Beginner's
Introduction to Computing at CARC
with JupyterHub

Matthew Fricke

Example ESACPE 2023

Version 0.1

Goals

- 1) Basic Linux Literacy
- 2) Using Compute Nodes with Slurm
- 3) Jupyterhub

I hope to spend about an hour on each with a 15 minute break.

Outline

- High Performance Computing Overview
- Logging in
- The BASH Shell
- The Slurm Job Scheduler
- Transferring data to and from CARC
- Accessing software and the module system
- Conda environments
- Jupyterhub

High Performance Computing

- What is high performance computing?
  - Really just means something that is a lot more powerful than your desktop or laptop.
  - Hardware:
    - That might mean more and faster processors to do the calculations more quickly (eg 400 CPUs instead of 4)
    - More RAM so you can work on bigger problems (3,000 GB instead of 8)
    - Bigger file systems so you can process larger datasets
    - More and bigger GPUs to accelerate your computations (12 GPUs at a time instead of 1)
  - People:
    - Someone else to manage the systems and keep them running and secure
    - Someone to answer your questions and help with problems

Logging into Wheeler

First login to the Linux **workstation** in front of you.

Use your CARC username and password.

Keven, Tannor, Viacheslav (Slava), and Jose can help you login if you have trouble.

This is an “important step” so don’t let me move on until you have logged in

Logging into Wheeler

First login to the **workstation** in front of you.

You will always login to CARC cluster remotely.

These clusters don’t even have monitors.

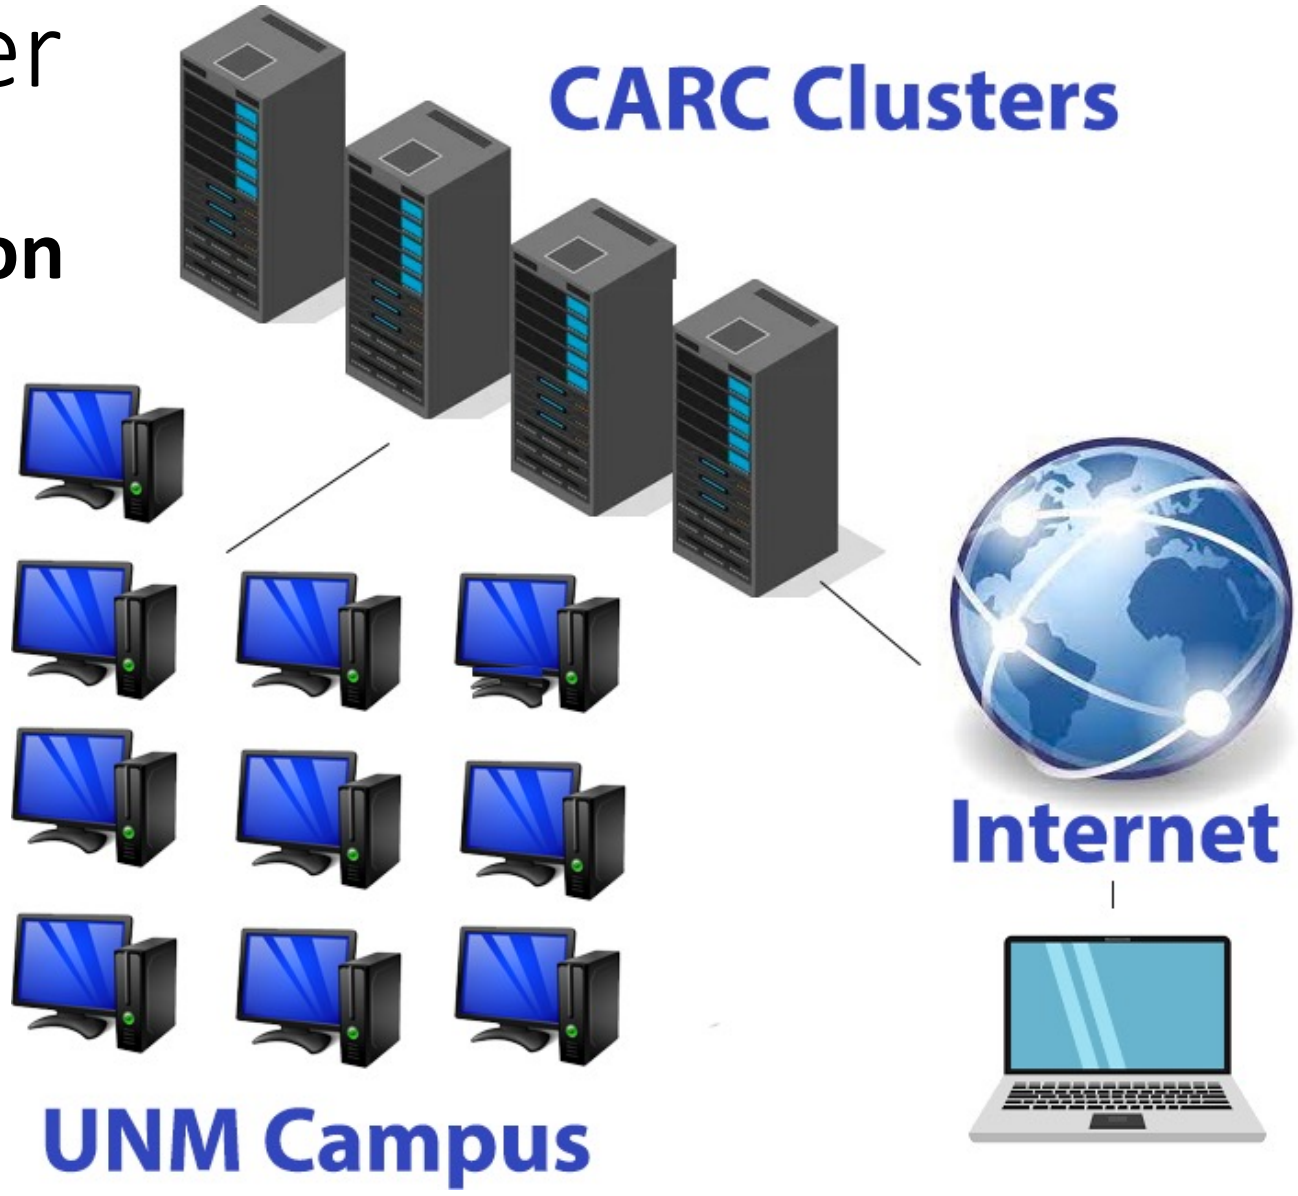

Logging into Wheeler

We are going to use a program called secure shell.

Secure shell (ssh) is now built into every major operating system (Windows, OSX, and Linux).

You don’t need third party programs like putty anymore.

Logging into Wheeler


```
ssh vanilla@wheeler.alliance.unm.edu
```

Should prompt you for a password…

Don’t let me move on until you are able to login.

