Exam LX0-103

PART I

Chapter

1

Exploring Linux Command-Line Tools

THE FOLLOWING EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ 103.1 Work on the command line
- ✓ 103.2 Process text streams using filters
- ✓ 103.4 Use streams, pipes, and redirects
- ✓ 103.7 Search text files using regular expressions

Linux borrows heavily from Unix, and Unix began as a text-based operating system (OS). Unix and Linux retain much of this heritage, which means to understand how to use and, especially administer Linux, you must understand at least the basics of its command-line tools. Using command-line tools requires the use of a shell. A *shell* is a program that accepts and interprets text-mode commands and provides an interface to the system.

This chapter begins with basic shell information, including the various shell programs available and the procedures for using them. From there, this chapter covers streams, pipes, and redirection, which you can use to move input and output between programs or between files and programs. These techniques are frequently combined with text processing using *filters*—commands you can use to manipulate text without the help of a conventional text editor. Sometimes you must manipulate text in an abstract way, using codes to represent several different types of text. This chapter, therefore, covers this topic as well.

Understanding Command-Line Basics

Before you do anything else with Linux, you should understand how to use a Linux shell. The shell allows you to enter commands as needed. Which commands can be entered depends on which shell program is running. Several of the available shell programs are briefly described.

In using shell commands, you should also understand shell *environment variables*, which are placeholders for data that may be useful to many programs. Finally, it is helpful to know how to get help with the shell commands you're trying to use.

Exploring Your Linux Shell Options

The shell to be used for entering commands is configured for each individual user, and Linux provides a range of available shells. A complete shell list would be quite long, but the following shells are among the more common choices:

`bash` The GNU Bourne Again Shell (`bash`) is based on the earlier Bourne shell for Unix but extends it in several ways. In Linux, `bash` is the most common default shell for user accounts, and it's the one emphasized in this book and on the exam.

`sh` The Bourne shell upon which `bash` is based goes by the name `sh`. It's not often used in Linux and the `sh` command is often a pointer to the `bash` shell or other shells.

`tcsh` This shell is based on the earlier C shell (`csh`). It's a fairly popular shell in some circles, but no major Linux distributions make it the default shell. Although it's similar to `bash` in many respects, some operational details differ. For instance, you don't assign environment variables the same way in `tcsh` as in `bash`.

`csh` The original C shell isn't used much on Linux, but if a user is familiar with `csh`, `tcsh` makes a good substitute.

`ksh` The Korn shell (`ksh`) was designed to take the best features of the Bourne shell and the C shell and extend them. It has a small but dedicated following among Linux users.

`zsh` The Z shell (`zsh`) takes shell evolution further than the Korn shell, incorporating features from earlier shells and adding still more.

In addition to these shells, dozens more obscure ones are available. In Linux, most users run `bash` because it is the most popular shell. Some other OSs use `csh` or `tcsh` as the default, so if your users have backgrounds on non-Linux Unix-like OSs, they may be more familiar with these other shells. You can change a user's default shell by editing their account, as described in Chapter 7, "Administering the System."

Be aware that there are two types of default shells. The *default interactive shell* is the shell program a user uses to enter commands, run programs from the command line, run shell scripts, and so on. The other default shell type is a default *system* shell. The *default system shell* is used by the Linux system to run system shell scripts, typically at startup.

The file `/bin/sh` is a pointer to the system's default system shell—normally `/bin/bash` for Linux. However, be aware that, on some distributions, the `/bin/sh` points to a different shell. For example, on Ubuntu, `/bin/sh` points to the dash shell, `/bin/dash`.

Using a Shell

Linux shell use is fairly straightforward for anybody who's used a text-mode OS before: You type a command, possibly including options to it, and the computer executes the command. For the most part, Linux commands are external—that is, they're programs that are separate from the shell.

A few commands are internal to the shell, though, and knowing the distinction can be important. You should also know some of the tricks that can make using the command shell easier—how to have the computer complete a long command or filename, retrieve a command you've recently run, or edit a command you've recently used (or haven't yet fully entered).

Starting a Shell

If you log into Linux using a text-mode login screen, you have logged into a virtual console terminal and, most likely, you'll be dropped directly into your default shell. The shell program is what presents the prompt and accepts subsequent commands.

If you log into Linux using a graphical user interface (GUI) login screen, you'll have to start a terminal emulator manually in order to reach your default shell. Some GUIs provide a menu option, such as `xterm` or `terminal`, to start a terminal emulator program. These programs enable you to run text-mode programs within Linux, and by default they come up running your shell. If you can't find such a menu option, look for a menu option that enables you to run an arbitrary command. Select it, and type **`xterm`** or **`konsole`** as the command name. This will launch a terminal emulator program that will run a shell.

Once you start a terminal or log into a virtual console terminal, the shell will provide you with a prompt for entering commands. Remember that the shell is a program providing you with an interface to the Linux system.

A good first command to try, `uname`, will show what operating system is being run:

```
$ uname
Linux
$
```

That's not too interesting. You can find out additional information by tacking on the `-a` option to the command. Be sure to include the necessary space between the command and the option:

```
$ uname -a
Linux server01.class.com 2.6.32-431.5.1.el6.x86_64 #1 SMP Wed Feb 12
00:41:43 UTC 2014 x86_64 x86_64 x86_64 GNU/Linux
$
```

The `uname -a` command provides a lot more information, including the current Linux kernel being used (`2.6.32`) as well as the system's hostname (`server01.class.com`). The `uname` command is an external command. The shell also provides internal commands. It's important to know the difference between the two command types, as explained in the next section.

Using Internal and External Commands

Internal commands are, as you might expect, built into the shell program. Thus they are also called *built-in commands*. Most shells offer a similar set of internal commands, but shell-to-shell differences do exist. Internal commands that you're likely to use enable you to perform some common tasks:

Change the Working Directory Whenever you're running a shell, you're working in a specific directory. The `cd` command changes the current working directory. For instance, typing **`cd /home/sally`** changes the current working directory to the `/home/sally` directory.

You can use shortcut characters with the `cd` command as well. The tilde (~) character is a useful shortcut; it stands for your home directory. Thus typing **`cd ~`** will have the same effect as typing **`cd /home/sally`** if your home directory is `/home/sally`.

Display the Working Directory The `pwd` command displays ("prints" to the screen) the current working directory. This command is helpful, especially after you have changed your working directory, to ensure you ended up in the right place.

Display a Line of Text The `echo` command displays the text you enter. For instance, typing **`echo Hello`** causes the system to display the string `Hello`. This may seem pointless, but it's useful in scripts (described in Chapter 9, "Writing Scripts, Configuring Email, and Using Databases"), and it can also be a good way to review the contents of environment variables (described later in this chapter, in the section "Using Environment Variables").

Time an Operation The `time` command times how long subsequent commands take to execute. For instance, typing **`time pwd`** tells you how long the system took to execute the `pwd` command. The time is displayed after the full command terminates. Three times are displayed: total execution time (aka real time), user CPU time, and system CPU time. The final two values tell you about CPU time consumed, which is likely to be much less than the total execution time.

Set Options In its most basic form, the `set` command displays a wide variety of options relating to `bash` shell operation. These options are formatted much like environment variables, but they aren't the same things. You can pass various options to `set` to have it affect a wide range of shell operations.

Terminate the Shell The `exit` and `logout` commands both terminate the shell. The `exit` command terminates any shell, but the `logout` command terminates only *login shells*. Login shells are shell programs that are launched automatically when you initiate a text-mode login as opposed to those that run in `xterm` windows or other terminal emulators.

The preceding list isn't complete. Later sections of this chapter and later chapters describe some additional internal commands. Consult your shell's documentation for a complete list of its internal commands.

You can quickly determine if a command is a built-in command by using the `type` command. Just enter the command `type` before the name of the command you wish to check:

```
$ type pwd
pwd is a shell builtin
$
$ type cd
cd is a shell builtin
$
$ type bash
bash is /bin/bash
$
```

Some of these internal commands are duplicated by external commands that do the same thing. But those external commands aren't always installed on all systems. You can see if there are internal commands with installed duplicate external commands by using the `-a` option on the `type` command:

```
$ type -a cd
cd is a shell builtin
```

```
$
$ type -a pwd
pwd is a shell builtin
pwd is /bin/pwd
$
```

You can see that on this system, there is no external `cd` command installed. However, it does have an external `pwd` command installed.

Keep in mind that even when external commands are installed, the internal command takes precedence. To access the external command, you must provide the complete external command path, as in typing **`/usr/bin/time`** rather than **`time`**.

Real World Scenario

Confusion over Internal and External Commands

When duplicate internal and external commands exist, they sometimes produce subtly different results or accept different options. These differences may occasionally cause problems if you are unaware of them. For example, the `time` built-in command returns slightly different results than the `/usr/bin/time` external command:

```
$ time pwd
/home/Christine

real    0m0.002s
user    0m0.002s
sys     0m0.001s
$
$ /usr/bin/time  pwd
/home/Christine
0.00user 0.00system 0:00.04elapsed 24%CPU
 (0avgtext+0avgdata 2336maxresident)k
56inputs+0outputs (1major+173minor)pagefaults 0swaps
$
```

As you can see, bash's internal `time` shows the time to execute the `pwd` command in a very nice format, while the external time command `/usr/bin/time` is not only a little sloppy in appearance, it also provides additional details. Be mindful of the potential behavior differences between internal and external commands.

When you type a command that's not recognized by the shell as one of its internal commands, the shell checks its path to find a program by that name to execute it. The *path* is a list of directories in which commands can be found. It's defined by the `$PATH` environment

variable, as described shortly in "Using Environment Variables." A typical user account has about half a dozen or so directories in its path. You can add and remove directories to the shell's path by changing the `$PATH` environment variable in a shell configuration file, as described in "Exploring Shell Configuration" later in this chapter.

You can run programs that aren't on the path by providing a complete path name on the command line. For instance, typing **`./myprog`** runs the `myprog` program in the current directory. Typing **`/home/arthur/thisprog`** runs the `thisprog` program in the `/home/arthur` directory.

