

MARUDHARKESARIJAINCOLLEGEFORWOMEN,VANIYAMBADI PG

DEPARTMENT OF COMPUTER APPLICATIONS

CLASS :I MCA

SUBJECT CODE :23PCA12

SUBJECTNAME:LINUXANDSHELLPROGRAMMING

SYLLABUS

Unit-II

More Structured Commands: Looping with for statement-Iterating with the until statement-Using the while statement-Combining loops-Redirecting loop output. **Handling User Input:** Passing parameters-Tracking parameters-Being shifty-Working with options-Standardizing options-Getting user input. **Script Control:** Handling signals-Running scripts in the background-Forbidding hang-ups -Controlling a Job-Modifying script priority-Automating script execution.

(Book-1,Chapters:13,14,and16)

More Structured Commands

IN THIS CHAPTER

Looping with the `for` statement

Iterating with the `until` statement

Using the `while` statement

Combining loops

Redirecting loop output

In the previous chapter, you saw how to manipulate the flow of a shell script program by checking the output of commands and the values of variables. In this chapter, we continue to look at structured commands that control the flow of your shell scripts. You'll see how you can perform repeating processes, commands that can loop through a set of commands until an indicated condition has been met. This chapter discusses and demonstrates the `for`, `while`, and `until` bash shell looping commands.

The `for` Command

Iterating through a series of commands is a common programming practice. Often, you need to repeat a set of commands until a specific condition has been met, such as processing all the files in a directory, all the users on a system, or all the lines in a text file.

The bash shell provides the `for` command to allow you to create a loop that iterates through a series of values. Each iteration performs a defined set of commands using one of the values in the series. Here's the basic format of the bash shell `for` command:

```
for var in list
do
    commands
done
```

You supply the series of values used in the iterations in the *list* parameter. You can specify the values in the list in several ways.

In each iteration, the variable `var` contains the current value in the list. The first iteration uses the first item in the list, the second iteration the second item, and so on until all the items in the list have been used.

The commands entered between the `do` and `done` statements can be one or more standard bash shell commands. Within the commands, the `$var` variable contains the current list item value for the iteration.

NOTE

If you prefer, you can include the `do` statement on the same line as the `for` statement, but you must separate it from the list items using a semicolon: `for var in list; do`.

We mentioned that there are several different ways to specify the values in the list. The following sections show the various ways to do that.

Reading values in a list

The most basic use of the `for` command is to iterate through a list of values defined within the `for` command itself:

```
$ cat test1
#!/bin/bash
# basic for command

for test in Alabama Alaska Arizona Arkansas California Colorado
do
    echo The next state is $test
done
$ ./test1
The next state is Alabama
The next state is Alaska
The next state is Arizona
The next state is Arkansas
The next state is California
The next state is Colorado
$
```

Each time the `for` command iterates through the list of values provided, it assigns the `$test` variable the next value in the list. The `$test` variable can be used just like any other script variable within the `for` command statements. After the last iteration, the `$test` variable remains valid throughout the remainder of the shell script. It retains the last iteration value (unless you change its value):

```
$ cat test1b
#!/bin/bash
```

```
#testingtheforvariableafterthelooping
fortestinAlabamaAlaskaArizonaArkansasCaliforniaColorado do
    echo"Thenextstateis$test"done
echo"Thelaststatewevisitedwas$test"
test=Connecticut
echo"Wait,nowwe'revisiting$test"
$./test1b
ThenextstateisAlabama The
next state is Alaska
ThenextstateisArizona
ThenextstateisArkansas
ThenextstateisCalifornia
ThenextstateisColorado
ThelaststatewevisitedwasColorado
Wait,nowwe'revisitingConnecticut
$
```

The `$test` variable retained its value and allowed us to change the value and use it outside of the `for` command loop, as any other variable would.

Reading complex values in a list

Things aren't always as easy as they seem with the `for` loop. There are times when you run into data that causes problems. Here's a classic example of what can cause problems for shell script programmers:

```
$catbadtest1
#!/bin/bash
#anotherexampleofhownottousetheforcommand

fortestinIdon'tknowifthis'llwork do
    echo"word:$test"done
ne
$./badtest1
word:I
word:dontknowifthisll
word:work
$
```

Ouch, that hurts. The shell saw the single quotation marks within the list values and attempted to use them to define a single data value, and it really messed things up in the process.

You have two ways to solve this problem:

- Use the escape character (the backslash) to escape the single quotation mark.
- Use double quotation marks to define the values that use single quotation marks.

Neither solution is all that fantastic, but each one helps solve the problem:

```
$ cat test2
#!/bin/bash
# another example of how not to use the for command

for test in I don\'t know if "this \'ll" work
do
    echo "word: $test"
done
$ ./test2
word: I
word: don't
word: know
word: if
word: this 'll
word: work
$
```

In the first problem value, you added the backslash character to escape the single quotation mark in the `don't` value. In the second problem value, you enclosed the `this 'll` value in double quotation marks. Both methods worked fine to distinguish the value.

Another problem you may run into is multi-word values. Remember that the `for` loop assumes that each value is separated with a space. If you have data values that contain spaces, you run into yet another problem:

```
$ cat badtest2
#!/bin/bash
# another example of how not to use the for command

for test in Nevada New Hampshire New Mexico New York North Carolina
do
    echo "Now going to $test"
done
$ ./badtest1
Now going to Nevada
Now going to New
Now going to Hampshire
Now going to New
Now going to Mexico
Now going to New
Now going to York
```

```
Now going to North
NowgoingtoCarolina
$
```

Oops, that's not exactly what we wanted. The `for` command separates each value in the list with a space. If there are spaces in the individual data values, you must accommodate them using double quotation marks:

```
$ cat test3
#!/bin/bash
# an example of how to properly define values

for test in Nevada "New Hampshire" "New Mexico" "New York"
do
    echo "Now going to $test"
done
$ ./test3
Now going to Nevada
Now going to New Hampshire
Now going to New Mexico
Now going to New York
$
```

Now the `for` command can properly distinguish between the different values. Also, notice that when you use double quotation marks around a value, the shell doesn't include the quotation marks as part of the value.

Reading a list from a variable

Often what happens in a shell script is that you accumulate a list of values stored in a variable and then need to iterate through the list. You can do this using the `for` command as well:

```
$ cat test4
#!/bin/bash
# using a variable to hold the list

list="Alabama Alaska Arizona Arkansas Colorado"
list="$list Connecticut"

for state in $list
do
    echo "Have you ever visited $state?"
done
$ ./test4
Have you ever visited Alabama?
Have you ever visited Alaska?
Have you ever visited Arizona?
```

```
HaveyouevervisitedArkansas?
HaveyouevervisitedColorado?
HaveyouevervisitedConnecticut?
$
```

The `$list` variable contains the standard text list of values to use for the iterations. Notice that the code also uses another assignment statement to add (or concatenate) an item to the existing list contained in the `$list` variable. This is a common method for adding text to the end of an existing text string stored in a variable.

Reading values from a command

Another way to generate values for use in the list is to use the output of a command. You use command substitution to execute any command that produces output and then use the output of the command in the `for` command:

```
$ cat test5
#!/bin/bash
# reading values from a file

file="states"

for state in $(cat $file) do
    echo "Visit beautiful $state"
done
$ cat states
Alabama
Alaska
Arizona
Arkansas
Colorado
Connecticut
Delaware
Florida
Georgia
$ ./test5
Visit beautiful Alabama
Visit beautiful Alaska
Visit beautiful Arizona
Visit beautiful Arkansas
Visit beautiful Colorado
Visit beautiful Connecticut
Visit beautiful Delaware
Visit beautiful Florida
Visit beautiful Georgia
$
```

This example uses the `cat` command in the command substitution to display the contents of the `states` file. Notice that the `states` file includes each state on a separate line, not separated by spaces. The `for` command still iterates through the output of the `cat` command one line at a time, assuming that each state is on a separate line. However, this doesn't solve the problem of having spaces in data. If you list a state with a space in it, the `for` command still takes each word as a separate value. There's a reason for this, which we look at in the next section.

NOTE

The `test5` code example assigned the filename to the variable using just the filename without a path. This requires that the file be in the same directory as the script. If this isn't the case, you need to use a full pathname (either absolute or relative) to reference the file location.

Changing the field separator

The cause of this problem is the special environment variable `IFS`, called the *internal field separator*. The `IFS` environment variable defines a list of characters the bash shell uses as field separators. By default, the bash shell considers the following characters as field separators:

- A space
- A tab
- A newline

If the bash shell sees any of these characters in the data, it assumes that you're restarting a new data field in the list. When working with data that can contain spaces (such as file-names), this can be annoying, as you saw in the previous script example.

To solve this problem, you can temporarily change the `IFS` environment variable values in your shell script to restrict the characters the bash shell recognizes as field separators. For example, if you want to change the `IFS` value to recognize only the newline character, you need to do this:

```
IFS=$'\n'
```

Adding this statement to your script tells the bash shell to ignore spaces and tabs in data values. Applying this technique to the previous script yields the following:

```
$ cat test5b #!/bin/bash
# reading values from a file

file="states"

IFS=$'\n'
```



```
for state in $(cat $file) do
    echo "Visit beautiful $state"
done
$ ./test5b
Visit beautiful Alabama
Visit beautiful Alaska
Visit beautiful Arizona
Visit beautiful Arkansas
Visit beautiful Colorado
Visit beautiful Connecticut
Visit beautiful Delaware
Visit beautiful Florida
Visit beautiful Georgia
Visit beautiful New York
Visit beautiful New Hampshire
Visit beautiful North Carolina
$
```

Now the shell script can use values in the list that contains spaces.

CAUTION

When working on long scripts, it's possible to change the IFS value in one place, and then forget about it and assume the default value elsewhere in the script. A safe practice to get into is to save the original IFS value before changing it and then restore it when you're finished.

This technique can be coded like this:

```
IFS_OLD=$IFS
IFS=$'\n'
<use the new IFS value in code> IFS=$IFS_OLD
```

This ensures that the IFS value is returned to the default value for future operations within the script.

Other excellent applications of the IFS environment variable are possible. Suppose you want to iterate through values in a file that are separated by a colon (such as in the `/etc/passwd` file). You just need to set the IFS value to a colon:

```
IFS=:
```

If you want to specify more than one IFS character, just string them together on the assignment line:

```
IFS=$'\n':; "
```

This assignment uses the newline, colon, semicolon, and double quotation mark characters as field separators. There's no limit to how you can parse your data using the IFS characters.

Reading a directory using wildcards

Finally, you can use the `for` command to automatically iterate through a directory of files. To do this, you must use a wildcard character in the file or pathname. This forces the shell to use *file globbing*. File globbing is the process of producing filenames or pathnames that match a specified wildcard character.

This feature is great for processing files in a directory when you don't know all the filenames:

```
$ cat test6
#!/bin/bash
# iterate through all the files in a directory

for file in /home/rich/test/* do

    if [-d "$file"]
    then
        echo "$file is a directory"
    fi
    f [ -f "$file" ]
    then
        echo "$file is a file"
    fi
done
$ ./test6
/home/rich/test/dirl is a directory
/home/rich/test/myprog.c is a file
/home/rich/test/myprog is a file
/home/rich/test/myscript is a file
/home/rich/test/newdir is a directory
/home/rich/test/newfile is a file
/home/rich/test/newfile2 is a file
/home/rich/test/testdir is a directory
/home/rich/test/testing is a file
/home/rich/test/testprog is a file
/home/rich/test/testprog.c is a file
$
```

The `for` command iterates through the results of the `/home/rich/test/*` listing. The code tests each entry using the `test` command (using the square bracket method) to see if it's a directory, using the `-d` parameter, or a file, using the `-f` parameter (See Chapter 12).

Notice in this example that we did something different in the `if` statement tests:

```
if [-d "$file"]
```

In Linux, it's perfectly legal to have directory and filenames that contain spaces. To accommodate that, you should enclose the `$file` variable in double quotation marks. If you don't, you'll get an error if you run into a directory or filename that contains spaces:

```
./test6:line6[:toomanyarguments
./test6:line9[:toomanyarguments
```

The `bash` shell interprets the additional words as arguments within the `test` command, causing an error.

You can also combine both the directory search method and the list method in the same `for` statement by listing a series of directory wildcards in the `for` command:

```
$ cat test7
#!/bin/bash
# iterating through multiple directories

for file in /home/rich/.b*/home/rich/badtest do
    if [-d "$file"] then
        echo "$file is a directory"
    elif [ -f "$file" ]
    then
        echo "$file is a file"
    else
        echo "$file doesn't exist"
    fi
done
$ ./test7
/home/rich/.backup.timestamp is a file
/home/rich/.bash_history is a file
/home/rich/.bash_logout is a file
/home/rich/.bash_profile is a file
/home/rich/.bashrc is a file
/home/rich/badtest doesn't exist
$
```

The `for` statement first uses file globbing to iterate through the list of files that result from the wildcard character; then it iterates through the next file in the list. You can combine any number of wildcard entries in the list to iterate through.

CAUTION

Notice that you can enter anything in the list data. Even if the file or directory doesn't exist, the `for` statement attempts to process whatever you place in the list. This can be a problem when working with files and directories. You have no way of knowing if you're trying to iterate through a nonexistent directory: It's always a good idea to test each file or directory before trying to process it.

TheC-StyleforCommand

If you've done any programming using the C programming language, you're probably surprised by the way the bash shell uses the `for` command. In the C language, a `for` loop normally defines a variable, which it then alters automatically during each iteration. Typically, programmers use this variable as a counter and either increment or decrement the counter by one in each iteration. The bash `for` command can also provide this functionality. This section shows you how to use a C-style `for` command in a bash shell script.

TheC languageforcommand

The C language `for` command has a specific method for specifying a variable, a condition that must remain true for the iterations to continue, and a method for altering the variable for each iteration. When the specified condition becomes false, the `for` loop stops. The condition equation is defined using standard mathematical symbols. For example, consider the following C language code:

```
for (i=0; i<10; i++)
{
    printf("The next number is %d\n", i);
}
```

This code produces a simple iteration loop, where the variable `i` is used as a counter. The first section assigns a default value to the variable. The middle section defines the condition under which the loop will iterate. When the defined condition becomes false, the `for` loop stops iterations. The last section defines the iteration process. After each iteration, the expression defined in the last section is executed. In this example, the `i` variable is incremented by one after each iteration.

