh
Phil Hughes
David Bandel
Boris Beletsky
Sean Dreilinger
Robert Kiesling
Evan Liebovitch
Henry Pierce

Names of all products herein are used for identification purposes only and are trademarks and/or registered trademarks of their respective owners. Specialized Systems Consultants, Inc., makes no claim of ownership or corporate association with the products or companies that own them.

Copyright 1992-1996 Matt Welsh

Copyright 1998 Specialized Systems Consultants, Inc (SSC)

P.O. Box 55549
Seattle, WA 98155-0549
USA
Phone: +1-206-782-7733
Fax: +1-206-782-7191
E-mail: ligs@ssc.com
URL: <http://www.ssc.com/>

Linux Installation and Getting Started is a free document; you may reproduce and/or modify it under the terms of version 2 (or, at your option, any later version) of the GNU General Public License as published by the Free Software Foundation.

This book is distributed in the hope it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details, in Appendix C.

The authors encourage wide distribution of this book for personal or commercial use, provided the above copyright notice remains intact and the method adheres to the provisions of the GNU General Public License (see Appendix C). In summary, you may copy and distribute this book free of charge or for a profit. No explicit permission is required from the author for reproduction of this book in any medium, physical or electronic.

Note, derivative works and translations of this document must be placed under the GNU General Public License, and the original copyright notice must remain intact. If you have contributed new material to this book, you must make the source code (e.g., LaTeX source) available for your revisions. Please make revisions and updates available directly to the document maintainers, Specialized Systems Consultants. This will allow for the merging of updates and provide consistent revisions to the Linux community.

If you plan to publish and distribute this book commercially, donations, royalties, and/or printed copies are greatly appreciated by the authors and the Linux Documentation Project. Contributing in this way shows your support for free software and the Linux Documentation Project. If you have questions or comments, please contact SSC.

3 Linux Tutorial

from: www.linuxdoc.org/LDP/gs/gs.html

3.1 Introduction.

If you're new to UNIX and Linux, you may be a bit intimidated by the size and apparent complexity of the system before you. This chapter does not go into great detail or cover advanced topics. Instead, we want you to hit the ground running.

We assume very little here about your background, except perhaps that you have some familiarity with personal computer systems, and MS-DOS. However, even if you're not an MS-DOS user, you should be able to understand everything here. At first glance, Linux looks a lot like MS-DOS—after all, parts of MS-DOS were modeled on the CP/M operating system, which in turn was modeled on UNIX. However, only the most superficial features of Linux resemble MS-DOS. Even if you're completely new to the PC world, this tutorial should help.

And, before we begin: *Don't be afraid to experiment.* The system won't bite you. You can't destroy anything by working on the system. Linux has built-in security features to prevent “normal” users from damaging files that are essential to the system. Even so, the worst thing that can happen is that you may delete some or all of your files and you'll have to re-install the system. So, at this point, you have nothing to lose.

3.2 Basic Linux concepts.

Linux is a multitasking, multiuser operating system, which means that many people can run many different applications on one computer at the same time. This differs from MS-DOS, where only one person can use the system at any one time. Under Linux, to identify yourself to the system, you must **log in**, which entails entering your **login name** (the name the system uses to identify you), and entering your **password**, which is your personal key for logging in to your account. Because only you know your password, no one else can log in to the system under your user name.

On traditional UNIX systems, the system administrator assigns you a user name and an initial password when you are given an account on the system. However, because in Linux *you* are the system administrator, you must set up your own account before you can log in. For the following discussions, we'll use the imaginary user name, “larry.”

In addition, each system has a **host name** assigned to it. It is this host name that gives your machine a name, gives it character and charm. The host name is used to identify individual machines on a network, but even if your machine isn't networked, it should have a host name. For our examples below, the system's host name is “mousehouse”.

3.2.1 Creating an account.

Before you can use a newly installed Linux system, you must set up a user account for yourself. It's usually not a good idea to use the root account for normal use; you should reserve the root account for running privileged commands and for maintaining the system as discussed below.

In order to create an account for yourself, log in as root and use the `useradd` or `adduser` command. See Section 4.6 for information on this procedure.

3.2.2 Logging in.

At login time, you'll see a prompt resembling the following:

```
mousehouse login:
```

Enter your user name and press the `Enter` key. Our hero, larry, would type:

```
mousehouse login: larry
Password:
```

Next, enter your password. The characters you enter won't be echoed to the screen, so type carefully. If you mistype your password, you'll see the message

```
Login incorrect
```

and you'll have to try again.

Once you have correctly entered the user name and password, you are officially logged in to the system, and are free to roam.

3.2.3 Virtual consoles.

The system's console is the monitor and keyboard connected directly to the system. (Because Linux is a multiuser operating system, you may have other terminals connected to serial ports on your system, but these would not be the console.) Linux, like some other versions of UNIX, provides access to **virtual consoles** (or VCs), that let you have more than one login session on the console at one time.

To demonstrate this, log in to your system. Next, press `Alt-F2`. You should see the login: prompt again. You're looking at the second virtual console. To switch back to the first VC, press `Alt-F1`. *Voila!* You're back to your first login session.

A newly-installed Linux system probably lets you to access only the first half-dozen or so VCs, by pressing `Alt-F1` through `Alt-F4`, or however many VCs are configured on your system. It is possible to enable up to 12 VCs—one for each function key on your keyboard. As you can see, use of VCs can be very powerful because you can work in several different sessions at the same time.

While the use of VCs is somewhat limiting (after all, you can look at only one VC at a time), it should give you a feel for the multiuser capabilities of Linux. While you're working on the first VC, you can switch over to the second VC and work on something else.

3.2.4 Shells and commands.

For most of your explorations in the world of Linux, you'll be talking to the system through a **shell**, a program that takes the commands you type and translates them into instructions to the operating system. This can be compared to the COMMAND.COM program under MS-DOS, which does essentially the same thing. A shell is just one interface to Linux. There are many possible interfaces—like the X Window System, which lets you run commands by using the mouse and keyboard.

As soon as you log in, the system starts the shell, and you can begin entering commands. Here's a quick example. Larry logs in and is waiting at the shell **prompt**.

```
mousehouse login: larry
Password: larry's password
Welcome to Mousehouse!
```

```
/home/larry#
```

The last line of this text is the shell's prompt, indicating that it's ready to take commands. (More on what the prompt itself means later.) Let's try telling the system to do something interesting:

```
/home/larry# make love
make: *** No way to make target 'love'. Stop.
/home/larry#
```

Well, as it turns out, `make` is the name of an actual program on the system, and the shell executed this program when given the command. (Unfortunately, the system was being unfriendly.)

This brings us to the burning question: What is a command? What happens when you type “`make love`”? The first word on the command line, “`make`”, is the name of the command to be executed. Everything else on the command line is taken as arguments to this command. Example:

```
/home/larry# cp foo bar
```

The name of this command is “`cp`”, and the arguments are “`foo`” and “`bar`”.

When you enter a command, the shell does several things. First, it checks the command to see if it is internal to the shell. (That is, a command which the shell knows how to execute itself. There are a number of these commands, and we'll go into them later.) The shell also checks to see if the command is an alias, or substitute name, for another command. If neither of these conditions apply, the shell looks for a program, on disk, having the specified name. If successful, the shell runs the program, sending the arguments specified on the command line.

In our example, the shell looks for a program called `make`, and runs it with the argument `love`. `Make` is a program often used to compile large programs, and takes as arguments the name of a “target” to compile. In the case of “`make love`”, we instructed `make` to compile the target `love`. Because `make` can't find a target by this name, it fails with a humorous error message, and returns us to the shell prompt.

What happens if we type a command to a shell and the shell can't find a program having the specified name? Well, we can try the following:

```
/home/larry# eat dirt!
eat: command not found
/home/larry#
```

Quite simply, if the shell can't find a program having the name given on the command line (here, “`eat`”), it prints an error message. You'll often see this error message if you mistype a command (for example, if you had typed “`mkae love`” instead of “`make love`”).

3.2.5 Logging out.

Before we delve much further, we should tell you how to log out of the system. At the shell prompt, use the command

```
/home/larry# exit
```

to log out. There are other ways of logging out, but this is the most foolproof one.

3.2.6 Changing your password.

You should also know how to change your password. The command `passwd` prompts you for your old password, and a new password. It also asks you to reenter the new password for validation. Be careful not to forget your password—if you do, you will have to ask the system administrator to reset it for you. (If you are the system administrator, see page .)