The `root` account should normally have a shorter path than ordinary user accounts. Typically, you'll omit directories that store GUI and other user-oriented programs from `root`'s path in order to discourage use of the `root` account for routine operations. This minimizes the risk of security breaches related to buggy or compromised binaries being run by `root`. Most important, `root`'s path should never include the current directory (`./`). Placing this directory in `root`'s path makes it possible for a local troublemaker to trick `root` into running replacements for common programs. Omitting the current directory from ordinary user paths is also generally a good idea. If this directory must be part of the ordinary user path, it should appear at the end of the path so that the standard programs take precedence over any replacement programs in the current directory.

Whether you need to enter the path or not for a command, the program file must be marked as executable. This is done via the execute bit that's stored with the file. Standard programs are marked as executable when they're installed, but if you need to adjust a program's executable status, you can do so with the `chmod` command, as described in Chapter 4, "Managing Files."

Performing Some Shell Command Tricks

Many users find typing commands to be tedious and error-prone. This is particularly true of slow or sloppy typists. For this reason, Linux shells include various tools that can help speed up operations. The first of these is *command completion*: Type part of a command or a filename (as an option to the command), and then press the Tab key. The shell tries to fill in the rest of the command or the filename. If just one command or filename matches the characters you've typed so far, the shell fills the rest of the command (or filename) for you and adds a space after it.

If the characters you've typed don't uniquely identify a command (or filename), the shell fills in what it can and then stops. Depending on the shell and its configuration, it may beep. If you press the Tab key again, the system responds by displaying the possible completions. You can then type another character or two and, if you haven't completed the command (or filename), press the Tab key again to have the process repeat.

The most fundamental Linux commands have fairly short names—`mv`, `ls`, `set`, and so on. However, some other commands are much longer, such as `traceroute` or

`service --status-all`. Filenames can also be quite lengthy—up to 255 characters on many filesystems. Thus command completion can save a lot of time when you're typing. It can also help you avoid typos.

The most popular Linux shells, including bash and tcsh, support command and filename completion. Some older shells, though, don't support this helpful feature.

Another useful shell shortcut is *history*. The shell history keeps a record of every command you type. If you've typed a long command recently and want to use it again or use a minor variant of it, you can pull the command out of the history.

There are several rather easy methods to retrieve commands. It comes down to determining the method you like best:

Retrieve a Command The simplest way to do this is to press the Up arrow key on your keyboard; this brings up the previous command. Pressing the Up arrow key repeatedly moves through multiple commands so you can find the one you want. If you overshoot, press the Down arrow key to move down the history. The Ctrl+P and Ctrl+N keystrokes double for the Up and Down arrow keys, respectively.

Search for a Command Press Ctrl+R to begin a backward (reverse) search, and begin typing characters that should be unique to the command you want to find. The characters you type need not be the ones that begin the command; they can exist anywhere in the command. You can either keep typing until you find the correct command or, after you've typed a few characters, press Ctrl+R repeatedly until you find the one you want.

The Ctrl+S keystroke is used to search forward in the command history. You can press the Ctrl+S keystroke *while* using the backward search. This reverses the history search from backward to forward. If you used a backward search and have passed by what you need, then this keystroke is useful.

If the Ctrl+S keystroke causes your terminal to hang, press Ctrl+Q to resume terminal operations. To keep your terminal from hanging when Ctrl+S is used, type **`stty -ixon`** at the command line.

In either event, if you can't find the command you want or if you change your mind and want to terminate the search, press Ctrl+G to do so.

Frequently, after finding a command in the history, you want to edit it. The bash shell, like many shells, provides editing features modeled after those of the Emacs editor:

Move within the Line Press Ctrl+A or Ctrl+E to move the cursor to the start or end of the line, respectively. The Left and Right arrow keys move within the line a character at a time. Ctrl+B and Ctrl+F do the same, moving backward and forward within a line. Pressing Ctrl plus the Left or Right arrow key moves backward or forward a word at a time, as does pressing Esc and then B or F.

Delete Text Pressing Ctrl+D or the Delete key deletes the character under the cursor. Pressing the Backspace key deletes the character to the left of the cursor. Pressing Ctrl+K deletes all text from the cursor to the end of the line. Pressing Ctrl+X and then Backspace deletes all of the text from the cursor to the beginning of the line.

Transpose Text Pressing Ctrl+T transposes the character before the cursor with the character under the cursor. Pressing Esc and then T transposes the two words immediately before (or under) the cursor.

Change Case Pressing Esc and then U converts text from the cursor to the end of the word to uppercase. Pressing Esc and then L converts text from the cursor to the end of the word to lowercase. Pressing Esc and then C converts the letter under the cursor (or the first letter of the next word) to uppercase, leaving the rest of the word unaffected.

Invoke an Editor You can launch a full-fledged editor to edit a command by pressing Ctrl+X followed by Ctrl+E. The `bash` shell attempts to launch the editor defined by the `$FCEDIT` or `$EDITOR` environment variable, or it launches Emacs as a last resort.

These editing commands are just the most useful ones supported by `bash`. In practice, you're likely to make heavy use of command and filename completion, the command history, and perhaps a few editing features.

If you prefer the vi editor to Emacs, you can use a vi-like `mode` in bash by typing **`set -o vi`**. (vi is described in Chapter 5, "Booting Linux and Editing Files.")

The `history` command provides an interface to view and manage the history. Typing **`history`** alone displays all of the commands in the history (typically the latest 500 commands).

To retrieve the last command in your shell history, type **`!!`** and press Enter. This will not only show you the command you recalled but execute it as well:

```
$ !!
type -a pwd
pwd is a shell builtin
pwd is /bin/pwd
$
```

You can execute a command by number via typing an exclamation mark followed by its number, as in **`!210`** to execute command 210. Typing **`history -c`** clears the history, which can be handy if you've recently typed commands you'd rather not have discovered by others, such as commands that include passwords.

The `bash` history is stored in the `.bash_history` file in your home directory. This is an ordinary plain-text file, so you can view it with a text editor or a command such as `less` (described later, in "Paging through Files with `less`").

Because your bash history is stored in a file, it can be examined by anybody who can read that file. Some commands enable you to type passwords or other sensitive data on the same line as the commands themselves, which can therefore be risky. The `~/.bash_history` file does *not* record what you type in response to other programs' prompts, just what you type at the bash prompt itself. Thus, if you have a choice, you should let commands that require passwords (or other sensitive data) prompt you to enter this data rather than enter such information as options to the command at the bash prompt.

In Exercise 1.1, you'll experiment with your shell's completion and command-editing tools.

EXERCISE 1.1

Editing Commands

To experiment with your shell's completion and command-editing tools, follow these steps:

1. Log in as an ordinary user.
2. Create a temporary directory by typing **`mkdir test`**. (Directory and file manipulation commands are described in more detail in Chapter 4.)
3. Change into the test directory by typing **`cd test`**.
4. Create a few temporary files by typing **`touch one two three`**. This command creates three empty files named `one`, `two`, and `three`.
5. Type **`ls -l t`** and, without pressing the Enter key, press the Tab key. The system may beep at you or display `two three`. If it doesn't display `two three`, press the Tab key again and it should do so. This reveals that either `two` or `three` is a valid completion to your command, because these are the two files in the test directory whose filenames begin with the letter `t`.
6. Type **h**, and again without pressing the Enter key, press the Tab key. The system should complete the command (`ls -l three`), at which point you can press the Enter key to execute it. (You'll see information on the file.)
7. Press the Up arrow key. You should see the `ls -l three` command appear on the command line.
8. Press Ctrl+A to move the cursor to the beginning of the line.
9. Press the Right arrow key once, and type **es** (without pressing the Enter key). The command line should now read `less -l three`.
10. Press the Right arrow key once, and press the Delete key three times. The command should now read `less three`. Press the Enter key to execute the command. (Note that you can do so even though the cursor isn't at the end of the line.) This invokes the `less` pager on the `three` file. (The `less` pager is described more fully later in

"Paging through Files with `less`.") Because this file is empty, you'll see a mostly empty screen.

11. Press the Q key to exit from the `less` pager.

Exploring Shell Configuration

Shells, like many Linux programs, are configured through files that hold configuration options in a plain-text format. The `bash` configuration files are actually `bash` shell scripts, which are described more fully in Chapter 9. A couple of examples of these configuration files are `~/.bashrc` and `/etc/profile`.

Even without knowing much about shell scripting, you can make simple changes to these files. Edit them in your favorite text editor, and change whatever needs changing. For instance, you can add directories to the `$PATH` environment variable, which takes a colon-delimited list of directories.

Be careful when changing your `bash` configuration files, particularly the global `bash` configuration files. Save a backup of the original file before making changes, and test your changes immediately by logging in using another virtual terminal. If you spot a problem, revert to your saved copy until you determine the problem's causes and create a working file.

Using Environment Variables

Environment variables are like variables in programming languages—they hold data to be referred to by the variable name. Environment variables differ from programs' internal variables in that they're part of the program's environment, and other programs, such as the shell, can modify this environment. Programs can rely on environment variables to set information that can apply to many different programs. For instance, many text-based programs need to know the capabilities of the terminal program you use. This information is conveyed in the `$TERM` environment variable, which is likely to hold a value such as `xterm` or `linux`. Programs that need to position the cursor, display color text, or perform other tasks that depend on terminal-specific capabilities can customize their output based on this information.

Chapter 9 describes environment variables and their manipulation in more detail. For the moment, you should know that you can set them in `bash` by using an assignment (`=`) operator followed by the `export` command. A fun environment variable to change is the `$PS1` variable. It modifies your shell prompt:

```
$
$ PS1="My New Prompt: "
My New Prompt: export PS1
My New Prompt:
```

You can combine these two commands into a single form:

```
My New Prompt: export PS1="Prompt: "
Prompt:
Prompt:
```

Either method sets the $PS1 environment variable to a new setting. When setting an environment variable, you omit the dollar sign, but subsequent references include a dollar sign to identify the environment variable as such. Thereafter, programs that need this information can refer to the environment variable. In fact, you can do so from the shell yourself using the echo command:

```
$ Prompt: echo $PS1
Prompt:
```

An echo of the $PS1 variable value can be a little confusing because it just shows your current prompt setting. However, you can get a better feel for displaying an environment variable by viewing the $PATH variable using echo:

```
Prompt: echo $PATH
/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:
/usr/local/sbin:/usr/sbin:/sbin:/home/Christine/bin
Prompt:
```

That's a little better. Remember, the $PATH environment variable provides the shell with a directory list to search when you're entering command or program names.

