

Assignment 0

Introduction to Linux

Due Date: Friday January 30, 2026 @ 11:59 pm
Version: 0.10
Note: Please report typographical errors, or unclear sections to the instructor

Objective

The purpose of this assignment is to create a local Linux Virtual Machine for use in future TCSS 422 projects and to gain experience using Ubuntu.

Here is an article describing Microsoft's adoption of Linux on the Azure cloud:

<https://www.zdnet.com/article/microsoft-developer-reveals-linux-is-now-more-used-on-azure-than-windows-server/>

For those new to Linux, before completing this assignment, please review the BASH shell and command line online tutorial found here:

Learn the BASH command line: <https://ryanstutorials.net/linuxtutorial/>

Task 1 – Install Oracle Virtual Box

Before creating a Linux virtual machine, we must install a **hypervisor**. Hypervisors, also known as a virtual machine monitors (VMMs), are a specialized layer of software, firmware, or hardware that enables creation of virtual machines on a physical host computer. Hypervisors contextualize each virtual machine to create a virtual computer by creating virtual hardware interfaces to mimic real components allowing VMs to interact and share the physical host computer's resources as if they were their own. The hypervisor contextualizes and virtualizes the VM's CPU, memory, virtual disks, and network interface card by sharing the underlying physical host hardware.

The default hypervisor for TCSS 422 is Oracle VirtualBox 7.x which can be downloaded from:

<https://www.virtualbox.org/wiki/Downloads>

Choose the appropriate platform package when downloading VirtualBox from the list for your host operating system:

Version 7.x of VirtualBox is now available for:

- Windows 10/11 hosts
- macOS / Intel and Apple Silicon hosts
- Linux distribution for Linux host
- Solaris hosts

UTM and Parallels are hypervisor alternatives if using macOS on Apple Silicon.

<https://mac.getutm.app/>

Parallels is not free, but there is a reduced cost student license:
<https://www.parallels.com/products/desktop/buy/?students>

An Apple Silicon port of Virtual Box 7.x for M1/M2/M3 macs is now available.
If anyone tests this, it would be appreciated to share your user experience with Virtual Box on Apple Silicon. For example, how is the performance/speed compared to UTM ?

Once you've downloaded VirtualBox, follow the instructions to complete the installation.

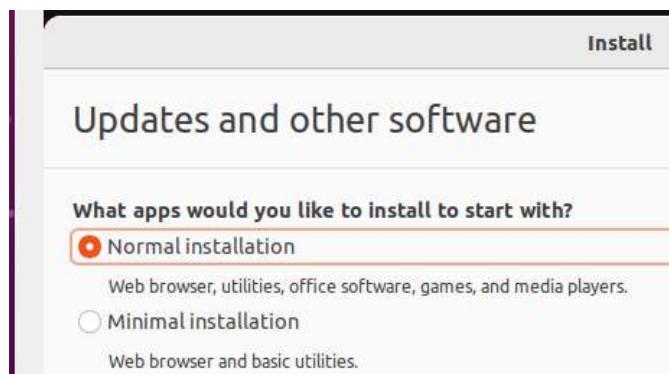
If you do not have access to a computer for Virtual Box and Ubuntu, the School of Engineering and Technology can create an Ubuntu 24.04 VirtualBox VM to support work in TCSS 422. If you haven't already requested a VM using the VM-survey, please contact the instructor by email.

Task 2 – Create Ubuntu 24.04 LTS Virtual Box VM

Once VirtualBox has been installed (or an alternative hypervisor), Ubuntu 24.04 VM can be installed on an Intel or ARM Mac/PC.

Ubuntu can be downloaded from: <https://www.ubuntu.com/downloads/desktop>

Please download and install Ubuntu the latest LTS version of Ubuntu 24.04. (Currently 24.04.3)
During the installation, a “normal installation” is recommended:



This option should include the Linux kernel headers required for kernel mode programming. Ubuntu 24.04 is a Debian based open source/free distribution of Linux, and is a leading commercial distribution of Linux. Some of the most widely used open-source Linux distributions include Ubuntu, Fedora, openSUSE, and Debian.