3.2.7 Files and directories.

Under most operating systems (including Linux), there is the concept of a **file**, which is just a bundle of information given a name (called a **filename**). Examples of files might be your history term paper, an e-mail message, or an actual program that can be executed. Essentially, anything saved on disk is saved in an individual file.

Files are identified by their file names. For example, the file containing your history paper might be saved with the file name `history-paper`. These names usually identify the file and its contents in some form that is meaningful to you. There is no standard format for file names as there is under MS-DOS and some other operating systems; in general, a file name can contain any character (except the `/` character—see the discussion of path names, below) and is limited to 256 characters in length.

With the concept of files comes the concept of **directories**. A directory is a collection of files. It can be thought of as a “folder” that contains many different files. Directories are given names, with which you can identify them. Furthermore, directories are maintained in a tree-like structure; that is, directories may contain other directories.

Consequently, you can refer to a file by its **path name**, which is made up of the filename, preceded by the name of the directory containing the file. For example, let’s say that Larry has a directory called `papers`, which contains three files: `history-final`, `english-lit`, and `masters-thesis`. Each of these three files contains information for three of Larry’s ongoing projects. To refer to the `english-lit` file, Larry can specify the file’s pathname, as in:

```
papers/english-lit
```

As you can see, the directory and filename are separated by a single slash (`/`). For this reason, filenames themselves cannot contain the `/` character. MS-DOS users will find this convention familiar, although in the MS-DOS world the backslash (`\`) is used instead.

As mentioned, directories can be nested within each other as well. For example, let’s say that there is another directory within `papers`, called `notes`. The `notes` directory contains the files `math-notes` and `cheat-sheet`. The pathname of the file `cheat-sheet` would be

```
papers/notes/cheat-sheet
```

Therefore, a path name is really like a path to the file. The directory that contains a given subdirectory is known as the **parent directory**. Here, the directory `papers` is the parent of the `notes` directory.

3.2.8 The directory tree.

Most Linux systems use a standard layout for files so that system resources and programs can be easily located. This layout forms a directory tree, which starts at the “`/`” directory, also known as the “root directory”. Directly underneath `/` are important subdirectories: `/bin`, `/etc`, `/dev`, and `/usr`, among others. These directories in turn contain other directories which contain system configuration files, programs, and so on.

In particular, each user has a **home directory**, which is the directory set aside for that user to store his or her files. In the examples above, all of Larry’s files (like `cheat-sheet` and `history-final`) are contained in Larry’s home directory. Usually, user home directories are contained under `/home`, and are named for the user owning that directory. Larry’s home directory is `/home/larry`.

The diagram on page shows a sample directory tree, which should give you an idea of how the directory tree on your system is organized.

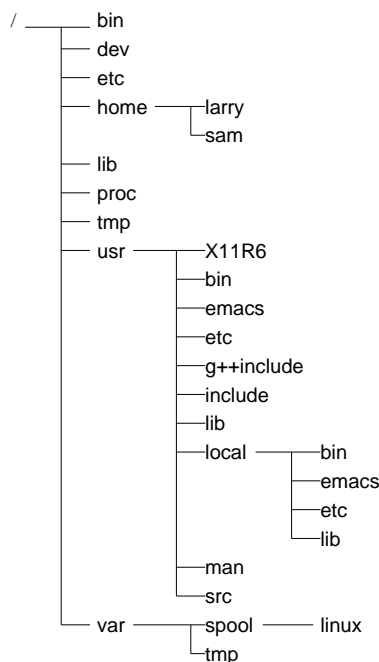


Figure 3.1: A typical (abridged) Linux directory tree.

3.2.9 The current working directory.

At any moment, commands that you enter are assumed to be relative to your **current working directory**. You can think of your working directory as the directory in which you are currently “located”. When you first log in, your working directory is set to your home directory—`/home/larry`, in our case. Whenever you refer to a file, you may refer to it in relationship to your current working directory, rather than specifying the full pathname of the file.

Here’s an example. Larry has the directory `papers`, and `papers` contains the file `history-final`. If Larry wants to look at this file, he can use the command

```
/home/larry# more /home/larry/papers/history-final
```

The `more` command simply displays a file, one screen at a time. However, because Larry’s current working directory is `/home/larry`, he can instead refer to the file *relative* to his current location by using the command

```
/home/larry# more papers/history-final
```

If you begin a filename (like `papers/final`) with a character other than `/`, you’re referring to the file in terms relative to your current working directory. This is known as a **relative path name**.

On the other hand, if you begin a file name with a `/`, the system interprets this as a full path name—that is, a path name that includes the entire path to the file, starting from the root directory, `/`. This is known as an absolute path name.

3.2.10 Referring to home directories.

Under both `tcsh` and `bash`¹ you can specify your home directory with the tilde character (`~`). For example, the command

¹`tcsh` and `bash` are two *shells* that run under Linux. The shell is a program that reads user commands and executes them; most Linux systems enable either `tcsh` or `bash` for new user accounts.

```
/home/larry# more ~/papers/history-final
```

is equivalent to

```
/home/larry# more /home/larry/papers/history-final
```

The shell replaces the `~` character with the name of your home directory.

You can also specify other user's home directories with the tilde character. The pathname `~karl/letters` translates to `/home/karl/letters` by the shell (if `/home/karl` is karl's home directory). Using a tilde is simply a shortcut; there is no directory named `~` — it's just syntactic sugar provided by the shell.

3.3 First steps into Linux.

Before we begin, it is important to know that all file and command names on a Linux system are case-sensitive (unlike operating systems such as MS-DOS). For example, the command `make` is very different from `Make` or `MAKE`. The same is true for file and directory names.

3.3.1 Moving around.

Now that you can log in, and you know how to refer to files using pathnames, how can you change your current working directory, to make life easier?

The command for moving around in the directory structure is `cd`, which is short for “change directory”. Many often-used Linux commands are two or three letters. The usage of the `cd` command is

```
cd directory
```

where *directory* is the name of the directory which you wish to become the current working directory.

As mentioned earlier, when you log in, you begin in your home directory. If Larry wanted to switch to the `papers` subdirectory, he'd use the command

```
/home/larry# cd papers
/home/larry/papers#
```

As you can see, Larry's prompt changes to reflect his current working directory (so he knows where he is). Now that he's in the `papers` directory, he can look at his history final with the command

```
/home/larry/papers# more history-final
```

Now, Larry is stuck in the `papers` subdirectory. To move back up to the next higher (or parent) directory, use the command

```
/home/larry/papers# cd ..
/home/larry#
```

(Note the space between the “`cd`” and the “`..`”.) Every directory has an entry named “`..`” which refers to the parent directory. Similarly, every directory has an entry named “`.`” which refers to itself. Therefore, the command

```
/home/larry/papers# cd .
```

gets us nowhere.

You can also use absolute pathnames with the `cd` command. To `cd` into Karl's home directory, we can use the command

```
/home/larry/papers# cd /home/karl
/home/karl#
```

Also, using `cd` with no argument will return you to your own home directory.

```
/home/karl# cd
/home/larry#
```

3.3.2 Looking at the contents of directories.

Now that you know how to move around directories, you might think, "So what?" Moving around directories is fairly useless by itself, so let's introduce a new command, `ls`. The `ls` command displays a listing of files and directories, by default from your current directory. For example:

```
/home/larry# ls
Mail
letters
papers
/home/larry#
```

Here we can see that Larry has three entries in his current directory: `Mail`, `letters`, and `papers`. This doesn't tell us much—are these directories or files? We can use the `-F` option of the `ls` command to get more detailed information.

```
/home/larry# ls -F
Mail/
letters/
papers/
/home/larry#
```

From the `/` appended to each filename, we know that these three entries are in fact subdirectories.

Using `ls -F` may also append `"*"` to the end of a filename in the resulting list which would indicate that the file is an **executable**, or a program which can be run. If nothing is appended to the filename using `ls -F`, the file is a "plain old file", that is, it's neither a directory nor an executable.

In general, each UNIX command may take a number of options in addition to other arguments. These options usually begin with a `"-"`, as demonstrated above with the `-F` option. The `-F` option tells `ls` to give more information about the type of the files involved—in this case, printing a `/` after each directory name.