Some environment variables, including the $PATH environment variable, are set automatically when you log in via the shell configuration files. If a program uses environment variables, its documentation should say so.

You can also view the entire environment by typing **env**. The result is likely to be several dozen lines of environment variables and their values. Chapter 9 describes what many of these variables are in more detail.

To delete an environment variable, use the unset command. The command takes the name of an environment variable (without the leading $ symbol) as an option. For instance, **unset PS1** removes the $PS1 environment variable. But if you do this, you will have no shell prompt!

Getting Help

Linux provides a text-based help system known as man. This command's name is short for *manual*, and its entries (its man pages) provide succinct summaries of what a command, file, or other feature does. For instance, to learn about man itself, you can type **man man**. The result is a description of the man command.

To peruse the manual pages for a particular command or topic, you type **man** followed by the command or topic as an option. For example, to read about the `export` command, you would type **man export** at the prompt. If you wanted to learn more about the shell built-in (internal) commands, you would type **man builtin** at the prompt.

The `man` utility uses the `less` pager by default to display information. This program displays text a page at a time. Press the spacebar to move forward a page, Esc followed by V to move back a page, the arrow keys to move up or down a line at a time, the slash (/) key to search for text, and so on. (Type **man less** to learn all the details, or consult the upcoming section "Paging through Files with `less`.") When you're done, press Q to exit `less` and the `man` page it's displaying.

You aren't stuck using the `less` pager with the `man` utility. You can change the pager by using the `-P` option. For example, if you decided to use the `more` pager instead to look up information on the `uname` command, you would type **man -P /bin/more uname** at the shell prompt.

Occasionally, the problem arises where you can't remember the exact name of a command to look up. The `man` utility has an option to help you here. You can use the `-k` option along with a keyword or two to search through the `man` pages:

```
$ man -k "system information"
dumpe2fs  (8)  - dump ext2/ext3/ext4 filesystem information
[…]
uname     (1)  - print system information
$
```

The returned information (shown as a partial listing above) can give you some clues as to your desired command name. Be aware that poor keyword choices may not produce the results you seek.

On some older Linux distributions, you may get no results from a `man` utility keyword search. This is most likely due to a missing whatis database. The *whatis database* contains a short description of each `man` page, and it is necessary for keyword searches. To create it or update it, type **makewhatis** at the prompt. You will need to do this as superuser, and it may take several minutes to run.

Linux `man` pages are organized into several sections, which are summarized in Table 1.1. Sometimes a single keyword has entries in multiple sections. For instance, `passwd` has entries under both section 1 and section 5. In most cases, `man` returns the entry in the lowest-numbered section, but you can force the issue by preceding the keyword by the section number. For instance, typing **man 5 passwd** returns information on the `passwd` file format rather than the `passwd` command.

TABLE 1.1 Manual sections

Section number	Description
1	Executable programs and shell commands
2	System calls provided by the kernel
3	Library calls provided by program libraries
4	Device files (usually stored in `/dev`)
5	File formats
6	Games
7	Miscellaneous (macro packages, conventions, and so on)
8	System administration commands (programs run mostly or exclusively by `root`)
9	Kernel routines

Some programs have moved away from `man` pages to `info` pages. The basic purpose of `info` pages is the same as that for `man` pages. However, `info` pages use a hypertext format so that you can move from section to section of the documentation for a program. Type **`info info`** to learn more about this system.

There are also pages specifically for the built-in (internal) commands called the `help` pages. To read the `help` pages for a particular built-in command, type **`help`** ***`command`***. For instance, to get help on the `pwd` command, type **`help pwd`** at the shell prompt. To learn more about how to use the `help` pages, type **`help help`** at the shell prompt.

The `man` pages, `info` pages, and `help` pages are usually written in a terse style. They're intended as reference tools, not tutorials! They frequently assume basic familiarity with the command, or at least with Linux in general. For more tutorial information, you must look elsewhere, such in books or on the Web.

Using Streams, Redirection, and Pipes

Streams, *redirection*, and *pipes* are some of the more powerful command-line tools in Linux. Linux treats the input to and output from programs as a stream, which is a data entity that can be manipulated. Ordinarily, input comes from the keyboard and output goes to the screen. You can redirect these input and output streams to come from or go to other sources, such as files. Similarly, you can pipe the output of one program as input into another program. These facilities can be great tools to tie together multiple programs.

Part of the Unix philosophy to which Linux adheres is, whenever possible, to do complex things by combining multiple simple tools. Redirection and pipes help in this task by enabling simple programs to be combined together in chains, each link feeding off the output of the preceding link.

Exploring File Descriptors

To begin understanding redirection and pipes, you must first understand the different *file descriptors*. Linux handles all objects as files. This includes a program's input and output stream. To identify a particular file object, Linux uses file descriptors:

Standard Input Programs accept keyboard input via *standard input*, abbreviated STDIN. Standard input's file descriptor is `0` (zero). In most cases, this is the data that comes into the computer from a keyboard.

Standard Output Text-mode programs send most data to their users via *standard output*, abbreviated STDOUT. Standard output is normally displayed on the screen, either in a full-screen text-mode session or in a GUI terminal emulator, such as an `xterm`. Standard output's file descriptor is `1` (one).

Standard Error Linux provides a second type of output stream, known as *standard error*, abbreviated STDERR. Standard error's file descriptor is `2` (two). This output stream is intended to carry high-priority information such as error messages. Ordinarily, standard error is sent to the same output device as standard output, so you can't easily tell them apart. You can redirect one independently of the other, though, which can be handy. For instance, you can redirect standard error to a file while leaving standard output going to the screen. This allows you to view the error messages at a later time.

Internally, programs treat STDIN, STDOUT, and STDERR just like data files—they open them, read from or write to the files, and close them when they're done. This is why the file descriptors are necessary and why they can be used in redirection.

Redirecting Input and Output

To redirect input or output, you use operators following the command, including any options it takes. For instance, to redirect the STDOUT of the `echo` command, you would type something like this:

```
$ echo $PATH 1> path.txt
$
$ cat path.txt
/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:
/usr/local/sbin:/usr/sbin:/sbin:/home/Christine/bin
$
```

The result is that the file `path.txt` contains the output of the command (in this case, the value of the `$PATH` environment variable). The operator used to perform this redirection was `>` and the file descriptor used to redirect STDOUT was `1` (one).

The `cat` command allows you to display a file's contents to STDOUT. It is described further in the section "Processing Text Using Filters" later in this chapter.

A nice feature of redirecting STDOUT is that you do not have to use its file descriptor, only the operator. Here's an example of leaving out the `1` (one) file descriptor, when redirecting STDOUT:

```
$ echo $PATH > another_path.txt
$
$ cat another_path.txt
/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:
/usr/local/sbin:/usr/sbin:/sbin:/home/Christine/bin
$
```

You can see that even without the STDOUT file descriptor, the output was redirected to a file. However, the redirection operator (`>`) was still needed.

You can also leave out the STDIN file descriptor when using the appropriate redirection operator. Redirection operators exist to achieve several effects, as summarized in Table 1.2.

TABLE 1.2 Common redirection operators

Redirection operator	Effect
`>`	Creates a new file containing standard output. If the specified file exists, it's overwritten. No file descriptor necessary.
`>>`	Appends standard output to the existing file. If the specified file doesn't exist, it's created. No file descriptor necessary.
`2>`	Creates a new file containing standard error. If the specified file exists, it's overwritten. File descriptor necessary.
`2>>`	Appends standard error to the existing file. If the specified file doesn't exist, it's created. File descriptor necessary.
`&>`	Creates a new file containing both standard output and standard error. If the specified file exists, it's overwritten. No file descriptors necessary.
`<`	Sends the contents of the specified file to be used as standard input. No file descriptor necessary.
`<<`	Accepts text on the following lines as standard input. No file descriptor necessary.
`<>`	Causes the specified file to be used for both standard input and standard output. No file descriptor necessary.

Most of these redirectors deal with output, both because there are two types of output (standard output and standard error) and because you must be concerned with what to do in case you specify a file that already exists. The most important input redirector is <, which takes the specified file's contents as standard input.

A common trick is to redirect standard output or standard error to `/dev/null`. This file is a device that's connected to nothing; it's used when you want to get rid of data. For instance, if the `whine` program is generating too many unimportant error messages, you can type **`whine 2> /dev/null`** to run it and discard its error messages.

One redirection operator that requires elaboration is the << operator. This operator implements something called a *here document*. A *here document* takes text from subsequent lines as standard input. Chances are you won't use this redirector on the command line. Subsequent lines *are* standard input, so there's no need to redirect them. Rather, you might use this command in a script to pass data to an interactive program. Unlike with most redirection operators, the text immediately following the << code isn't a filename; instead, it's a word that's used to mark the end of input. For instance, typing **`someprog << EOF`** causes `someprog` to accept input until it sees a line that contains *only* the string `EOF` (without even a space following it).

Some programs that take input from the command line expect you to terminate input by pressing Ctrl+D. This keystroke corresponds to an end-of-file marker using the American Standard Code for Information Interchange (ASCII).

Piping Data between Programs

Programs can frequently operate on other programs' outputs. For instance, you might use a text-filtering command (such as the ones described shortly in "Processing Text Using Filters") to manipulate text output by another program. You can do this with the help of redirection operators: send the first program's standard output to a file, and then redirect the second program's standard input to read from that file. This method is awkward, though, and it involves the creation of a file that you might easily overlook, leading to unnecessary clutter on your system.