The bash shell also supports a version of the `for` loop that looks similar to the C-style `for` loop, although it does have some subtle differences, including a couple of things that will confuse shell script programmers. Here's the basic format of the C-style bash `for` loop:

```
for ((variableassignment; condition; iterationprocess))
```

The format of the C-style `for` loop can be confusing for bash shell script programmers, because it uses C-style variable references instead of the shell-style variable references. Here's what a C-style `for` command looks like:

```
for ((a=1; a<10; a++))
```

Notice that there are a couple of things that don't follow the standard bash shell `for` method:

- The assignment of the variable value can contain spaces.
- The variable in the condition isn't preceded with a dollar sign.
- The equation for the iteration process doesn't use the `expr` command format.

The shell developers created this format to more closely resemble the C-style `for` command. Although this is great for C programmers, it can throw even expert shell programmers into a tizzy. Be careful when using the C-style `for` loop in your scripts.

Here's an example of using the C-style `for` command in a bash shell program:

```
$ cat test8
#!/bin/bash
#testing the C-style for loop

for ((i=1;i<=10;i++)) do
    echo "The next number is $i" done
$ ./test8
The next number is 1
The next number is 2
The next number is 3
The next number is 4
The next number is 5
The next number is 6
The next number is 7
The next number is 8
The next number is 9
The next number is 10
$
```

The `for` loop iterates through the commands using the variable defined in the `for` loop (the letter `i` in this example). In each iteration, the `$i` variable contains the value assigned in the `for` loop. After each iteration, the loop iteration process is applied to the variable, which in this example, increments the variable by one.

Using multiple variables

The C-style `for` command also allows you to use multiple variables for the iteration. The loop handles each variable separately, allowing you to define a different iteration process for each variable. Although you can have multiple variables, you can define only one condition in the `for` loop:

```
$ cat test9
#!/bin/bash
#multiple variables

for ((a=1,b=10;a<=10;a++,b--))
do
    echo "$a-$b" done
$ ./test9
```

```
1-10
2-9
3-8
4-7
5-6
6-5
7-4
8-3
9-2
10-1
$
```

The `a` and `b` variables are each initialized with different values, and different iteration processes are defined. While the loop increases the `a` variable, it decreases the `b` variable for each iteration.

The `while` Command

The `while` command is somewhat of a cross between the `if-then` statement and the `for` loop. The `while` command allows you to define a command to test and then loop through a set of commands for as long as the defined test command returns a zero exit status. It tests the `test` command at the start of each iteration. When the `test` command returns a non-zero exit status, the `while` command stops executing the set of commands.

Basic `while` format

Here's the format of the `while` command:

```
while testcommand
do
    othercommands
done
```

The `test` command defined in the `while` command is the exact same format as in `if-then` statements (see Chapter 12). As in the `if-then` statement, you can use any normal `bash` shell command, or you can use the `test` command to test for conditions, such as variable values.

The key to the `while` command is that the exit status of the `test` command specified must change, based on the commands run during the loop. If the exit status never changes, the `while` loop will get stuck in an infinite loop.

The most common use of the `test` command is to use brackets to check a value of a shell variable that's used in the loop commands:

```
$cattest10
#!/bin/bash
```

```
#whilecommandtest

var1=10
while[$var1-gt0] do
    echo$var1
    var1=$((var1-1))
done
$./test10
10
9
8
7
6
5
4
3
2
1
$
```

The `while` command defines the test condition to check for each iteration:

```
while[$var1-gt0]
```

As long as the test condition is true, the `while` command continues to loop through the commands defined. Within the commands, the variable used in the test condition must be modified, or you'll have an infinite loop. In this example, we use shell arithmetic to decrease the variable value by one:

```
var1=$((var1-1))
```

The `while` loop stops when the test condition is no longer true.

Using multiple test commands

The `while` command allows you to define multiple test commands on the `while` statement line. Only the exit status of the last test command is used to determine when the loop stops. This can cause some interesting results if you're not careful. Here's an example of what we mean:

```
$cat test11 #!/bin/bash
#testing a multi-command while loop var1=10

while echo $var1
```

```
        [ $var1 -ge 0 ]
do
    echo "This is inside the loop" $var1
    var1=$((var1 - 1))
done
$ ./test11
10
This is inside the loop 9
This is inside the loop 8
This is inside the loop 7
This is inside the loop 6
This is inside the loop 5
This is inside the loop 4
This is inside the loop 3
This is inside the loop 2
This is inside the loop 1
This is inside the loop 0
This is inside the loop
-1
$
```

Pay close attention to what happened in this example. Two test commands were defined in the `while` statement:

```
while echo $var1
[ $var1 -ge 0 ]
```

The first test simply displays the current value of the `var1` variable. The second test uses brackets to determine the value of the `var1` variable. Inside the loop, an `echo` statement displays a simple message, indicating that the loop was processed. Notice when you run the example how the output ends:

```
This is inside the loop
-1
$
```

The `while` loop executed the `echo` statement when the `var1` variable was equal to zero and then decreased the `var1` variable value. Next, the test commands were executed for

the next iteration. The `echo` test command was executed, displaying the value of the `var1` variable, which is now less than zero. It's not until the shell executes the `test` test command that the `while` loop terminates.

This demonstrates that in a multi-command `while` statement, all the test commands are executed in each iteration, including the last iteration when the last test command fails. Be careful of this. Another thing to be careful of is how you specify the multiple test commands. Note that each test command is on a separate line!

The `until` Command

The `until` command works in exactly the opposite way from the `while` command. The `until` command requires that you specify a test command that normally produces a non-zero exit status. As long as the exit status of the test command is non-zero, the bash shell executes the commands listed in the loop. When the test command returns a zero exit status, the loop stops.

As you would expect, the format of the `until` command is:

```
until test commands
do
    other commands
done
```

Similar to the `while` command, you can have more than one *test command* in the `until` command statement. Only the exit status of the last command determines if the bash shell executes the *other commands* defined.

The following is an example of using the `until` command:

```
$ cat test12 #!/bin/bash
# using the until command

var1=100

until [ $var1 -eq 0 ] do
    echo $var1
    var1=$((var1-25)) done
$ ./test12
100
75
50
25
$
```

This example tests the `var1` variable to determine when the `until` loop should stop. As soon as the value of the variable is equal to zero, the `until` command stops the loop. The same caution as for the `while` command applies when you use multiple test commands with the `until` command:

```
$ cat test13
#!/bin/bash
# using the until command

var1=100

until echo $var1
    [ $var1 -eq 0 ]
do
    echo Inside the loop: $var1
    var1=$(( $var1 - 25 ))
done
$ ./test13
100
Inside the loop: 100 75
Inside the loop: 75 50
Inside the loop: 50 25
Inside the loop: 25 0
$
```

The shell executes the test commands specified and stops only when the last command is true.

Nesting Loops

A loop statement can use any other type of command within the loop, including other loop commands. This is called a *nested loop*. Care should be taken when using nested loops, because you're performing an iteration within an iteration, which multiplies the number of times commands are being run. If you don't pay close attention to this, it can cause problems in your scripts.

Here's a simple example of nesting a `for` loop inside another `for` loop:

```
$ cat test14
#!/bin/bash
```

```
#nestingforloops

for((a=1;a<=3;a++)) do
    echo"Startingloop$a:"
    for((b=1;b<=3;b++)) do
        echo "    Insideloop:$b"
    done
done
$ ./test14
Startingloop1:
    Insideloop:1
    Insideloop:2
    Insideloop:3
Startingloop2:
    Insideloop:1
    Insideloop:2
    Insideloop:3
Startingloop3:
    Insideloop:1
    Insideloop:2
    Insideloop:3
$
```

The nested loop (also called the *inner loop*) iterates through its values for each iteration of the outer loop. Notice that there's no difference between the `do` and `done` commands for the two loops. The bash shell knows when the first `done` command is executed that it refers to the inner loop and not the outer loop.

The same applies when you mix loop commands, such as placing a `for` loop inside a `while` loop:

```
$ cat test15 #!/bin/bash
#placing a for loop inside a while loop var1=5

while[$var1-ge0] do
    echo"Outerloop:$var1"
    for((var2=1;$var2<3;var2++)) do
        var3=$((var1*var2))
        echo"Innerloop:$var1*$var2=$var3"
    done
    var1=$((var1-1))
done
$ ./test15
```

```
Outerloop:5
  Innerloop: 5 * 1 = 5
  Innerloop: 5 * 2 = 10
Outerloop:4
  Innerloop: 4 * 1 = 4
  Innerloop: 4 * 2 = 8
Outerloop:3
  Innerloop: 3 * 1 = 3
  Innerloop: 3 * 2 = 6
Outerloop:2
  Innerloop: 2 * 1 = 2
  Innerloop: 2 * 2 = 4
Outerloop:1
  Innerloop: 1 * 1 = 1
  Innerloop: 1 * 2 = 2
Outerloop:0
  Innerloop: 0 * 1 = 0
  Innerloop: 0 * 2 = 0
$
```

Again, the shell distinguished between the `do` and `done` commands of the inner `for` loop from the same commands in the outer `while` loop.

If you really want to test your brain, you can even combine `until` and `while` loops:

```
$ cat test16
#!/bin/bash
# using until and while loops

var1=3

until [ $var1 -eq 0 ] do
  echo "Outerloop: $var1"
  var2=1
  while [ $var2 -lt 5 ] do
    var3=$((echo "scale=4; $var1/$var2" | bc))
    echo "          Innerloop: $var1/$var2=$var3"
    var2=$(( $var2 + 1 ))
  done
  var1=$(( $var1 - 1 ))
done
$ ./test16
Outerloop:3
  Inner loop: 3 / 1 = 3.0000
  Inner loop: 3 / 2 = 1.5000
  Inner loop: 3 / 3 = 1.0000
  Inner loop: 3 / 4 = .7500
```

```
Outerloop:2
  Innerloop: 2 / 1 = 2.0000
  Innerloop: 2 / 2 = 1.0000
  Innerloop: 2 / 3 = .6666
  Innerloop: 2 / 4 = .5000
Outerloop:1
  Innerloop:1/1=1.0000
  Innerloop:1/2=.5000
  Innerloop:1/3=.3333
  Innerloop:1/4=.2500
$
```

The outer `until` loop starts with a value of 3 and continues until the value equals 0. The inner `while` loop starts with a value of 1 and continues as long as the value is less than 5. Each loop must change the value used in the test condition, or the loop will get stuck infinitely.

Looping on File Data

Often, you must iterate through items stored inside a file. This requires combining two of the techniques covered:

- Using nested loops
- Changing the `IFS` environment variable

By changing the `IFS` environment variable, you can force the `for` command to handle each line in the file as a separate item for processing, even if the data contains spaces. After you've extracted an individual line in the file, you may have to loop again to extract data contained within it.

The classic example of this is processing data in the `/etc/passwd` file. This requires that you iterate through the `/etc/passwd` file line by line and then change the `IFS` variable value to a colon so you can separate the individual components in each line.

The following is an example of doing just that:

```
#!/bin/bash
#changing the IFS value

IFS_OLD=$IFS
IFS=$'\n'
for entry in $(cat /etc/passwd) do
  echo "Values in $entry -"
  IFS=:
  for value in $entry
```

```
do
    echo "    $value"
done
done
$
```

This script uses two different IFS values to parse the data. The first IFS value parses the individual lines in the `/etc/passwd` file. The inner `for` loop next changes the IFS value to the colon, which allows you to parse the individual values within the `/etc/passwd` lines.

When you run this script, you get output something like this:

```
Values in rich:x:501:501:RichBlum:/home/rich:/bin/bash- rich
x
501
501
RichBlum
/home/rich
/bin/bash
Values in katie:x:502:502:KatieBlum:/home/katie:/bin/bash- katie
x
506
509
KatieBlum
/home/katie
/bin/bash
```

The inner loop parses each individual value in the `/etc/passwd` entry. This is also a great way to process comma-separated data, a common way to import spreadsheet data.

Controlling the Loop

You might be tempted to think that after you start a loop, you're stuck until the loop finishes all its iterations. This is not true. A couple of commands help us control what happens inside of a loop:

- The `break` command
- The `continue` command

Each command has a different use in how to control the operation of a loop. The following sections describe how you can use these commands to control the operation of your loops.

Thebreakcommand

The `break` command is a simple way to escape a loop in progress. You can use the `break` command to exit any type of loop, including `while` and `until` loops.

You can use the `break` command in several situations. This section shows each of these methods.

Breaking out of a single loop

When the shell executes a `break` command, it attempts to break out of the loop that's currently processing:

```
$ cat test17 #!/bin/bash
#breaking out of a for loop

for var1 in 1 2 3 4 5 6 7 8 9 10 do
    if [ $var1 -eq 5 ] then
        break
    fi
    echo "Iteration number: $var1"
done
echo "The for loop is completed"
$ ./test17
Iteration number: 1
Iteration number: 2
Iteration number: 3
Iteration number: 4
The for loop is completed
$
```

The `for` loop should normally have iterated through all the values specified in the list. However, when the `if-then` condition was satisfied, the shell executed the `break` command, which stopped the `for` loop.

This technique also works for `while` and `until` loops:

```
$ cat test18 #!/bin/bash
#breaking out of a while loop var1=1

while [ $var1 -lt 10 ] do
    if [ $var1 -eq 5 ]
```

```
        then
            break
        fi
        echo "Iteration:$var1"var1=$
        [ $var1 + 1 ]
    done
    echo "Thewhileloopiscompleted"
$ ./test18
Iteration: 1
Iteration: 2
Iteration: 3
Iteration: 4
Thewhileloopiscompleted
$
```

Thewhileloopterminatedwhentheif-thenconditionwasmets,executingthebreak command.

Breakingoutofaninnerloop

Whenyou'reworkingwithmultipleloops,thebreakcommandautomaticallyterminates the innermost loop you're in:

```
$cat test19
#!/bin/bash
#breakingoutofaninnerloop

for ((a=1;a<4;a++)) do
    echo "Outerloop:$a"
    for ((b=1;b<100;b++)) do
        if [$b-eq5]
        then
            break
        fi
        echo "    Innerloop:$b"
    done
done
$ ./test19
Outerloop:1
    Innerloop:1
    Innerloop:2
    Innerloop:3
    Innerloop:4
Outerloop:2
    Innerloop:1
    Innerloop:2
    Innerloop:3
```



```
        Innerloop:4
Outerloop:3
        Innerloop:1
        Innerloop:2
        Innerloop:3
        Innerloop:4
$
```

The `for` statement in the inner loop specifies to iterate until the `b` variable is equal to 100. However, the `if-then` statement in the inner loop specifies that when the `b` variable value is equal to 5, the `break` command is executed. Notice that even though the inner loop is terminated with the `break` command, the outer loop continues working as specified.

