```
# Change directory to your home directory.
# Remember, ~ means your home directory
cd ~

# Check to see your current working directory
pwd

# You should be in /home/PUNetID
# Let's see what time it is
date

# Let's see the files and directories in your home directory.
ls

# Let's see all the files, including those that start with a dot
ls -a

# Let's see some more information about those files and directories
# The -l option (dash ell) displays the long entry for the file or
# directory
ls -l

# Let's see who else has a home directory on this machine
# the .. (period period) means to look in the directory above the
# current directory
# In this case, that is the directory /home.
ls -l ..

# In this view, you also see the permissions of the directories.
# your home directory is marked as
# drwx------
# which means d: directory
# rwx: the owner can read, write, and execute,
# ---: (middle three dashes) the group has no permissions
# ---: (rightmost three dashes) the world has no permissions
# These permissions allow you access to your home directory but
# no one else has access. You can notice the group is listed in the
# middle of the line, users.

# Check to make sure you are still in your home directory. Using
# ls does not move you even if you use the .. option.
```

_____

```
# Let's make a new directory inside of our home directory
mkdir cs300_test

# Check to make sure the new directory is listed
```

_____

```
# Let's move into that directory.
# Change Directory
cd cs300_test

# Confirm your new current working directory
_________________

# If you ever need an explanation of a command, use the man command
# this will explain the command and list all the options available.
man ls
(then press enter at the question, q to quit)

# Let's copy some files from zeus to the local machine for the
# rest of the exercise.
# wget will retrieve a file from a web server.
wget zeus.cs.pacificu.edu/chadd/LinuxTest.tar.gz

# The fetch command is similar to wget. The curl command is
# also similar. Not all systems have all three commands.

# Copy this LinuxTest.tar.gz file to your home directory on zeus.
# we will use Secure Copy (scp) to do this.
# Note the : at the end!
scp LinuxTest.tar.gz punetid@zeus.cs.pacificu.edu:

# Connect to zeus to make sure the file transferred up correctly
ssh punetid@zeus.cs.pacificu.edu

# Make sure your prompt says zeus, otherwise you are still on your
# local machine.

# Confirm the file exists on Zeus
_______________________________

# Rename (move the file on Zeus)

mv LinuxTest.tar.gz LinuxTest_PUNETID.tar.gz

# use the Submit script [*see the note at the end about the submit script]
# to submit the file you just created for grading.

submit cs300f18 LinuxTest_PUNETID.tar.gz

# Did you get a success message and a receipt?

ls -al LinuxTest_PUNETID.tar.gz.cs300f18.receipt

# The submit script is how you will submit most of your assignments this
# semester.  The receipt plus your submitted file is is your proof you
# submitted your assignment.  Don't alter either file.
# validate your receipt

checkReceipt LinuxTest_PUNETID.tar.gz cs300f18 LinuxTest_PUNETID.tar.gz.cs300f18.receipt
# submit and checkReceipt only exist on Zeus.
```

```
# Use the exit command to disconnect from zeus.
exit

# Make sure your prompt says the name of your local machine and
# not zeus!
# tar.gz files are compressed archives.
# Uncompress and extract the files:
# tar: the command
# x: extract
# v: verbose list to the screen the files extracted
# z: compressed (the .gz file extension requires this)
# f: file (the filename must follow directly after the f option)
tar xvzf LinuxTest.tar.gz

# Change directory to the newly created LinuxTest directory
_________________

# List all the files in that directory
___________________

# make the welcome.sh file executable
chmod u+x welcome.sh

# execute welcome.sh  The ./ says to look in your current directory
# for welcome.sh rather than search $PATH
./welcome.sh

# what directories are in your PATH?  These are directories Linux will
# search when you try to run an executable command
echo $PATH

# A number of text files ( .txt) should be listed.
# Let's look at one of them
# cat displays the contents of a text file to the screen
cat CS150.txt

# If you just want to see the first two lines in a file:
head -n 2 CS150.txt

# If you just want to see the last two lines in a file:
tail -n 2 CS150.txt

# A good way to edit the file is to use nano.
# The menu is at the bottom of the screen. ^ means the control key.
# Control-X exits nano
nano CS150.txt

# pico is similar to nano.
# vi is more powerful than pico or nano, but takes a while to learn.
# google vi cheat sheet if you are interested in vi.

# There are many words in that file. Let's count how many
# lines, words, and characters are in that file. The wc command
```

```
# will do that.
wc CS150.txt

# How many lines are in the file?
__________________

# How many characters are in the file?
_________________

# You can also print just the number of lines, words, or characters
wc -l
wc -w
wc -m

# Let's check how many files contain the word data.
grep data *

# Grep explained:
# grep : the command
# data : the pattern to look for
# * : the files to check (* is a wildcard meaning every file)

# Let's find all the text files (.txt) that exist in the
# LinuxTest directory and its subdirectories
find . -name '*.txt'

# Find explained:
# find : the command
# . : Start looking in the current directory.
# -name '*txt' : look for file names that match the pattern *.txt
# the single quotes are important!

# Let's save that list of text files into a new file:
# the > will take the output of the command to the left, and write
# that output to the file on the right.
find . -name '*.txt' > listOfFiles.txt

# Let's make sure that new file got created
_________________________________

# Show the contents of the new file listOfFiles.txt
_________________________________

# Let's find all the text files again
_________________________________

# Notice that listOfFiles.txt is now listed.
# Let's save the list of text files to a new file, named
# newListOfFiles.txt
_________________________________

# Let's see how the two list of text files differ.
# The diff command compares two text files and displays the
```

```
# difference

diff listOfFiles.txt newListOfFiles.txt

# The output contains an > which means that line exists in the
# file listed on the left, newListOfFiles.txt, and not in the file
# listed on the right, listOfFiles.txt.
# Let's count how many files we found:
# The pipe, or vertical bar, takes the output from the command on the
# left and uses that output as input to the command on the right.
find . -name '*.txt' | wc -l

# Let's list information about each of the text files:
# the xargs command will take each of the lines from the
# input (the command on the left) and pass that line
# to the command on the right (in this case, ls l)

# Let's find all the text files that contain the string CS

find . -name '*.txt' | xargs grep CS

# Let's move into the documents directory to work with those files
_____________________

# Display the contents of the file preamble.txt
_____________________

# Display the contents of the file declaration.txt
_____________________

# This file declaration.txt is really long and scrolls by too fast.
# Let's slowly scroll through the file.
# the less command shows you a screen full of text at a time
# press space for the next screen. press enter for the next line.
# press q to quit

cat declaration.txt | less

# The more command is similar
cat declaration.txt | more

# You don't need to use cat
more declaration.txt
less declaration.txt

# Compiler Options
# Don't forget the gcc compiler options:
# -Wall    : show all warnings
# -g       : include debug symbols
# -c       : compile only (produce object file)
# -o  file  : name the output filename
# -S       : produce assembly
# -fverbose-asm : produce verbose assembly
```

```
gcc -o main_asm.S -S main.c
gcc -o main_asm.S -fverbose-asm -S main.c
gcc -o main.o -c -g -Wall main.c
gcc -o main -g main.o
# Sort

\ls CS*.txt | sort -r

# the capital R option to sort is quite different. Run the following
# command twice.

\ls CS*.txt | sort -R

What does -R do?
_____________________________________

# Open up another terminal
# Run top to show all the processes currently running
top

# restrict top to show only your processes by typing u in the window
# displaying top and typing your punetid

# Press q in the window displaying top to quit top. Do not close the
# terminal.

# Type ps to see all the processes running in your current window
ps

# The sleep command will sleep for some number of seconds
sleep 3

# Putting an & at the end of a command will run the command in the
# background and allow you to type more commands.

sleep 120 &

# Check all the running processes again, you should see sleep running
# (or sleeping)
__________

# Open top again in the other terminal
# restrict top to show only your processes
_________________________

# Go back to the original terminal and launch sleep again with the &
sleep 600 &

# Look at top. The PID or process id is listed at the left. Sleep
# should be listed.
# Find the process id.
# Go back to the other terminal and kill the sleep process
kill PID
```

```
# Some processes don’t die when you use kill alone. Sometimes you
# need to do the following. 9 stands for SIGKILL:
kill -9 PID

# Check out the following man page.
# You are looking at the definition of kill from section 1p.
# SIGKILL is listed around line 70.
man -s 1p kill

# A few more commands.
# Let’s check how much space is left on the hard drive
# You should be able to see how much space is available on /home
df -h

# Another useful command is screen.
# Let’s first check to make sure screen is available in the lab.
# The which command will tell you where the command is on the
# file system, if it exists.
which screen

which bob
# should fail and tell you in which directories which looked for bob

# Screen exists, good.
# screen is especially useful if you
# need to run a long running command on a remote machine, such as
# zeus. You would log into zeus and then start screen on zeus.
screen -S test
# This creates a shell session that won’t die if you get
# disconnected. Also, you can reconnect to this shell session from
# another location to check
# the progress of your work or make changes.
# In a new terminal, connect to the screen you just created:
screen -x test
ls
ls al

# you should see the same output being shown in both terminals.
# everything is happening exactly once, the output is just
# copied to both places.
# Let’s disconnect from screen in one terminal.
Control-A D

# In the other terminal, let’s write a little loop
# The bash shell is powerful, you can program right in the shell!
# The following is a while loop that will sleep for 3 seconds,
# display information about the hard drive, and then display the
# date.  Remember, the semi-colon is just a command separator
# The loop is equivalent to while(true) in C++ so it will
# never terminate. Later, we will use Control-C
to kill the loop.
while [ true ] ; do sleep 3; df -h ; date ; done

# Let’s disconnect from screen in this terminal.
```