The solution is to use data pipes (aka pipelines). A pipe redirects the first program's standard output to the second program's standard input, and it is denoted by a vertical bar (|):

```
$ first | second
```

For instance, suppose that *`first`* generates some system statistics, such as system uptime, CPU use, number of users logged in, and so on. This output might be lengthy, so you want to trim it a bit. You might therefore use *`second`*, which could be a script or command that echoes from its standard input only the information in which you're interested. (The `grep` command, described in "Using `grep`," is often used in this role.)

Pipes can be used in sequences of arbitrary length:

```
$ first | second | third | fourth | fifth | sixth [...]
```

Another redirection tool often used with pipes is the `tee` command. This command splits standard input so that it's displayed on standard output and in as many files as you specify. Typically, `tee` is used in conjunction with data pipes so that a program's output can be both stored and viewed immediately. For instance, to view and store the output of the `echo $PATH` command, you might type this:

```
$ echo $PATH | tee path.txt
/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:
/usr/local/sbin:/usr/sbin:/sbin:/home/Christine/bin
$
$ cat path.txt
/usr/lib64/qt-3.3/bin:/usr/local/bin:/bin:/usr/bin:
/usr/local/sbin:/usr/sbin:/sbin:/home/Christine/bin
$
```

Notice that not only were the results of the command displayed to STDOUT, but they were also redirected to the `path.txt` file by the `tee` command. Ordinarily, `tee` overwrites any files whose names you specify. If you want to append data to these files, pass the `-a` option to `tee`.

Generating Command Lines

Sometimes you'll find yourself needing to conduct an unusual operation on your Linux server. For instance, suppose you want to remove every file in a directory tree that belongs to a certain user. With a large directory tree, this task can be daunting!

The usual file-deletion command, `rm` (described in more detail in Chapter 4), doesn't provide an option to search for and delete every file that matches a specific criterion. One command that can do the search portion is `find` (also described in more detail in Chapter 4). This command displays all of the files that match the criteria you provide. If you could combine the output of `find` to create a series of command lines using `rm`, the task would be solved. This is precisely the purpose of the `xargs` command.

The `xargs` command builds a command from its standard input. The basic syntax for this command is as follows:

```
xargs [options] [command [initial-arguments]]
```

The *`command`* is the command you want to execute, and *`initial-arguments`* is a list of arguments you want to pass to the command. The *`options`* are `xargs` options; they aren't passed to *`command`*. When you run `xargs`, it runs *`command`* once for every word passed to it on standard input, adding that word to the argument list for *`command`*. If you want to pass multiple options to the command, you can protect them by enclosing the group in quotation marks.

For instance, consider the task of deleting several files that belong to a particular user. You can do this by piping the output of `find` to `xargs`, which then calls `rm`:

```
# find / -user Christine | xargs -d "\n" rm
```

The first part of this command (`find / -user Christine`) finds all of the files in directory tree (`/`) and its subdirectories that belong to user `Christine`. (Since you are looking through the entire directory tree, you need superuser privileges for this to work properly.) This list is then piped to `xargs`, which adds each input value to its own `rm` command. Problems can arise if filenames contain spaces because by default `xargs` uses both spaces and newlines as item delimiters. The `-d "\n"` option tells `xargs` to use only newlines as delimiters, thus avoiding this problem in this context. (The `find` command separates each found filename with a newline.)

It is important to exercise caution when using the `rm` command with superuser privileges. This is especially true when piping the files to delete into the `rm` command. You could easily delete the wrong files unintentionally.

A tool that's similar to `xargs` in many ways is the backtick (`` ` ``), which is a character to the left of the 1 key on most keyboards. The backtick is *not* the same as the single quote character (`'`), which is located to the right of the semicolon (`;`) on most keyboards.

Text within backticks is treated as a separate command whose results are substituted on the command line. For instance, to delete those user files, you can type the following command:

```
# rm `find ./ -user Christine`
```

The backtick solution works fine in some cases, but it breaks down in more complex situations. The reason is that the output of the backtick-contained command is passed to the command it precedes as if it had been typed at the shell. By contrast, when you use `xargs`, it runs the command you specify (`rm` in these examples) once for each of the input items. What's more, you can't pass options such as `-d "\n"` to a backtick. Thus these two examples will work the same in many cases, but not in all of them.

Use of the backtick is falling out of favor because backticks are so often confused with single quotation marks. In several shells, you can use `$()` instead. For instance, the backtick example used in the preceding example would be changed to

```
# rm $(find ./ -user Christine)
```

This command works just as well, and it is much easier to read and understand.

Processing Text Using Filters

In keeping with Linux's philosophy of providing small tools that can be tied together via pipes and redirection to accomplish more complex tasks, many simple commands to manipulate text are available. These commands accomplish tasks of various types, such as combining files, transforming the data in files, formatting text, displaying text, and summarizing data.

Many of the following descriptions include input-file specifications. In most cases, you can omit these input-file specifications, in which case the utility reads from standard input instead.

File-Combining Commands

The first text-filtering commands are those used to combine two or more files into one file. Three important commands in this category are `cat`, `join`, and `paste`, which join files end to end based on fields in the file or by merging on a line-by-line basis.

Combining Files with *cat*

The `cat` command's name is short for *concatenate*, and this tool does just that: It links together an arbitrary number of files end to end and sends the result to standard output. By combining `cat` with output redirection, you can quickly combine two files into one:

```
$ cat first.txt second.txt > combined.txt
$
$ cat first.txt
Data from first file.
$
$ cat second.txt
Data from second file.
$
$ cat combined.txt
Data from first file.
Data from second file.
$
```

Although `cat` is officially a tool for combining files, it's also commonly used to display the contents of a short file to STDOUT. If you type only one filename as an option, `cat` displays that file. This is a great way to review short files; but for long files, you're better off using a full-fledged pager command, such as `more` or `less`.

You can add options to have `cat` perform minor modifications to the files as it combines them:

Display Line Ends If you want to see where lines end, add the `-E` or `--show-ends` option. The result is a dollar sign (`$`) at the end of each line.

Number Lines The `-n` or `--number` option adds line numbers to the beginning of every line. The `-b` or `--number-nonblank` option is similar, but it numbers only lines that contain text.

Minimize Blank Lines The `-s` or `--squeeze-blank` option compresses groups of blank lines down to a single blank line.

Display Special Characters The `-T` or `--show-tabs` option displays tab characters as `^I`. The `-v` or `--show-nonprinting` option displays most control and other special characters using carat (`^`) and `M-` notations.

The `tac` command is similar to `cat`, but it reverses the order of lines in the output:

```
$ cat combined.txt
Data from first file.
Data from second file.
$
$ tac combined.txt
Data from second file.
Data from first file.
$
```

Joining Files by Field with *join*

The `join` command combines two files by matching the contents of specified fields within the files. Fields are typically space-separated entries on a line. However, you can specify another character as the field separator with the `-t` *char* option, where *char* is the character you want to use. You can cause `join` to ignore case when performing comparisons by using the `-i` option.

The effect of `join` may best be understood through a demonstration. Consider Listing 1.1 and Listing 1.2, which contain data on telephone numbers. Listing 1.1 shows the names associated with those numbers, and Listing 1.2 shows whether the numbers are listed or unlisted.

Listing 1.1: Demonstration file containing telephone numbers and names

```
555-2397 Beckett, Barry
555-5116 Carter, Gertrude
555-7929 Jones, Theresa
555-9871 Orwell, Samuel
```

Listing 1.2: Demonstration file containing telephone number listing status

```
555-2397 unlisted
555-5116 listed
555-7929 listed
555-9871 unlisted
```

You can display the contents of both files using `join`:

```
$ join listing1.1.txt listing1.2.txt
555-2397 Beckett, Barry unlisted
555-5116 Carter, Gertrude listed
555-7929 Jones, Theresa listed
555-9871 Orwell, Samuel unlisted
```

By default, `join` uses the first field as the one to match across files. Because Listing 1.1 and Listing 1.2 both place the phone number in this field, it's the key field in the output. You can specify another field by using the `-1` or `-2` option to indicate the join field for the first or second file, respectively. For example, type **`join -1 3 -2 2 cameras.txt lenses.txt`** to join using the third field in `cameras.txt` and the second field in `lenses.txt`. The `-o` *`FORMAT`* option enables more complex specifications for the output file's format. You can consult the `man` page for `join` for even more details.

The `join` command can be used at the core of a set of simple customized database-manipulation tools using Linux text-manipulation commands. It's very limited by itself, though. For instance, it requires its two files to have the same ordering of lines. (You can use the `sort` command to ensure this is so.)

Merging Lines with *paste*

The `paste` command merges files line by line, separating the lines from each file with tabs, as shown in the following example, using Listings 1.1 and 1.2 again:

```
$ paste listing1.1.txt listing1.2.txt
555-2397 Beckett, Barry      555-2397 unlisted
555-5116 Carter, Gertrude    555-5116 listed
555-7929 Jones, Theresa      555-7929 listed
555-9871 Orwell, Samuel      555-9871 unlisted
```

You can use `paste` to combine data from files that aren't keyed with fields suitable for use by `join`. Of course, to be meaningful, the files' line numbers must be exactly equivalent. Alternatively, you can use `paste` as a quick way to create a two-column output of textual data; however, the alignment of the second column may not be exact if the first column's line lengths aren't exactly even.

File-Transforming Commands

Many of Linux's text-manipulation commands are aimed at transforming the contents of files. These commands don't actually change files' contents but instead send the changed

files' contents to standard output. You can then pipe this output to another command or redirect it into a new file.

An important file-transforming command is sed. This command is very complex and is covered later in this chapter in "Using sed."

Converting Tabs to Spaces with *expand*

Sometimes text files contain tabs but programs that need to process the files don't cope well with tabs. In such a case, you may want to convert tabs to spaces. The expand command does this.

By default, expand assumes a tab stop every eight characters. You can change this spacing with the `-t` *num* or `--tabs=`*num* option, where *num* is the tab spacing value.

Displaying Files in Octal with *od*

Some files aren't easily displayed in ASCII. For example, most graphics files, audio files, and so on use non-ASCII characters that look like gibberish. Worse, these characters can do strange things to your display if you try to view such a file with cat or a similar tool. For instance, your font may change, or your console may begin beeping uncontrollably. Nonetheless, you may sometimes want to display such files, particularly if you want to investigate the structure of a data file.