If you give `ls` a directory name, the system will print the contents of that directory.

```
/home/larry# ls -F papers
english-lit
history-final
masters-thesis
notes/
/home/larry#
```

Or, for a more interesting listing, let's see what's in the system's `/etc` directory.

```
/home/larry# ls /etc
Images      ftpusers    lpc          rc.new       shells
adm         getty       magic        rc0.d        startcons
bcheckrc    gettydefs   motd         rc1.d        swapoff
brc         group       mount        rc2.d        swapon
brc         inet        mtab         rc3.d        syslog.conf
csh.cshrc   init        mtools       rc4.d        syslog.pid
csh.login   init.d      pac          rc5.d        syslogd.reload
default     initrunlvl  passwd       rmt          termcap
disktab     inittab     printcap     rpc          umount
fdprm       inittab.old profile       rpcinfo      update
fstab       issue       psdatabase   securetty    utmp
ftppaccess  lilo        rc           services     wtmp

/home/larry#
```

If you're a MS-DOS user, you may notice that the filenames can be longer than 8 characters, and can contain periods in any position. You can even use more than one period in a filename.

Let's move to the top of the directory tree, and then down to another directory with the commands

```
/home/larry# cd ..
/home# cd ..
/# cd usr
/usr# cd bin
/usr/bin#
```

You can also move into directories in one step, as in `cd /usr/bin`.

Try moving around various directories, using `ls` and `cd`. In some cases, you may run into the foreboding "Permission denied" error message. This is simply UNIX security kicking in: in order to use the `ls` or `cd` commands, you must have permission to do so. We talk more about this in section 3.10 File Permissions.

3.3.3 Creating new directories.

It's time to learn how to create directories. This involves the use of the `mkdir` command. Try the following:

```
/home/larry# mkdir foo
/home/larry# ls -F
Mail/
foo/
letters/
papers/
/home/larry# cd foo
/home/larry/foo# ls
/home/larry/foo#
```

Congratulations! You made a new directory and moved into it. Since there aren't any files in this new directory, let's learn how to copy files from one place to another.

3.3.4 Copying files.

To copy files, use the command `cp`, as shown here:

```
/home/larry/foo# cp /etc/termcap .
/home/larry/foo# cp /etc/shells .
/home/larry/foo# ls -F
shells      termcap
/home/larry/foo# cp shells bells
/home/larry/foo# ls -F
bells      shells      termcap
/home/larry/foo#
```

The `cp` command copies the files listed on the command line to the file or directory given as the last argument. Notice that we use “.” to refer to the current directory.

3.3.5 Moving files.

The `mv` command moves files, rather than copying them. The syntax is very straightforward:

```
/home/larry/foo# mv termcap sells
/home/larry/foo# ls -F
bells      sells      shells
/home/larry/foo#
```

Notice that the `termcap` file has been renamed `sells`. You can also use the `mv` command to move a file to a completely new directory.

- ◇ **Note:** `mv` and `cp` will overwrite a destination file having the same name without asking you. Be careful when you move a file into another directory. There may already be a file having the same name in that directory, which you’ll overwrite!

3.3.6 Deleting files and directories.

You now have an ugly rhyme developing with the use of the `ls` command. To delete a file, use the `rm` command, which stands for “remove”, as shown here:

```
/home/larry/foo# rm bells sells
/home/larry/foo# ls -F
shells
/home/larry/foo#
```

We’re left with nothing but shells, but we won’t complain. Note that `rm` by default won’t prompt you before deleting a file—so be careful.

A related command to `rm` is `rmdir`. This command deletes a directory, but only if the directory is empty. If the directory contains any files or subdirectories, `rmdir` will complain.

3.3.7 Looking at files.

The commands `more`, `less`, and `cat` are used for viewing the contents of files. `more` displays a file, one screenful at a time, while `cat` displays the whole file at once.

To look at the file `shells`, use the command

```
/home/larry/foo# more shells
```

In case you're interested what `shells` contains, it's a list of valid shell programs on your system. On most systems, this includes `/bin/sh`, `/bin/bash`, and `/bin/csh`. We'll talk about these different types of shells later.

While using `more`, press `Space` to display the next page of text, and `b` to display the previous page. There are other commands available in `more` as well, these are just the basics. Pressing `q` will quit `more`.

Quit `more` and try `cat /etc/termcap`. The text will probably fly by too quickly for you to read it all. The name “`cat`” actually stands for “`concatenate`”, which is the real use of the program. The `cat` command can be used to concatenate the contents of several files and save the result to another file. This will be seen again in section 3.14.1.

3.3.8 Getting online help.

Almost every UNIX system, including Linux, provides a facility known as **manual pages**. These manual pages contain online documentation for system commands, resources, configuration files and so on.

The command used to access manual pages is `man`. If you're interested in learning about other options of the `ls` command, you can type

```
/home/larry# man ls
```

and the manual page for `ls` will be displayed.

Unfortunately, most manual pages are written for those who already have some idea of what the command or resource does. For this reason, manual pages usually contain only the technical details of the command, without much explanation. However, manual pages can be an invaluable resource for jogging your memory if you forget the syntax of a command. Manual pages will also tell you about commands that we don't cover in this book.

I suggest that you try `man` for the commands that we've already gone over and whenever I introduce a new command. Some of these commands won't have manual pages, for several reasons. First, the manual pages may not have been written yet. (The Linux Documentation Project is responsible for manual pages under Linux as well. We are gradually accumulating most of the manual pages available for the system.) Second, the the command might be an internal shell command, or an alias (discussed in Section 3.2.4), which would not have a manual page of its own. One example is `cd`, which is an internal shell command. The shell itself actually processes the `cd`—there is no separate program that implements this command.

3.4 Accessing MS-DOS files.

If, for some twisted and bizarre reason, you want to access files from MS-DOS, it's easily done under Linux.

The usual way to access MS-DOS files is to mount an MS-DOS partition or floppy under Linux, allowing you to access the files directly through the file system. For example, if you have an MS-DOS floppy in `/dev/fd0`, the command

```
# mount -t msdos /dev/fd0 /mnt
```

will mount it under `/mnt`. See Section 4.8.4 for more information on mounting floppies.

You can also mount an MS-DOS partition of your hard drive for access under Linux. If you have an MS-DOS partition on `/dev/hda1`, the command

```
# mount -t msdos /dev/hda1 /mnt
```

mounts it. Be sure to **umount** the partition when you're done using it. You can have a MS-DOS partition automatically mounted at boot time if you include the entry in `/etc/fstab`. See Section 4.4 for details. The following line in `/etc/fstab` will mount an MS-DOS partition on `/dev/hda1` on the directory `/dos`.

```
/dev/hda1    /dos    msdos    defaults
```

You can also mount the VFAT file systems that are used by Windows 95:

```
# mount -t vfat /dev/hda1 /mnt
```

This allows access to the long filenames of Windows 95. This only applies to partitions that actually have the long filenames stored. You can't mount a normal FAT16 file system and use this to get long filenames.

The Mtools software may also be used to access MS-DOS files. The commands `mcd`, `mdir`, and `mcopy` all behave like their MS-DOS counterparts. If you install Mtools, there should be manual pages available for these commands.

Accessing MS-DOS files is one thing; running MS-DOS programs is another. There is an MS-DOS Emulator under development for Linux; it is widely available, and included in most distributions. It can also be retrieved from a number of locations, including the various Linux FTP sites listed in Appendix B. The MS-DOS Emulator is reportedly powerful enough to run a number of applications, including WordPerfect, from Linux. However, Linux and MS-DOS are vastly different operating systems. The power of any MS-DOS emulator under UNIX is limited. In addition, a Microsoft Windows emulator that runs under X Windows is under development.

3.5 Summary of basic UNIX commands.

This section introduces some of the most useful basic commands of a UNIX system, including those that are covered in the previous section.

Note that options usually begin with “-”, and in most cases you can specify more than one option with a single “-”. For example, rather than use the command `ls -l -F`, you can use `ls -lF`.

Rather than listing every option for each command, we only present useful or important commands at this time. In fact, most of these commands have many options that you'll never use. You can use `man` to see the manual pages for each command, which list all of the available options.

Also note that many of these commands take as arguments a list of files or directories, denoted in this table by “`file1 ...fileN`”. For example, the `cp` command takes as arguments a list of files to copy, followed by the destination file or directory. When copying more than one file, the destination must be a directory.

`cd` Change the current working directory

Syntax: `cd directory`

Where *directory* is the directory which you want to change to. (“.” refers to the current directory, “..” the parent directory. If no directory is specified it defaults to your home directory.)