Control-A D

```
# Take a deep breath.
# Let's reconnect:
screen -x test

# Your previous output should show up and you should see that the
# loop is still running.

# Press control-C to kill the loop

# Typing exit will terminate the screen.
```

QUESTIONS TO ANSWER

Perform the following tasks at the Linux command line. Record the commands you ran in a text file: CS300_Linux_PUNETID.txt and use the submit script to turn in your assignment by **Aug 31, 2018, 6 pm.** Number each answer and make sure your name is at the top of the file. **You will submit two separate files for this assignment.**

1. Move to your home directory

2. Attempt to move to Chadd's home directory (username: chadd).

3. Show all of the files and directories in your home directory, except those that start with a dot.

4. Make a directory called CS300_Linux_Practice

5. Move into CS300_Linux_Practice

6. Use nano to create a file named WhoIAm.txt On the first line type CS300. On the second line type your name and PUNetID. Save the file and exit nano.

7. Determine how many characters are in WhoIAm.txt and save that information to the file CharCount.txt. Write one chain of commands to solve this. Hint: > | <

8. Copy the file WhoAmI_CS300.txt from Chadd's web directory on zeus.cs.pacificu.edu to the current directory (CS300_Linux_Practice)

9. Determine how the file WhoAmI_CS300.txt differs from WhoAmI.txt

10. Move up one directory.

11. Create the file CS300_Linux_Practice_PUNETID.tar.gz by tarring up the CS300_Linux_Practice file.

12. Copy the file CS300_Linux_Practice_PUNETID.tar.gz to your home directory on zeus.

13. Log on to zeus.

14. Use the submit script to submit CS300_Linux_Practice_PUNETID.tar.gz

15. Check to make sure you received a valid receipt.