In such a case, od (whose name stands for *octal dump*) can help. It displays a file in an unambiguous format—octal (base 8) numbers by default. For instance, consider Listing 1.2 as parsed by od:

```
$ od listing1.2.txt
0000000 032465 026465 031462 033471 072440 066156 071551 062564
0000020 005144 032465 026465 030465 033061 066040 071551 062564
0000040 005144 032465 026465 034467 034462 066040 071551 062564
0000060 005144 032465 026465 034071 030467 072440 066156 071551
0000100 062564 005144
0000104
```

The first field on each line is an index into the file in octal. For instance, the second line begins at octal 20 (16 in base 10) bytes into the file. The remaining numbers on each line represent the bytes in the file. This type of output can be difficult to interpret unless you're well versed in octal notation and perhaps in the ASCII code.

Although od is nominally a tool for generating octal output, it can generate many other output formats, such as hexadecimal (base 16), decimal (base 10), and even ASCII with escaped control characters. Consult the man page for od for details on creating these variants.

Sorting Files with *sort*

Sometimes you'll create an output file that you want sorted. To do so, you can use a command that's called, appropriately enough, `sort`. This command can sort in several ways, including the following:

Ignore Case Ordinarily, `sort` sorts by ASCII value, which differentiates between uppercase and lowercase letters. The `-f` or `--ignore-case` option causes `sort` to ignore case.

Month Sort The `-M` or `--month-sort` option causes the program to sort by three-letter month abbreviation (`JAN` through `DEC`).

Numeric Sort You can sort by number by using the `-n` or `--numeric-sort` option.

Reverse Sort Order The `-r` or `--reverse` option sorts in reverse order.

Sort Field By default, `sort` uses the first field as its sort field. You can specify another field with the `-k` *`field`* or `--key=`*`field`* option. (The *`field`* can be two numbered fields separated by commas, to sort on multiple fields.)

As an example, suppose you wanted to sort Listing 1.1 by first name. You could do so like this:

```
$ sort -k 3 listing1.1.txt
555-2397 Beckett, Barry
555-5116 Carter, Gertrude
555-9871 Orwell, Samuel
555-7929 Jones, Theresa
```

The `sort` command supports a large number of additional options, many of them quite exotic. Consult `sort`'s `man` page for details.

Breaking a File into Pieces with *split*

The `split` command can split a file into two or more files. Unlike most of the text-manipulation commands described in this chapter, this command requires you to enter an output filename or, more precisely, an output filename prefix, to which is added an alphabetic code. You must also normally specify how large you want the individual files to be:

Split by Bytes The `-b` *`size`* or `--bytes=`*`size`* option breaks the input file into pieces of *`size`* bytes. This option can have the usually undesirable consequence of splitting the file mid-line.

Split by Bytes in Line-Sized Chunks You can break a file into files of no more than a specified size without breaking lines across files by using the `-C=`*`size`* or `--line-bytes=`*`size`* option. (Lines will still be broken across files if the line length is greater than *`size`*.)

Split by Number of Lines The `-l` *`lines`* or `--lines=`*`lines`* option splits the file into chunks with no more than the specified number of lines.

As an example, consider breaking Listing 1.1 into two parts by number of lines:

```
$ split -l 2 listing1.1.txt numbers
```

The result is two files, `numbersaa` and `numbersab`, which together hold the original contents of `listing1.1.txt`.

If you don't specify any defaults (as in **`split listing1.1.txt`**), the result is output files split into 1,000-line chunks, with names beginning with `x` (`xaa`, `xab`, and so on). If you don't specify an input filename, `split` uses standard input.

Translating Characters with *tr*

The `tr` command changes individual characters from standard input. Its syntax is as follows:

```
tr [options] SET1 [SET2]
```

You specify the characters you want replaced in a group (*SET1*) and the characters with which you want them to be replaced as a second group (*SET2*). Each character in *SET1* is replaced with the one at the equivalent position in *SET2*. Here's an example using Listing 1.1:

```
$ tr BCJ bc < listing1.1.txt
555-2397 beckett, barry
555-5116 carter, Gertrude
555-7929 cones, Theresa
555-9871 Orwell, Samuel
```

The `tr` command relies on standard input, which is the reason for the input redirection (<) in this example. This is the only way to pass the command a file.

This example translates some, but not all, of the uppercase characters to lowercase. Note that *SET2* in this example was shorter than *SET1*. The result is that `tr` substitutes the last available letter from *SET2* for the missing letters. In this example, the `J` in `Jones` became a `c`. The `-t` or `--truncate-set1` option causes `tr` to truncate *SET1* to the size of *SET2* instead.

Another `tr` option is `-d`, which causes the program to delete the characters from *SET1*. When using `-d`, you omit *SET2* entirely.

The `tr` command also accepts a number of shortcuts, such as `[:alnum:]` (all numbers and letters), `[:upper:]` (all uppercase letters), `[:lower:]` (all lowercase letters), and `[:digit:]` (all digits). You can specify a range of characters by separating them with dashes (`-`), as in `A-M` for characters between `A` and `M`, inclusive. Consult `tr`'s `man` page for a complete list of these shortcuts.

Converting Spaces to Tabs with *unexpand*

The `unexpand` command is the logical opposite of `expand`; it converts multiple spaces to tabs. This can help compress the size of files that contain many spaces and can be helpful if a file is to be processed by a utility that expects tabs in certain locations.

Like expand, unexpand accepts the `-t` *num* or `--tabs=`*num* option, which sets the tab spacing to once every *num* characters. If you omit this option, unexpand assumes a tab stop every eight characters.

Deleting Duplicate Lines with *uniq*

The uniq command removes duplicate lines. It's most likely to be useful if you've sorted a file and don't want duplicate items. For instance, suppose you want to summarize Shakespeare's vocabulary. You might create a file with all of the Bard's works, one word per line. You can then sort this file using sort and pass it through uniq. Using a shorter example file containing the text `to be or not to be, that is the question` (one word per line), the result looks like this:

```
$ sort shakespeare.txt | uniq
be
is
not
or
question
that
the
to
```

Note that the words to and be, which appeared in the original file twice, appear only once in the uniq-processed version.

File-Formatting Commands

The next three commands—fmt, nl, and pr—reformat the text in a file. The first of these is designed to reformat text files, such as when a program's README documentation file uses lines that are too long for your display. The nl command numbers the lines of a file, which can be helpful in referring to lines in documentation or correspondence. Finally, pr is a print-processing tool; it formats a document in pages suitable for printing.

Reformatting Paragraphs with *fmt*

Sometimes text files arrive with outrageously long line lengths, irregular line lengths, or other problems. Depending on the difficulty, you may be able to cope simply by using an appropriate text editor or viewer to read the file. If you want to clean up the file a bit, though, you can do so with fmt. If called with no options (other than the input filename, if you're not having it work on standard input), the program attempts to clean up paragraphs, which it assumes are delimited by two or more blank lines or by changes in indentation. The new paragraph formatting defaults to paragraphs that are no more than 75 characters wide. You can change this with the `-`*width*, `-w` *width*, and `--width=`*width* options, which set the line length to *width* characters.

Numbering Lines with *nl*

As described earlier, in "Combining Files with `cat`," you can number the lines of a file with that command. The `cat` line-numbering options are limited, though, if you need to do complex line numbering. The `nl` command is the tool to use in this case. In its simplest form, you can use `nl` alone to accomplish much the same goal as `cat -b` achieves: numbering all the non-blank lines in a file. You can add many options to `nl` to achieve various special effects:

Body Numbering Style You can set the numbering style for the bulk of the lines with the `-b` *`style`* or `--body-numbering=`*`style`* option, where *`style`* is a style format code, described shortly.

Header and Footer Numbering Style If the text is formatted for printing and has headers or footers, you can set the style for these elements with the `-h` *`style`* or `--header-numbering=`*`style`* option for the header and `-f` *`style`* or `--footer-numbering=`*`style`* option for the footer.

Page Separator Some numbering schemes reset the line numbers for each page. You can tell `nl` how to identify a new page with the `-d=`*`code`* or `--section-delimiter=`*`code`* option, where *`code`* is a code for the character that identifies the new page.

Line-Number Options for New Pages Ordinarily, `nl` begins numbering each new page with line 1. If you pass the `-p` or `--no-renumber` option, though, it doesn't reset the line number with a new page.

Number Format You can specify the numbering format with the `-n` *`format`* or `--number-format=`*`format`* option, where *`format`* is `ln` (left justified, no leading zeros), `rn` (right justified, no leading zeros), or `rz` (right justified with leading zeros).

The body, header, and footer options enable you to specify a numbering style for each of these page elements, as described in Table 1.3.

TABLE 1.3 Styles used by `nl`

Style code	Description
`t`	The default behavior is to number lines that aren't empty. You can make this default explicit by using a style code of `t`.
`a`	This style code causes all lines to be numbered, including empty lines.
`n`	This style code causes all line numbers to be omitted, which may be desirable for headers or footers.
`p`*`REGEXP`*	This option causes only lines that match the specified regular expression (*`REGEXP`*) to be numbered. Regular expressions are described later in "Using Regular Expressions."

As an example, suppose you've created a script, `buggy`, but you find that it's not working as you expect. When you run it, you get error messages that refer to line numbers, so you want to create a version of the script with lines that are numbered for easy reference. You can do so by calling `nl` with the option to number all lines, including blank lines (`-b a`):

```
$ nl -b a buggy > numbered-buggy.txt
```

Because the input file doesn't have any explicit page delimiters, the output will be numbered in a single sequence. The `nl` command doesn't try to impose its own page-length limits.

The `numbered-buggy.txt` file created by this command isn't useful as a script because of the line numbers that begin each line. You can, however, load it into a text editor or display it with a pager such as `less` to view the text and see the line numbers along with the commands they contain.