Example: `cd ../foo` sets the current directory up one level, then back down to `foo`.

<code>ls</code>	<p>Displays information about the named files and directories.</p> <p>Syntax: <code>ls files</code></p> <p>Where <i>files</i> consists of the filenames or directories to list. The most commonly used options are <code>-F</code> (to display the file type), and <code>-l</code> (to give a “long” listing including file size, owner, permissions, and so on).</p> <p>Example: <code>ls -lF /home/larry</code> displays the contents of the directory <code>/home/larry</code>.</p>
<code>cp</code>	<p>Copies one or more files to another file or directory.</p> <p>Syntax: <code>cp files destination</code></p> <p>Where <i>files</i> lists the files to copy, and <i>destination</i> is the destination file or directory.</p> <p>Example: <code>cp ../frog joe</code> copies the file <code>../frog</code> to the file or directory <code>joe</code>.</p>
<code>mv</code>	<p>Moves one or more files to another file or directory. This command does the equivalent of a copy followed by the deletion of the original file. you can use this to rename files, like in the MS-DOS command <code>RENAME</code>.</p> <p>Syntax: <code>mv files destination</code></p> <p>Where <i>files</i> lists the files to move, and <i>destination</i> is the destination file or directory.</p> <p>Example: <code>mv ../frog joe</code> moves the file <code>../frog</code> to the file or directory <code>joe</code>.</p>
<code>rm</code>	<p>Deletes files. Note that when you delete a file under UNIX, they are unrecoverable (unlike MS-DOS, where you can usually “undelete” the file).</p> <p>Syntax: <code>rm files</code></p> <p>Where <i>files</i> describes the filenames to delete.</p> <p>The <code>-i</code> option prompts for confirmation before deleting the file.</p> <p>Example: <code>rm -i /home/larry/joe /home/larry/frog</code> deletes the files <code>joe</code> and <code>frog</code> in <code>/home/larry</code>.</p>
<code>mkdir</code>	<p>Creates new directories.</p> <p>Syntax: <code>mkdir dirs</code></p> <p>Where <i>dirs</i> are the directories to create.</p> <p>Example: <code>mkdir /home/larry/test</code> creates the directory <code>test</code> in <code>/home/larry</code>.</p>
<code>rmdir</code>	<p>Deletes empty directories. When using <code>rmdir</code>, the current working directory must not be within the directory to be deleted.</p> <p>Syntax: <code>rmdir dirs</code></p> <p>Where <i>dirs</i> defines the directories to delete.</p> <p>Example: <code>rmdir /home/larry/papers</code> deletes the directory <code>/home/larry/papers</code>, if empty.</p>
<code>man</code>	<p>Displays the manual page for the given command or resource (that is, any system utility that isn’t a command, such as a library function.)</p> <p>Syntax: <code>man command</code></p> <p>Where <i>command</i> is the name of the command or resource to get help on.</p> <p>Example: <code>man ls</code> gives help on the <code>ls</code> command.</p>

<code>more</code>	<p>Displays the contents of the named files, one screenful at a time. Syntax: <code>more files</code> Where <i>files</i> lists the files to display. Example: <code>more papers/history-final</code> displays the file <code>papers/history-final</code>.</p>
<code>cat</code>	<p>Officially used to concatenate files, <code>cat</code> is also used to display the contents of a file on screen. Syntax: <code>cat files</code> Where <i>files</i> lists the files to display. Example: <code>cat letters/from-mdw</code> displays the file <code>letters/from-mdw</code>.</p>
<code>echo</code>	<p>Displays the given arguments on the screen. Syntax: <code>echo args</code> Where <i>args</i> lists arguments to echo. Example: <code>echo "Hello world"</code> displays the string "Hello world".</p>
<code>grep</code>	<p>Display every line in one or more files that match the given pattern. Syntax: <code>grep pattern files</code> Where <i>pattern</i> is a regular expression pattern, and <i>files</i> lists the files to search. Example: <code>grep loomer /etc/hosts</code> displays every line in the file <code>/etc/hosts</code> that contains the pattern "loomer".</p>

3.6 Exploring the file system.

A **file system** is the collection of files and the hierarchy of directories on a system. The time has now come to escort you around the file system.

You now have the skills and the knowledge to understand the Linux file system, and you have a roadmap. (Refer to Figure 3.1).

First, change to the root directory (`cd /`), and then enter `ls -F` to display a listing of its contents. You'll probably see the following directories²: `bin`, `dev`, `etc`, `home`, `install`, `lib`, `mnt`, `proc`, `root`, `tmp`, `user`, `usr`, and `var`.

Now, let's take a look at each of these directories.

`/bin` `/bin` is short for "binaries" or executables, where many essential system programs reside. Use `ls -F /bin` to list the files here. If you look down the list you may see a few commands that you recognize, such as `cp`, `ls`, and `mv`. These are the actual programs for these commands. When you use the `cp` command, for example, you're running the program `/bin/cp`.

Using `ls -F`, you'll see that most (if not all) of the files in `/bin` have an asterisk("*") appended to their filenames. This indicates that the files are executables, as described in Section 3.3.2.

²You may see others, and you might not see all of them. Every release of Linux differs in some respects.

`/dev`

The “files” in `/dev` are **device files**—they access system devices and resources like disk drives, modems, and memory. Just as your system can read data from a file, it can also read input from the mouse by accessing `/dev/mouse`.

Filenames that begin with `fd` are floppy disk devices. `fd0` is the first floppy disk drive, and `fd1` is the second. You may have noticed that there are more floppy disk devices than the two listed above: these represent specific types of floppy disks. For example, `fd1H1440` access high-density, 3.5” diskettes in drive 1.

The following is a list of some of the most commonly used device files. Even though you may not have some of the physical devices listed below, chances are that you’ll have drivers in `/dev` for them anyway.

- `/dev/console` refers to the system’s console—that is, the monitor connected directly to your system.
- The various `/dev/ttyS` and `/dev/cua` devices are used for accessing serial ports. `/dev/ttyS0` refers to “COM1” under MS-DOS. The `/dev/cua` devices are “callout” devices, and used with a modem.
- Device names beginning with `hd` access hard drives. `/dev/hda` refers to the *whole* first hard disk, while `/dev/hda1` refers to the first *partition* on `/dev/hda`.
- Device names that begin with `sd` are SCSI drives. If you have a SCSI hard drive, instead of accessing it through `/dev/hda`, you would access `/dev/sda`. SCSI tapes are accessed via `st` devices, and SCSI CD-ROM via `sr` devices.
- Device names that begin with `lp` access parallel ports. `/dev/lp0` is the same as “LPT1” in the MS-DOS world.
- `/dev/null` is used as a “black hole”—data sent to this device is gone forever. Why is this useful? Well, if you wanted to suppress the output of a command appearing on your screen, you could send that output to `/dev/null`. We’ll talk more about this later.
- Devices whose names are `/dev/tty` followed by a number refer to the “virtual consoles” on your system (accessed by pressing `Alt-F1`, `Alt-F2`, and so on). `/dev/tty1` refers to the first VC, `/dev/tty2` refers to the second, and so on.
- Device names beginning with `/dev/pty` are **pseudo-terminals**, which are used to provide a “terminal” to remote login sessions. For example, if your machine is on a network, incoming `telnet` logins would use one of the `/dev/pty` devices.

<code>/etc</code>	<code>/etc</code> contains a number of miscellaneous system configuration files. These include <code>/etc/passwd</code> (the user database), <code>/etc/rc</code> (the system initialization script), and so on.
<code>/sbin</code>	<code>/sbin</code> contains essential system binaries that are used for system administration.
<code>/home</code>	<code>/home</code> contains user's home directories. For example, <code>/home/larry</code> is the home directory for the user "larry". On a newly install system, there may not be any users in this directory.
<code>/lib</code>	<code>/lib</code> contains shared library images , which are files that contain code which many programs share in common. Rather than each program using its own copy of these shared routines, they are all stored in one common place, in <code>/lib</code> . This makes executable files smaller, and saves space on your system.
<code>/proc</code>	<code>/proc</code> supports a "virtual file system", where the files are stored in memory, not on disk. These "files" refer to the various processes running on the system, and let you get information about the programs and processes that are running at any given time. This is discussed in more detail in Section 3.12
<code>/tmp</code>	Many programs store temporary information and data in a file that is deleted when the program has finished executing. The standard location for these files is in <code>/tmp</code> .
<code>/usr</code>	<code>/usr</code> is a very important directory which contains subdirectories that contain some of the most important and useful programs and configuration files used on the system.