Logging into Wheeler

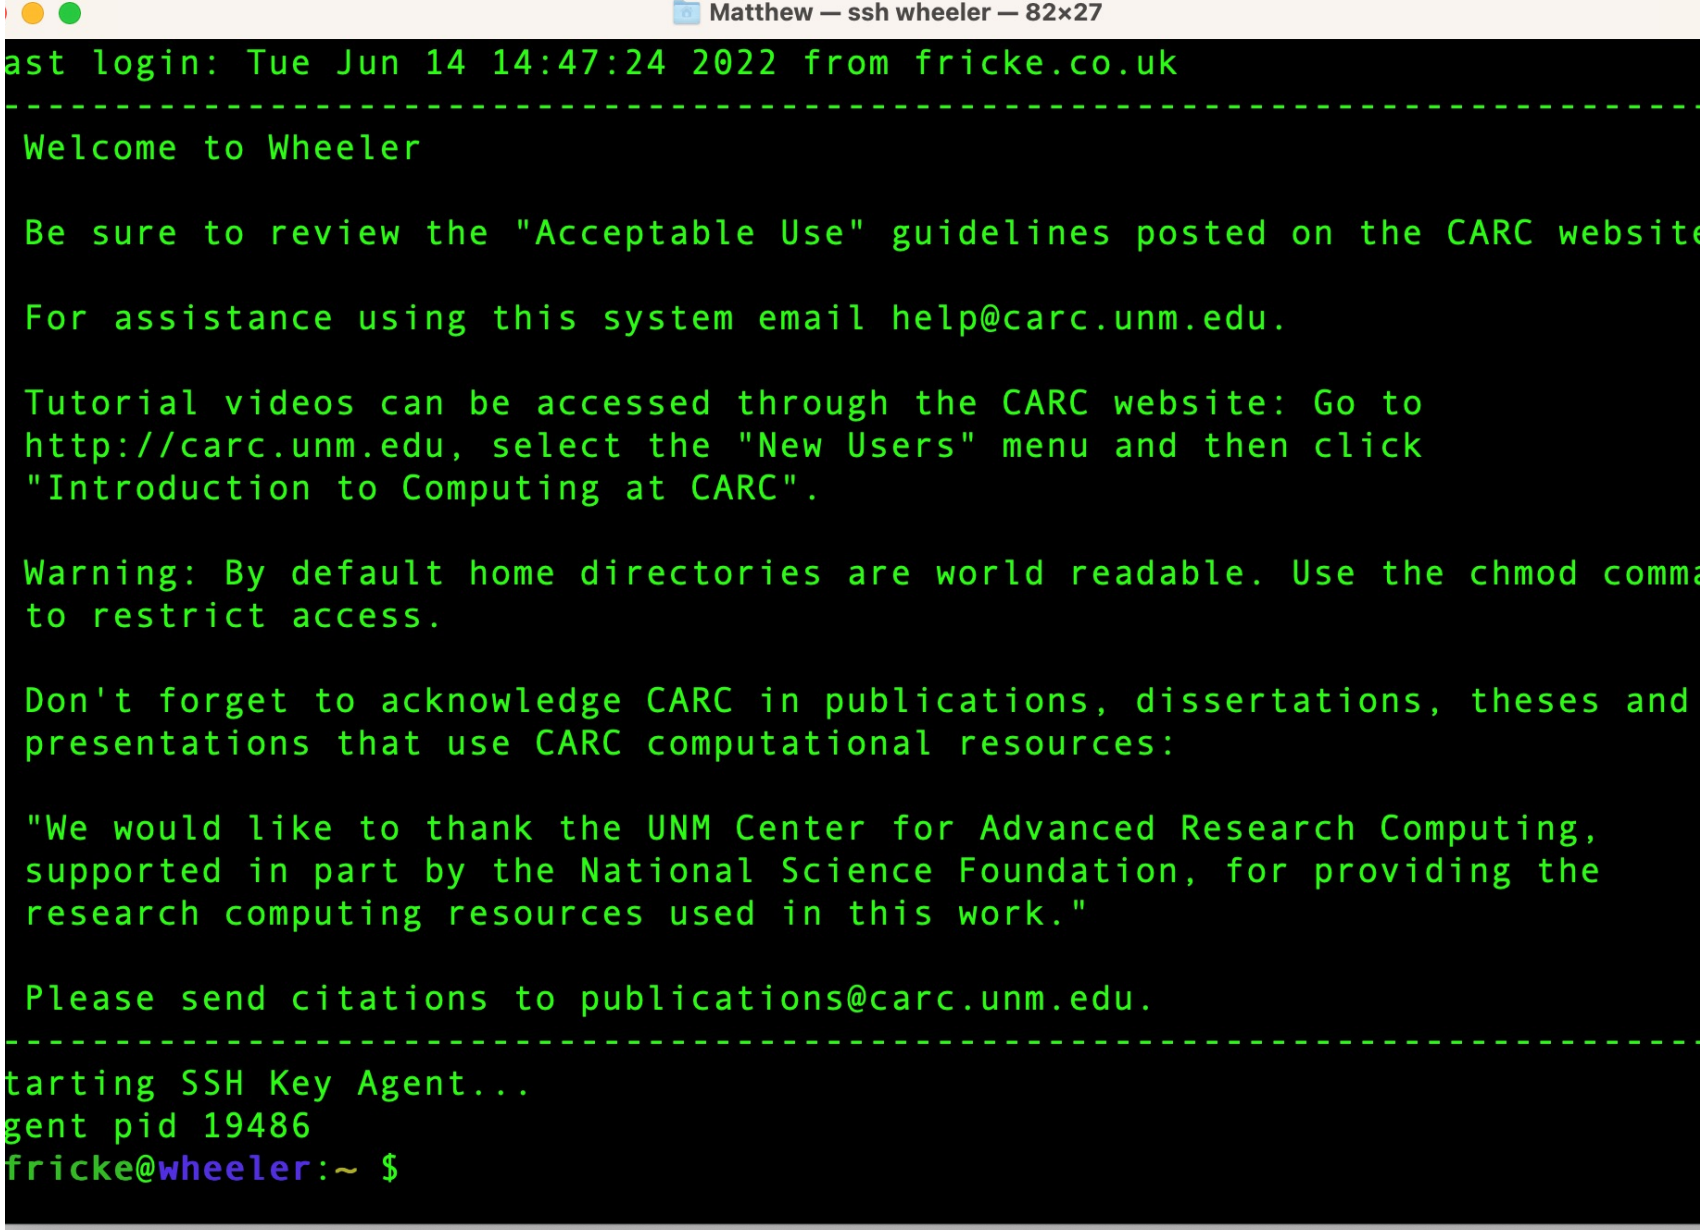

Please enter the following command

cp –r /projects/shared/workshops/beginner/mystuff ~/

We will come help you if you have any trouble.

(Later I will go over what this command does)

Linux and the BASH Shell

- The Kernel manages access to the hardware in a computer.
- An Operating System (OS) is the Kernel plus useful programs provided by the OS.
- The “shell” is the outermost layer of the OS.
- It is where the user interacts with the OS.