Preparing a File for Printing with *pr*

If you want to print a plain-text file, you may want to prepare it with headers, footers, page breaks, and so on. The `pr` command was designed to do this. In its most basic form, you pass the command a file:

```
$ pr myfile.txt
```

The result is text formatted for printing on a line printer—that is, `pr` assumes an 80-character line length in a monospaced font. Of course, you can also use `pr` in a pipe, either to accept input piped from another program or to pipe its output to another program. (The recipient program might be `lpr`, which is used to print files, as described in Chapter 6, "Configuring the X Window System, Localization, and Printing.")

By default, `pr` creates output that includes the original text with headers, which lists the current date and time, the original filename, and the page number. You can tweak the output format in a variety of ways, including the following:

Generate Multicolumn Output Passing the `-`*`numcols`* or `--columns=`*`numcols`* option creates output with *`numcols`* columns. For example, if you typed **`pr -3 myfile.txt`**, the output would be displayed in three columns. Note that `pr` doesn't reformat text; if lines are too long, they're truncated or run over into multiple columns.

Generate Double-Spaced Output The `-d` or `--double-space` option causes double-spaced output from a single-spaced file.

Use Form Feeds Ordinarily, `pr` separates pages by using a fixed number of blank lines. This works fine if your printer uses the same number of lines that `pr` expects. If you have problems with this issue, you can pass the `-F`, `-f`, or `--form-feed` option, which causes `pr` to output a form-feed character between pages. This works better with some printers.

Set Page Length The `-l` *`lines`* or `--length=`*`lines`* option sets the length of the page in lines.

Set the Header Text The `-h` *`text`* or `--header=`*`text`* option sets the text to be displayed in the header, replacing the filename. To specify a multi-word string, enclose it in quotes, as in `--header="My File"`. The `-t` or `--omit-header` option omits the header entirely.

Set Left Margin and Page Width The `-o` *`chars`* or `--indent=`*`chars`* option sets the left margin to *`chars`* characters. This margin size is added to the page width, which defaults to 72 characters and can be explicitly set with the `-w` *`chars`* or `--width` *`chars`* option.

These options are just the beginning; `pr` supports many more options, which are described in its `man` page. As an example of `pr` in action, consider printing a double-spaced and numbered version of a configuration file (say, `/etc/profile`) for your reference. You can do this by piping together `cat` and its `-n` option to generate a numbered output, `pr` and its `-d` option to double-space the result, and `lpr` to print the file:

```
$ cat -n /etc/profile | pr -d | lpr
```

The result should be a printout that might be handy for taking notes on the configuration file. One caveat, though: If the file contains lines that approach or exceed 80 characters in length, the result can be single lines that spill across two lines. The result will be disrupted page boundaries. As a workaround, you can set a somewhat short page length with `-l` and use `-f` to ensure that the printer receives form feeds after each page:

```
$ cat -n /etc/profile | pr -dfl 50 | lpr
```

The `pr` command is built around assumptions about printer capabilities that were reasonable in the early 1980s. It's still useful today, but you might prefer to look into GNU Enscript (`www.codento.com/people/mtr/genscript/`). This program has many of the same features as `pr`, but it generates PostScript output that can take better advantage of modern printer features.

File-Viewing Commands

Sometimes you just want to view a file or part of a file. A few commands can help you accomplish this goal without loading the file into a full-fledged editor.

As described earlier, the `cat` command is also handy for viewing short files.

Viewing the Starts of Files with *head*

Sometimes all you need to do is see the first few lines of a file. This may be enough to identify what a mystery file is, for instance; or you may want to see the first few entries of a log file to determine when that file was started. You can accomplish this goal with the `head`

command, which echoes the first 10 lines of one or more files to standard output. (If you specify multiple filenames, each one's output is preceded by a header to identify it.) You can modify the amount of information displayed by `head` in two ways:

Specify the Number of Bytes The `-c` *num* or `--bytes=`*num* option tells `head` to display *num* bytes from the file rather than the default 10 lines.

Specify the Number of Lines You can change the number of lines displayed with the `-n` *num* or `--lines=`*num* option.

Viewing the Ends of Files with *tail*

The `tail` command works just like `head`, except that `tail` displays the *last* 10 lines of a file. (You can use the `-c` or `--bytes`, and `-n` or `--lines` options to change the amount of data displayed, just as with `head`.) This command is useful for examining recent activity in log files or other files to which data may be appended.

The `tail` command supports several options that aren't present in `head` and that enable the program to handle additional duties, including the following:

Track a File The `-f` or `--follow` option tells `tail` to keep the file open and to display new lines as they're added. This feature is helpful for tracking log files because it enables you to see changes as they're made to the file.

Stop Tracking on Program Termination The `--pid=`*pid* option tells `tail` to terminate tracking (as initiated by `-f` or `--follow`) once the process with a process ID (PID) of *pid* terminates. (PIDs are described in more detail in Chapter 2, "Managing Software.")

Some additional options provide more obscure capabilities. Consult `tail`'s `man` page for details.

You can combine `head` with `tail` to display or extract portions of a file. For instance, suppose you want to display lines 11 through 15 of a file, `sample.txt`. You can extract the first 15 lines of the file with `head` and then display the last five lines of that extraction with `tail`. The final command would be **`head -n 15 sample.txt | tail -n 5`**.

Paging through Files with *less*

The `less` command's name is a joke; it's a reference to the `more` command, which was an early file pager. The idea was to create a better version of `more`, so the developers called it `less` ("less is more").

The idea behind `less` (and `more`, for that matter) is to enable you to read a file a screen at a time. When you type **`less` *`filename`***, the program displays the first few lines of *filename*. You can then page back and forth through the file:

- Pressing the spacebar moves forward through the file a screen at a time.
- Pressing Esc followed by V moves backward through the file a screen at a time.

- The Up and Down arrow keys move up or down through the file a line at a time.
- You can search the file's contents by pressing the slash (/) key followed by the search term. For instance, typing **/portable** finds the first occurrence of the string portable after the current position. Typing a slash followed by the Enter key moves to the next occurrence of the search term. Typing **n** alone repeats the search forward, while typing **N** alone repeats the search backward.
- You can search backward in the file by using the question mark (?) key rather than the slash key.
- You can move to a specific line by typing **g** followed by the line number, as in **g50** to go to line 50.
- When you're done, type **q** to exit from the program.

Unlike most of the programs described here, less can't be readily used in a pipe, except as the final command in the pipe. In that role, though, less is very useful because it enables you to examine lengthy output conveniently.

Although less is quite common on Linux systems and is typically configured as the default text pager, some Unix-like systems use more in this role. Many of less's features, such as the ability to page backward in a file, don't work in more.

One additional less feature can be handy: Typing **h** displays less's internal help system. This display summarizes the commands you may use, but it's long enough that you must use the usual less paging features to view it all! When you're done with the help screens, just type **q** as if you were exiting from viewing a help document with less. This action will return you to your original document.

File-Summarizing Commands

The final text-filtering commands described here are used to summarize text in one way or another. The cut command takes segments of an input file and sends them to standard output, while the wc command displays some basic statistics on the file.

Extracting Text with *cut*

The cut command extracts portions of input lines and displays them on standard output. You can specify what to cut from input lines in several ways:

By Byte The -b *list* or --bytes=*list* option cuts the specified list of bytes from the input file. (The format of *list* is described shortly.)

By Character The -c *list* or --characters=*list* option cuts the specified list of characters from the input file. In practice, this method and the by-byte method usually produce identical results. (If the input file uses a multibyte encoding system, though, the results won't be identical.)

By Field The `-f` *`list`* or `--fields=`*`list`* option cuts the specified list of fields from the input file. By default, a field is a tab-delimited section of a line, but you can change the delimiting character with the `-d` *`char`*, `--delim=`*`char`*, or `--delimiter=`*`char`* option, where *`char`* is the character you want to use to delimit fields. Ordinarily, `cut` echoes lines that don't contain delimiters. Including the `-s` or `--only-delimited` option changes this behavior so that the program doesn't echo lines that don't contain the delimiter character.

Many of these options take a *`list`* option, which is a way to specify multiple bytes, characters, or fields. You make this specification by number. It can be a single number (such as `4`), a closed range of numbers (such as `2-4`), or an open range of numbers (such as `-4` or `4-`). In this final case, all bytes, characters, or fields from the beginning of the line to the specified number (or from the specified number to the end of the line) are included in the list.

The `cut` command is frequently used in scripts to extract data from some other command's output. For instance, suppose you're writing a script and the script needs to know the hardware address of your Ethernet adapter. This information can be obtained from the `ifconfig` command (described in more detail in Chapter 8, "Configuring Basic Networking"):

```
$ ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:0C:76:96:A3:73
          inet addr:192.168.1.3  Bcast:192.168.1.255
Mask:255.255.255.0
          inet6 addr: fe80::20c:76ff:fe96:a373/64 Scope:Link
          UP BROADCAST NOTRAILERS RUNNING MULTICAST MTU:1500
Metric:1
          RX packets:7127424 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5273519 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:6272843708 (5982.2 Mb)  TX bytes:1082453585 (1032.3 Mb)
          Interrupt:10 Base address:0xde00
```

Unfortunately, most of this information is extraneous for the desired purpose. The hardware address is the 6-byte hexadecimal number following `HWaddr`. To extract that data, you can combine `grep` (described shortly in "Using `grep`") with `cut` in a pipe:

```
$ ifconfig eth0 | grep HWaddr | cut -d " " -f 11
00:0C:76:96:A3:73
```

Of course, in a script, you would probably assign this value to a variable or otherwise process it through additional pipes. (Chapter 9 describes scripts in more detail.)

Obtaining a Word Count with *wc*

The `wc` command produces a word count (that's where it gets its name), as well as line and byte counts, for a file:

```
$ wc file.txt
  308  2343  15534 file.txt
```

This file contains 308 lines (or more precisely, 308 newline characters), 2,343 words, and 15,534 bytes. You can limit the output to the newline count, the word count, the byte count, or a character count with the `--lines` (`-l`), `--words` (`-w`), `--bytes` (`-c`), or `--chars` (`-m`) option, respectively. You can also learn the maximum line length with the `--max-line-length` (`-L`) option.

For an ordinary ASCII file, the character and byte counts will be identical. These values may diverge for files that use multibyte character encodings.