The various directories described above are essential for the system to operate, but most of the items found in `/usr` are optional. However, it is these optional items that make the system useful and interesting. Without `/usr`, you'd have a boring system that supports only programs like `cp` and `ls`. `/usr` contains most of the larger software packages and the configuration files that accompany them.

<code>/usr/X11R6</code>	<code>/usr/X11R6</code> contains the X Window System, if you installed it. The X Window System is a large, powerful graphical environment that provides a large number of graphical utilities and programs, displayed in "windows" on your screen. If you're at all familiar with the Microsoft Windows or Macintosh environments, X Windows will look familiar. The <code>/usr/X11R6</code> directory contains all of the X Windows executables, configuration files, and support files. This is covered in more details in another Chapter.
<code>/usr/bin</code>	<code>/usr/bin</code> is the real warehouse for software on any Linux system, containing most of the executables for programs not found in other places, like <code>/bin</code>
<code>/usr/etc</code>	Just as <code>/etc</code> contains essential miscellaneous system programs and configuration files, <code>/usr/etc</code> contains miscellaneous utilities and files that, in general, are not essential to the system.

<code>/usr/include</code>	<code>/usr/include</code> contains include files for the C compiler. These files (most of which end in <code>.h</code> for “header”), declare data structure names, subroutines, and constants used when writing programs in C. Files in <code>/usr/include/sys</code> are generally used when programming on the UNIX system level. If you are familiar with the C programming language, here you’ll find header files like <code>stdio.h</code> , which declare functions like <code>printf()</code> .
<code>/usr/g++-include</code>	<code>/usr/g++-include</code> contains include files for the C++ compiler (much like <code>/usr/include</code>).
<code>/usr/lib</code>	<code>/usr/lib</code> contains the “stub” and “static” library equivalents for the files found in <code>/lib</code> . When compiling a program, the program is “linked” with the libraries found in <code>/usr/lib</code> , which then directs the program to look in <code>/lib</code> when it needs the actual code in the library. In addition, various other programs store configuration files in <code>/usr/lib</code> .
<code>/usr/local</code>	<code>/usr/local</code> is much like <code>/usr</code> —it contains various programs and files not essential to the system, but which make the system fun and exciting. In general, programs in <code>/usr/local</code> are specialized for your system—consequently, <code>/usr/local</code> differs greatly between Linux systems.
<code>/usr/man</code>	This directory contains manual pages. There are two subdirectories in it for every manual page “section” (use the command <code>man man</code> for details). For example, <code>/usr/man/man1</code> contains the source (that is, the unformatted original) for manual pages in section 1, and <code>/usr/man/cat1</code> contains the formatted manual pages for section 1.
<code>/usr/src</code>	<code>/usr/src</code> contains the source code (the uncompiled instructions) for various programs on your system. The most important directory here is <code>/usr/src/linux</code> , which contains the source code for the Linux kernel.
<code>/var</code>	<code>/var</code> holds directories that often change in size or tend to grow. Many of those directories used to reside <code>/usr</code> , but since those who support Linux are trying to keep it relatively unchangeable, the directories that change often have been moved to <code>/var</code> . Some Linux distributions maintain their software package databases in directories under <code>/var</code> .
<code>/var/log</code>	<code>/var/log</code> contains various files of interest to the system administrator, specifically system logs, which record errors or problems with the system. Other files record logins to the system as well as failed login attempts. This will be covered in Chapter 4.
<code>/var/spool</code>	<code>/var/spool</code> contains files which are “spooled” to another program. For example, if your machine is connected to a network, incoming mail is stored in <code>/var/spool/mail</code> until you read or delete it. Outgoing or incoming news articles are in <code>/var/spool/news</code> , and so on.

3.7 Types of shells.

As mentioned before, Linux is a multitasking, multiuser operating system. Multitasking is very useful, and once you understand it, you’ll use it all of the time. Before long, you’ll run programs in the background, switch between tasks, and pipeline programs together to achieve complicated results with a single command.

Many of the features we’ll cover in this section are features provided by the shell itself. Be careful not to confuse Linux (the actual operating system) with a shell—a shell is just an interface

to the underlying system. The shell provides functionality in addition to Linux itself.

A shell is not only an interpreter for the interactive commands you type at the prompt, but also a powerful programming language. It lets you to write **shell scripts**, to “batch” several shell commands together in a file. If you know MS-DOS you’ll recognize the similarity to “batch files”. Shell scripts are a very powerful tool, that will let you automate and expand your use of Linux.

There are several types of shells in the Linux world. The two major types are the “Bourne shell” and the “C shell”. The Bourne shell uses a command syntax like the original shell on early UNIX systems, like System III. The name of the Bourne shell on most Linux systems is `/bin/sh` (where **sh** stands for “shell”). The C shell (not to be confused with sea shell) uses a different syntax, somewhat like the programming language C, and on most Linux systems is named `/bin/csh`.

Under Linux, several variations of these shells are available. The two most commonly used are the Bourne Again Shell, or “Bash” (`/bin/bash`), and “Tcsh” (`/bin/tcsh`). **bash** is a form of the Bourne shell that includes many of the advanced features found in the C shell. Because **bash** supports a superset of the Bourne shell syntax, shell scripts written in the standard Bourne shell should work with **bash**. If you prefer to use the C shell syntax, Linux supports **tcsh**, which is an expanded version of the original C shell.

The type of shell you decide to use is mostly a religious issue. Some folks prefer the Bourne shell syntax with the advanced features of **bash**, and some prefer the more structured C shell syntax. As far as normal commands such as `cp` and `ls` are concerned, the shell you use doesn’t matter—the syntax is the same. Only when you start to write shell scripts or use advanced features of a shell do the differences between shell types begin to matter.

As we discuss the features of the various shells, we’ll note differences between Bourne and C shells. However, for the purposes of this manual most of those differences are minimal. (If you’re really curious at this point, read the `man` pages for **bash** and **tcsh**).

3.8 Wildcards.

A key feature of most Linux shells is the ability to refer to more than one file by name using special characters. These **wildcards** let you refer to, say, all file names that contain the character “n”.

The wildcard “*” specifies any character or string of characters in a file name. When you use the character “*” in a file name, the shell replaces it with all possible substitutions from file names in the directory you’re referencing.

Here’s a quick example. Suppose that Larry has the files `frog`, `joe`, and `stuff` in his current directory.

```
/home/larry# ls
frog      joe      stuff
/home/larry#
```

To specify all files containing the letter “o” in the filename, use the command

```
/home/larry# ls *o*
frog      joe
/home/larry#
```

As you can see, each instance of “*” is replaced with all substitutions that match the wildcard from filenames in the current directory.

The use of “*” by itself simply matches all filenames, because all characters match the wildcard.

```
/home/larry# ls *
frog      joe      stuff
/home/larry#
```

Here are a few more examples:

```
/home/larry# ls f*
frog
/home/larry# ls *ff
stuff
/home/larry# ls *f*
frog      stuff
/home/larry# ls s*f
stuff
/home/larry#
```

The process of changing a “*” into a series of filenames is called **wildcard expansion** and is done by the shell. This is important: an individual command, such as `ls`, *never* sees the “*” in its list of parameters. The shell expands the wildcard to include all filenames that match. So, the command

```
/home/larry# ls *o*
```

is expanded by the shell to

```
/home/larry# ls frog joe
```

One important note about the “*” wildcard: it does *not* match file names that begin with a single period (“.”). These files are treated as **hidden** files—while they are not really hidden, they don’t show up on normal `ls` listings and aren’t touched by the use of the “*” wildcard.

Here’s an example. We mentioned earlier that each directory contains two special entries: “.” refers to the current directory, and “..” refers to the parent directory. However, when you use `ls`, these two entries don’t show up.

```
/home/larry# ls
frog      joe      stuff
/home/larry#
```

If you use the `-a` switch with `ls`, however, you can display filenames that begin with “.”. Observe:

```
/home/larry# ls -a
.      ..      .bash.profile      .bashrc      frog      joe      stuff
/home/larry#
```

The listing contains the two special entries, “.” and “..”, as well as two other “hidden” files—`.bash_profile` and `.bashrc`. These two files are startup files used by `bash` when `larry` logs in. They are described in a later section.