For Apple Silicon users, be sure to download the 64-bit ARM (ARMv8/AArch64) desktop image for M1 Mac (Apple Silicon) from:

<https://cdimages.ubuntu.com/ubuntu/daily-live/current/>

The recommended configuration for your virtual machine's disk space is a minimum of 25 GB. 40+ GB is recommended if planning to use the Ubuntu VM after TCSS 422. Previously, students with 10GB or less have run out of disk space before the end of the quarter leading to headaches where the VM had to be reinstalled with a larger disk.

Desktop / Laptop "Host" Physical Memory	Recommended VM memory	Close all applications while running VM
Less than 16 GB	4 GB	YES
16 GB to 24 GB	4 GB to 6 GB	YES
24 GB to 31 GB	5 GB to 7 GB	Recommended
32 GB	6 GB to 8 GB	Recommended
33 GB to 48 GB	8 GB+	Recommended
48 GB to 56 GB	10 to 12 GB+	Recommended
More than 56 GB	12 GB +	Optional

(OLD) Installing Ubuntu 22.04 on Windows 10 Oracle VirtualBox: <https://youtu.be/zHwFtxJsog>

Two videos on how to install Ubuntu 24.04 on Windows 11 Oracle VirtualBox:

<https://youtu.be/DhVjgI57Ino>

<https://youtu.be/8uYjoBh21IU>

Video on how to install Ubuntu 24.04 using UTM on M1 Mac:

<https://youtu.be/JrNS3brSnmA>

Video on how to install Ubuntu 24.04 using VMware Fusion on M1 Mac:

<https://youtu.be/kDosGTdwqO0>

Another option is Parallels on M1 Mac, but it is not free, but there is a student edition:

<https://www.parallels.com/landingpage/pd/education/>

(OLD) Video on how to install Ubuntu 22.04 on Mac with Parallels*: <https://youtu.be/1vht7h3EQtc>

* - note for Mac users, UTM or Parallels is currently recommended for virtual machines

Documentation on Installing Virtual Box (free) on Intel Macs (<=2020):

<https://medium.com/tech-lounge/how-to-install-ubuntu-on-mac-using-virtualbox-3a26515aa869>

<https://help.cyberstart.com/help/installing-and-setting-up-virtualbox-intel-mac-users>

Task 3 – Install Virtual Box Linux Guest Additions (Virtual Box Only)

After installing Ubuntu, install the [Linux Guest Additions](#):

"Guest Additions" enable sharing of the Host OS Hard Disk (e.g. Windows, Mac) with the Virtual Machine. Guest Additions also enable sharing of the clipboard, and provide mouse pointer integration. Please spend the time installing the Guest Additions. Things includes a two-way clipboard where you can copy-and-paste data between the host and guest VM seamlessly. You'll be happy that you did!

How-To Video – Installing the Linux Guest Additions:

<https://youtu.be/zdkl16oAS1k> (old video)

Instructions including how to install Virtual Box Guest Additions:

When working with Virtual Machines, the base operating system (e.g. Windows 10/11) on your laptop that hosts the virtual machine is called the **host operating system**. The operating system used by the VM is called the **guest operating system**.

Many people will have the following configuration in TCSS 422:
Host operating system = Windows 10/11 or Mac OSX
Guest operating system (VM) = Ubuntu 24.04 LTS

Download the Virtual Box Guest Additions ISO file.

This will be called “VBoxGuestAdditions_(version-number).iso”, where (version-number) is the version of Virtual Box installed on your host computer. (i.e. your laptop, desktop, etc.)