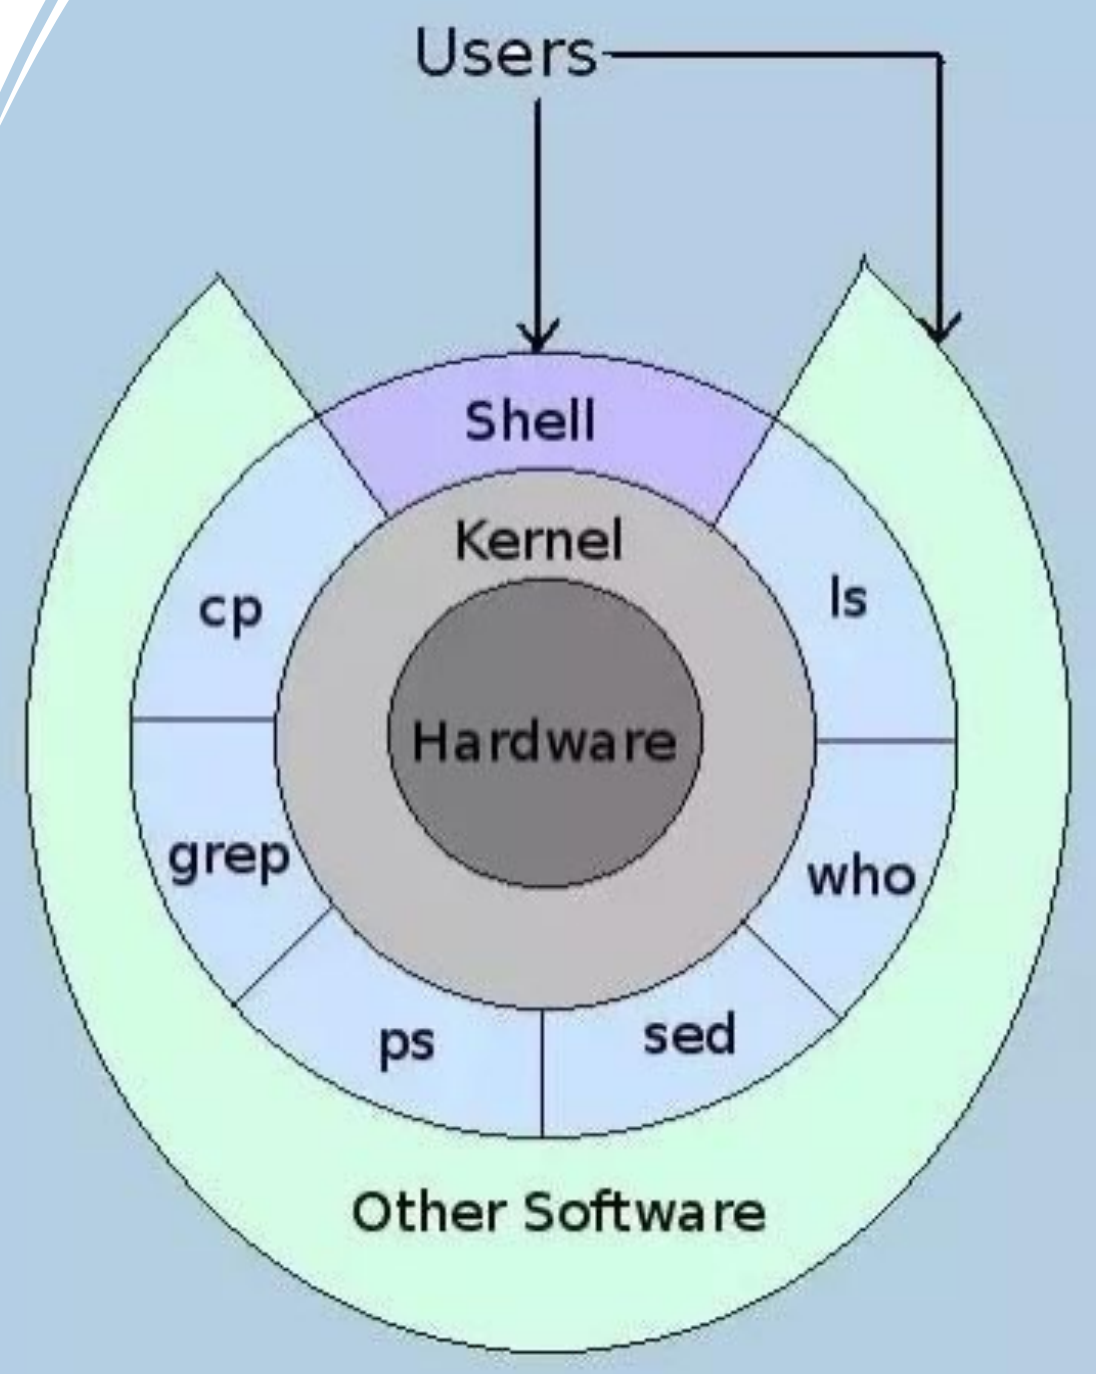

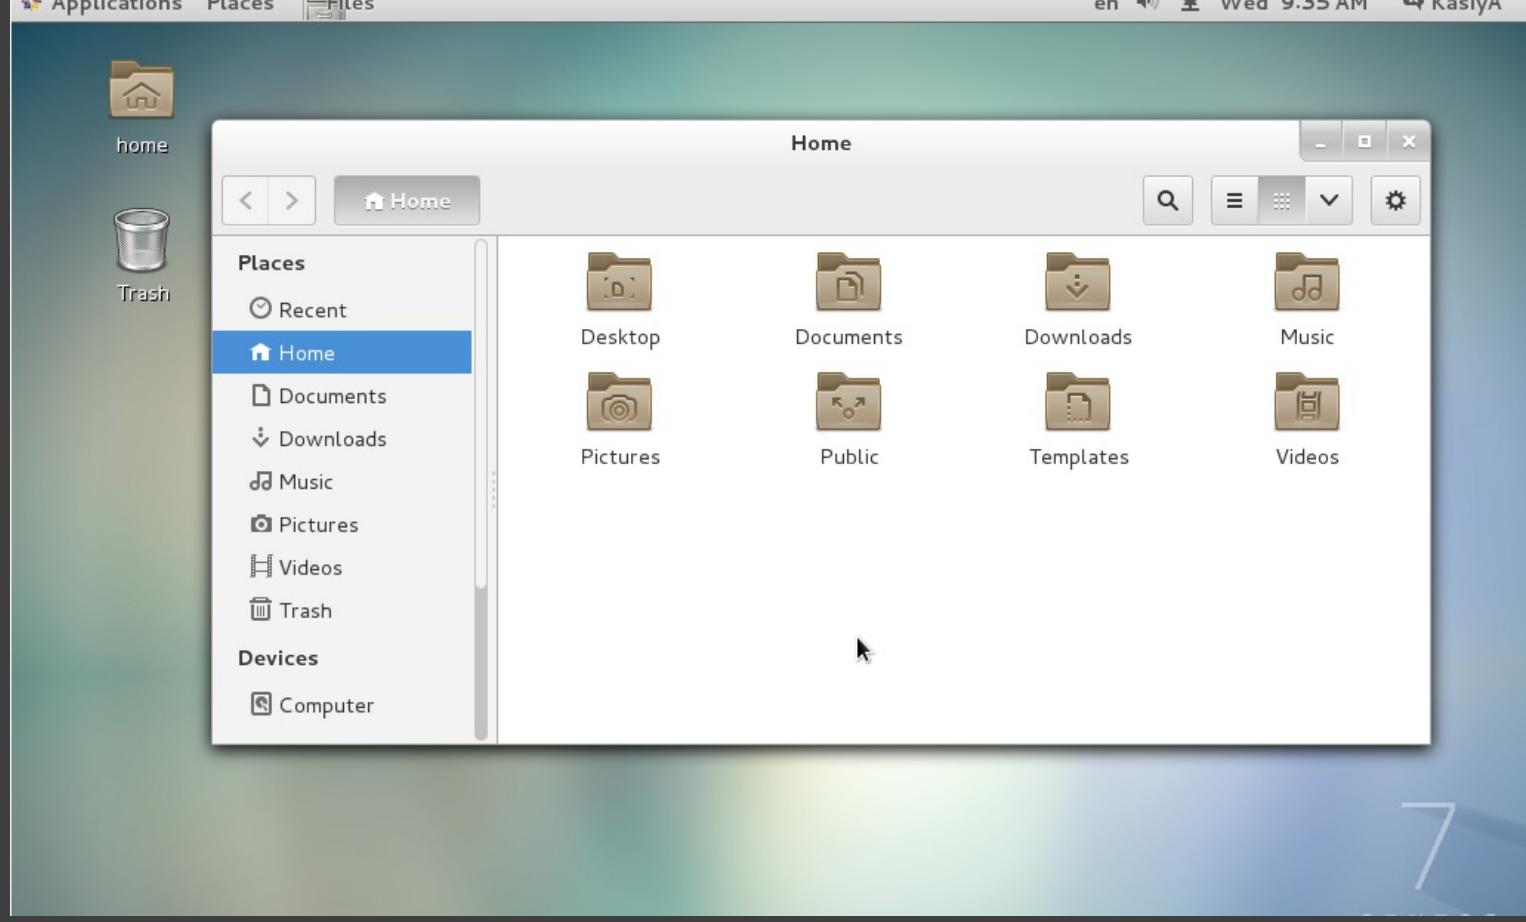

Graphical Shells (GUIs)

Logging into Wheeler

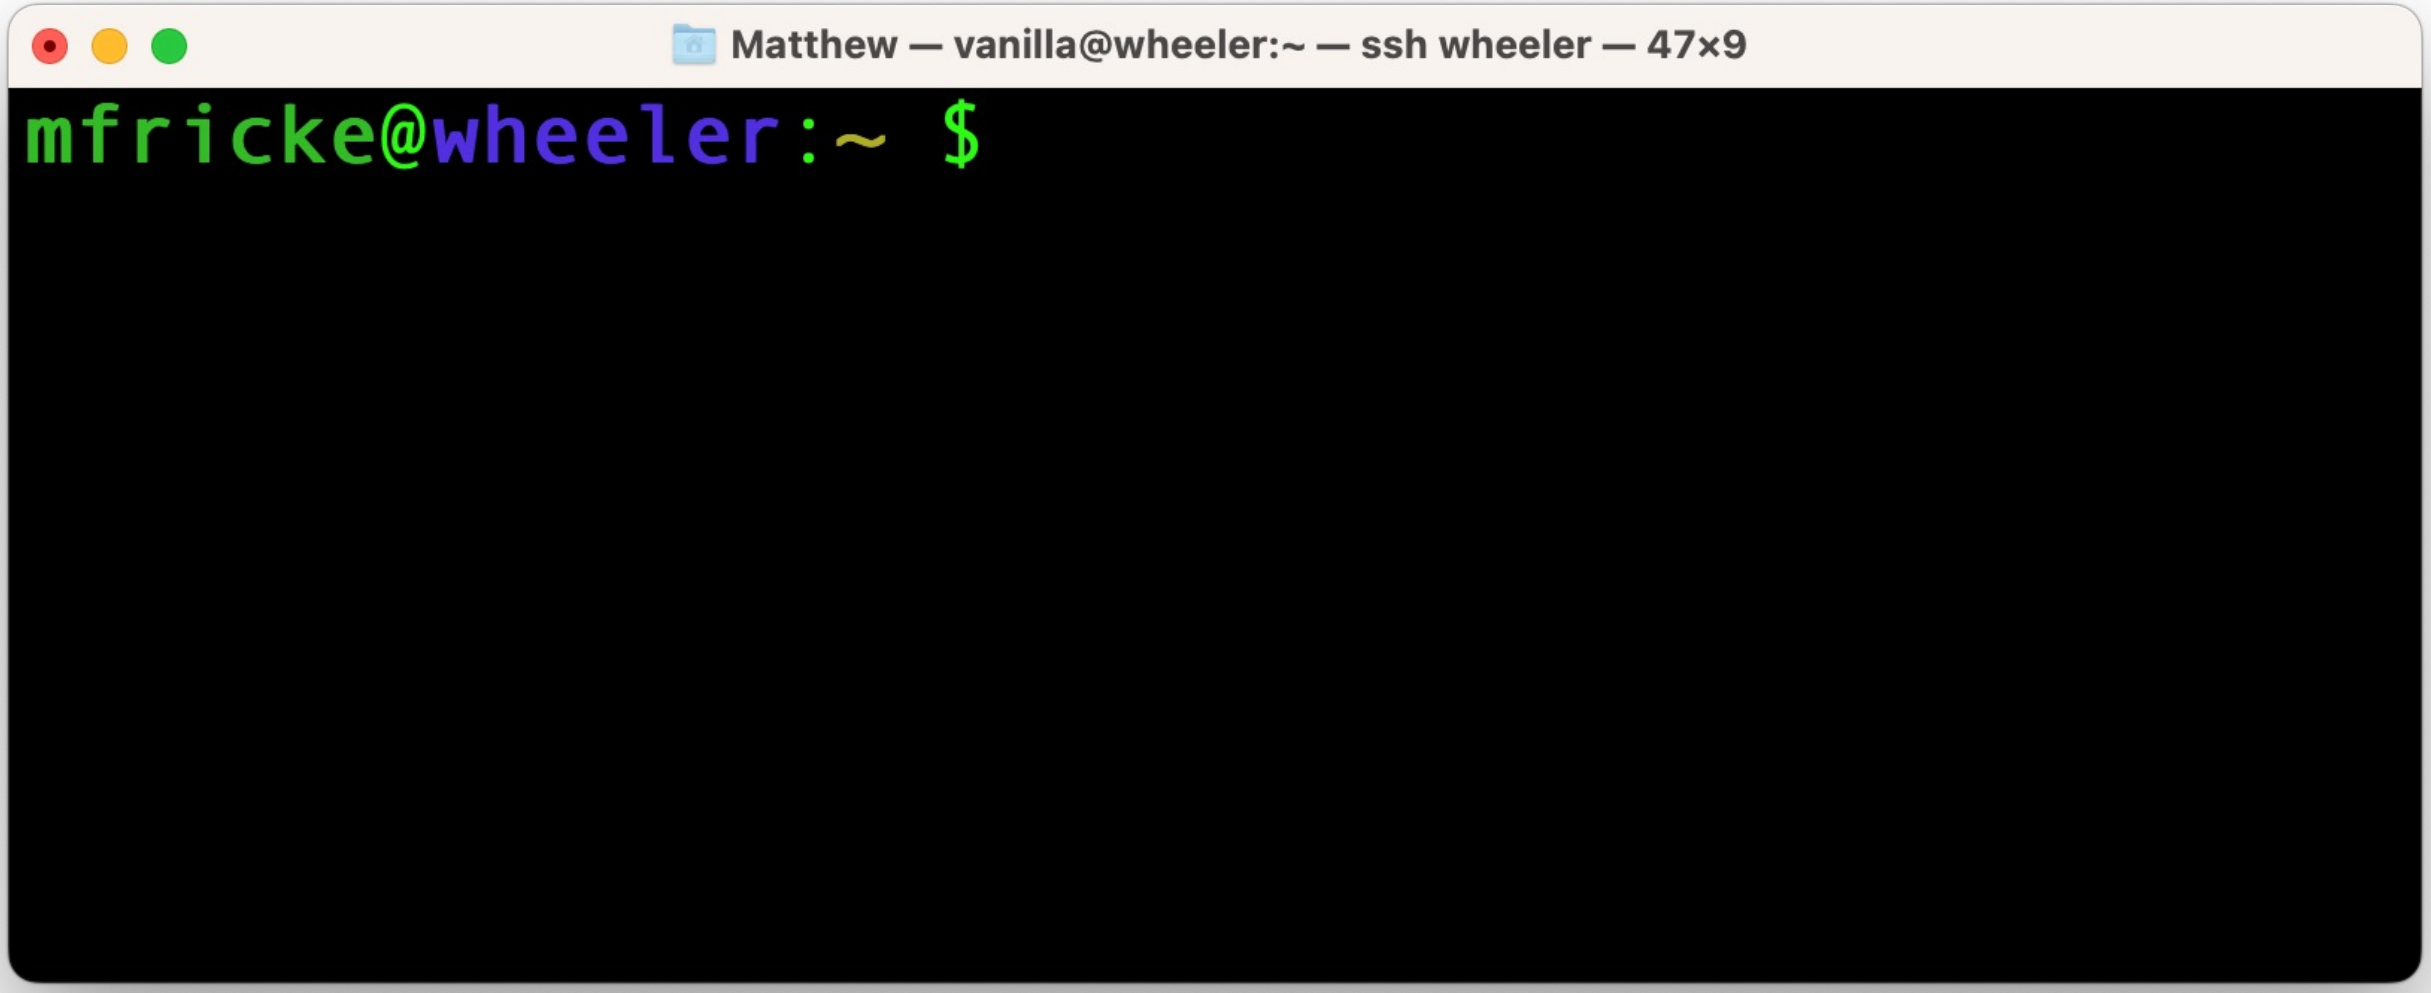

Linux and the BASH Shell

The Borne-Again Shell (BASH)

Written in 1976 by Stephen Bourne for UNIX version 7.