Note that when you use the “*” wildcard, none of the filenames beginning with “.” are displayed.

```
/home/larry# ls *
frog      joe      stuff
/home/larry#
```

This is a safety feature: if the “*” wildcard matched filenames beginning with “.”, it would also match the directory names “.” and “..”. This can be dangerous when using certain commands.

Another wildcard is “?”. The “?” wildcard expands to only a single character. Thus, “`ls ?`” displays all one-character filenames. And “`ls termca?`” would display “`termcap`” but not “`termcap.backup`”. Here’s another example:

```

/home/larry# ls j?e
joe
/home/larry# ls f??g
frog
/home/larry# ls ???f
stuff
/home/larry#

```

As you can see, wildcards lets you specify many files at one time. In the previous command summary, we said that the `cp` and `mv` commands actually can copy or move more than one file at a time. For example,

```

/home/larry# cp /etc/s* /home/larry

```

copies all filenames in `/etc` beginning with “s” to the directory `/home/larry`. The format of the `cp` command is really

```

cp files destination

```

where *files* lists the filenames to copy, and *destination* is the destination file or directory. `mv` has an identical syntax.

If you are copying or moving more than one file, the *destination* must be a directory. You can only copy or move a single file to another file.

3.9 Linux plumbing.

3.9.1 Standard input and standard output.

Many Linux commands get input from what is called **standard input** and send their output to **standard output** (often abbreviated as `stdin` and `stdout`). Your shell sets things up so that standard input is your keyboard, and standard output is the screen.

Here’s an example using the `cat` command. Normally, `cat` reads data from all of the files specified by the command line, and sends this data directly to `stdout`. Therefore, using the command

```

/home/larry/papers# cat history-final masters-thesis

```

displays the contents of the file `history-final` followed by `masters-thesis`.

However, if you don’t specify a filename, `cat` reads data from `stdin` and sends it back to `stdout`. Here’s an example:

```

/home/larry/papers# cat
Hello there.
Hello there.
Bye.
Bye.
Ctrl-D
/home/larry/papers#

```

Each line that you type is immediately echoed back by `cat`. When reading from standard input, you indicate the input is “finished” by sending an EOT (end-of-text) signal, in general, generated by pressing Ctrl-D.

Here’s another example. The `sort` command reads lines of text (again, from `stdin`, unless you specify one or more filenames) and sends the sorted output to `stdout`. Try the following.

```
/home/larry/papers# sort
bananas
carrots
apples
[Ctrl-D]
apples
bananas
carrots
/home/larry/papers#
```

Now we can alphabetize our shopping list... isn't Linux useful?

3.9.2 Redirecting input and output.

Now, let's say that you want to send the output of `sort` to a file, to save our shopping list on disk. The shell lets you **redirect** standard output to a filename, by using the “`>`” symbol. Here's how it works:

```
/home/larry/papers# sort > shopping-list
bananas
carrots
apples
[Ctrl-D]
/home/larry/papers#
```

As you can see, the result of the `sort` command isn't displayed, but is saved to the file named `shopping-list`. Let's look at this file:

```
/home/larry/papers# cat shopping-list
apples
bananas
carrots
/home/larry/papers#
```

Now you can sort your shopping list, and save it, too! But let's suppose that you are storing the unsorted, original shopping list in the file `items`. One way of sorting the information and saving it to a file would be to give `sort` the name of the file to read, in lieu of standard input, and redirect standard output as we did above, as follows:

```
/home/larry/papers# sort items > shopping-list
/home/larry/papers# cat shopping-list
apples
bananas
carrots
/home/larry/papers#
```

However, there's another way to do this. Not only can you redirect standard output, you can redirect standard *input* as well, using the “`<`” symbol.

```
/home/larry/papers# sort < items
apples
bananas
carrots
/home/larry/papers#
```

Technically, `sort < items` is equivalent to `sort items`, but lets you demonstrate the following point: `sort < items` behaves as if the data in the file `items` was typed to standard input. The shell handles the redirection. `sort` wasn't given the name of the file (`items`) to read; as far as `sort` is concerned, it still reads from standard input as if you had typed the data from your keyboard.

This introduces the concept of a **filter**. A filter is a program that reads data from standard input, processes it in some way, and sends the processed data to standard output. Using redirection, standard input and standard output can be referenced from files. As mentioned above, `stdin` and `stdout` default to the keyboard and screen respectively. `sort` is a simple filter. It sorts the incoming data and sends the result to standard output. `cat` is even simpler. It doesn't do anything with the incoming data, it simply outputs whatever is given to it.

3.9.3 Using pipes.

We already demonstrated how to use `sort` as a filter. However, these examples assume that you have data stored in a file somewhere or are willing to type the data from the standard input yourself. What if the data that you wanted to sort came from the output of another command, like `ls`?

The `-r` option to `sort` sorts the data in reverse-alphabetical order. If you want to list the files in your current directory in reverse order, one way to do it is as follows:

```
/home/larry/papers# ls
english-list
history-final
masters-thesis
notes
/home/larry/papers#
```

Now redirect the output of the `ls` command into a file called `file-list`:

```
/home/larry/papers# ls > file-list
/home/larry/papers# sort -r file-list
notes
masters-thesis
history-final
english-list
/home/larry/papers#
```

Here, you save the output of `ls` in a file, and then run `sort -r` on that file. But this is unwieldy and uses a temporary file to save the data from `ls`.

The solution is **pipelining**. This is a shell feature that connects a string of commands via a “pipe.” The `stdout` of the first command is sent to the `stdin` of the second command. In this case, we want to send the `stdout` of `ls` to the `stdin` of `sort`. Use the “|” symbol to create a pipe, as follows:

```
/home/larry/papers# ls | sort -r
notes
```

```
masters-thesis
history-final
english-list
/home/larry/papers#
```

This command is shorter and easier to type.

Here's another useful example, the command

```
/home/larry/papers# ls /usr/bin
```

displays a long list of files, most of which fly past the screen too quickly for you to read. So, let's use `more` to display the list of files in `/usr/bin`.

```
/home/larry/papers# ls /usr/bin | more
```

Now you can page down the list of files at your leisure.

But the fun doesn't stop here! You can pipe more than two commands together. The command `head` is a filter that displays the first lines from an input stream (in this case, input from a pipe). If you want to display the last filename in alphabetical order in the current directory, use commands like the following:

```
/home/larry/papers# ls | sort -r | head -1
notes
/home/larry/papers#
```

where `head -1` displays the first line of input that it receives (in this case, the stream of reverse-sorted data from `ls`).

3.9.4 Non-destructive redirection of output.

Using “>” to redirect output to a file is destructive: in other words, the command

```
/home/larry/papers# ls > file-list
```

overwrites the contents of the file `file-list`. If instead, you redirect with the symbol “>>”, the output is appended to (added to the end of) the named file instead of overwriting it. For example,

```
/home/larry/papers# ls >> file-list
```

appends the output of the `ls` command to `file-list`.

Keep in mind that redirection and pipes are features of the shell—which supports the use of “>”, “>>” and “|”. It has nothing to do with the commands themselves.

3.10 File permissions.

3.10.1 Concepts of file permissions.

Because there is typically more than one user on a Linux system, Linux provides a mechanism known as **file permissions**, which protect user files from tampering by other users. This mechanism lets files and directories be “owned” by a particular user. For example, because Larry created the files in his home directory, Larry owns those files and has access to them.

Linux also lets files be shared between users and groups of users. If Larry desired, he could cut off access to his files so that no other user could access them. However, on most systems the default is to allow other users to read your files but not modify or delete them in any way.

Every file is owned by a particular user. However, files are also owned by a particular **group**, which is a defined group of users of the system. Every user is placed into at least one group when that user's account is created. However, the system administrator may grant the user access to more than one group.

Groups are usually defined by the type of users who access the machine. For example, on a university Linux system users may be placed into the groups **student**, **staff**, **faculty** or **guest**. There are also a few system-defined groups (like **bin** and **admin**) which are used by the system itself to control access to resources—very rarely do actual users belong to these system groups.

Permissions fall into three main divisions: read, write, and execute. These permissions may be granted to three classes of users: the owner of the file, the group to which the file belongs, and to all users, regardless of group.

Read permission lets a user read the contents of the file, or in the case of directories, list the contents of the directory (using **ls**). Write permission lets the user write to and modify the file. For directories, write permission lets the user create new files or delete files within that directory. Finally, execute permission lets the user run the file as a program or shell script (if the file is a program or shell script). For directories, having execute permission lets the user **cd** into the directory in question.