For example, if using Virtual Box 7.2.4 on your host computer, the file to download is:
https://download.virtualbox.org/virtualbox/7.2.4/VBoxGuestAdditions_7.2.4.iso

If using another version of Virtual Box, search for the version, and download the file from here:
<https://download.virtualbox.org/virtualbox/>

Once you have downloaded the ISO file, follow the instructions for installing the guest additions onto your Ubuntu 24.04 Virtual Box VM as in the articles:

<https://linuxconfig.org/installing-virtualbox-guest-additions-on-ubuntu-24-04>

<https://www.linuxtechi.com/install-virtualbox-guest-additions-on-ubuntu-24-04/>

Please do yourself a favor, and do not go the entire quarter without installing the guest additions.

After installing the Guest Additions, be sure to enable the **shared clipboard** on any VMs you create on Virtual Box. Select the VM under Virtual Box, and select “Settings”. Under “General”, go to the “Advanced” tab, and enable the “Bidirectional” clipboard. This will allow you to copy-and-paste between your host (laptop) and guest (VM). Also, on the left-hand-side, select “Shared Folders”. Add a new shared folder by clicking on the folder icon with a green-plus sign on the right. One or more folders from your host (laptop) can be shared with the guest (VM). Shared folders often will then appear under “/media” on the Ubuntu filesystem. These file system mounts can be viewed using the “`df -h`” Linux command. Note: you typically need to be “superuser” to access these file system mounts. To become superuser type “`sudo bash`” and then “`cd /media/{dir-name}`”. Take note that modifications as super user will use the “root” user and not your typical user account. Review Linux file permissions including the commands `chmod`, `chown`, and `chgrp` to learn how to change file and group ownership as well as file permissions. “`mv`” and “`cp`” can move and copy files from `/media` to other locations.

For Ubuntu VMs (guests), the shared clipboard requires Guest Additions to be installed.

Task 4 – Become familiar with Linux supported by generative AI

For this task, the objective is to discover **TWO** different Linux commands that provide answers to each of the questions. In Linux, there are notoriously **many** ways to query the same data or perform the same operations. Using your favorite Generative AI tool, such as Chat GPT, prompt the AI to provide two commands. Alternatively, traditional web search can also be used.

For this assignment, we will specifically prompt the AI **NOT** to use the “cut” or “awk” commands to parse the command output to provide only a precise answer.

In your prompts/searches, include the text:
“do not use awk or cut to answer the question(s)”

Use of “grep” is permitted – it is not necessary to provide any prompt information regarding grep.

Ideally the command provides some raw output, so it is possible to see the answer to the question with some context. Awk and cut parse the output so only the exact answer is provided. The problem with this is context is lost, so it is not possible to see the overall information the command produces – and so you learn less about how the command(s) work.

After finding two Linux commands for each question, your next task is to time them using the Linux ‘**time**’ command. The Linux ‘time’ command reports the real, user, and kernel time required to execute the command. ‘Real’ time is the elapsed time (i.e. wall-clock or watch time). ‘User’ time is the time the CPU spends executing instructions in user mode (unprivileged mode). ‘Kernel’ time is the time the CPU spends executing instructions in kernel mode (privileged mode). Your goal is to run both commands on your Ubuntu VM, show the full output, and identify which command is faster (has the lower ‘real’ time). If the commands appear to have similar time, then you will report “the runtime of the commands is too close to choose a winner”. It may be necessary to run each command 10-20x or more to see if one is clearly faster.

GENERATIVE AI PITFALL: It is possible that Generative AI will provide incorrect commands. You are responsible for ensuring the commands execute correctly and there are no incorrect commands which fail to correctly answer the question. One approach is to generate more than 2 commands. Try generating 3, 4, or more and compare their output. Using consensus, it should be possible to determine the correct answer. Note that some answers are not static, but change constantly. If you encounter errors, you can work with the LLM to correct them, or simply find a working command on your own to substitute in the Generative AI’s output.