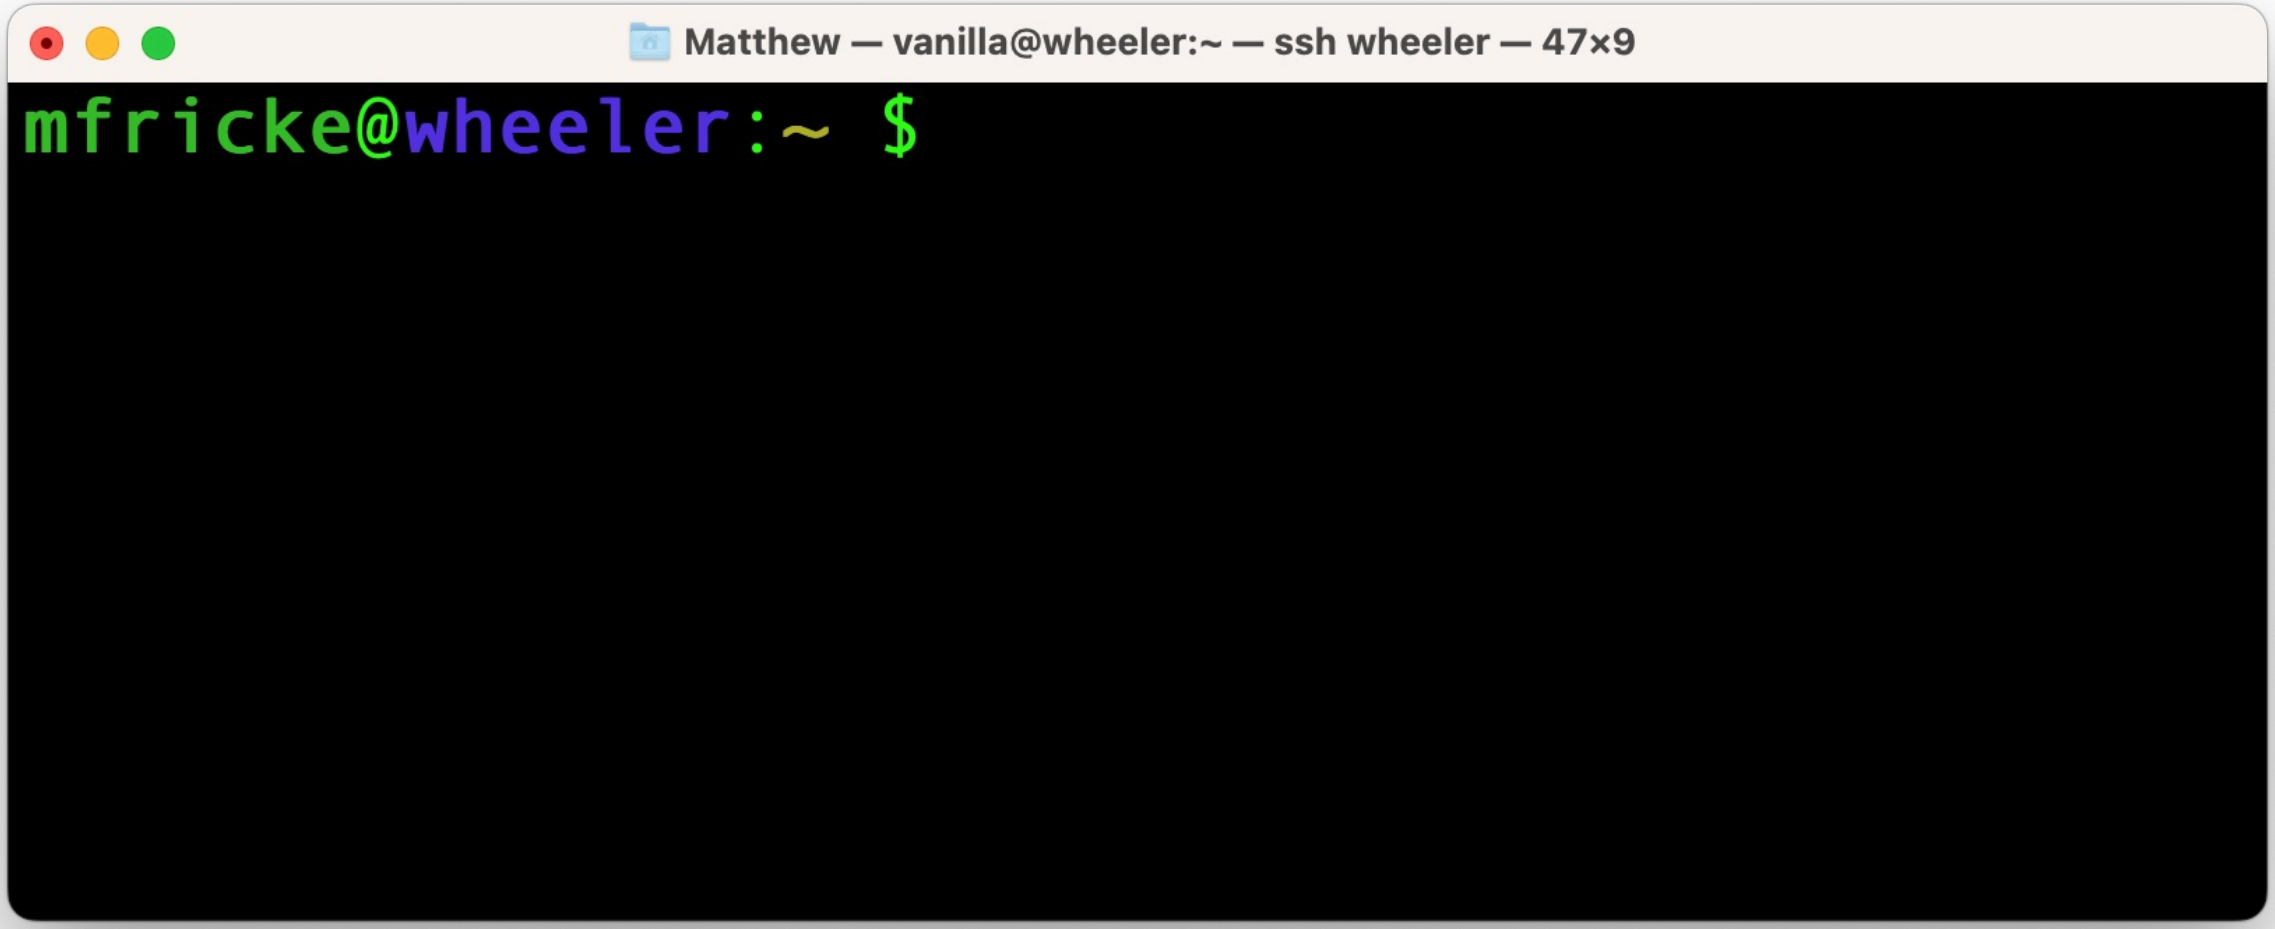

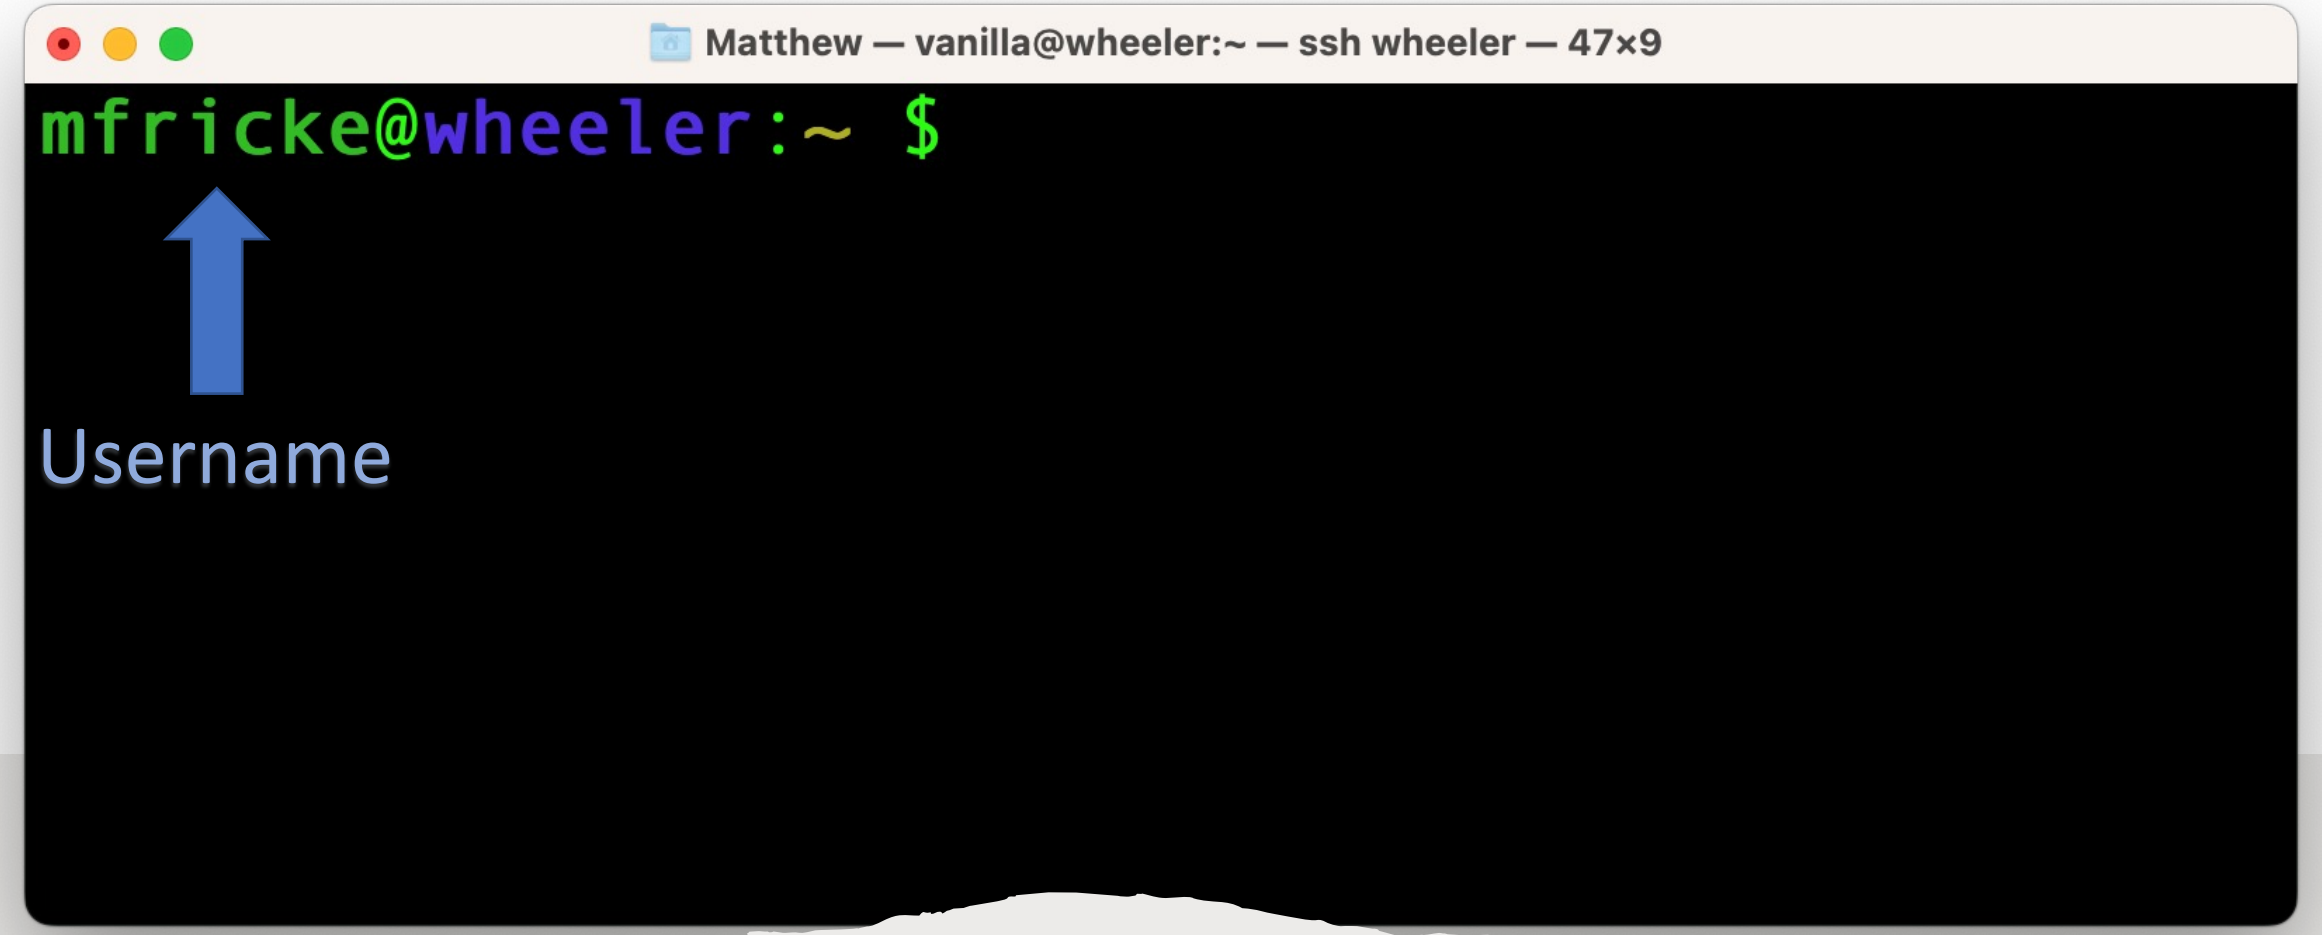

Understanding the BASH prompt…