3.10.2 Interpreting file permissions.

Let's look at an example that demonstrates file permissions. Using the **ls** command with the **-l** option displays a “long” listing of the file, including file permissions.

```
/home/larry/foo# ls -l stuff
-rw-r--r--  1 larry  users          505 Mar 13 19:05 stuff

/home/larry/foo#
```

The first field in the listing represents the file permissions. The third field is the owner of the file (**larry**) and the fourth field is the group to which the file belongs (**users**). Obviously, the last field is the name of the file (**stuff**). We'll cover the other fields later.

This file is owned by **larry**, and belongs to the group **users**. The string **-rw-r--r--** lists, in order, the permissions granted to the file's owner, the file's group, and everybody else.

The first character of the permissions string (“-”) represents the type of file. A “-” means that this is a regular file (as opposed to a directory or device driver). The next three characters (“**rw-**”) represent the permissions granted to the file's owner, **larry**. The “**r**” stands for “read” and the “**w**” stands for “write”. Thus, **larry** has read and write permission to the file **stuff**.

As mentioned, besides read and write permission, there is also “execute” permission—represented by an “**x**”. However, a “-” is listed here in place of an “**x**”, so **Larry** doesn't have execute permission on this file. This is fine, as the file **stuff** isn't a program of any kind. Of course, because **Larry** owns the file, he may grant himself execute permission for the file if he so desires. (This will be covered shortly.)

The next three characters, (“**r--**”), represent the group's permissions on the file. The group that owns this file is **users**. Because only an “**r**” appears here, any user who belongs to the group **users** may read this file.

The last three characters, also (“**r--**”), represent the permissions granted to every other user on the system (other than the owner of the file and those in the group **users**). Again, because only an “**r**” is present, other users may read the file, but not write to it or execute it.

Here are some other examples of permissions:

<code>-rwxr-xr-x</code>	The owner of the file may read, write and execute the file. Users in the file's group, and all other users, may read and execute the file.
<code>-rw-----</code>	The owner of the file may read and write the file. No other user can access the file.
<code>-rwxrwxrwx</code>	All users may read, write, and execute the file.

3.10.3 Permissions Dependencies.

The permissions granted to a file also depend on the permissions of the directory in which the file is located. For example, even if a file is set to `-rwxrwxrwx`, other users cannot access the file unless they have read and execute access to the directory in which the file is located. For example, if Larry wanted to restrict access to all of his files, he could set the permissions to his home directory `/home/larry` to `-rwx-----`. In this way, no other user has access to his directory, and all files and directories within it. Larry doesn't need to worry about the individual permissions on each of his files.

In other words, to access a file at all, you must have execute access to all directories along the file's pathname, and read (or execute) access to the file itself.

Typically, users on a Linux system are very open with their files. The usual set of permissions given to files is `-rw-r--r--`, which lets other users read the file but not change it in any way. The usual set of permissions given to directories is `-rwxr-xr-x`, which lets other users look through your directories, but not create or delete files within them.

However, many users wish to keep other users out of their files. Setting the permissions of a file to `-rw-----` will prevent any other user from accessing the file. Likewise, setting the permissions of a directory to `-rwx-----` keeps other users out of the directory in question.

3.10.4 Changing permissions.

The command `chmod` is used to set the permissions on a file. Only the owner of a file may change the permissions on that file. The syntax of `chmod` is

```
chmod {a,u,g,o}{+,-}{r,w,x} filenames
```

Briefly, you supply one or more of **a**ll, **u**ser, **g**roup, or **o**ther. Then you specify whether you are adding rights (+) or taking them away (-). Finally, you specify one or more of **r**ead, **w**rite, and **e**xecute. Some examples of legal commands are:

<code>chmod a+r stuff</code>	gives all users read access to the file
<code>chmod +r stuff</code>	Same as above—if none of a , u , g , or o is specified, a is assumed.
<code>chmod og-x stuff</code>	Remove execute permission from users other than the owner
<code>chmod u+rwx stuff</code>	Let the owner of the file read, write, and execute the file.
<code>chmod o-rwx stuff</code>	Remove read, write, and execute permission from users other than the owner and users in the file's group

3.11 Managing file links.

Links let you give a single file more than one name. Files are actually identified by the system by their **inode number**, which is just the unique file system identifier for the file. A directory is actually a listing of inode numbers with their corresponding filenames. Each filename in a directory is a **link** to a particular inode.

3.11.1 Hard links.

The `ln` command is used to create multiple links for one file. For example, let's say that you have a file called `foo` in a directory. Using `ls -i`, you can look at the inode number for this file.

```
/home/larry# ls -i foo
22192 foo
/home/larry#
```

Here, `foo` has an inode number of 22192 in the file system. You can create another link to `foo`, named `bar`, as follows:

```
/home/larry# ln foo bar
```

With `ls -i`, you see that the two files have the same inode.

```
/home/larry# ls -i foo bar
22192 bar 22192 foo
/home/larry#
```

Now, specifying either `foo` or `bar` will access the same file. If you make changes to `foo`, those changes appear in `bar` as well. For all purposes, `foo` and `bar` are the same file.

These links are known as **hard links** because they create a direct link to an inode. Note that you can hard-link files only when they're on the same file system; symbolic links (see below) don't have this restriction.

When you delete a file with `rm`, you are actually only deleting one link to a file. If you use the command

```
/home/larry# rm foo
```

then only the link named `foo` is deleted, `bar` will still exist. A file is only truly deleted on the system when it has no links to it. Usually, files have only one link, so using the `rm` command deletes the file. However, if a file has multiple links to it, using `rm` will delete only a single link; in order to delete the file, you must delete all links to the file.

The command `ls -l` displays the number of links to a file (among other information).

```
/home/larry# ls -l foo bar
-rw-r--r-- 2 root root 12 Aug 5 16:51 bar
-rw-r--r-- 2 root root 12 Aug 5 16:50 foo
/home/larry#
```

The second column in the listing, "2", specifies the number of links to the file.

As it turns out, a directory is actually just a file containing information about link-to-inode associations. Also, every directory contains at least two hard links: `."` (a link pointing to itself), and `.."` (a link pointing to the parent directory). The root directory (`/`) `."` link just points back to `/`. (In other words, the parent of the root directory is the root directory itself.)

3.11.2 Symbolic links.

Symbolic links, or **symlinks**, are another type of link, which are different from hard links. A symbolic link lets you give a file another name, but doesn't link the file by inode.

The command `ln -s` creates a symbolic link to a file. For example, if you use the command

```
/home/larry# ln -s foo bar
```

you will create a symbolic link named `bar` that points to the file `foo`. If you use `ls -i`, you'll see that the two files have different inodes, indeed.

```
/home/larry# ls -i foo bar
22195 bar  22192 foo
/home/larry#
```

However, using `ls -l`, we see that the file `bar` is a symlink pointing to `foo`.

```
/home/larry# ls -l foo bar
lrwxrwxrwx  1 root    root          12 Aug  5 16:51 bar -> foo
-rw-r--r--  1 root    root          12 Aug  5 16:50 foo
/home/larry#
```

The file permissions on a symbolic link are not used (they always appear as `lrwxrwxrwx`). Instead, the permissions on the symbolic link are determined by the permissions on the target of the symbolic link (in our example, the file `foo`).

Functionally, hard links and symbolic links are similar, but there are differences. For one thing, you can create a symbolic link to a file that doesn't exist; the same is not true for hard links. Symbolic links are processed by the kernel differently than are hard links, which is just a technical difference but sometimes an important one. Symbolic links are helpful because they identify the file they point to; with hard links, there is no easy way to determine which files are linked to the same inode.

Links are used in many places on the Linux system. Symbolic links are especially important to the shared library images in `/lib`.

3.12 Job control.

3.12.1 Jobs and processes.

Job control is a feature provided by many shells (including `bash` and `tcsh`) that let you control multiple running commands, or **jobs**, at once. Before we can delve much further, we need to talk about **processes**.