Summary of prompt requirements:

- Do not use awk or cut to parse Linux command output. Use of grep is permitted. Include in your prompt(s): “do not use awk or cut to answer the question(s)”
- Provide the 2 commands in a bash script called a0.sh.
- For each question, there should be a header line that describes the question.
- The two commands should be labeled as ‘Command #1’ and ‘Command #2’ in the output. Use the echo command to print this text. (see example output below)
- The questions should be numbered in the output, and there should be blank lines between the questions.

Questions

Question 1 – How many total **processes** are present shortly after a Ubuntu Linux computer boots up?

Required: use a command that simply outputs the answer and exits. Use a command that does not have a graphical or text-based interface to answer the question.

Question 2 – How many total **threads** are present shortly after the Ubuntu Linux computer boots up?

Required: use a command that simply outputs the answer and exits. Use a command that does not have a graphical or text-based interface to answer the question.

Question 3 – What is the version number of the Linux kernel installed on a Ubuntu Linux computer?

Question 4 – What is the model name of the CPU(s) on a Ubuntu Linux computer?

Question 5 – What is the total size of the memory swap space (virtual memory) in MB on a Ubuntu Linux computer?

Question 6 – What is size of the root disk partition in MB?

In Linux, the root partition is always mounted at “/”. A mount point is the directory or location in the file system where an I/O device has been mounted. The mount point is used to access the device through a file system.

Question 7 – What is the total number of inodes on the root filesystem?

If unfamiliar with what an inode is, look up the definition and how to display the number of free/used inodes on Linux/Ubuntu.

Question 8* – What is the HTTP status code that is returned from www.google.com?

The status code indicates whether the address is reachable via the network using the application-layer. In this case HTTP is our application layer protocol used to test network connectivity to google.

Question 9* - What is the interface name of the network interface device used to route the ICMP ping packets to www.google.com?

Question 10 – Identify the file system type of the “/” root partition?

* - These questions require networking to be configured on the Ubuntu VM.

BONUS QUESTIONS

These questions should be answered after running your a0.sh script. These questions will replace any missed questions from above. After running the Linux commands above, answer the follow-up questions.

Question 11 – As an extension to Question 4, look up and report details regarding your CPU using the wikipedia or wikichip websites, or use Intel/AMD product specification webpages. (1) Identify how many CPU cores the CPU has. (2) Identify the release date of the CPU (If the information is available), and (3) identify the original retail price.

Question 12 – As an extension to Question 10, briefly describe the file system type (1-2 sentences). Using the Linux manual pages, look up the file system type. The manual pages identify the Linux kernel version when specific features were added to the file system. From the manual page for the file system, identify the name of two features added to the filesystem since July 2017. You need to determine the Linux kernel version in July 2017, and identify two features released after this date. Using an Internet search engine, search for the two file system features you identified and write a short description regarding what each is.

Note: If you are using Apple Silicon as the host computer, the Ubuntu VM won't identify the model name of the CPU. You can try and infer your Apple CPU model from “About this Mac” on MacOS.

See: https://en.wikipedia.org/wiki/Apple_M1 or https://en.wikipedia.org/wiki/Apple_M2 (etc)

What to Submit

Produce **FOUR** files for submission to Canvas:

File #1: BASH SCRIPT (a0.sh file)

For the assignment, submit a BASH script that captures the list of commands used to answer each of the questions. Number each question and command clearly.

How to create a simple bash script:

Create a “bash” script which runs each of the two commands used to answer the questions. Be sure to number the questions in the output, and use the echo command to output the first and second command to the console before running it. Follow the format of the example below.

Use a text editor such as “gedit”, “vi”, “pico”, or “nano” to create and make changes to the script, “a0.sh”:

Example Script:

Question #1: What is the command to show the user’s current working directory?

```
# Question 1 - What is a command to show the user's current working
directory?

echo "1. Command to show the current working directory"
echo "Command 1: pwd"
time pwd
echo

echo "Command 2: echo \$PWD"
time echo $PWD
echo...
```

To run the script assign the script to have execute permission

Give the “user” (u) “execute” (x) permission with chmod (u+x):

```
$ chmod u+x a0.sh
```

Note the script example above runs the commands and also times how long they take.