Breaking out of an outer loop

There may be times when you're in an inner loop but need to stop the outer loop. The `break` command includes a single command-line parameter value:

```
break n
```

where *n* indicates the level of the loop to break out of. By default, *n* is 1, indicating to break out of the current loop. If you set *n* to a value of 2, the `break` command stops the next level of the outer loop:

```
$ cat test20 #!/bin/bash
#breaking out of an outer loop

for ((a=1;a<4;a++)) do
    echo "Outerloop:$a"
    for ((b=1;b<100;b++)) do
        if [ $b -gt 4 ] then
            break 2
        fi
        echo "    Innerloop:$b"
    done
done
$ ./test20
Outerloop:1
    Innerloop:1
    Innerloop:2
    Innerloop:3
    Innerloop:4
$
```

Now when the shell executes the `break` command, the outer loop stops.

The `continue` command

The `continue` command is a way to prematurely stop processing commands inside of a loop but not terminate the loop completely. This allows you to set conditions within a loop where the shell won't execute commands. Here's a simple example of using the `continue` command in a `for` loop:

```
$ cat test21
#!/bin/bash
# using the continue command

for ((var1=1; var1<15; var1++)) do
    if [ $var1 -gt 5 ] && [ $var1 -lt 10 ] then
        continue
    fi
    echo "Iteration number: $var1"
done
$ ./test21
Iteration number: 1
Iteration number: 2
Iteration number: 3
Iteration number: 4
Iteration number: 5
Iteration number: 10
Iteration number: 11
Iteration number: 12
Iteration number: 13
Iteration number: 14
$
```

When the conditions of the `if-then` statement are met (the value is greater than 5 and less than 10), the shell executes the `continue` command, which skips the rest of the commands in the loop, but keeps the loop going. When the `if-then` condition is no longer met, things return to normal.

You can use the `continue` command in `while` and `until` loops, but be extremely careful with what you're doing. Remember that when the shell executes the `continue` command, it skips the remaining commands. If you're incrementing your test condition variable in one of those conditions, bad things happen:

```
$ cat badtest3
#!/bin/bash
# improperly using the continue command in a while loop var1=0

while echo "while iteration: $var1"
```

[illegible]

You'll want to make sure you redirect the output of this script to the `more` command so you can stop things. Everything seems to be going just fine until the `if-then` condition is met, and the shell executes the `continue` command. When the shell executes the `continue` command, it skips the remaining commands in the `while` loop. Unfortunately, that's where the `$var1` counter variable that is tested in the `while` test command is incremented. That means that the variable isn't incremented, as you can see from the continually displaying output.

As with the `break` command, the `continue` command allows you to specify what level of loop to continue with a command line parameter:

continuen

where `end` defines the loop level to continue. Here's an example of continuing an outer for loop:

```
$cat test22
#!/bin/bash
#continuing an outer loop

for ((a=1;a<=5;a++)) do
    echo "Iteration$a:"
    for ((b=1;b<3;b++)) do
        if [[ $a -gt 2 ] && [ $a -lt 4 ] then
            continue 2
        fi
        var3=$((a*b))
        echo "    The result of $a * $b is $var3"
    done
done
$ ./test22
Iteration1:
    The result of 1 * 1 is 1
    The result of 1 * 2 is 2
Iteration2:
    The result of 2 * 1 is 2
    The result of 2 * 2 is 4
Iteration3:
Iteration4:
    The result of 4 * 1 is 4
    The result of 4 * 2 is 8
Iteration5:
    The result of 5 * 1 is 5
    The result of 5 * 2 is 10
$
```

The `if-then` statement:

```
if [[ $a -gt 2 ] && [ $a -lt 4 ] then
    continue 2
fi
```

uses the `continue` command to stop processing the commands inside the loop but continue the outer loop. Notice in the script output that the iteration for the value 3 doesn't process any inner loop statements, because the `continue` command stopped the processing, but it continues with the outer loop processing.

ProcessingtheOutputofaLoop

Finally, you can either pipe or redirect the output of a loop within your shell script. You do this by adding the processing command to the end of the `done` command:

```
for file in /home/rich/* do
    if [-d "$file"] then
        echo "$file is a directory"
    elif
        echo "$file is a file"
    fi
done > output.txt
```

Instead of displaying the results on the monitor, the shell redirects the results of the `for` command to the file `output.txt`.

Consider the following example of redirecting the output of a `for` command to a file:

```
$ cat test23 #!/bin/bash
# redirecting the for output to a file

for ((a=1;a<10;a++)) do
    echo "The number is $a"
done > test23.txt
echo "The command is finished."
$ ./test23
The command is finished.
$ cat test23.txt
The number is 1
The number is 2
The number is 3
The number is 4
The number is 5
The number is 6
The number is 7
The number is 8
The number is 9
$
```

The shell creates the file `test23.txt` and redirects the output of the `for` command only to the file. The shell displays the `echo` statement after the `for` command just as normal.

This same technique also works for piping the output of a loop to another command:

```
$ cat test24 #!/bin/bash
```

```
#pipingalooptoanothercommand

forstatein"NorthDakota"ConnecticutIllinoisAlabamaTennessee do
    echo"$stateisthenextplacetogo"done |
sort
echo"This completesourtravels"
$./test24
Alabamaisthenextplacetogo
Connecticutisthenextplacetogo
Illinoisisthenextplacetogo
NorthDakotaisthenextplacetogo
Tennesseeisthenextplacetogo
This completesourtravels
$
```

The state values aren't listed in any particular order in the `for` command list. The output of the `for` command is piped to the `sort` command, which changes the order of the `for` command output. Running the script indeed shows that the output was properly sorted within the script.

Practical Examples

Now that you've seen how to use the different ways to create loops in shell scripts, let's look at some practical examples of how to use them. Looping is a common way to iterate through data on the system, whether it's files in folders or data contained in a file. Here are a couple of examples that demonstrate using simple loops to work with data.

Finding executable files

When you run a program from the command line, the Linux system searches a series of folders looking for that file. Those folders are defined in the `PATH` environment variable. If you want to find out just what executable files are available on your system for you to use, just scan all the folders in the `PATH` environment variable. That may take some time to do manually, but it's a breeze working out a small shell script to do that.

The first step is to create a `for` loop to iterate through the folders stored in the `PATH` environment variable. When you do that, don't forget to set the `IFS` separator character:

```
IFS=:
for folder in $PATH
do
```

Now that you have the individual folders in the `$folder` variable, you can use another `for` loop to iterate through all the files inside that particular folder:

```
for file in $folder/* do
```

The last step is to check whether the individual files have the executable permission set, which you can do using the `if-test` feature:

```
if [-x$file] then
    echo"    $file"
fi
```

And there you have it! Putting all the pieces together into a script looks like this:

```
$cat test25
#!/bin/bash
#finding files in the PATH

IFS=:
for folder in $PATH do
    echo"$folder:"
    for file in $folder/* do
        if [-x$file] then
            echo"    $file"
        fi
    done
done
$
```

When you run the code, you get a listing of the executable files that you can use from the command line:

```
$. /test25 | more
/usr/local/bin:
/usr/bin:
    /usr/bin/Mail
    /usr/bin/Thunar
    /usr/bin/X
    /usr/bin/Xorg
    /usr/bin/[
    /usr/bin/a2p
    /usr/bin/abiword
    /usr/bin/ac
    /usr/bin/activation-client
    /usr/bin/addr2line
...
```

The output shows all the executable files found in all the folders defined in the `PATH` environment variable, which is quite a few!

Creating multiple user accounts

The goal of shell scripts is to make life easier for the system administrator. If you happen to work in an environment with lots of users, one of the most boring tasks can be creating new user accounts. Fortunately, you can use the `while` loop to make your job a little easier!

Instead of having to manually enter `useradd` commands for every new user account you need to create, you can place the new user accounts in a text file and create a simple shell script to do that work for you. The format of the text file that we'll use looks like this:

```
userid,username
```

The first entry is the `userid` you want to use for the new user account. The second entry is the full name of the user. The two values are separated by a comma, making this a comma-separated file format, or `.csv`. This is a very common file format used in spreadsheets, so you can easily create the user account list in a spreadsheet program and save it in `.csv` format for your shell script to read and process.

To read the file data, we're going to use a little shell scripting trick. We'll actually set the `IFS` separator character to a comma as the test part of the `while` statement. Then to read the individual lines, we'll use the `read` command. That looks like this:

```
while IFS=',' read -r userid name
```

The `read` command does the work of moving onto the next line of text in the `.csv` text file, so we don't need another loop to do that. The `while` command exits when the `read` command returns a `FALSE` value, which happens when it runs out of lines to read in the file. Tricky!

To feed the data from the file into the `while` command, you just use a redirection symbol at the end of the `while` command.

Putting everything together results in this script:

```
$ cat test26
#!/bin/bash
# process new user accounts

input="users.csv"
while IFS=',' read -r userid name do
    echo "adding $userid"
    useradd -c "$name" -m $userid
done <"$input"
$
```

The `$input` variable points to the data file and is used as the redirect data for the `while` command. The `users.csv` file looks like this:

```
$ cat users.csv
rich, Richard Blum
```



```
christine,ChristineBresnahan
barbara,Barbara Blum
tim,Timothy Bresnahan
$
```

To run the problem, you must be the root user account, because the `useradd` command requires root privileges:

```
# ./test26
adding rich
adding christine
adding barbara
adding tim
#
```

Then by taking a quick look at the `/etc/passwd` file, you can see that the accounts have been created:

```
#tail/etc/passwd
rich:x:1001:1001:Richard Blum:/home/rich:/bin/bash
christine:x:1002:1002:Christine Bresnahan:/home/christine:/bin/bash
barbara:x:1003:1003:Barbara Blum:/home/barbara:/bin/bash
tim:x:1004:1004:Timothy Bresnahan:/home/tim:/bin/bash
#
```

Congratulations, you've saved yourself a lot of time in adding user accounts!

Summary

Looping is an integral part of programming. The bash shell provides three looping commands that you can use in your scripts.

The `for` command allows you to iterate through a list of values, either supplied within the command line, contained in a variable, or obtained by using file globbing, to extract file and directory names from a wildcard character.


The `while` command provides a method to loop based on the condition of a command, using either ordinary commands or the `test` command, which allows you to test conditions of variables. As long as the command (or condition) produces a zero exit status, the `while` loop continues to iterate through the specified set of commands.

The `until` command also provides a method to iterate through commands, but it bases its iterations on a command (or condition) producing a non-zero exit status. This feature allows you to set a condition that must be met before the iteration stops.

You can combine loops in shell scripts, producing multiple layers of loops. The bash shell provides the `continue` and `break` commands, which allow you to alter the flow of the normal loop process based on different values within the loop.

The bash shell also allows you to use standard command redirection and piping to alter the output of a loop. You can use redirection to redirect the output of a loop to a file or piping to redirect the output of a loop to another command. This provides a wealth of features with which you can control your shell script execution.

The next chapter discusses how to interact with your shell script user. Often, shell scripts aren't completely self-contained. They require some sort of external data that must be supplied at the time you run them. The next chapter discusses different methods with which you can provide real-time data to your shell scripts for processing.



Handling User Input

IN THIS CHAPTER

Passing parameters

Tracking parameters

Being shift-y

Working with options

Standardizing options

Getting user input

So far you've seen how to write scripts that interact with data, variables, and files on the Linux system. Sometimes, you need to write a script that has to interact with the person running the script. The bash shell provides a few different methods for retrieving data from people, including command line parameters (data values added after the command), command line options (single-letter values that modify the behavior of the command), and the capability to read input directly from the keyboard. This chapter discusses how to incorporate these different methods into your bash shell scripts to obtain data from the person running your script.

Passing Parameters

The most basic method of passing data to your shell script is to use *command line parameters*. Command line parameters allow you to add data values to the command line when you execute the script:

```
$ ./addem1030
```

This example passes two command line parameters (10 and 30) to the script `addem`. The script handles the command line parameters using special variables. The following sections describe how to use command line parameters in your bash shell scripts.

Reading parameters

The bash shell assigns special variables, called *positional parameters*, to all of the command line parameters entered. This includes the name of the script the shell is executing. The positional parameter variables are standard numbers, with *\$0* being the script's name, *\$1* being the first parameter, *\$2* being the second parameter, and so on, up to *\$9* for the ninth parameter.

Here's a simple example of using one command line parameter in a shell script:

```
$ cat test1.sh
#!/bin/bash
# using one command line parameter #
factorial=1
for ((number=1; number<=$1; number++)) do
    factorial=$((factorial*number)) done
echo The factorial of $1 is $factorial
$
$ ./test1.sh 5
The factorial of 5 is 120
$
```

You can use the *\$1* variable just like any other variable in the shell script. The shell script automatically assigns the value from the command line parameter to the variable; you don't need to do anything with it.

If you need to enter more command line parameters, each parameter must be separated by a space on the command line:

```
$ cat test2.sh
#!/bin/bash
# testing two command line parameters #
total=$(( $1 * $2 ))
echo The first parameter is $1.
echo The second parameter is $2.
echo The total value is $total.
$
$ ./test2.sh 2 5
The first parameter is 2.
The second parameter is 5.
The total value is 10.
$
```

The shell assigns each parameter to the appropriate variable.

In the preceding example, the command line parameters used were both numerical values. You can also use text strings in the command line:

```
$ cat test3.sh
#!/bin/bash
#testing string parameters #
echo Hello $1, glad to meet you.
$
$ ./test3.sh Rich
Hello Rich, glad to meet you.
$
```

The shell passes the string value entered into the command line to the script. However, you'll have a problem if you try to do this with a text string that contains spaces:

```
$ ./test3.sh Rich Blum
Hello Rich, glad to meet you.
$
```

Remember that each of the parameters is separated by a space, so the shell interpreted the space as just separating the two values. To include a space as a parameter value, you must use quotation marks (either single or double quotation marks):

```
$ ./test3.sh 'Rich Blum'
Hello Rich Blum, glad to meet you.
$
$ ./test3.sh "Rich Blum"
Hello Rich Blum, glad to meet you.
$
```

NOTE

The quotation marks used when you pass text strings as parameters are not part of the data. They just delineate the beginning and the end of the data.

If your script needs more than nine command line parameters, you can continue, but the variable names change slightly. After the ninth variable, you must use braces around the variable number, such as `${10}`. Here's an example of doing that:

```
$ cat test4.sh
#!/bin/bash
#handling lots of parameters #
total=${${10}}*${11}
echo The tenth parameter is ${10}
echo The eleventh parameter is ${11} echo
The total is $total
```

```
$  
$ ./test4.sh123456789101112  
The tenth parameter is 10  
Theeleventhparameteris11 The  
total is 110  
$
```

This technique allows you to add as many command line parameters to your scripts as you could possibly need.