Every time you run a program, you start what is called a process. The command `ps` displays a list of currently running processes, as shown here:

```
/home/larry# ps

  PID TT STAT  TIME COMMAND
   24  3  S    0:03 (bash)
  161  3  R    0:00 ps

/home/larry#
```

The PID listed in the first column is the **process ID**, a unique number given to every running process. The last column, `COMMAND`, is the name of the running command. Here, we're looking only at the processes which Larry himself is currently running. (There are many other processes running on the system as well—"ps -aux" lists them all.) These are `bash` (Larry's shell), and the `ps` command itself. As you can see, `bash` is running concurrently with the `ps` command. `bash` executed `ps` when Larry typed the command. After `ps` has finished running (after the table of processes is displayed), control is returned to the `bash` process, which displays the prompt, ready for another command.

A running process is also called a *job*. The terms *process* and *job* are interchangeable. However, a process is usually referred to as a "job" when used in conjunction with **job control**—a feature of the shell that lets you switch between several independent jobs.

In most cases users run only a single job at a time—whatever command they last typed to the shell. However, using job control, you can run several jobs at once, and switch between them as needed.

How might this be useful? Let’s say you are editing a text file and want to interrupt your editing and do something else. With job control, you can temporarily suspend the editor, go back to the shell prompt and start to work on something else. When you’re done, you can switch back to the editor and be back where you started, as if you didn’t leave the editor. There are many other practical uses of job control.

3.12.2 Foreground and background.

Jobs can either be in the **foreground** or in the **background**. There can only be one job in the foreground at a time. The foreground job is the job with which you interact—it receives input from the keyboard and sends output to your screen, unless, of course, you have redirected input or output, as described starting in section 3.9). On the other hand, jobs in the background do not receive input from the terminal—in general, they run along quietly without the need for interaction.

Some jobs take a long time to finish and don’t do anything interesting while they are running. Compiling programs is one such job, as is compressing a large file. There’s no reason why you should sit around being bored while these jobs complete their tasks; just run them in the background. While jobs run in the background, you are free to run other programs.

Jobs may also be **suspended**. A suspended job is a job that is temporarily stopped. After you suspend a job, you can tell the job to continue in the foreground or the background as needed. Resuming a suspended job does not change the state of the job in any way—the job continues to run where it left off.

Suspending a job is not equal to interrupting a job. When you **interrupt** a running process (by pressing the interrupt key, which is usually `Ctrl-C`)³, the process is killed, for good. Once the job is killed, there’s no hope of resuming it. You’ll must run the command again. Also, some programs trap the interrupt, so that pressing `Ctrl-C` won’t immediately kill the job. This is to let the program perform any necessary cleanup operations before exiting. In fact, some programs don’t let you kill them with an interrupt at all.

3.12.3 Backgrounding and killing jobs.

Let’s begin with a simple example. The command `yes` is a seemingly useless command that sends an endless stream of y’s to standard output. (This is actually useful. If you piped the output of `yes` to another command which asked a series of yes and no questions, the stream of y’s would confirm all of the questions.)

Try it out:

```
/home/larry# yes
y
y
y
y
y
```

The y’s will continue *ad infinitum*. You can kill the process by pressing the interrupt key, which is usually `Ctrl-C`. So that we don’t have to put up with the annoying stream of y’s, let’s redirect the standard output of `yes` to `/dev/null`. As you may remember, `/dev/null` acts as a “black hole”

³You can set the interrupt key with the `stty` command

for data. Any data sent to it disappears. This is a very effective method of quieting an otherwise verbose program.

```
/home/larry# yes > /dev/null
```

Ah, much better. Nothing is printed, but the shell prompt doesn't come back. This is because `yes` is still running, and is sending those inane `y`'s to `/dev/null`. Again, to kill the job, press the interrupt key.

Let's suppose that you want the `yes` command to continue to run but wanted to get the shell prompt back so that you can work on other things. You can put `yes` into the background, allowing it to run, without need for interaction.

One way to put a process in the background is to append an "&" character to the end of the command.

```
/home/larry# yes > /dev/null &
[1] 164
/home/larry#
```

As you can see, the shell prompt has returned. But what is this "[1] 164"? And is the `yes` command really running?

The "[1]" represents the **job number** for the `yes` process. The shell assigns a job number to every running job. Because `yes` is the one and only job we're running, it is assigned job number 1. The "164" is the process ID, or PID, number given by the system to the job. You can use either number to refer to the job, as you'll see later.

You now have the `yes` process running in the background, continuously sending a stream of `y`'s to `/dev/null`. To check on the status of this process, use the internal shell command `jobs`.

```
/home/larry# jobs
[1]+  Running                  yes >/dev/null &
/home/larry#
```

Sure enough, there it is. You could also use the `ps` command as demonstrated above to check on the status of the job.

To terminate the job, use the `kill` command. This command takes either a job number or a process ID number as an argument. This was job number 1, so using the command

```
/home/larry# kill %1
```

kills the job. When identifying the job with the job number, you must prefix the number with a percent ("%") character.

Now that you've killed the job, use `jobs` again to check on it:

```
/home/larry# jobs
[1]+  Terminated              yes >/dev/null &
/home/larry#
```

The job is in fact dead, and if you use the `jobs` command again nothing should be printed.

You can also kill the job using the process ID (PID) number, displayed along with the job ID when you start the job. In our example, the process ID is 164, so the command

```
/home/larry# kill 164
```

is equivalent to

```
/home/larry# kill %1
```

You don't need to use the “%” when referring to a job by its process ID.

3.12.4 Stopping and restarting jobs.

There is another way to put a job into the background. You can start the job normally (in the foreground), **stop** the job, and then restart it in the background.

First, start the **yes** process in the foreground, as you did before:

```
/home/larry# yes > /dev/null &
```

Again, because **yes** is running in the foreground, you shouldn't get the shell prompt back.

Now, rather than interrupt the job with `Ctrl-C`, **suspend** the job. Suspending a job doesn't kill it: it only temporarily stops the job until you restart it. To do this, press the suspend key, which is usually `Ctrl-Z`.

```
/home/larry# yes > /dev/null &
```

```
Ctrl-Z
```

```
[1]+  Stopped                yes >/dev/null &
```

```
/home/larry#
```

While the job is suspended, it's simply not running. No CPU time is used for the job. However, you can restart the job, which causes the job to run again as if nothing ever happened. It will continue to run where it left off.

To restart the job in the foreground, use the **fg** command (for “foreground”).

```
/home/larry# fg
```

```
yes >/dev/null
```

The shell displays the name of the command again so you're aware of which job you just put into the foreground. Stop the job again with `Ctrl-Z`. This time, use the **bg** command to put the job into the background. This causes the command to run just as if you started the command with “&” as in the last section.

```
/home/larry# bg
```

```
[1]+ yes >/dev/null &
```

```
/home/larry#
```

And you have your prompt back. **Jobs** should report that **yes** is indeed running, and you can kill the job with **kill** as we did before.

How can you stop the job again? Using `Ctrl-Z` won't work, because the job is in the background. The answer is to put the job in the foreground with **fg**, and then stop it. As it turns out, you can use **fg** on either stopped jobs or jobs in the background.

There is a big difference between a job in the background and a job that is stopped. A stopped job is not running—it's not using any CPU time, and it's not doing any work (the job still occupies system memory, although it may have been swapped out to disk). A job in the background is running and using memory, as well as completing some task while you do other work.

However, a job in the background may try to display text on your terminal, which can be annoying if you're trying to work on something else. For example, if you used the command

```
/home/larry# yes &
```

without redirecting `stdout` to `/dev/null`, a stream of `y`'s would be displayed on your screen, without any way for you to interrupt it. (You can't use `Ctrl-C` to interrupt jobs in the background.) In order to stop the endless `y`'s, use the `fg` command to bring the job to the foreground, and then use `Ctrl-C` to kill it.

Another note. The `fg` and `bg` commands normally affect the job that was last stopped (indicated by a "+" next to the job number when you use the `jobs` command). If you are running multiple jobs at once, you can put jobs in the foreground or background by giving the job ID as an argument to `fg` or `bg`, as in

```
/home/larry# fg %2
```

(to put job number 2 into the foreground), or

```
/home/larry# bg %3
```

(to put job number 3 into the background). You can't use process ID numbers with `fg` or `bg`.

Furthermore, using the job number alone, as in

```
/home/larry# %2
```

is equivalent to

```
/home/larry# fg %2
```

Just remember that using job control is a feature of the shell. The `fg`, `bg`, and `jobs` commands are internal to the shell. If for some reason you use a shell that doesn't support job control, don't expect to find these commands available.

In addition, there are some aspects of job control that differ between `bash` and `tcsh`. In fact, some shells don't provide job control at all—however, most shells available for Linux do.