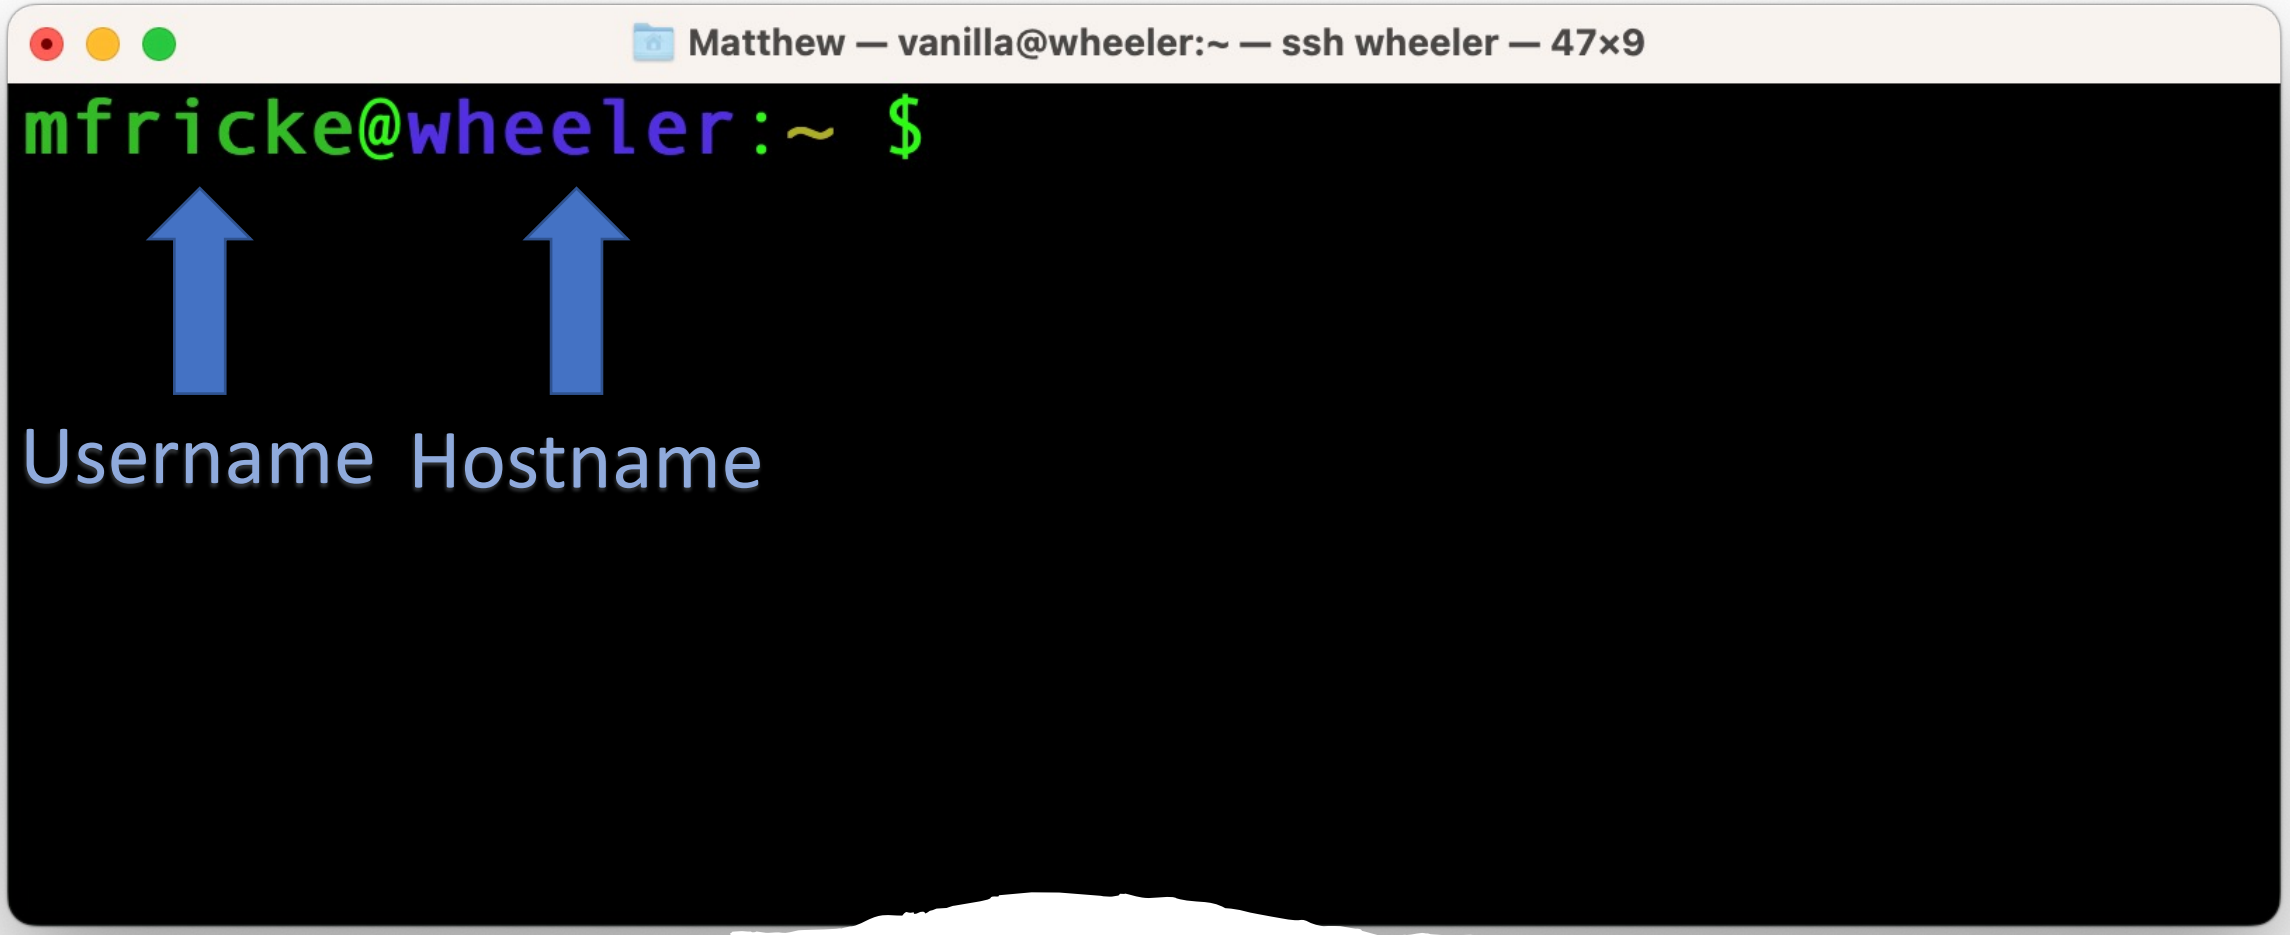

Understanding the BASH prompt…

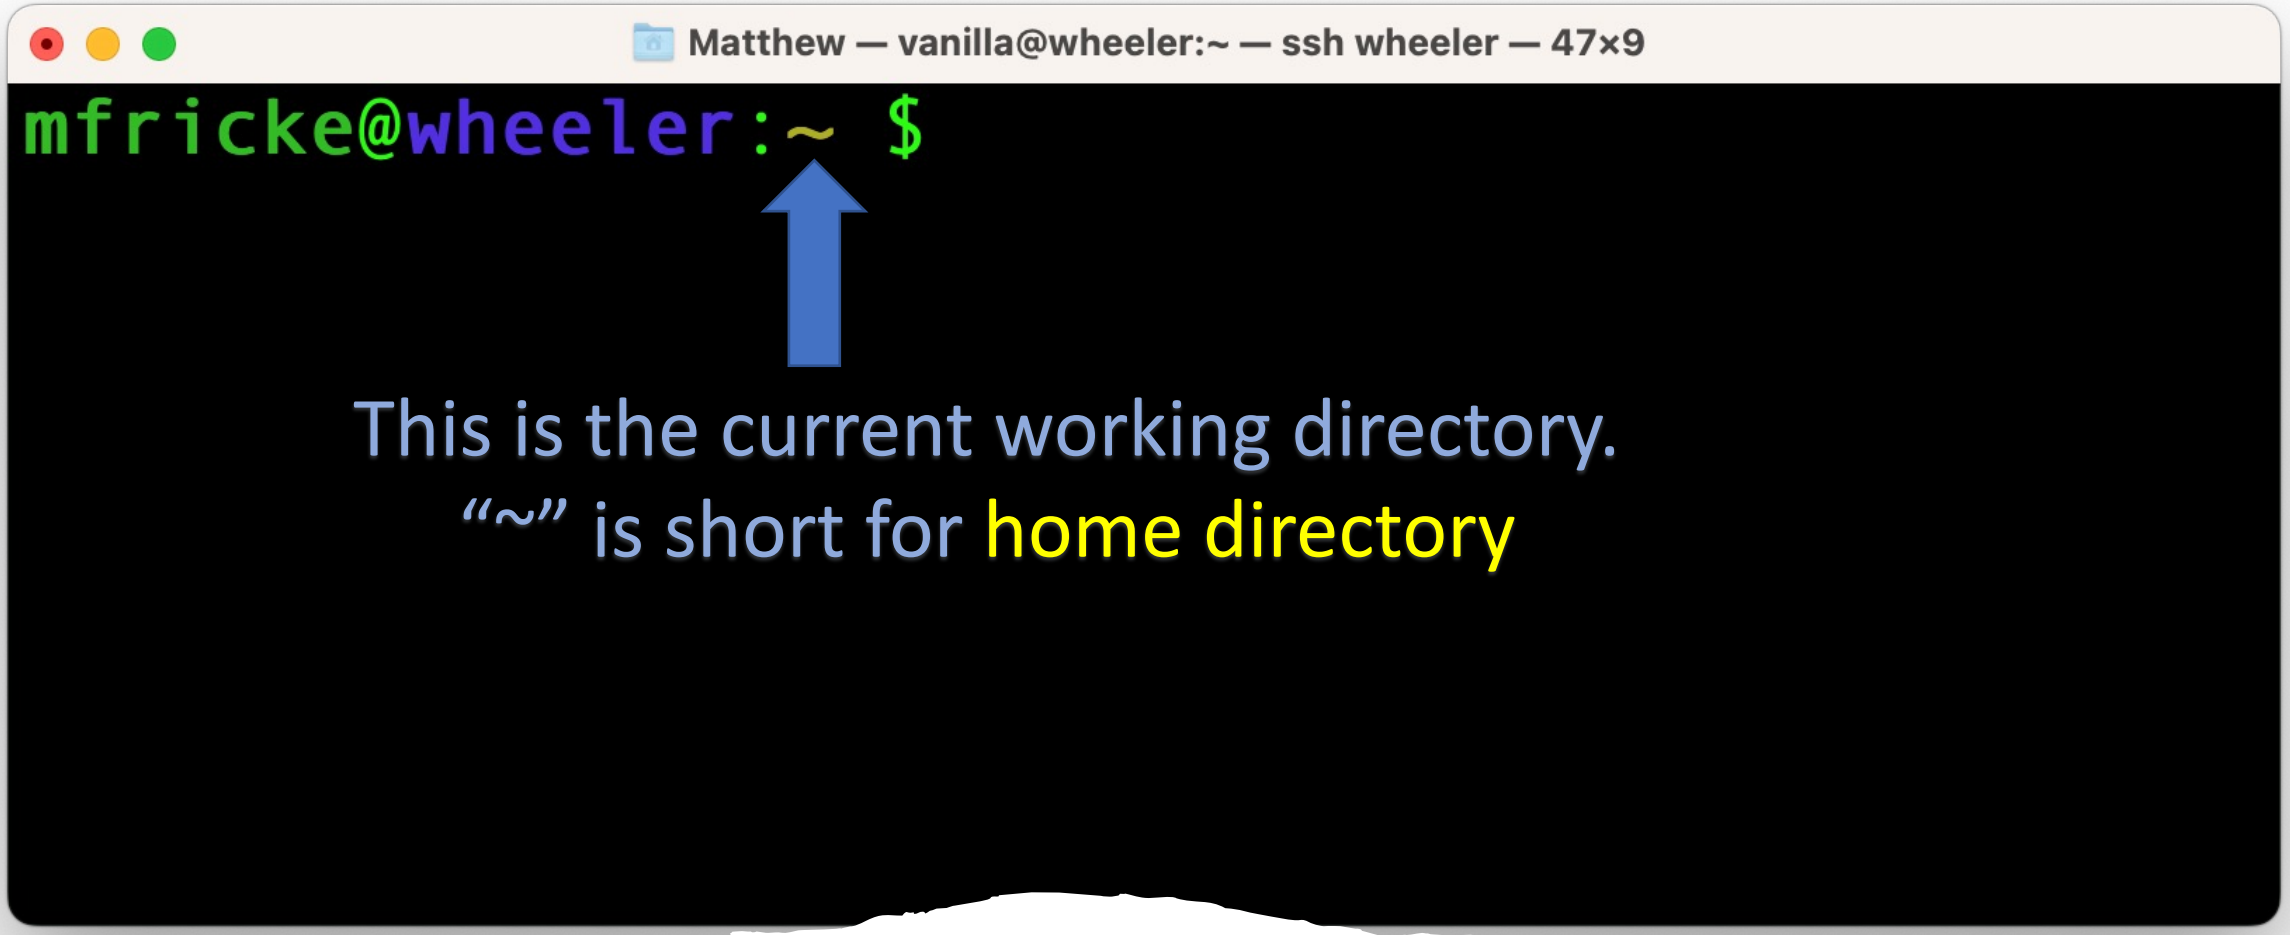