Readingthescriptname

You can use the `$0` parameter to determine the script name the shell started from the command line. This can come in handy if you're writing a utility that can have multiple functions.

```
$ cat test5.sh  
#!/bin/bash  
#Testingthe$0parameter #  
echoThezeroparameterissetto:$0 #  
$  
$ bash test5.sh  
Thezeroparameterissetto:test5.sh  
$
```

However, there is a potential problem. When using a different command to run the shell script, the command becomes entangled with the script name in the `$0` parameter:

```
$ ./test5.sh  
Thezeroparameterissetto:./test5.sh  
$
```

There is another potential problem. When the actual string passed is the full script path, and not just the script's name, the `$0` variable gets set to the full script path and name:

```
$ bash/home/Christine/test5.sh  
Thezeroparameterissetto:/home/Christine/test5.sh  
$
```

If you want to write a script that performs different functions based on just the script's name, you'll have to do a little work. You need to be able to strip off whatever path is used to run the script. Also, you need to be able to remove any entangled commands from the script.

Fortunately, there's a handy little command available that does just that. The `basename` command returns just the script's name without the path:

```
$ cat test5b.sh  
#!/bin/bash
```

```
#Usingbasenamewiththe$0parameter #
name=$(basename$0)
echo
echoThescriptnameis:$name #
$bash/home/Christine/test5b.sh

Thescriptnameis:test5b.sh
$
$./test5b.sh

Thescriptnameis:test5b.sh
$
```

Now that's much better. You can use this technique to write scripts that perform different functions based on the script name used. Here's a simple example:

```
$cattest6.sh
#!/bin/bash
#TestingaMulti-functionscript #
name=$(basename$0)
#
if[$name="addem"]
then
    total=$(( $1+$2 ))
#
elif[$name="multem"] then
    total=$(( $1*$2 ))
fi
#
echo
echoThecalculatedvalueis$total #
$
$cptest6.shaddem
$chmodu+xaddem
$
$ln-stest6.shmultem
$
$ls-l*em
-rwxrw-r--.1ChristineChristine224Jun3023:50addem
lrwxrwxrwx.1ChristineChristine      8Jun3023:50multem->test6.sh
$
$./addem25

Thecalculatedvalueis7
```



```
$  
$ ./multem25  
  
Thecalculatedvalueis10  
$
```

The example creates two separate filenames from the `test6.sh` script, one by just copying the file to a new script (`addem`) and the other by using a symbolic link (see Chapter 3) to create the new script (`multem`). In both cases, the script determines the script's base name and performs the appropriate function based on that value.

Testing parameters

Be careful when using command line parameters in your shell scripts. If the script is run without the parameters, bad things can happen:

```
$ ./addem2  
./addem: line 8: 2+: syntax error: operand expected (error token is  
"")  
Thecalculatedvalueis  
$
```

When the script assumes there is data in a parameter variable, and no data is present, most likely you'll get an error message from your script. This is a poor way to write scripts. Always check your parameters to make sure the data is there before using it:

```
$ cat test7.sh  
#!/bin/bash  
#testing parameters before use #  
if [-n "$1"] then  
    echo Hello $1, glad to meet you. else  
    echo "Sorry, you did not identify yourself."  
fi  
$  
$ ./test7.sh Rich  
Hello Rich, glad to meet you.  
$  
$ ./test7.sh  
Sorry, you did not identify yourself.  
$
```

In this example, the `-n` test evaluation was used to check for data in the `$1` command line parameter. In the next section, you'll learn another way to check command line parameters.

UsingSpecialParameterVariables

A few special bash shell variables track command line parameters. This section describes what they are and how to use them.

Countingparameters

As you saw in the last section, you should verify command line parameters before using them in your script. For scripts that use multiple command line parameters, this checking can get tedious.

Instead of testing each parameter, you can count how many parameters were entered on the command line. The bash shell provides a special variable for this purpose.

The special `$#` variable contains the number of command line parameters included when the script was run. You can use this special variable anywhere in the script, just like a normal variable:

```
$ cat test8.sh
#!/bin/bash
# getting the number of parameters #
echo There were $# parameters supplied.
$
$ ./test8.sh
There were 0 parameters supplied.
$
$ ./test8.sh 12345
There were 5 parameters supplied.
$
$ ./test8.sh 12345678910
There were 10 parameters supplied.
$
$ ./test8.sh "Rich Blum"
There were 1 parameters supplied.
$
```

Now you have the ability to test the number of parameters present before trying to use them:

```
$ cat test9.sh
#!/bin/bash
# Testing parameters
#
if [ $# -ne 2 ] then
    echo
```

PartII:ShellScriptingBasics

```
        echoUsage:test9.shab
        echo
    else
        total=$((1+$2))
        echo
        echoThetotalis$total
        echo
    fi
#
$
$bashtest9.sh

Usage:test9.shab

$bashtest9.sh10

Usage:test9.shab

$bashtest9.sh1015

Thetotalis25

$
```

The `if-then` statement uses the `-ne` evaluation to perform a numeric test of the command line parameters supplied. If the correct number of parameters isn't present, an error message displays showing the correct usage of the script.

This variable also provides a cool way of grabbing the last parameter on the command line without having to know how many parameters were used. However, you need to use a little trick to get there.

If you think this through, you might think that because the `$#` variable contains the value of the number of parameters, using the variable `${$#}` would represent the last command line parameter variable. Try that and see what happens:

```
$catbadtest1.sh
#!/bin/bash
#testinggrabbinglastparameter #
echoThelastparameterwas${$#}
$
$./badtest1.sh10
Thelastparameterwas15354
$
```

Wow, what happened? Obviously, something went wrong. It turns out that you can't use the dollar sign within the braces. Instead, you must replace the dollar sign with an exclamation mark. Odd, but it works:

```
$ cat test10.sh
#!/bin/bash
#Grabbing the last parameter #
params=$#
echo
echo The last parameter is $params
echo The last parameter is ${!#} echo
#
$
$ bashtest10.sh 12345

The last parameter is 5
The last parameter is 5

$
$ bashtest10.sh

The last parameter is 0
The last parameter is test10.sh

$
```

Perfect. This script also assigned the `$#` variable value to the variable `params` and then used that variable within the special command line parameter variable format as well. Both versions worked. It's also important to notice that, when there weren't any parameters on the command line, the `$#` value was zero, which is what appears in the `params` variable, but the `${!#}` variable returns the script name used on the command line.

Grabbing all the data

In some situations you want to grab all the parameters provided on the command line. Instead of having to mess with using the `$#` variable to determine how many parameters are on the command line and having to loop through all of them, you can use a couple of other special variables.

The `$*` and `$@` variables provide easy access to all your parameters. Both of these variables include all the command line parameters within a single variable.

The `$*` variable takes all the parameters supplied on the command line as a single word. The word contains each of the values as they appear on the command line. Basically, instead of treating the parameters as multiple objects, the `$*` variable treats them all as one parameter.

The `$@` variable, on the other hand, takes all the parameters supplied on the command line as separate words in the same string. It allows you to iterate through the values, separating out each parameter supplied. This is most often accomplished using the `for` command.

PartII:ShellScriptingBasics

It can easily get confusing to figure out how these two variables operate. Let's look at the difference between the two:

```
$ cat test11.sh
#!/bin/bash
#testing$*and$@ #
echo
echo "Using the \${*} method: ${*}"
echo "Using the \${@} method: ${@}"
$
$ ./test11.sh rich bar barakatie jessica

Using the ${*} method: rich bar barakatie jessica

Using the ${@} method: rich bar barakatie jessica
$
```

Notice that on the surface, both variables produce the same output, showing all the command line parameters provided at once.

The following example demonstrates where the differences are:

```
$ cat test12.sh
#!/bin/bash
#testing$*and$@ #
echo
count=1
#
for param in "$*"
do
    echo "\${*} Parameter # $count = $param"
    count=$((count + 1))
done
#
echo
count=1
#
for param in "$@"
do
    echo "\${@} Parameter # $count = $param"
    count=$((count + 1))
done
$
$ ./test12.sh rich bar barakatie jessica

${*} Parameter #1=rich bar barakatie jessica

${@} Parameter #1=rich
```

```
$@Parameter#2=barbara
$@Parameter#3=katie
$@Parameter#4=jessica
$
```

Now we're getting somewhere. By using the `for` command to iterate through the special variables, you can see how they each treat the command line parameters differently. The `$*` variable treated all the parameters as a single parameter, while the `$@` variable treated each parameter separately. This is a great way to iterate through command line parameters.

Being Shifty

Another tool you have in your bash shell tool belt is the `shift` command. The bash shell provides the `shift` command to help you manipulate command line parameters. The `shift` command literally shifts the command line parameters in their relative positions.

When you use the `shift` command, it moves each parameter variable one position to the left by default. Thus, the value for variable `$3` is moved to `$2`, the value for variable `$2` is moved to `$1`, and the value for variable `$1` is discarded (note that the value for variable `$0`, the program name, remains unchanged).

This is another great way to iterate through command line parameters, especially if you don't know how many parameters are available. You can just operate on the first parameter, shift the parameters over, and then operate on the first parameter again.

Here's a short demonstration of how this works:

```
$ cat test13.sh
#!/bin/bash
#demonstrating the shift command
echo
count=1
while [ -n "$1" ] do
    echo "Parameter#$count=$1"count
    count=$((count + 1))
    shift
done
$ ./test13.sh rich barbara katie jessica
```

```
Parameter #1 = rich
Parameter#2=barbara
Parameter #3 = katie
Parameter#4=jessica
$
```

The `script` performs a `while` loop, testing the length of the first parameter's value. When the first parameter's length is zero, the loop ends. After testing the first parameter, the `shift` command is used to shift all the parameters one position.

TIP

Be careful when working with the `shift` command. When a parameter is shifted out, its value is lost and can't be recovered.

Alternatively, you can perform a multiple location shift by providing a parameter to the `shift` command. Just provide the number of places you want to shift:

```
$ cat test14.sh
#!/bin/bash
#demonstrating a multi-position shift #
echo
echo "The original parameters: $*"
shift
2
echo "Here's the new first parameter: $1"
$
$ ./test14.sh 12345

The original parameters: 12345
Here's the new first parameter: 3
$
```

By using values in the `shift` command, you can easily skip over parameters you don't need.

Working with Options

If you've been following along in the book, you've seen several bash commands that provide both parameters and options. *Options* are single letters preceded by a dash that alter the behavior of a command. This section shows three methods for working with options in your shell scripts.

Finding your options

On the surface, there's nothing all that special about command line options. They appear on the command line immediately after the script name, just the same as command line parameters. In fact, if you want, you can process command line options the same way you process command line parameters.

Processing simple options

In the `test13.sh` script earlier, you saw how to use the `shift` command to work your way down the command line parameters provided with the script program. You can use this same technique to process command line options.

As you extract each individual parameter, use the `case` statement (see Chapter 12) to determine when a parameter is formatted as an option:

```
$ cat test15.sh
#!/bin/bash
# extracting command line options as parameters #
echo
while [-n "$1"] do
    case "$1" in
        -a) echo "Found the -a option";;
        -b) echo "Found the -b option";;
        -c) echo "Found the -c option";;
        *) echo "$1 is not an option";; esac
    shift
done
$
$ ./test15.sh -a -b -c -d

Found the -a option
Found the -b option
Found the -c option
-d is not an option
$
```

The `case` statement checks each parameter for valid options. When one is found, the appropriate commands are run in the `case` statement.

This method works, no matter in what order the options are presented on the command line:

```
$ ./test15.sh -d -c -a

-d is not an option
Found the -c option
Found the -a option
$
```

The `case` statement processes each option as it finds it in the command line parameters. If any other parameters are included on the command line, you can include commands in the catch-all part of the `case` statement to process them.

Separating options from parameters

Often you'll run into situations where you'll want to use both options and parameters for a shell script. The standard way to do this in Linux is to separate the two with a special character code that tells the script when the options are finished and when the normal parameters start.

For Linux, this special character is the double dash (--). The shell uses the double dash to indicate the end of the option list. After seeing the double dash, your script can safely process the remaining command line parameters as parameters and not options.

To check for the double dash, simply add another entry in the case statement:

```
$ cat test16.sh
#!/bin/bash
# extracting options and parameters
echo
while [-n "$1"] do
    case "$1" in
        -a) echo "Found the -a option";;
        -b) echo "Found the -b option";;
        -c) echo "Found the -c option";;
        --) shift
            break;;
        *) echo "$1 is not an option";; esac
    shift
done
#
count=1
for param in $@ do
    echo "Parameter # $count: $param"
    count=$((count + 1))
done
$
```

This script uses the `break` command to break out of the `while` loop when it encounters the double dash. Because we're breaking out prematurely, we need to ensure that we stick in another `shift` command to get the double dash out of the parameter variables.

For the first test, try running the script using a normal set of options and parameters:

```
$ ./test16.sh -c -a -b test1 test2 test3
```

```
Found the -c option
Found the -a option
Found the -b option
test1 is not an option
```

```
test2isnotanoption
test3isnotanoption
$
```

The results show that the script assumed that all the command line parameters were options when it processed them. Next, try the same thing, only this time using the doubledash to separate the options from the parameters on the command line:

```
$ ./test16.sh -c -a -b --test1test2test3
```

```
Foundthe-coption
Foundthe-aoption
Foundthe-boption
Parameter#1:test1
Parameter#2:test2
Parameter#3:test3
$
```

When the script reaches the double dash, it stops processing options and assumes that any remaining parameters are command line parameters.

Processing options with values

Some options require an additional parameter value. In these situations, the command line looks something like this:

```
$ ./testing.sh -a test1 -b -c -d test2
```

Your script must be able to detect when your command line option requires an additional parameter and be able to process it appropriately. Here's an example of how to do that:

```
$ cat test17.sh
#!/bin/bash
#extractingcommandlineoptionsandvalues echo
while [-n "$1"] do
    case "$1" in
        -a) echo "Foundthe-aoption";;
        -b) param="$2"
            echo "Foundthe-boption, with parameter value $param" shift ;;
        -c) echo "Foundthe-coption";;
        --) shift
            break;;
        *) echo "$1 is not an option";; esac
    shift
done
```

```
#
count=1
for param in "$@"
do
    echo "Parameter # $count: $param"
    count=$((count + 1))
done
$
$ ./test17.sh -a -b test1 -d

Found the -a option
Found the -b option, with parameter value test1
-disnotanoption
$
```

In this example, the `case` statement defines three options that it processes. The `-b` option also requires an additional parameter value. Because the parameter being processed is `$1`, you know that the additional parameter value is located in `$2` (because all the parameters are shifted after they are processed). Just extract the parameter value from the `$2` variable. Of course, because we used two parameter spots for this option, you also need to set the `shift` command to shift one additional position.