Using Regular Expressions

Many Linux programs employ *regular expressions*, which are tools for describing or matching patterns in text. Regular expressions are similar in principle to the wildcards that can be used to specify multiple filenames. At their simplest, regular expressions can be plain text without adornment. However, certain characters are used to denote patterns. Because of their importance, regular expressions are described in the following section.

Two programs that make heavy use of regular expressions, `grep` and `sed`, are also covered. These programs search for text within files and permit editing of files from the command line, respectively.

Understanding Regular Expressions

Two forms of regular expression are common: basic and extended. Which form you must use depends on the program. Some accept one form or the other, but others can use either type, depending on the options passed to the program. (Some programs use their own minor or major variants on either of these classes of regular expression.) The differences between basic and extended regular expressions are complex and subtle, but the fundamental principles of both are similar.

The simplest type of regular expression is an alphabetic string, such as `Linux` or `HWaddr`. These regular expressions match any string of the same size or longer that contains the regular expression. For instance, the `HWaddr` regular expression matches `HWaddr`, `This is the HWaddr`, and `The HWaddr is unknown`. The real strength of regular expressions comes in the use of nonalphabetic characters, which activate advanced matching rules:

Bracket Expressions Characters enclosed in square brackets (`[]`) constitute bracket expressions, which match any one character within the brackets. For instance, the regular expression `b[aeiou]g` matches the words `bag`, `beg`, `big`, `bog`, and `bug`.

Range Expressions A range expression is a variant of a bracket expression. Instead of listing every character that matches, range expressions list the start and end points separated by a dash (`-`), as in `a[2-4]z`. This regular expression matches `a2z`, `a3z`, and `a4z`.

Any Single Character The dot (`.`) represents any single character except a newline. For instance, `a.z` matches `a2z`, `abz`, `aQz`, or any other three-character string that begins with `a` and ends with `z`.

Start and End of Line The carat (`^`) represents the start of a line, and the dollar sign (`$`) denotes the end of a line.

Repetition Operators A full or partial regular expression may be followed by a special symbol to denote how many times a matching item must exist. Specifically, an asterisk (`*`) denotes zero or more occurrences, a plus sign (`+`) matches one or more occurrences, and a question mark (`?`) specifies zero or one match. The asterisk is often combined with the dot (as in `.*`) to specify a match with any substring. For instance, `A.*Lincoln` matches any string that contains `A` and `Lincoln`, in that order—`Abe Lincoln` and `Abraham Lincoln` are just two possible matches.

Multiple Possible Strings The vertical bar (`|`) separates two possible matches; for instance, `car|truck` matches either `car` or `truck`.

Parentheses Ordinary parentheses (`()`) surround subexpressions. Parentheses are often used to specify how operators are to be applied; for example, you can put parentheses around a group of words that are concatenated with the vertical bar to ensure that the words are treated as a group, any one of which may match, without involving surrounding parts of the regular expression.

Escaping If you want to match one of the special characters, such as a dot, you must *escape* it—that is, precede it with a backslash (`\`). For instance, to match a computer hostname (say, `twain.example.com`), you must escape the dots, as in `twain\.example\.com`.

The preceding descriptions apply to extended regular expressions. Some details are different for basic regular expressions. In particular, the `?`, `+`, `|`, `(`, and `)` symbols lose their special meanings. To perform the tasks handled by these characters, some programs, such as `grep`, enable you to recover the functions of these characters by escaping them (say, using `\|` instead of `|`). Whether you use basic or extended regular expressions depends on which form the program supports. For programs such as `grep`, which support both, you can use either. Which form you choose is mostly a matter of personal preference.

You can get more help on regular expressions at the command-line by typing **`man 7 regex`**. The certification objectives list this particular man page as `regex(7)`.

Regular expression rules can be confusing, particularly when you're first introduced to them. Some examples of their use, in the context of the programs that use them, will help. The next couple of sections provide such examples.

Using *grep*

The `grep` command is extremely useful. It searches for files that contain a specified string and returns the name of the file and (if it's a text file) a line of context for that string. The basic `grep` syntax is as follows:

```
grep [options] regexp [files]
```

The *regexp* is a regular expression, as just described. The grep command supports a large number of options. Some of the common options enable you to modify the way the program searches files:

Count Matching Lines Instead of displaying context lines, grep displays the number of lines that match the specified pattern if you use the -c or --count option.

Specify a Pattern Input File The -f *file* or --file=*file* option takes pattern input from the specified file rather than from the command line.

Ignore Case You can perform a search that isn't case sensitive, rather than the default case-sensitive search, by using the -i or --ignore-case option.

Search Recursively The -r or --recursive option searches in the specified directory and all subdirectories rather than simply the specified directory. You can use rgrep rather than specify this option.

Use a Fixed Strings Pattern If you want to turn off the grep command's use of regular expressions and use basic pattern searching instead, you can use the -F or --fixed-strings option. Alternatively, you can use fgrep rather than grep. Either way, the characters in the basic pattern string are treated literally. For example, $ is treated literally as a $ and not as a regular expression.

Use an Extended Regular Expression The grep command interprets *regexp* as a basic regular expression by default. To use an extended regular expression, you can pass the -E or --extended-regexp option. Alternatively, you can call egrep rather than grep. This variant command uses extended regular expressions by default.

A simple example of grep uses a regular expression with no special components:

```
$ grep -r eth0 /etc/*
```

This example finds all the files in /etc that contain the string eth0 (the identifier for the first wired Ethernet device on most Linux distributions). Because the example includes the -r option, it searches recursively, so files in subdirectories of /etc are examined in addition to those in /etc itself. For each matching text file, the line that contains the string is printed.

Some files in /etc can't be read by ordinary users. Thus if you type this command as a non-root user, you'll see some error messages relating to grep's inability to open files.

Suppose you want to locate all the files in /etc that contain the string eth0 or eth1. You can enter the following command, which uses a bracket expression to specify both variant devices:

```
$ grep eth[01] /etc/*
```

A still more complex example searches all files in /etc that contain the hostname twain.example.com or bronto.pangaea.edu and, later on the same line, the number

127. This task requires using several of the regular expression features. Expressed using extended regular expression notation, the command looks like this:

```
$ grep -E "(twain\.example\.com|bronto\.pangaea\.edu).*127" /etc/*
```

This command illustrates another feature you may need to use: shell quoting. Because the shell uses certain characters, such as the vertical bar and the asterisk, for its own purposes, you must enclose certain regular expressions in quotes lest the shell attempt to parse the regular expression and pass a modified version of what you type to `grep`.

You can use `grep` in conjunction with commands that produce a lot of output in order to sift through that output for the material that's important to you. (Several examples throughout this book use this technique.) For example, suppose you want to find the process ID (PID) of a running `xterm`. You can use a pipe to send the result of a `ps` command (described in Chapter 2) through `grep`:

```
# ps ax | grep xterm
```

The result is a list of all running processes called `xterm`, along with their PIDs. You can even do this in series, using `grep` to restrict further the output on some other criterion, which can be useful if the initial pass still produces too much output.

Using *sed*

The `sed` command directly modifies a file's contents, sending the changed file to standard output. Its syntax can take one of two forms:

```
sed [options] -f script-file [input-file]
sed [options] script-text [input-file]
```

In either case, *`input-file`* is the name of the file you want to modify. (Modifications are temporary unless you save them in some way, as illustrated shortly.) The script (*`script-text`* or the contents of *`script-file`*) is the set of commands you want `sed` to perform. When you pass a script directly on the command line, the *`script-text`* is typically enclosed in single quote marks. Table 1.4 summarizes a few `sed` commands that you can use in its scripts.

TABLE 1.4 Common `sed` commands

Command	Addresses	Meaning
`=`	0 or 1	Display the current line number.
`a\`*`text`*	0 or 1	Append *`text`* to the file.
`i\`*`text`*	0 or 1	Insert *`text`* into the file.

Command	Addresses	Meaning
`r` *filename*	0 or 1	Append text from *filename* into the file.
`c\`*text*	Range	Replace the selected range of lines with the provided *text*.
`s/`*regexp*`/`*replacement*	Range	Replace text that matches the regular expression (*regexp*) with *replacement*.
`w` *filename*	Range	Write the current pattern space to the specified file.
`q`	0 or 1	Immediately quit the script, but print the current pattern space.
`Q`	0 or 1	Immediately quit the script.

Table 1.4 is incomplete; sed is quite complex, and this section merely introduces this tool.

The Addresses column of Table 1.4 requires elaboration: sed commands operate on addresses, which are line numbers. Commands may take no addresses, in which case they operate on the entire file. If one address is specified, they operate on the specified line. If two addresses (a range) are specified, the commands operate on that range of lines, inclusive.

In operation, sed looks something like this:

```
$ sed 's/2012/2013/' cal-2012.txt > cal-2013.txt
```

This command processes the input file, `cal-2012.txt`, using sed's `s` command to replace the first occurrence of `2012` on each line with `2013`. (If a single line may have more than one instance of the search string, you must perform a global search by appending `g` to the command string, as in `s/2012/2013/g`.) By default, sed sends the modified file to standard output, so this example uses redirection to send the output to `cal-2013.txt`. The idea in this example is to convert a file created for the year 2012 quickly so that it can be used in 2013. If you don't specify an input filename, sed works from standard input, so it can accept the output of another command as its input.

Although it's conceptually simple, sed is a very complex tool; even a modest summary of its capabilities would fill a chapter. You can consult its man page for basic information, but to understand sed fully, you may want to consult a book that tackles this tough subject, such as our book *Linux Command Line and Shell Scripting Bible*, *3rd Edition* (Wiley, 2015).

Certain sed commands, including the substitution command, are also used in vi, which is described more fully in Chapter 5.

Real World Scenario

Doing One Thing in Many Ways

As you become experienced with Linux and compare notes with other Linux administrators, you may find that the way you work is different from the way others work. This is because Linux often provides multiple methods to solve certain problems. For instance, ASCII text files use certain characters to encode the end of a line. Unix (and Linux) use a single line feed character (ASCII 0x0a, sometimes represented as `\n`), whereas DOS and Windows use the combination of a carriage return (ASCII 0x0d or `\r`) and a line feed. When moving ASCII files between computers, you may need to convert from one form to the other. How can you do this?