Understanding the BASH prompt…

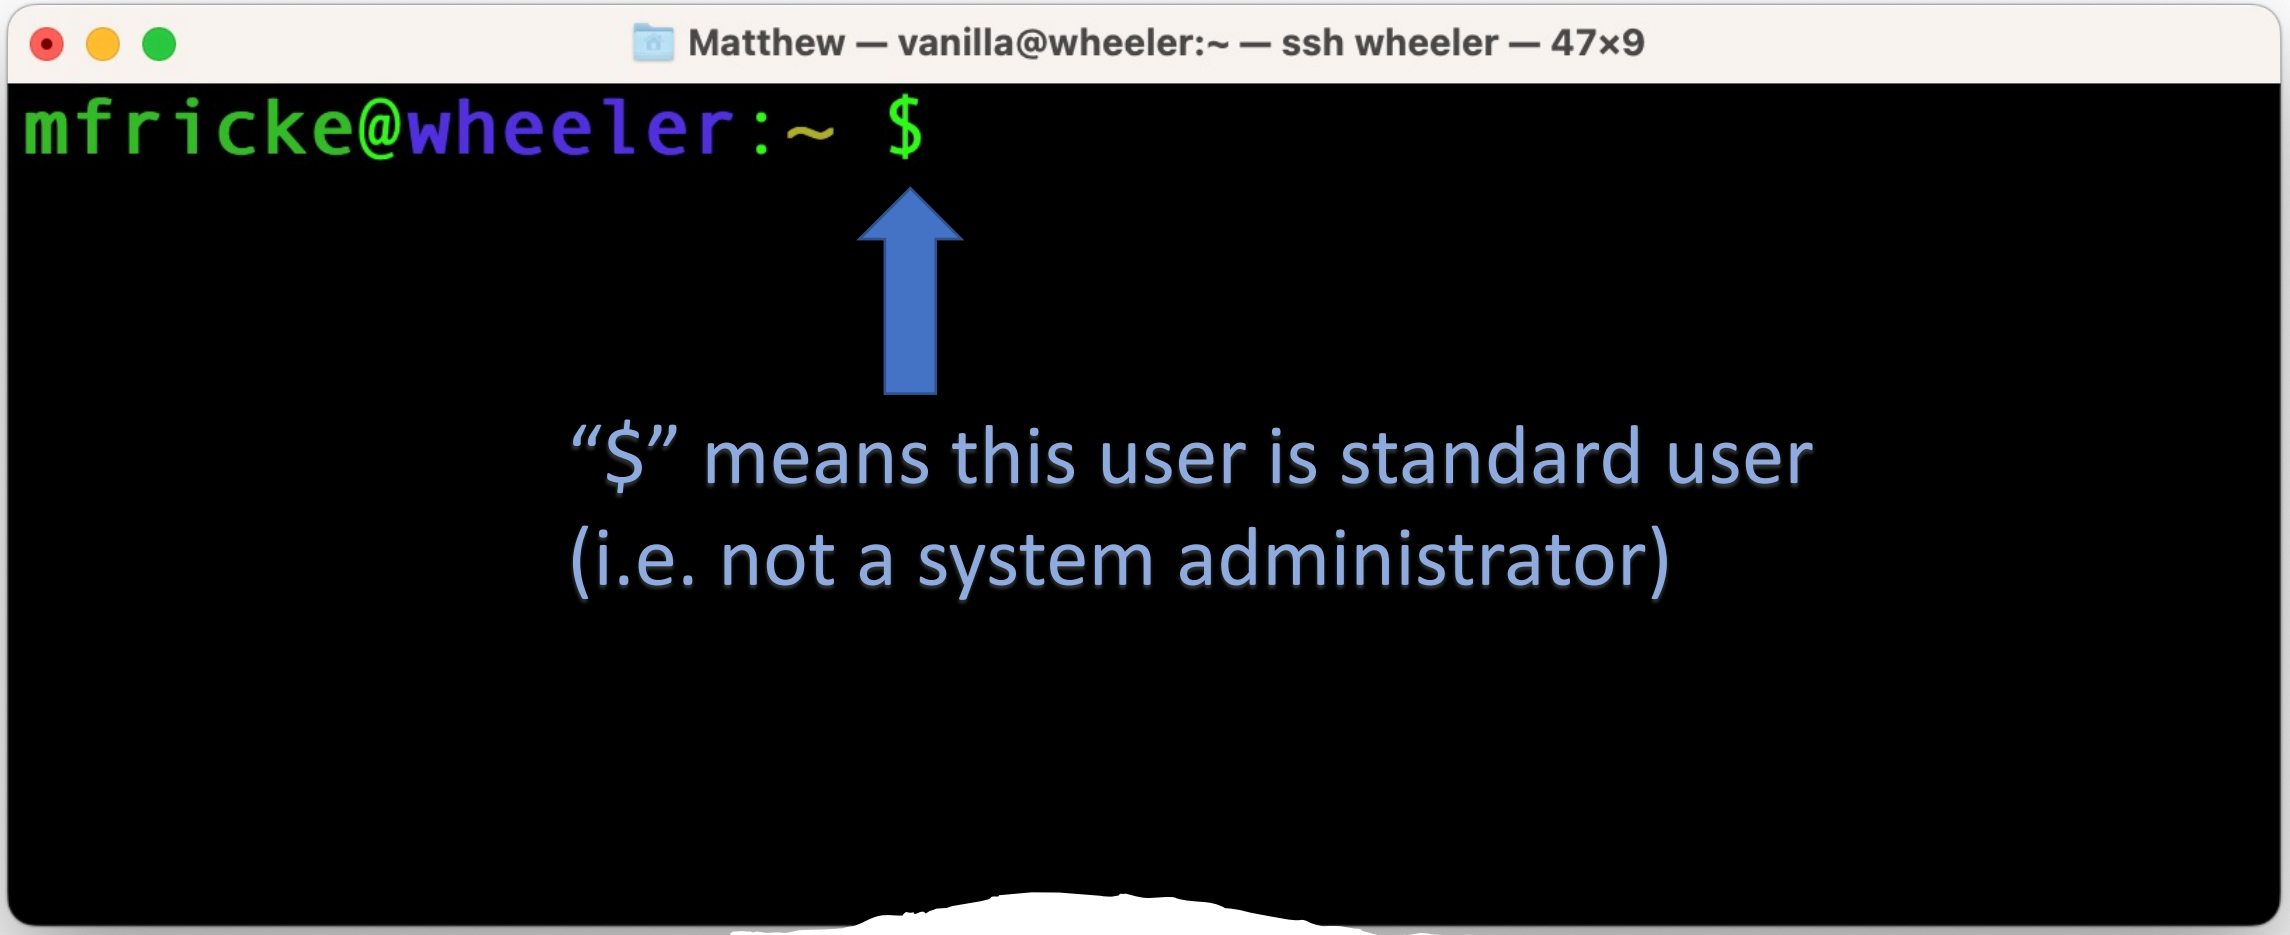

Understanding the BASH prompt…