Just as with the basic feature, this process works no matter what order you place the options in (just remember to include the appropriate option parameter with the each option):

```
$ ./test17.sh -b test1 -a -d
Found the -b option, with parameter value test1
Found the -a option
-disnotanoption
$
```

Now you have the basic ability to process command line options in your shell scripts, but there are limitations. For example, this doesn't work if you try to combine multiple options in one parameter:

```
$ ./test17.sh -ac
-ac is not an option
$
```

It is a common practice in Linux to combine options, and if your script is going to be user-friendly, you'll want to offer this feature for your users as well. Fortunately, there's another method for processing options that can help you.

Using the `getopt` command

The `getopt` command is a great tool to have handy when processing command line options and parameters. It reorganizes the command line parameters to make parsing them in your script easier.

Looking at the command format

The `getopt` command can take a list of command line options and parameters, in any form, and automatically turn them into the proper format. It uses the following command format:

```
getopt optstring parameters
```

The *optstring* is the key to the process. It defines the valid option letters that can be used in the command line. It also defines which option letters require a parameter value.

First, list each command line option letter you're going to use in your script in the *optstring*. Then place a colon after each option letter that requires a parameter value. The `getopt` command parses the supplied parameters based on the *optstring* you define.

Tip

A more advanced version of the `getopt` command, called `getopts` (notice it is plural), is available. The `getopts` command is covered later in this chapter. Because of their nearly identical spelling, it's easy to get these two commands confused. Be careful!

Here's a simple example of how `getopt` works:

```
$getoptab:cd-a-btest1-cdtest2test3
-a-btest1-c-d--test2test3
$
```

The *optstring* defines four valid option letters, *a*, *b*, *c*, and *d*. A colon (*:*) is placed behind the letter *b* in order to require option *b* to have a parameter value. When the `getopt` command runs, it examines the provided parameter list (`-a-btest1-cdtest2test3`) and parses it based on the supplied *optstring*. Notice that it automatically separated the `-cd` options into two separate options and inserted the double dash to separate the additional parameters on the line.

If you specify a parameter option not in the *optstring*, by default the `getopt` command produces an error message:

```
$getoptab:cd-a-btest1-cdetest2test3
getopt: invalid option --e
-a-btest1-c-d--test2test3
$
```

If you prefer to just ignore the error messages, use `getopt` with the `-q` option:

```
$getopt-qab:cd-a-btest1-cdetest2test3
-a-b'test1'-c-d--'test2''test3'
$
```

Note that the `getopt` command options must be listed before the *optstring*. Now you should be ready to use this command in your scripts to process command line options.

Usinggetoptinyourscripts

You can use the `getopt` command in your scripts to format any command line options or parameters entered for your script. It's a little tricky, however, to use.

The trick is to replace the existing command line options and parameters with the formatted version produced by the `getopt` command. The way to do that is to use the `set` command.

You saw the `set` command back in Chapter 6. The `set` command works with the different variables in the shell.

One of the `set` command options is the double dash (`--`). The double dash instructs `set` to replace the command line parameter variables with the values on the `set` command's command line.

The trick then is to feed the original script command line parameters to the `getopt` command and then feed the output of the `getopt` command to the `set` command to replace the original command line parameters with the nicely formatted ones from `getopt`. This looks something like this:

```
set--$(getopt-qab:cd"$@")
```

Now the values of the original command line parameter variables are replaced with the output from the `getopt` command, which formats the command line parameters for us.

Using this technique, we can now write scripts that handle our command line parameters for us:

```
$cat test18.sh
#!/bin/bash
#Extract command line options&values with getopt #
set--$(getopt-qab:cd"$@") #
echo
while [-n "$1"] do
    case "$1" in
        -a) echo "Found the -a option";;
        -b) param="$2"
            echo "Found the -b option, with parameter value $param" shift ;;
        -c) echo "Found the -c option";;
        --) shift
            break;;
        *) echo "$1 is not an option";; esac
    shift
done
```

```
done
#
count=1
for param in "$@"
do
    echo "Parameter#$count:$param"
    count=$((count + 1))
done
#
$
```

You'll notice this is basically the same script as in `test17.sh`. The only thing that changed is the addition of the `getopt` command to help format our command line parameters.

Now when you run the script with complex options, things work much better:

```
$. /test18.sh -ac

Found the -a option
Found the -c option
$
```

And of course, all the original features work just fine as well:

```
$. /test18.sh -a-b test1 -cd test2 test3 test4

Found the -a option
Found the -b option, with parameter value 'test1'
Found the -c option
Parameter#1: 'test2'
Parameter#2: 'test3'
Parameter#3: 'test4'
$
```

Now things are looking pretty fancy. However, there's still one small bug that lurks in the `getopt` command. Check out this example:

```
$. /test18.sh -a-b test1 -cd "test2 test3" test4

Found the -a option
Found the -b option, with parameter value 'test1'
Found the -c option
Parameter#1: 'test2
Parameter#2: test3'
Parameter#3: 'test4'
$
```

The `getopt` command isn't good at dealing with parameter values with spaces and quotation marks. It interpreted the space as the parameter separator, instead of following the

double quotation marks and combining the two values into one parameter. Fortunately, this problem has another solution.

Advancingtogetopts

The `getopts` command (notice that it is plural) is built into the bash shell. It looks much like its `getopt` cousin, but has some expanded features.

Unlike `getopt`, which produces one output for all the processed options and parameters found in the command line, the `getopts` command works on the existing shell parameter variables sequentially.

It processes the parameters it detects in the command line one at a time each time it's called. When it runs out of parameters, it exits with an exit status greater than zero. This makes it great for using in loops to parse all the parameters on the command line.

Here's the format of the `getopts` command:

```
getopts optstring variable
```

The *optstring* value is similar to the one used in the `getopt` command. Valid option letters are listed in the *optstring*, along with a colon if the option letter requires a parameter value. To suppress error messages, start the *optstring* with a colon. The `getopts` command places the current parameter in the *variable* defined in the command line.

The `getopts` command uses two environment variables. The `OPTARG` environment variable contains the value to be used if an option requires a parameter value. The `OPTIND` environment variable contains the value of the current location within the parameter list where `getopts` left off. This allows you to continue processing other command line parameters after finishing the options.

Let's look at a simple example that uses the `getopts` command:

```
$ cat test19.sh
#!/bin/bash
# simple demonstration of the getopts command #
echo
while getopts :ab:c opt do
    case "$opt" in
        a) echo "Found the -a option";;
        b) echo "Found the -b option, with value $OPTARG";;
        c) echo "Found the -c option";;
        *) echo "Unknown option: $opt";; esac
    done
done
$
```

```
$ ./test19.sh -abtest1 -c

Found the -a option
Found the -b option, with value test1
Found the -c option
$
```

The `while` statement defines the `getopts` command, specifying what command line options to look for, along with the variable name (*opt*) to store them in for each iteration.

You'll notice something different about the `case` statement in this example. When the `getopts` command parses the command line options, it strips off the leading dash, so you don't need leading dashes in the `case` definitions.

The `getopts` command offers several nice features. For starters, you can include spaces in your parameter values:

```
$ ./test19.sh -b "test1 test2" -a

Found the -b option, with value test1 test2
Found the -a option
$
```

Another nice feature is that you can run the option letter and the parameter value together without a space:

```
$ ./test19.sh -abtest1

Found the -a option
Found the -b option, with value test1
$
```

The `getopts` command correctly parsed the `test1` value from the `-b` option. In addition, the `getopts` command bundles any undefined option it finds in the command line into a single output, the question mark:

```
$ ./test19.sh -d

Unknown option: ?
$
$ ./test19.sh -acde

Found the -a option
Found the -c option
Unknown option: ?
Unknown option: ?
$
```

Any option letter not defined in the *optstring* value is sent to your code as a question mark.

PartII:ShellScriptingBasics

The `getopts` command knows when to stop processing options and leave the parameters for you to process. As `getopts` processes each option, it increments the `OPTIND` environment variable by one. When you've reached the end of the `getopts` processing, you can use the `OPTIND` value with the `shift` command to move to the parameters:

```
$cat test20.sh
#!/bin/bash
# Processing options & parameters with getopts #
echo
while getopts :ab:cd opt do
    case "$opt" in
        a) echo "Found the -a option";;
        b) echo "Found the -b option, with value $OPTARG";;
        c) echo "Found the -c option";;
        d) echo "Found the -d option";;
        *) echo "Unknown option: $opt";; esac
    done
#
shift $[${OPTIND}-1] #
echo
count=1
for param in "$@" do
    echo "Parameter $count: $param"
    count=$((count + 1))
done
#
$
$ ./test20.sh -a -b test1 -d test2 test3 test4

Found the -a option
Found the -b option, with value test1
Found the -d option

Parameter 1: test2
Parameter 2: test3
Parameter 3: test4
$
```

Now you have a full-featured command line option and parameter processing utility you can use in all your shell scripts!

StandardizingOptions

When you create your shell script, obviously you're in control of what happens. It's completely up to you as to which letter options you select to use and how you select to use them.

However, a few letter options have achieved a somewhat standard meaning in the Linux world. If you leverage these options in your shell script, your scripts will be more user-friendly.

Table 14-1 shows some of the common meanings for command-line options used in Linux.

TABLE 14-1 Common Linux Command-Line Options

Option	Description
-a	Shows all objects
-c	Produces a count
-d	Specifies a directory
-e	Expands an object
-f	Specifies a file to read data from
-h	Displays a help message for the command
-i	Ignores text case
-l	Produces a long format version of the output
-n	Uses a non-interactive (batch) mode
-o	Specifies an output file to redirect all output to
-q	Runs in quiet mode
-r	Processes directories and files recursively
-s	Runs in silent mode
-v	Produces verbose output
-x	Excludes an object
-y	Answers yes to all questions

You'll probably recognize most of these option meanings just from working with the various `bash` commands throughout the book. Using the same meaning for your options helps users interact with your script without having to worry about manuals.

GettingUserInput

Although providing command line options and parameters is a great way to get data from your script users, sometimes your script needs to be more interactive. Sometimes you need to ask a question while the script is running and wait for a response from the person running your script. The bash shell provides the `read` command just for this purpose.

Readingbasics

The `read` command accepts input either from standard input (such as from the keyboard) or from another file descriptor. After receiving the input, the `read` command places the data into a variable. Here's the `read` command at its simplest:

```
$ cat test21.sh
#!/bin/bash
#testing the read command #
echo -n "Enter your name: "
read name
echo "Hello $name, welcome to my program."
$
$ ./test21.sh
Enter your name: Rich Blum
Hello Rich Blum, welcome to my program.
$
```

That's pretty simple. Notice that the `echo` command that produced the prompt uses the `-n` option. This suppresses the newline character at the end of the string, allowing the script user to enter data immediately after the string, instead of on the next line. This gives your scripts a more form-like appearance.

In fact, the `read` command includes the `-p` option, which allows you to specify a prompt directly in the `read` command line:

```
$ cat test22.sh
#!/bin/bash
#testing the read -p option #
read -p "Please enter your age: " age
days=$((age * 365))
echo "That makes you over $days days old!"
$
$ ./test22.sh
Please enter your age: 10
That makes you over 3650 days old!
$
```

You'll notice in the first example that when a name was entered, the `read` command assigned both the first name and last name to the same variable. The `read` command assigns all data entered at the prompt to a single variable, or you can specify multiple variables. Each data value entered is assigned to the next variable in the list. If the list of variables runs out before the data does, the remaining data is assigned to the last variable:

```
$ cat test23.sh
#!/bin/bash
#entering multiple variables #
read -p "Enter your name: " first last
echo "Checking data for $last, $first..."
$
$ ./test23.sh
Enter your name: Rich Blum
Checking data for Blum, Rich...
$
```

You can also specify no variables on the `read` command line. If you do that, the `read` command places any data it receives in the special environment variable `REPLY`:

```
$ cat test24.sh
#!/bin/bash
#Testing the REPLY Environment variable #
read -p "Enter your name: " echo
echo Hello $REPLY, welcome to my program. #
$
$ ./test24.sh
Enter your name: Christine

Hello Christine, welcome to my program.
$
```

The `REPLY` environment variable contains all the data entered in the input, and it can be used in the shell script as any other variable.

Timeout

Be careful when using the `read` command. Your script may get stuck waiting for the script user to enter data. If the script must go on regardless of whether any data was entered, you can use the `-t` option to specify a timer. The `-t` option specifies the number of seconds for the `read` command to wait for input. When the timer expires, the `read` command returns a non-zero exit status:

```
$ cat test25.sh
#!/bin/bash
```

PartII:ShellScriptingBasics

```
#timingthedataentry #
ifread-t5-p"Pleaseenteryourname:"name then
    echo"Hello$name,welcometomyscript" else
    echo
    echo"Sorry,tooslow!"
fi
$
$ ./test25.sh
Pleaseenteryourname:Rich
HelloRich,welcometomyscript
$
$ ./test25.sh
Pleaseenteryourname:
Sorry,tooslow!
$
```

Because the `read` command exits with a non-zero exit status if the timer expires, it's easy to use the standard structured statements, such as an `if-then` statement or a `while` loop to track what happened. In this example, when the timer expires, the `if` statement fails, and the shell executes the commands in the `else` section.

Instead of timing the input, you can also set the `read` command to count the input characters. When a preset number of characters has been entered, it automatically exits, assigning the entered data to the variable:

```
$ cat test26.sh
#!/bin/bash
#gettingjustonecharacterofinput #
read-n1-p"Doyouwanttocontinue[Y/N]?"answer case
$answer in
Y|y) echo
    echo"fine,continueon...";;
N|n) echo
    echoOK,goodbye
    exit;;
esac
echo"Thisistheendofthescript"
$
$ ./test26.sh
Doyouwanttocontinue[Y/N]?Y
fine,continueon...
Thisistheendofthescript
$
$ ./test26.sh
```

```
Doyouwanttocontinue[Y/N]?n
OK, goodbye
$
```

This example uses the `-n` option with the value of 1, instructing the `read` command to accept only a single character before exiting. As soon as you press the single character to answer, the `read` command accepts the input and passes it to the variable. You don't need to press the Enter key.