One solution is to use a special-purpose program, such as `dos2unix` or `unix2dos`. You could type **`dos2unix file.txt`** to convert `file.txt` from DOS-style to Unix-style ASCII, for instance. This is usually the simplest solution, but not all distributions have these utilities installed by default or even available to install.

Another approach is to use `tr`. For instance, to convert from DOS style to Unix style, you might type this:

```
$ tr -d \\r < dosfile.txt > unixfile.txt
```

This approach won't work when converting from Unix style to DOS style, though. For that, you can use `sed`:

```
sed s/$/"\r"/ unixfile.txt > dosfile.txt
```

Variants on both the `tr` and `sed` commands exist. For instance, sometimes the quotes around `\r` may be omitted from the `sed` command; whether they're required depends on your shell and its configuration.

Yet another approach is to load the file into a text editor and then save it using different file-type settings. (Not all editors support such changes, but some do.)

Many other examples exist of multiple solutions to a problem. Sometimes one solution stands out above others as being superior, but at other times the differences may be subtle, or each approach may have merit in particular situations. Thus it's best to be at least somewhat familiar with many of the alternatives, such as the options described throughout this book.

Summary

The command line is the key to Linux. Even if you prefer GUI tools to text-mode tools, understanding text-mode commands is necessary to fully manage a Linux system. This task begins with the shell, which accepts commands you type and displays the results of those commands. In addition, shells support linking programs together via pipes and redirecting programs' input and output. These features enable you to perform complex tasks using simple tools by having each program perform its own small part of the task. This technique is frequently used with Linux text filters, which manipulate text files in various ways—sorting text by fields, merging multiple files, and so on.

Exam Essentials

Summarize features that Linux shells offer to speed up command entry. The command history often enables you to retrieve an earlier command that's similar or identical to the one you want to enter. Tab completion reduces typing effort by letting the shell finish long command names or filenames. Command-line editing lets you edit a retrieved command or change a typo before committing the command.

Describe the purpose of the `man` command. The `man` command displays the manual page for the keyword (command, filename, system call, or other feature) that you type. This documentation provides succinct summary information that's useful as a reference to learn about exact command options or features.

Explain the purpose of environment variables. Environment variables store small pieces of data—program options, information about the computer, and so on. This information can be read by programs and used to modify program behavior in a way that's appropriate for the current environment.

Describe the difference between standard output and standard error. Standard output carries normal program output, whereas standard error carries high-priority output, such as error messages. The two can be redirected independently of one another.

Explain the purpose of pipes. Pipes tie programs together by feeding the standard output from the first program into the second program's standard input. They can be used to link together a series of simple programs to perform more complex tasks than any one of the programs could manage.

Describe the filter commands. The various simple filter commands allow the manipulation of text. These commands accomplish tasks of various types, such as combining files, transforming the data in files, formatting text, displaying text, and summarizing data.

Summarize the structure of regular expressions. Regular expressions are strings that describe other strings. They can contain normal alphanumeric characters, which match the exact same characters in the string they are describing, as well as several special symbols and symbol sets that match multiple different characters. The combination is a powerful pattern-matching tool used by many Linux programs.

Review Questions

1. You type a command into `bash` and pass a long filename to it, but after you enter the command, you receive a `File not found` error message because of a typo in the filename. How might you proceed?
   - **A.** Retype the command, and be sure you type the filename correctly, letter by letter.
   - **B.** Retype the command, but press the Tab key after typing a few letters of the long filename to ensure that the filename is entered correctly.
   - **C.** Press the Up arrow key, and use `bash`'s editing features to correct the typo.
   - **D.** Any of the above.
   - **E.** None of the above.
2. Which of the following commands is implemented as an internal command in `bash`?
   - **A.** `cat`
   - **B.** `less`
   - **C.** `tee`
   - **D.** `sed`
   - **E.** `echo`
3. You type **`echo $PROC`**, and the computer replies `Go away`. What does this mean?
   - **A.** No currently running processes are associated with your shell, so you may log out without terminating them.
   - **B.** The remote computer `PROC` isn't accepting connections; you should contact its administrator to correct the problem.
   - **C.** Your computer is handling too many processes; you must kill some of them to regain control of the computer.
   - **D.** Your central processing unit (CPU) is defective and must be replaced as soon as possible.
   - **E.** You, one of your configuration files, or a program you've run has set the `$PROC` environment variable to `Go away`.
4. What does the `pwd` command accomplish?
   - **A.** It prints the name of the working directory.
   - **B.** It changes the current working directory.
   - **C.** It prints wide displays on narrow paper.
   - **D.** It parses web page URLs for display.
   - **E.** It prints the terminal's width in characters.
5. What is the surest way to run a program (say, `myprog`) that's located in the current working directory?
   - **A.** Type **`./`** followed by the program name: **`./myprog`**.
   - **B.** Type the program name alone: **`myprog`**.

C. Type **run** followed by the program name: **run myprog**.

D. Type **/.** followed by the program name: **/.myprog**.

E. Type the program name followed by an ampersand (&): **myprog &**.

6. How does man display information by default on most Linux systems?

A. Using a custom X-based application

B. Using the Firefox web browser

C. Using the info browser

D. Using the vi editor

E. Using the less pager

7. You want to store the standard output of the ifconfig command in a text file (file.txt) for future reference, and you want to wipe out any existing data in the file. You do *not* want to store standard error in this file. How can you accomplish these goals?

A. **ifconfig < file.txt**

B. **ifconfig >> file.txt**

C. **ifconfig > file.txt**

D. **ifconfig | file.txt**

E. **ifconfig 2> file.txt**

8. What is the effect of the following command?

```
$ myprog &> input.txt
```

A. Standard error to myprog is taken from input.txt.

B. Standard input to myprog is taken from input.txt.

C. Standard output and standard error from myprog are written to input.txt.

D. All of the above.

E. None of the above.

9. How many commands can you pipe together at once?

A. 2

B. 3

C. 4

D. 16

E. >16

10. You want to run an interactive script, gabby, which produces a lot of output in response to the user's inputs. To facilitate future study of this script, you want to copy its output to a file. How might you do this?

A. **gabby > gabby-out.txt**

B. **gabby | tee gabby-out.txt**

C. **gabby < gabby-out.txt**

D. **gabby &> gabby-out.txt**

E. **gabby `gabby-out.txt`**

11. A text-mode program, `verbose`, prints a lot of bogus "error" messages to standard error. How might you get rid of those messages while still interacting with the program?

A. **`verbose | quiet`**

B. **`verbose &> /dev/null`**

C. **`verbose 2> /dev/null`**

D. **`verbose > junk.txt`**

E. **`quiet-mode verbose`**

12. How do the > and >> redirection operators differ?

A. The > operator creates a new file or overwrites an existing one; the >> operator creates a new file or appends to an existing one.

B. The > operator creates a new file or overwrites an existing one; the >> operator appends to an existing file or issues an error message if the specified file doesn't exist.

C. The > operator redirects standard output; the >> operator redirects standard error.

D. The > operator redirects standard output; the >> operator redirects standard input.

E. The > operator writes to an existing file but fails if the file doesn't exist; the >> operator writes to an existing file or creates a new one if it doesn't already exist.

13. What program would you use to display the end of a configuration file?

A. `uniq`

B. `cut`

C. `tail`

D. `wc`

E. `fmt`

14. What is the effect of the following command?

```
$ pr report.txt | lpr
```

A. The file `report.txt` is formatted for printing and sent to the `lpr` program.

B. The files `report.txt` and `lpr` are combined together into one file and sent to standard output.

C. Tabs are converted to spaces in `report.txt`, and the result is saved in `lpr`.

D. The file `report.txt` is printed, and any error messages are stored in the file `lpr`.

E. None of the above.

15. Which of the following commands will number the lines in `aleph.txt`? (Select three.)

A. **`fmt aleph.txt`**

B. **`nl aleph.txt`**

C. **`cat -b aleph.txt`**

D. **`cat -n aleph.txt`**

E. **`od -nl aleph.txt`**

16. You have a data file, `data.txt`, to be processed by a particular program. However, the program cannot handle data separated by tabs. The `data.txt` file's data is separated by a tab stop at every eight characters. What command should you use before processing the data file with the program?

    A. `od data.txt > data1.txt`

    B. `expand data.txt >> data.txt`

    C. `fmt --remove-tabs data.txt`

    D. `expand data.txt > data1.txt`

    E. `unexpand -t 8 data.txt`

17. Which of the following commands will change all occurrences of `dog` in the file `animals.txt` to `mutt` in the screen display?

    A. `sed -s "dog" "mutt" animals.txt`

    B. `grep -s "dog||mutt" animals.txt`

    C. `sed 's/dog/mutt/g' animals.txt`

    D. `cat animals.txt | grep -c "dog" "mutt"`

    E. `fmt animals.txt | cut 'dog' > 'mutt'`

18. You've received an ASCII text file (`longlines.txt`) that uses no carriage returns within paragraphs but two carriage returns between paragraphs. The result is that your preferred text editor displays each paragraph as a very long line. How can you reformat this file so that you can more easily edit it (or a copy)?

    A. `sed 's/Ctrl-M/NL/' longlines.txt`

    B. `fmt longlines.txt > longlines2.txt`

    C. `cat longlines.txt > longlines2.txt`

    D. `pr longlines.txt > longlines2.txt`

    E. `grep longlines.txt > longlines2.txt`

19. Which of the following commands will print lines from the file `world.txt` that contain matches to `changes` and `changed`?

    A. `grep change[ds] world.txt`

    B. `sed change[d-s] world.txt`

    C. `od "change'd|s'" world.txt`

    D. `cat world.txt changes changed`

    E. `find world.txt "change(d|s)"`

20. Which of the following regular expressions will match the strings `dog`, `dug`, and various other strings but not `dig`?

    A. `d.g`

    B. `d[ou]g`

    C. `d[o-u]g`

    D. `di*g`

    E. `d.ig`