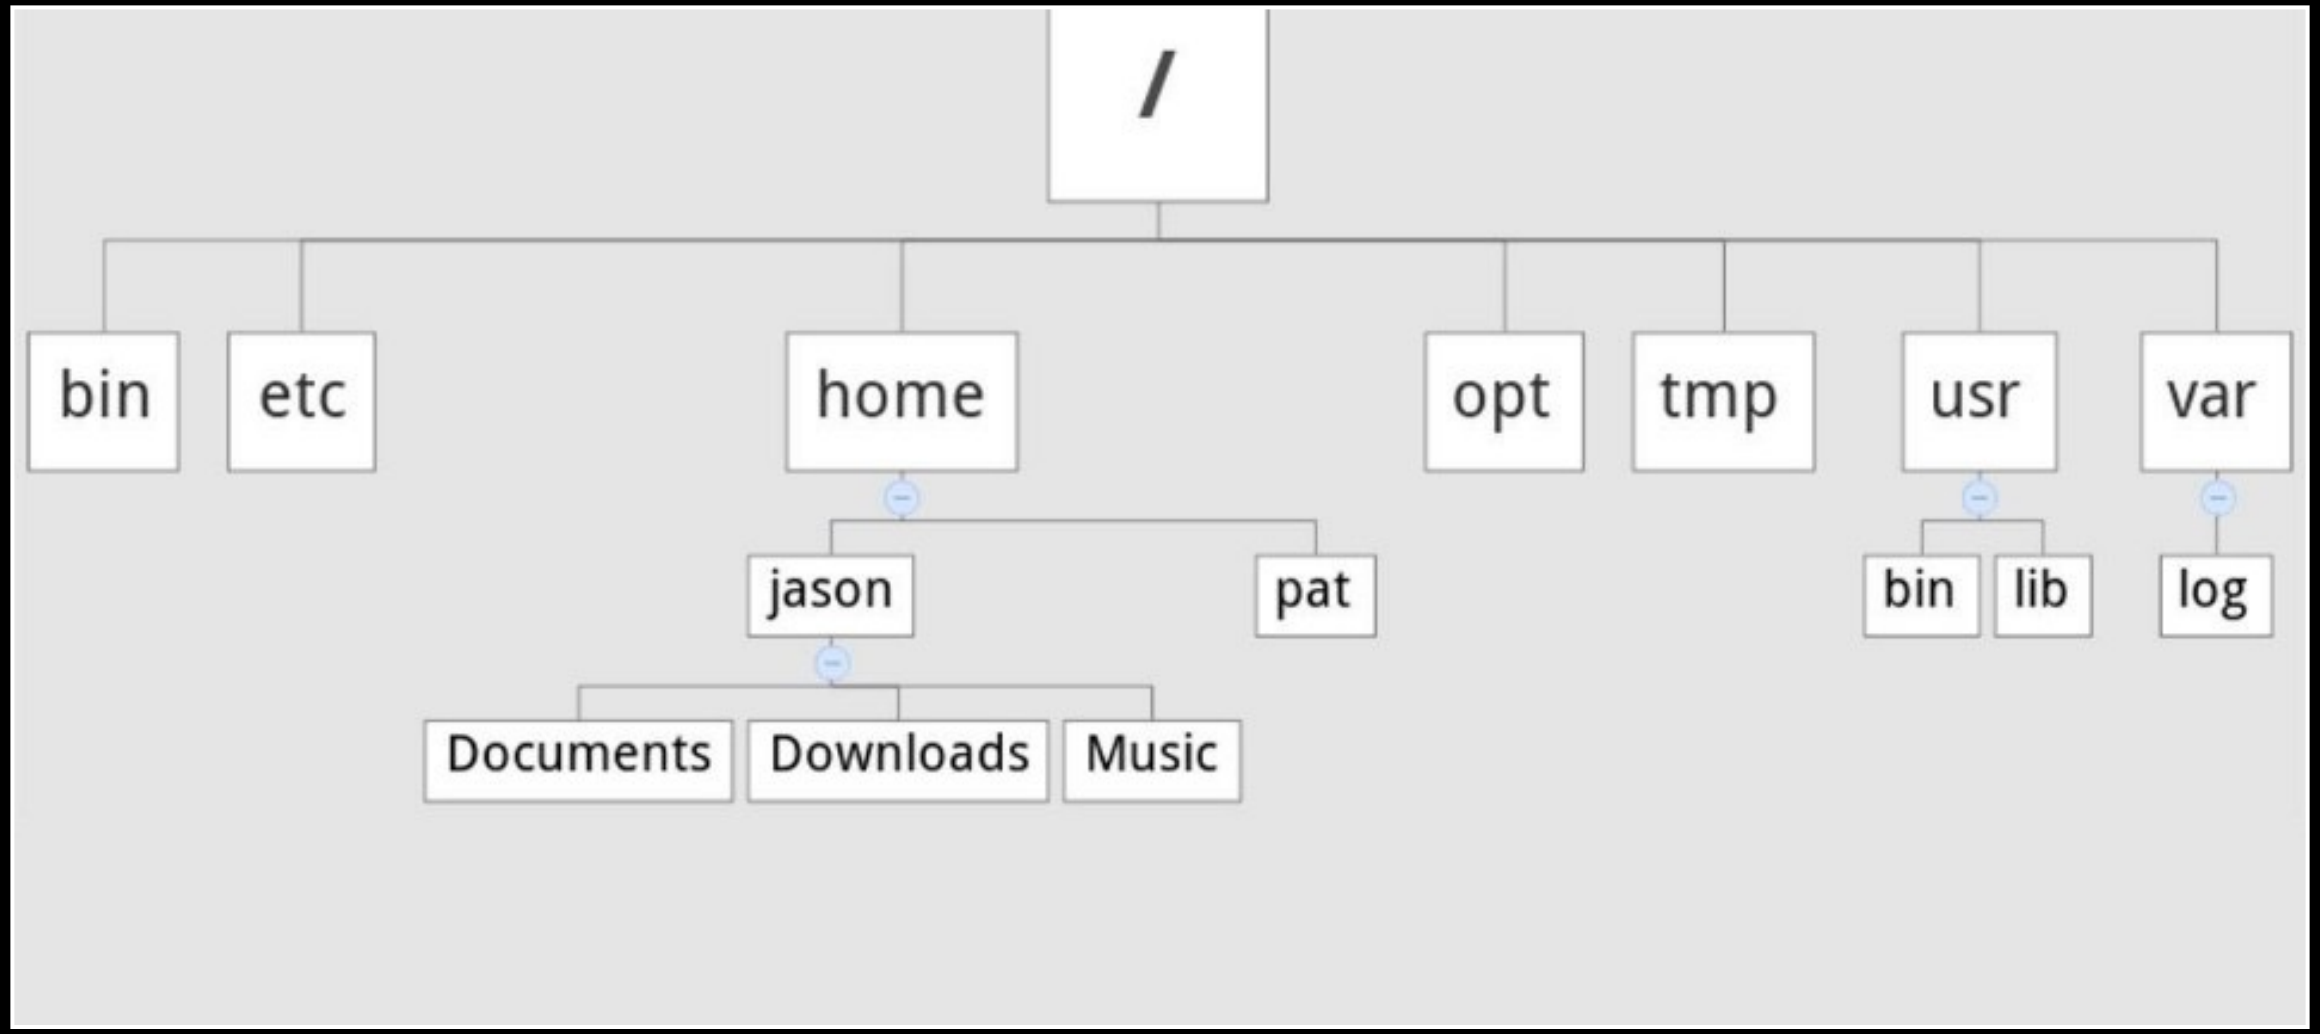

Example Filesystem Tree

Figuring out where you are