Reading with `no display`

Sometimes you need input from the script user, but you don't want that input to display on the monitor. The classic example is when entering passwords, but there are plenty of other types of data that you need to hide.

The `-s` option prevents the data entered in the `read` command from being displayed on the monitor; actually, the data is displayed, but the `read` command sets the text color to the same as the background color. Here's an example of using the `-s` option in a script:

```
$ cat test27.sh
#!/bin/bash
# hiding input data from the monitor #
read -s -p "Enter your password: " pass
echo "Is your password really $pass?"
$
$ ./test27.sh
Enter your password:
Is your password really T3st1ng?
$
```

The data typed at the input prompt doesn't appear on the monitor but is assigned to the variable for use in the script.

Reading from a file

Finally, you can also use the `read` command to read data stored in a file on the Linux system. Each call to the `read` command reads a single line of text from the file. When no more lines are left in the file, the `read` command exits with a non-zero exit status.

The tricky part is getting the data from the file to the `read` command. The most common method is to pipe the result of the `cat` command of the file directly to a `while` command that contains the `read` command. Here's an example:

```
$ cat test28.sh
#!/bin/bash
```

```
#readingdatafromafile #
count=1
cat test | while read line
do
    echo "Line$count:$line"
    count=$((count + 1))
done
echo "Finished processing the file"
$
$ cat test
The quick brown dog jumps over the lazy fox.
This is a test, this is only a test.
O Romeo, Romeo! Wherefore art thou Romeo?
$
$ ./test28.sh
Line1: The quick brown dog jumps over the lazy fox.
Line2: This is a test, this is only a test.
Line3: O Romeo, Romeo! Wherefore art thou Romeo?
Finished processing the file
$
```

The `while` command loop continues processing lines of the file with the `read` command, until the `read` command exits with a non-zero exit status.

Summary

This chapter showed three methods for retrieving data from the script user. Command line parameters allow users to enter data directly on the command line when they run the script. The script uses positional parameters to retrieve the command line parameters and assign them to variables.

The `shift` command allows you to manipulate the command line parameters by rotating them within the positional parameters. This command allows you to easily iterate through the parameters without knowing how many parameters are available.

You can use three special variables when working with command line parameters. The shell sets the `$#` variable to the number of parameters entered on the command line. The `$*` variable contains all the parameters as a single string, and the `$@` variable contains all the parameters as separate words. These variables come in handy when you're trying to process long parameter lists.

Besides parameters, your script users can use command line options to pass information to your script. Command line options are single letters preceded by a dash. Different options can be assigned to alter the behavior of your script.

The bash shell provides three ways to handle command line options.

The first way is to handle them just like command line parameters. You can iterate through the options using the positional parameter variables, processing each option as it appears on the command line.

Another way to handle command line options is with the `getopt` command. This command converts command line options and parameters into a standard format that you can process in your script. The `getopt` command allows you to specify which letters it recognizes as options and which options require an additional parameter value. The `getopt` command processes the standard command line parameters and outputs the options and parameters in the proper order.

The final method for handling command line options is via the `getopts` command (note that it's plural). The `getopts` command provides more advanced processing of the command line parameters. It allows for multi-value parameters, along with identifying options not defined by the script.

An interactive method to obtain data from your script users is the `read` command. The `read` command allows your script to query users for information and wait. The `read` command places any data entered by the script user into one or more variables, which you can use within the script.

Several options are available for the `read` command that allow you to customize the data input into your script, such as using hidden data entry, applying timed data entry, and requesting a specific number of input characters.

In the next chapter, we look further into how bash shell scripts output data. So far, you've seen how to display data on the monitor and redirect it to a file. Next, we explore a few other options that you have available not only to direct data to specific locations but also to direct specific types of data to specific locations. This will help make your shell scripts look professional!

ScriptControl

IN THIS CHAPTER

Handling signals

Running scripts in the background

Forbidding hang-ups

Controlling a Job

Modifying script priority

Automating script execution

As you start building advanced scripts, you'll probably wonder how to run and control them on your Linux system. So far in this book, the only way we've run scripts is directly from the command line interface in real-time mode. This isn't the only way to run scripts in Linux.

Quite a few options are available for running your shell scripts. There are also options for controlling your scripts. Various control methods include sending signals to your script, modifying a script's priority, and switching the run mode while a script is running. This chapter examines the different ways you can control your shell scripts.

Handling Signals

Linux uses signals to communicate with processes running on the system. Chapter 4 described the different Linux signals and how the Linux system uses these signals to stop, start, and kill processes. You can control the operation of your shell script by programming the script to perform certain commands when it receives specific signals.

Signaling the bash shell

There are more than 30 Linux signals that can be generated by the system and applications. Table 16-1 lists the most common Linux system signals that you'll run across in your shell script writing.

TABLE 16-1 **Linux Signals**

Signal	Value	Description
1	SIGHUP	Hangsup the process
2	SIGINT	Interrupts the process
3	SIGQUIT	Stopsthe process
9	SIGKILL	Unconditionally terminates the process
15	SIGTERM	Terminates the process if possible
17	SIGSTOP	Unconditionally stops, but doesn't terminate, the process
18	SIGTSTP	Stops or pauses the process, but doesn't terminate
19	SIGCONT	Continues a stopped process

By default, the bash shell ignores any `SIGQUIT` (3) and `SIGTERM` (15) signals it receives (so an interactive shell cannot be accidentally terminated). However, the bash shell does not ignore any `SIGHUP` (1) and `SIGINT` (2) signals it receives.

If the bash shell receives a `SIGHUP` signal, such as when you leave an interactive shell, it exits. Before it exits, however, it passes the `SIGHUP` signal to any processes started by the shell, including any running shell scripts.

With a `SIGINT` signal, the shell is just interrupted. The Linux kernel stops giving the shell processing time on the CPU. When this happens, the shell passes the `SIGINT` signal to any processes started by the shell to notify them of the situation.

As you probably have noticed, the shell passes these signals on to your shell script program for processing. However, a shell script's default behavior does not govern these signals, which may have an adverse effect on the script's operation. To avoid this situation, you can program your script to recognize signals and perform commands to prepare the script for the consequences of the signal.

Generating signals

The bash shell allows you to generate two basic Linux signals using key combinations on the keyboard. This feature comes in handy if you need to stop or pause a runaway script.

Interrupting a process

The `Ctrl+C` key combination generates a `SIGINT` signal and sends it to any processes currently running in the shell. You can test this by running a command that normally takes a long time to finish and pressing the `Ctrl+C` key combination:

```
$ sleep 100
^C
$
```

The `Ctrl+C` key combination sends a `SIGINT` signal, which simply stops the current process running in the shell. The `sleep` command pauses the shell's operation for the specified number of seconds and returns the shell prompt. By pressing the `Ctrl+C` key combination before the time passed, the `sleep` command terminated prematurely.

Pausing a process

Instead of terminating a process, you can pause it in the middle of whatever it's doing. Sometimes, this can be a dangerous thing (for example, if a script has a file lock open on a crucial system file), but often it allows you to peek inside what a script is doing without actually terminating the process.

The `Ctrl+Z` key combination generates a `SIGTSTP` signal, stopping any processes running in the shell. Stopping a process is different than terminating the process. Stopping the process leaves the program in memory and able to continue running from where it left off. In the "Controlling the Job" section later in this chapter, you learn how to restart a process that's been stopped.

When you use the `Ctrl+Z` key combination, the shell informs you that the process has been stopped:

```
$ sleep 100
^Z
[1]+ Stopped                  sleep100
$
```

The number in the square brackets is the *job number* assigned by the shell. The shell refers to each process running in the shell as a *job* and assigns each job a unique job number within the current shell. It assigns the first started process job number 1, the second job number 2, and so on.

If you have a stopped job assigned to your shell session, `bash` warns you if you try to exit the shell:

```
$ sleep 100
^Z
[1]+ Stopped                  sleep100
$ exit
exit
There are stopped jobs.
$
```

You can view the stopped jobs using the `ps` command:

```
$ sleep 100
^Z
[1]+ Stopped                  sleep100
$
$ ps -l
```

```
F S UID    PID  PPID  C PRI NI ADDR SZ WCHAN  TTY          TIME CMD
0 S 501    2431  2430  0  80  0 -27118 wait   pts/0 00:00:00 bash
0 T 501    2456  2431  0  80  0 -25227 signal pts/0 00:00:00 sleep
0 R 501    2458  2431  0  80  0 -27034 -      pts/0 00:00:00 ps
$
```

In the `S` column (process state), the `ps` command shows the stopped job's state as `T`. This indicates the command is either being traced or is stopped.

If you really want to exit the shell with a stopped job still active, just type the `exit` command again. The shell exits, terminating the stopped job. Alternately, now that you know the PID of the stopped job, you can use the `kill` command to send a `SIGKILL` signal to terminate it:

```
$kill-92456
$
[1]+Killed                  sleep100
$
```

When you kill the job, initially you don't get any response. However, the next time you do something that produces a shell prompt (such as pressing the Enter key), you'll see a message indicating that the job was killed. Each time the shell produces a prompt, it also displays the status of any jobs that have changed states in the shell. After you kill a job, the next time you force the shell to produce a prompt, it displays a message showing that the job was killed while running.

Trapping signals

Instead of allowing your script to leave signals ungoverned, you can trap them when they appear and perform other commands. The `trap` command allows you to specify which Linux signals your shell script can watch for and intercept from the shell. If the script receives a signal listed in the `trap` command, it prevents it from being processed by the shell and instead handles it locally.

The format of the `trap` command is:

```
trap commands signals
```

On the `trap` command line, you just list the commands you want the shell to execute, along with a space-separated list of signals you want to trap. You can specify the signals either by their numeric value or by their Linux signal name.

Here's a simple example of using the `trap` command to capture the `SIGINT` signal and govern the script's behavior when the signal is sent:

```
$cat test1.sh
#!/bin/bash
#Testingsignaltrapping #
```

```
trap"echo'Sorry!IhavetrappedCtrl-C'"SIGINT #
echoThisistestscript #
count=1
while[$count-le10] do
    echo"Loop#$count"s
    sleep 1
    count=$((count+1)) done
#
echo"Thisistheendofthetestscript"#
```

The `trap` command used in this example displays a simple text message each time it detects the `SIGINT` signal. Trapping this signal makes this script impervious to the user attempting to stop the program by using the bash shell keyboard `Ctrl+C` command:

```
$ ./test1.sh
Thisistestscript
Loop #1
Loop#2
Loop#3
Loop#4
Loop#5
^CSorry!IhavetrappedCtrl-C
Loop #6
Loop#7
Loop#8
^CSorry!IhavetrappedCtrl-C
Loop #9
Loop#10
Thisistheendofthetestscript
$
```

Each time the `Ctrl+C` key combination was used, the script executed the `echo` statement specified in the `trap` command instead of not managing the signal and allowing the shell to stop the script.

Trapping script exit

Besides trapping signals in your shell script, you can trap them when the shell script exits. This is a convenient way to perform commands just as the shell finishes its job.

To trap the shell script exiting, just add the `EXIT` signal to the `trap` command:

```
$ cat test2.sh
#!/bin/bash
```

PartII:ShellScriptingBasics

```
#Trappingthescriptexit #
trap"echoGoodbye..."EXIT #
count=1
while[$count-le5] do
    echo"Loop#$count"s
    leep 1
    count=${count+1} done
#
$
$ ./test2.sh
Loop#1
Loop#2
Loop#3
Loop#4
Loop #5
Goodbye...
$
```

When the script gets to the normal exit point, the trap is triggered, and the shell executes the command you specify on the `trap` command line. The `EXIT` trap also works if you prematurely exit the script:

```
$ ./test2.sh
Loop#1
Loop#2
Loop#3
^C Goodbye...
$
```

Because the `SIGINT` signal isn't listed in the `trap` command list, when the `Ctrl+C` key combination is used to send that signal, the script exits. However, before the script exits, because the `EXIT` is trapped, the shell executes the `trap` command.

Modifying or removing a trap

To handle traps differently in various sections of your shell script, you simply reissue the `trap` command with new options:

```
$ cat test3.sh
#!/bin/bash
# Modifying a set trap #
trap"echo'Sorry...Ctrl-C is trapped.'"SIGINT
```

```

#
count=1
while[$count-le5] do
    echo"Loop#$count"s
    leep 1
    count=${count+1} done
#
trap"echo'Imodifiedthetrap!'"SIGINT #
count=1
while[$count-le5] do
    echo"SecondLoop#$count"sleee
    1
    count=${count+1} done
#
$

```

After the signal trap is modified, the script manages the signal or signals differently. However, if a signal is received before the trap is modified, the script processes it per the original `trap` command:

```

$ ./test3.sh
Loop#1
Loop#2
Loop#3
^C Sorry...Ctrl-C is trapped. Loop
#4
Loop#5
SecondLoop#1
SecondLoop#2
^C Imodifiedthetrap!
Second Loop #3
SecondLoop#4
SecondLoop#5
$

```

You can also remove a set trap. Simply add two dashes after the `trap` command and a list of the signals you want to return to default behavior:

```

$ cat test3b.sh
#!/bin/bash
# Removing a set trap #
trap"echo'Sorry...Ctrl-C is trapped.'"SIGINT #

```


PartII:ShellScriptingBasics

```
count=1
while[$count-le5] do
    echo"Loop#$count"s
    leep 1
    count=${count+1} done
#
#Removethetrap
trap -- SIGINT
echo"Ijustremovedthetrap"#
count=1
while[$count-le5] do
    echo"SecondLoop#$count"sleee
    1
    count=${count+1} done
#
$./test3b.sh
Loop#1
Loop#2
Loop#3
Loop#4
Loop#5
Ijustremovedthetrap
Second Loop #1
SecondLoop#2
SecondLoop#3
^C
$
```

TIP

You can use a single dash instead of a double dash after the `trap` command to return signals to their default behavior. Both the single and double dash work properly.

After the signal trap is removed, the script handles the `SIGINT` signal in its default manner, terminating the script. However, if a signal is received before the trap is removed, the script processes it per the original `trap` command:

```
$./test3b.sh
Loop#1
Loop#2
Loop#3
^CSorry...Ctrl-Cistrapped. Loop
#4
```

```
Loop#5
Ijustremovedthetrap
Second Loop #1
SecondLoop#2
^C
$
```

In this example, the first Ctrl+C key combination was used to attempt to terminate the script prematurely. Because the signal was received before the trap was removed, the script executed the command specified in the trap. After the script executed the trap removal, then Ctrl+C could prematurely terminate the script.