File #2: BASH SCRIPT OUTPUT (a0.out file)

Next, capture the output of the script file to a text file, and submit this text file to Canvas.

Using I/O redirection, capture the bash script output to a text file:

```
$ ./a0.sh > a0.out 2>&1
```

“2>&1” redirects error messages written to the ‘standard error stream’ to the ‘standard output stream’ so that all of the information is written to the same file.

Check the output for proper formatting by printing to the screen:

```
$ cat a0.out
```

1. Command to show the current working directory

```

Command 1: pwd
/home/fred

real 0m0.000s
user 0m0.000s
sys 0m0.000s

Command 2: echo $PWD
/home/fred

real 0m0.000s
user 0m0.000s
sys 0m0.000s

```

Note that the output for ‘time’ goes to stderr not stdout. Therefore it is necessary to redirect the 2nd file stream to the 1st file stream using the argument “2>&1”.

File #3: ANSWERS FILE (a0_answers.txt file)

Create a text file called “a0_answers.txt”.

In this file, you are to identify the large language models (LLMs) used for the assignment (if any). If the same LLM is used for every question, then just identify the full LLM model name. If different LLMs are used, then identify the full LLM model name used for each question.

Next answer each Linux question by interpreting the best results from the output of the Linux commands. Then you are to report if command 1 or command 2 is faster -- *or if there is a tie* --.

```

$ gedit a0_answers.txt

[Add text to answer each question]

LLM: for all questions, ChatGPT, based on GPT-5.2

1. /home/fred
faster command: tie (both commands have the same runtime)

2. (q2 answer)
faster command: Command 1

3. (q3 answer)
faster command: Command 2
...

[Save the file, Exit GEDIT.]

```

File #4: LLM PROMPT FILE (a0_prompt.txt file)

If you used one or more LLMs in this assignment, then please submit your LLM prompts in a file called a0_prompt.txt. If you did not use LLMs (Generative AI) then submit a file with a brief description of the tools and techniques used to support answering the questions.

To submit the assignment, upload the FOUR files to Canvas:

1- **BASH SCRIPT:** (e.g. a0.sh)

Contains the bash script with the commands to answer each of the questions

2- **BASH SCRIPT OUTPUT** (e.g. a0.out)

Output captured by running a0.sh script, and redirecting the output using “>” to the a0.out output file.

3- **ANSWERS FILE** (e.g. a0_answers.txt) File provides your interpretation of the answers to each of the questions and also your interpretation of which command is faster. Your answers must be based on output from your bash script output. (a0.out).

4- **PROMPT FILE**(e.g. a0_prompt.txt) file captures LLM prompts used to complete the assignment. If LLMs are not used, this file describes any tools and techniques used to support completion of the assignment.

Grading

This assignment will be scored out of 40 points. (40/40)=100%

Each question is worth 4 points: one point is for including command output and timing values for each question that provides the answer in a0.out, one point is for answering the question correctly in a0_answers.txt, and one point is for including two valid commands for each question in a0.sh. One additional point is for providing the LLM prompt for the question, or if LLMs are not used, for describing the tools and techniques.

Two questions (11 & 12) are counted as bonus questions. They are not extra credit questions, but they will replace any missed question in 1 through 10. Answering bonus questions enables a 100% score on the assignment while skipping, or incorrectly answering up to 2 questions.

Please create, format, and submit four text files to Canvas with the specified names to submit the assignment:

File #1: a0.sh

File #2: a0.out

File #3: a0_answers.txt

File #4: a0_prompt.txt

Change History

Version	Date	Change
0.1	01/13/2026	Original Version