

## System Administration Decal

Intermediate Lab #1

February 20, 2008

### Introduction

This week, we learned about the shell and various typical UNIX commands we can use within the shell. Additionally, we explored some special features of the shell, including output and input redirection, pipes, and substitution. In this lab, you will practice these skills while learning even more about the shell, in a series of themed exercises.

Please email your responses to [joshk.decal@triplehelix.org](mailto:joshk.decal@triplehelix.org).

### I. Beginner Lab

Please do exercises 7-9 and 11-12 from the beginner's lab. You will need to read up on the *chmod* and *chown* commands for #11, which we did not cover in lecture, and of course, you can do that using *man*.

### II. The *wget* command

The *wget* command is used to retrieve web pages and other files hosted on web servers on a command-line interface. Familiarize yourself with this tool by issuing *man wget* from a terminal to read its manual page. Then, answer these questions (the man page is useful for many of these):

1. If you run *wget* on a URL, you will notice that it prints a lot of output, including a text-based progress bar. Come up with **at least two ways** to inhibit its output, either using the capabilities of the shell or by passing options to *wget*.
2. Figure out how to make *wget* print the contents of a URL to standard output. Use this option to create a **pipe sequence** that searches [www.cnn.com](http://www.cnn.com) for the term 'sports'.
3. What is special about the way *wget* prints its informational status messages / progress bar? **Hint:** try using output redirection to hide one or more output channels.

### III. Exit codes

In UNIX environments, all programs report whether they have failed or succeeded when they terminate using an **exit code**. You might think of it this way: running a program is like calling a function in a piece of code, and checking to see what its return value is.

1. Find a way to display the exit code of the program that was last run in the shell using Google (or your favorite search engine.) Using the Web to find answers is instrumental even for seasoned sysadmins; it's impossible to know *everything* off the top of your head. In your answer, include the query you used and the web page from which you got the answer.
2. Try using *wget* to download a URL that does not exist. Using your answer to #1, determine what exit codes are commonly used for success and failure.

#### IV. Job Control

The shell has the ability to manage many running programs at once, and stop and continue them at your whim. In this section of the lab, you will learn how to control this power. To begin, use your knowledge from exercise II to start downloading a large file from somewhere. (A few megabytes is enough – you just need to buy yourself time.)

1. While it is downloading, hit Control and Z at the same time. (From now on we'll abbreviate this to ^Z; similarly Control + C will be ^C.) Clearly describe what happens.
2. Hitting ^Z will return you to a prompt. Describe the output of the *jobs* command, and then try using the *fg* or *bg* commands. Figure out what all three of these do and how they pertain to job control. You shouldn't need a search engine to do this. Just experiment!
3. Use this procedure to run at least three programs at once. Supply the output of the *jobs* command when this is happening.
4. Now start your download again, but this time interrupt it using ^C. Compare the behavior of ^C and ^Z.
5. Every time we have run programs up until now, they have started in the foreground (i.e. you are not returned to a prompt immediately.) Find a way to start a program in the background.