Running Scripts in Background Mode

Sometimes, running a shell script directly from the command line interface is inconvenient. Some scripts can take a long time to process, and you may not want to tie up the command line interface waiting. While the script is running, you can't do anything else in your terminal session. Fortunately, there's a simple solution to that problem.

When you use the `ps` command, you see a whole bunch of different processes running on the Linux system. Obviously, all these processes are not running on your terminal monitor. This is called running processes in the *background*. In background mode, a process runs without being associated with a `STDIN`, `STDOUT`, and `STDERR` on a terminal session (see Chapter 15).

You can exploit this feature with your shell scripts as well, allowing them to run behind the scenes and not lock up your terminal session. The following sections describe how to run your scripts in background mode on your Linux system.

Running in the background

Running a shell script in background mode is a fairly easy thing to do. To run a shell script in background mode from the command line interface, just place an ampersand symbol (&) after the command:

```
$ cat test4.sh
#!/bin/bash
# Test running in the background #
count=1
while [ $count -le 10 ] do
    sleep 1
    count=$((count+1)) done
```

PartII:ShellScriptingBasics

```
#
$
$ ./test4.sh&
[1] 3231
$
```

When you place the ampersand symbol after a command, it separates the command from the bash shell and runs it as a separate background process on the system. The first thing that displays is the line:

```
[1] 3231
```

The number in the square brackets is the job number assigned by the shell to the background process. The next number is the Process ID (PID) the Linux system assigns to the process. Every process running on the Linux system must have a unique PID.

As soon as the system displays these items, a new command line interface prompt appears. You are returned to the shell, and the command you executed runs safely in background mode. At this point, you can enter new commands at the prompt.

When the background process finishes, it displays a message on the terminal:

```
[1]      Done                  ./test4.sh
```

This shows the job number and the status of the job (Done), along with the command used to start the job.

Be aware that while the background process is running, it still uses your terminal monitor for `STDOUT` and `STDERR` messages:

```
$ cat test5.sh
#!/bin/bash
# Test running in the background without output #
echo "Start the test script" count=1
while [ $count -le 5 ] do
    echo "Loop # $count"s
    sleep 5
    count=$((count+1)) done
#
echo "Test script is complete" #
$
$ ./test5.sh&
[1] 3275
```

```
$Startthetestscript
Loop #1
Loop#2
Loop#3
Loop#4
Loop#5
Testscriptiscomplete

[1]   Done                  ./test5.sh
$
```

You'll notice from the example that the output from the `test5.sh` script displays. The output intermixes with the shell prompt, which is why `Startthetestscript` appears next to the `$` prompt.

You can still issue commands while this output is occurring:

```
$ ./test5.sh&
[1] 3319
$Startthetestscript
Loop #1
Loop#2
Loop#3
lsmyprog*
myprogmyprog.c
$Loop#4
Loop#5
Testscriptiscomplete

[1]+Done                  ./test5.sh
$$
```

While the `test5.sh` script is running in the background, the command `lsmyprog*` was entered. The script's output, the typed command, and the command's output all intermixed with each other's output display. This can be confusing! It is a good idea to redirect `STDOUT` and `STDERR` for scripts you will be running in the background (Chapter 15) to avoid this messy output.

Running multiple background jobs

You can start any number of background jobs at the same time from the command line prompt:

```
$ ./test6.sh&
[1] 3568
$ThisIsTestScript#1

$ ./test7.sh&
```

```
[2]3570
$ThisisTestScript#2

$./test8.sh&
[3]3573
$And...anotherTestscript

$./test9.sh&
[4]3576
$Then...therewasonemoretestscript

$
```

Each time you start a new job, the Linux system assigns it a new job number and PID. You can see that all the scripts are running using the `ps` command:

```
$ps
  PID TTY          TIME CMD
 2431 pts/0    00:00:00 bash
 3568 pts/0    00:00:00 test6.sh
 3570 pts/0    00:00:00 test7.sh
 3573 pts/0    00:00:00 test8.sh
 3574 pts/0    00:00:00 sleep
 3575 pts/0    00:00:00 sleep
 3576 pts/0    00:00:00 test9.sh
 3577 pts/0    00:00:00 sleep
 3578 pts/0    00:00:00 sleep
 3579 pts/0    00:00:00 ps
$
```

You must be careful when using background processes from a terminal session. Notice in the output from the `ps` command that each of the background processes is tied to the terminal session (pts/0) terminal. If the terminal session exits, the background process also exits.

NOTE

Earlier in this chapter we mentioned that when you attempt to exit a terminal session, a warning is issued if there are stopped processes. However, with background processes, only some terminal emulators remind you that a background job is running, before you attempt to exit the terminal session.

If you want your script to continue running in background mode after you have logged off the console, there's something else you need to do. The next section discusses that process.

Running Scripts without a Hang-Up

Sometimes, you may want to start a shell script from a terminal session and let the script run in background mode until it finishes, even if you exit the terminal session. You can do this by using the `nohup` command.

The `nohup` command runs another command blocking any `SIGHUP` signals that are sent to the process. This prevents the process from exiting when you exit your terminal session.

The format used for the `nohup` command is as follows:

```
$nohup./test1.sh&
[1]3856
$nohup:ignoringinputandappendingoutputto'nohup.out'

$
```

As with a normal background process, the shell assigns the command a job number, and the Linux system assigns a PID number. The difference is that when you use the `nohup` command, the script ignores any `SIGHUP` signals sent by the terminal session if you close the session.

Because the `nohup` command disassociates the process from the terminal, the process loses the `STDOUT` and `STDERR` output links. To accommodate any output generated by the command, the `nohup` command automatically redirects `STDOUT` and `STDERR` messages to a file, called `nohup.out`.

NOTE

If you run another command using `nohup`, the output is appended to the existing `nohup.out` file. Be careful when running multiple commands from the same directory, because all the output is sent to the same `nohup.out` file, which can get confusing.

The `nohup.out` file contains all the output that would normally be sent to the terminal monitor. After the process finishes running, you can view the `nohup.out` file for the output results:

```
$catnohup.out
This is a test script
Loop 1
Loop 2
Loop 3
Loop 4
Loop 5
Loop 6
Loop 7
Loop 8
Loop 9
Loop 10
This is the end of the test script
$
```

The output appears in the `nohup.out` file just as if the process ran on the command line.

Controlling the Job

Earlier in this chapter, you saw how to use the Ctrl+C key combination to stop a job running in the shell. After you stop a job, the Linux system lets you either kill or restart it. You can kill the process by using the `kill` command. Restarting a stopped process requires that you send it a `SIGCONT` signal.

The function of starting, stopping, killing, and resuming jobs is called *job control*. With job control, you have full control over how processes run in your shell environment. This section describes the commands used to view and control jobs running in your shell.

Viewing jobs

The key command for job control is the `jobs` command. The `jobs` command allows you to view the current jobs being handled by the shell:

```
$ cat test10.sh
#!/bin/bash
# Test job control
#
echo "Script Process ID: $$"
count=1
while [ $count -le 10 ] do
    echo "Loop # $count"
    sleep 10
    count=$((count+1))
done
#
echo "End of script..."
$
```

The script uses the `$$` variable to display the PID that the Linux system assigns to the script; then it goes into a loop, sleeping for 10 seconds at a time for each iteration.

You can start the script from the command line interface and then stop it using the Ctrl+Z key combination:

```
$ ./test10.sh
Script Process ID: 1897
Loop #1
Loop #2
^Z
[1]+ Stopped                  ./test10.sh
$
```

Using the same script, another job is started as a background process, using the ampersand symbol. To make life a little easier, the output of that script is redirected to a file so it doesn't appear on the screen:

```
$ ./test10.sh>test10.out&
[2]1917
$
```

The `jobs` command enables you to view the jobs assigned to the shell. The `jobs` command shows both the stopped and the running jobs, along with their job numbers and the commands used in the jobs:

```
$ jobs
[1]+ Stopped                  ./test10.sh
[2]- Running                  ./test10.sh>test10.out&
$
```

You can view the various jobs' PIDs by adding the `-l` parameter (lowercase L) on the `jobs` command:

```
$ jobs -l
[1]+1897 Stopped              ./test10.sh
[2]-1917 Running              ./test10.sh>test10.out&
$
```

The `jobs` command uses a few different command line parameters, as shown in Table 16-2.

TABLE 16-2 The `jobs` Command Parameters

Parameter	Description
-l	List the PID of the process along with the job number
-n	List only jobs that have changed their status since the last notification from the shell
-p	List only the PIDs of the jobs
-r	List only the running jobs
-s	List only stopped jobs

You probably noticed the plus and minus signs in the `jobs` command output. The job with the plus sign is considered the default job. It would be the job referenced by any job control commands if a job number wasn't specified in the command line.

The job with the minus sign is the job that would become the default job when the current default job finishes processing. There will be only one job with the plus sign and one job with the minus sign at any time, no matter how many jobs are running in the shell.

The following is an example showing how the next job in line takes over the default status, when the default job is removed. Three separate processes are started in the background. The `jobs` command listing shows the three processes, their PID, and their status. Note that the default process (the one listed with the plus sign) is the last process started, job #3.

```
$ ./test10.sh>test10a.out&
[1]1950
$ ./test10.sh>test10b.out&
[2]1952
$ ./test10.sh>test10c.out&
[3]1955
$
$ jobs-l
[1] 1950Running                ./test10.sh>test10a.out&
[2]-1952Running                ./test10.sh>test10b.out&
[3]+1955Running                ./test10.sh>test10c.out&
$
```

Using the `kill` command to send a `SIGHUP` signal to the default process causes the job to terminate. In the next `jobs` listing, the job that previously had the minus sign now has the plus sign and is the default job:

```
$ kill1955
$
[3]+Terminated                ./test10.sh>test10c.out
$
$ jobs-l
[1]-1950Running                ./test10.sh>test10a.out&
[2]+1952Running                ./test10.sh>test10b.out&
$
$ kill1952
$
[2]+Terminated                ./test10.sh>test10b.out
$
$ jobs-l
[1]+1950Running                ./test10.sh>test10a.out&
$
```

Although changing a background job to the default process is interesting, it doesn't seem very useful. In the next section, you learn how to use commands to interact with the default process using no PID or job number.

Restartingstoppedjobs

Under bash job control, you can restart any stopped job as either a background process or a foreground process. A foreground process takes over control of the terminal you're working on, so be careful about using that feature.

To restart a job in background mode, use the `bg` command:

```
$ ./test11.sh
^Z
[1]+ Stopped                  ./test11.sh
$
$bg
[1]+ ./test11.sh&
$
$jobs
[1]+ Running                  ./test11.sh&
$
```

Because the job was the default job, indicated by the plus sign, only the `bg` command was needed to restart it in background mode. Notice that no PID is listed when the job is moved into background mode.

If you have additional jobs, you need to use the job number along with the `bg` command:

```
$ ./test11.sh
^Z
[1]+ Stopped                  ./test11.sh
$
$ ./test12.sh
^Z
[2]+ Stopped                  ./test12.sh
$
$bg2
[2]+ ./test12.sh&
$
$jobs
[1]+ Stopped                  ./test11.sh
[2]- Running                  ./test12.sh&
$
```

The command `bg2` was used to send the second job into background mode. Notice that when the `jobs` command was used, it listed both jobs with their status, even though the default job is not currently in background mode.

To restart a job in foreground mode, use the `fg` command, along with the job number:

```
$fg2
./test12.sh
This is the script 'send...'
$
```

Because the job is running in foreground mode, the command line interface prompt does not appear until the job finishes.

BeingNice

In a multitasking operating system (which Linux is), the kernel is responsible for assigning CPU time for each process running on the system. The *schedulingpriority* is the amount of CPU time the kernel assigns to the process relative to the other processes. By default, all processes started from the shell have the same scheduling priority on the Linux system.

The scheduling priority is an integer value, from -20 (the highest priority) to +19 (the lowest priority). By default, the bash shell starts all processes with a scheduling priority of 0.

Tip

It's confusing to remember that -20, the lowest value, is the highest priority and 19, the highest value, is the lowest priority. Just remember the phrase, "Nice guys finish last." The "nicer" or higher you are in value, the lower your chance of getting the CPU.

Sometimes, you want to change the priority of a shell script, either lowering its priority so it doesn't take as much processing power away from other processes or giving it a higher priority so it gets more processing time. You can do this by using the `nice` command.

Using the `nice` command

The `nice` command allows you to set the scheduling priority of a command as you start it. To make a command run with less priority, just use the `-n` command line option for `nice` to specify a new priority level:

```
$ nice -n 10 ./test4.sh > test4.out &
[1] 4973
$
$ ps -p 4973 -o pid,ppid,ni,cmd
  PID PPID NI CMD
  4973 472 10 /bin/bash ./test4.sh
$
```

Notice that you must use the `nice` command on the same line as the command you are starting. The output from the `ps` command confirms that the `nice` value (column `NI`) has been set to 10.

The `nice` command causes the script to run at a lower priority. However, if you try to increase the priority of one of your commands, you might be in for a surprise:

```
$ nice -n -10 ./test4.sh > test4.out &
[1] 4985
$ nice: cannot set niceness: Permission denied

[1]+ Done                               nice -n -10 ./test4.sh > test4.out
$
```

The `nice` command prevents normal system users from increasing the priority of their commands. Notice that the job does run, even though the attempt to raise its priority with the `nice` command failed.

You don't have to use the `-n` option with the `nice` command. You can simply type the priority preceded by a dash:

```
$ nice -10 ./test4.sh > test4.out &
[1] 4993
$
$ ps -p 4993 -o pid,ppid,ni,cmd
    PID PPID NIC CMD
    4993 4721 10 /bin/bash ./test4.sh
$
```

However, this can get confusing when the priority is a negative number, because you must have a double-dash. It's best just to use the `-n` option to avoid confusion.

Using the `renice` command

Sometimes, you'd like to change the priority of a command that's already running on the system. That's what the `renice` command is for. It allows you to specify the PID of a running process to change its priority:

```
$ ./test11.sh &
[1] 5055
$
$ ps -p 5055 -o pid,ppid,ni,cmd
    PID PPID NIC CMD
    5055 4721  0 /bin/bash ./test11.sh
$
$ renice -n 10 -p 5055
5055: oldpriority 0, newpriority 10
$
$ ps -p 5055 -o pid,ppid,ni,cmd
    PID PPID NIC CMD
    5055 4721 10 /bin/bash ./test11.sh
$
```

The `renice` command automatically updates the scheduling priority of the running process. As with the `nice` command, the `renice` command has some limitations:

- You can only `renice` processes that you own.
- You can only `renice` your processes to a lower priority.
- The root user can `renice` any process to any priority.

If you want to fully control running processes, you must be logged in as the root account or use the `sudo` command.

Running Like Clockwork

When you start working with scripts, you may want to run a script at a preset time, usually at a time when you're not there. The Linux system provides a couple of ways to run a script at a preselected time: the `at` command and the `crontable`. Each method uses a different technique for scheduling when and how often to run scripts. The following sections describe each of these methods.

Scheduling a job using the `at` command

The `at` command allows you to specify a time when the Linux system will run a script. The `at` command submits a job to a queue with directions on when the shell should run the job. The `atd` daemon, `atd`, runs in the background and checks the job queue for jobs to run. Most Linux distributions start this daemon automatically at boot time.

The `atd` daemon checks a special directory on the system (usually `/var/spool/at`) for jobs submitted using the `at` command. By default, the `atd` daemon checks this directory every 60 seconds. When a job is present, the `atd` daemon checks the time the job is set to be run. If the time matches the current time, the `atd` daemon runs the job.

The following sections describe how to use the `at` command to submit jobs to run and how to manage these jobs.

Understanding the `at` command format

The basic `at` command format is pretty simple:

```
at [-f filename] time
```

By default, the `at` command submits input from `STDIN` to the queue. You can specify a filename used to read commands (your script file) using the `-f` parameter.

The `time` parameter specifies when you want the Linux system to run the job. If you specify a time that has already passed, the `at` command runs the job at that time on the next day.

You can get pretty creative with how you specify the time. The `at` command recognizes lots of different time formats:

- A standard hour and minute, such as `10:15`
- An AM/PM indicator, such as `10:15PM`
- A specific named time, such as `now`, `noon`, `midnight`, or `teatime` (`4PM`)

In addition to specifying the time to run the job, you can also include a specific date, using a few different date formats:

- A standard date format, such as MMDDYY, MM/DD/YY, or DD.MM.YY
- A text date, such as Jul 4 or Dec 25, with or without the year
- A time increment:
 - Now+25 minutes
 - 10:15 PM tomorrow
 - 10:15+7 days

When you use the `at` command, the job is submitted into a *jobqueue*. The job queue holds the jobs submitted by the `at` command for processing. There are 26 different job queues available for different priority levels. Job queues are referenced using lowercase letters, *a* through *z*, and uppercase letters *A* through *Z*.

NOTE

A few years ago, the `batch` command was another method that allowed a script to be run at a later time. The `batch` command was unique because you could schedule a script to run when the system was at a lower usage level. However, nowadays, the `batch` command is just simply a script, `/usr/bin/batch`, that calls the `at` command and submits your job to the `batch` queue.

The higher alphabetically the job queue, the lower the priority (higher `nice` value) the job will run under. By default, `at` jobs are submitted to the `at` job queue. If you want to run a job at a lower priority, you can specify a different queue letter using the `-q` parameter.

Retrieving job output

When the job runs on the Linux system, there's no monitor associated with the job. Instead, the Linux system uses the e-mail address of the user who submitted the job as `STDOUT` and `STDERR`. Any output destined to `STDOUT` or `STDERR` is mailed to the user via the mail system.

Here's a simple example using the `at` command to schedule a job to run on a CentOS distribution:

```
$ cat test13.sh
#!/bin/bash
# Test using at command #
echo "This script ran at $(date +%B%d, %T)"
echo "sleep 5"
echo "This is the script's end..." #
```

```
$at-ftest13.shnow
job7at2015-07-1412:38
$
```

The `at` command displays the job number assigned to the job along with the time the job is scheduled to run. The `-f` option tells what script file to use and the `now` time designation directs `at` to run the script immediately.

Using e-mail for the `at` command's output is inconvenient at best. The `at` command sends e-mail via the `sendmail` application. If your system does not use `sendmail`, you won't get any output! Therefore, it's best to redirect `STDOUT` and `STDERR` in your scripts (see Chapter 15) when using the `at` command, as the following example shows:

```
$cat test13b.sh
#!/bin/bash
#Test using at command #
echo "This script ran at $(date +%B%d,%T)" >> test13b.out
echo >> test13b.out
sleep 5
echo "This is the script's send..." >> test13b.out #
$
$at-M-f test13b.shnow
job8at2015-07-1412:48
$
$cat test13b.out
This script ran at July 14, 12:48:18

This is the script's send...
$
```

If you don't want to use e-mail or redirection with `at`, it is best to add the `-M` option to suppress any output generated by jobs using the `at` command.

Listing pending jobs

The `atq` command allows you to view what jobs are pending on the system:

```
$at-M-f test13b.sh teatime
job17at2015-07-1416:00
$
$at-M-f test13b.sh tomorrow
job18at2015-07-1513:03
$
$at-M-f test13b.sh 13:30
job19at2015-07-1413:30
$
$at-M-f test13b.shnow
```

```
job20at2015-07-1413:03
$
$atq
20      2015-07-1413:03=Christine
18      2015-07-1513:03aChristine
17      2015-07-1416:00aChristine
19      2015-07-1413:30aChristine
$
```

The job listing shows the job number, the date and time the system will run the job, and the job queue the job is stored in.

Removing jobs

After you know the information about what jobs are pending in the job queues, you can use the `atrm` command to remove a pending job:

```
$atq
18      2015-07-15 13:03 a Christine
17      2015-07-14 16:00 a Christine
19      2015-07-14 13:30 a Christine
$
$atrm 18
$
$atq
17      2015-07-14 16:00 a Christine
19      2015-07-14 13:30 a Christine
$
```

Just specify the job number you want to remove. You can only remove jobs that you submit for execution. You can't remove jobs submitted by others.

Scheduling regular scripts

Using the `at` command to schedule a script to run at a preset time is great, but what if you need that script to run at the same time every day or once a week or once a month? Instead of having to continually submit `at` jobs, you can use another feature of the Linux system.

The Linux system uses the `cron` program to allow you to schedule jobs that need to run on a regular basis. The `cron` program runs in the background and checks special tables, called *cron tables*, for jobs that are scheduled to run.

Looking at the crontable

The `crontable` uses a special format for allowing you to specify when a job should be run. The format for the `crontable` is:

```
min hour day of month month day of week command
```


Part II: Shell Scripting Basics

The `crontable` allows you to specify entries as specific values, ranges of values (such as 1–5), or as a wildcard character (the asterisk). For example, if you want to run a command at 10:15 on every day, you would use this `crontable` entry:

```
1510***command
```

The wildcard character used in the `dayofmonth`, `month`, and `dayofweek` fields indicates that `cron` will execute the command every day of every month at 10:15. To specify a command to run at 4:15 PM every Monday, you would use the following:

```
1516**1command
```

You can specify the `dayofweek` entry as either a three-character text value (mon, tue, wed, thu, fri, sat, sun) or a numeric value, with 0 being Sunday and 6 being Saturday.

Here's another example: to execute a command at 12 noon on the first day of every month, you would use the following format:

```
00121**command
```

The `dayofmonth` entry specifies a date value (1–31) for the month.

NOTE

The astute reader might be wondering just how you would be able to set a command to execute on the last day of every month because you can't set the `dayofmonth` value to cover every month. This problem has plagued Linux and Unix programmers, and has spawned quite a few different solutions. A common method is to add an `if-then` statement that uses the `date` command to check if tomorrow's date is 01:

```
0012***if[`date+%d-dtomorrow`=01];then;command
```

This check every day at 12 noon to see if it's the last day of the month, and if so, `cron` runs the command.

The command list must specify the full command pathname or shell script to run. You can add any command line parameters or redirection symbols you like, as a regular command line:

```
1510***/home/rich/test4.sh>test4out
```

The `cron` program runs the script using the user account that submitted the job. Thus, you must have the proper permissions to access the command and output files specified in the command listing.

Building the crontable

Each system user can have their own `crontable` (including the root user) for running scheduled jobs. Linux provides the `crontab` command for handling the `crontable`. To list an existing `crontable`, use the `-l` parameter:

```
$crontab-l
nocrontabforrich
$
```

By default, each user's crontable file doesn't exist. To add an entry to your crontable, use the `-e` parameter. When you do that, the `crontab` command starts a text editor (see Chapter 10) with the existing crontable (or an empty file if it doesn't yet exist).

Viewing crond directories

When you create a script that has less precise execution time needs, it is easier to use one of the pre-configured `cron` script directories. There are four basic directories: hourly, daily, monthly, and weekly.

```
$ls/etc/cron.*ly
/etc/cron.daily:
cups          makewhatis.cronprelink      tmpwatch
logrotate locate.cron      readahead.cron

/etc/cron.hourly:
0anacron

/etc/cron.monthly:
readahead-monthly.cron

/etc/cron.weekly:
$
```

Thus, if you have a script that needs to be run one time per day, just copy the script to the `daily` directory and `cron` executes it each day.

Looking at the anacron program

The only problem with the `cron` program is that it assumes that your Linux system is operational 24 hours a day, 7 days a week. Unless you're running Linux in a server environment, this may not necessarily be true.

If the Linux system is turned off at the time a job is scheduled to run in the crontable, the job doesn't run. The `cron` program doesn't retroactively run missed jobs when the system is turned back on. To resolve this issue, many Linux distributions also include the `anacron` program.

If `anacron` determines that a job has missed a scheduled running, it runs the job as soon as possible. This means that if your Linux system is turned off for a few days, when it starts back up, any jobs scheduled to run during the time it was off are automatically run.

This feature is often used for scripts that perform routine log maintenance. If the system is always off when the scripts should run, the log files would never get trimmed and could

grow to undesirable sizes. With `anacron`, you're guaranteed that the log files are trimmed at least each time the system is started.

The `anacron` program deals only with programs located in the `crond` directories, such as `/etc/cron.monthly`. It uses timestamps to determine if the jobs have been run at the proper scheduled interval. A timestamp file exists for each `crond` directory and is located in `/var/spool/anacron`:

```
$sudo cat /var/spool/anacron/cron.monthly
20150626
$
```

The `anacron` program has its own table (usually located at `/etc/anacrontab`) to check the job directories:

```
$sudo cat /etc/anacrontab
# /etc/anacrontab: configuration file for anacron

# See anacron(8) and anacrontab(5) for details.

SHELL=/bin/sh
PATH=/sbin:/bin:/usr/sbin:/usr/bin
MAILTO=root
# the maximal random delay added to the basedelay of the jobs  RANDOM_DELAY=45
# the jobs will be started during the following hour only
START_HOURS_RANGE=3-22

#period in days      delay in minutes   job-identifier  command
1          5          cron.daily      nicerun-parts/etc/cron.daily
7          25         cron.weekly      nicerun-parts/etc/cron.weekly
@monthly 45         cron.monthly     nicerun-parts/etc/cron.monthly
$
```

The basic format of the `anacrontab` is slightly different from that of the `crontab`:

period delay identifier command

The period entry defines how often the jobs should be run, specified in days. The `anacron` program uses this entry to check against the jobs' timestamp file. The delay entry specifies how many minutes after the system starts the `anacron` program should run missed scripts. The command entry contains the `run-parts` program and a `cronscript` directory name. The `run-parts` program is responsible for running any script in the directory passed to it.

Notice that `anacron` does not run the scripts located in `/etc/cron.hourly`. This is because the `anacron` program does not deal with scripts that have execution time needs of less than daily.

The identifier entry is a unique non-blank character string — for example, `cron-weekly`. It is used to uniquely identify the job in log messages and error e-mails.

Starting scripts with a new shell

The ability to run a script every time a user starts a new bash shell (even just when a specific user starts a bash shell) can come in handy. Sometimes, you want to set shell features for a shell session or just ensure that a specific file has been set.

Recall the startup files run when a user logs into the bash shell (covered in detail in Chapter 6). Also, remember that not every distribution has all the startup files. Essentially, the first file found in the following ordered list is run and the rest are ignored:

- `$HOME/.bash_profile`
- `$HOME/.bash_login`
- `$HOME/.profile`

Therefore, you should place any scripts you want to run at login time in the first file listed.

The bash shell runs the `.bashrc` file any time a new shell is started. You can test this by adding a simple `echo` statement to the `.bashrc` file in your home directory and starting a new shell:

```
$ cat .bashrc
#.bashrc

# Source global definitions
if [ -f /etc/bashrc ]; then
    . /etc/bashrc
fi

# User specific aliases and functions
echo "I'm in a new shell!"
$
$ bash
I'm in a new shell!
$
$ exit
exit
$
```

The `.bashrc` file is also typically run from one of the bash startup files. Because the `.bashrc` file runs both when you log into the bash shell and when you start a bash shell, if you need a script to run in both instances, place your shell scripts in this file.

Summary

The Linux system allows you to control your shell scripts by using signals. The bash shell accepts signals and passes them on to any process running under the shell process. Linux signals allow you to easily kill a runaway process or temporarily pause a long-running process.

You can use the `trap` statement in your scripts to catch signals and perform commands. This feature provides a simple way to control whether a user can interrupt your script while it's running.

By default, when you run a script in a terminal session shell, the interactive shell is suspended until the script completes. You can cause a script or command to run in background mode by adding an ampersand sign (&) after the command name. When you run a script or command in background mode, the interactive shell returns, allowing you to continue entering more commands. Any background processes run using this method are still tied to the terminal session. If you exit the terminal session, the background processes also exit.

To prevent this from happening, use the `nohup` command. This command intercepts any signals intended for the command that would stop it — for example, when you exit the terminal session. This allows scripts to continue running in background mode even if you exit the terminal session.

When you move a process to background mode, you can still control what happens to it. The `jobs` command allows you to view processes started from the shell session. After you know the job ID of a background process, you can use the `kill` command to send Linux signals to the process or use the `fg` command to bring the process back to the foreground in the shell session. You can suspend a running foreground process by using the `Ctrl+Z` key combination and place it back in background mode, using the `bg` command.

The `nice` and `renice` commands allow you to change the priority level of a process. By giving a process a lower priority, you allow the CPU to allocate less time to it. This comes in handy when running long processes that can take lots of CPU time.

In addition to controlling processes while they're running, you can also determine when a process starts on the system. Instead of running a script directly from the command line interface prompt, you can schedule the process to run at an alternative time. You can accomplish this in several different ways. The `at` command enables you to run a script once at a preset time. The `cron` program provides an interface that can run scripts at a regularly scheduled interval.

Finally, the Linux system provides script files for you to use for scheduling your scripts to run whenever a user starts a new bash shell. Similarly, the startup files, such as `.bashrc`, are located in every user's home directory to provide a location to place scripts and commands that run with a new shell.

In the next chapter, we look at how to write script functions. Script functions allow you to write code blocks once and then use them in multiple locations throughout your script.