in the filesystem…

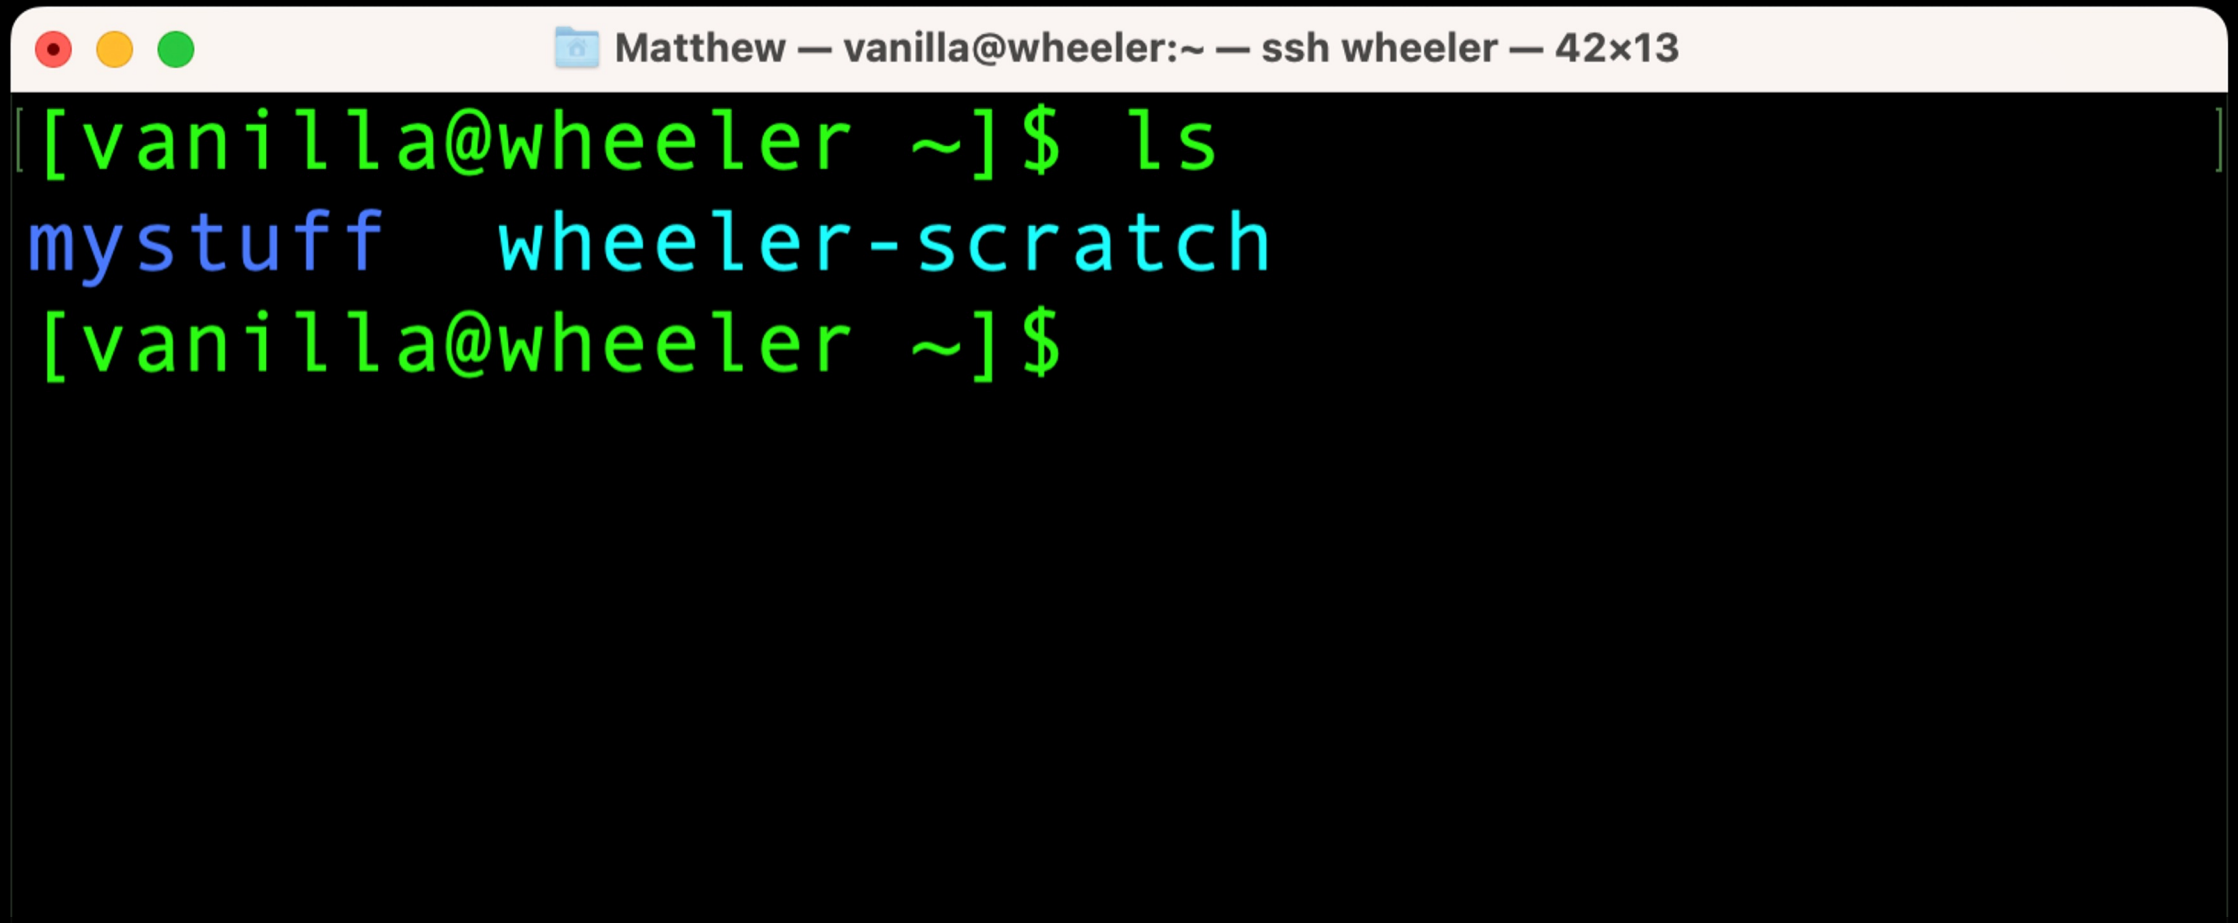

Figuring out where you are in the filesystem…

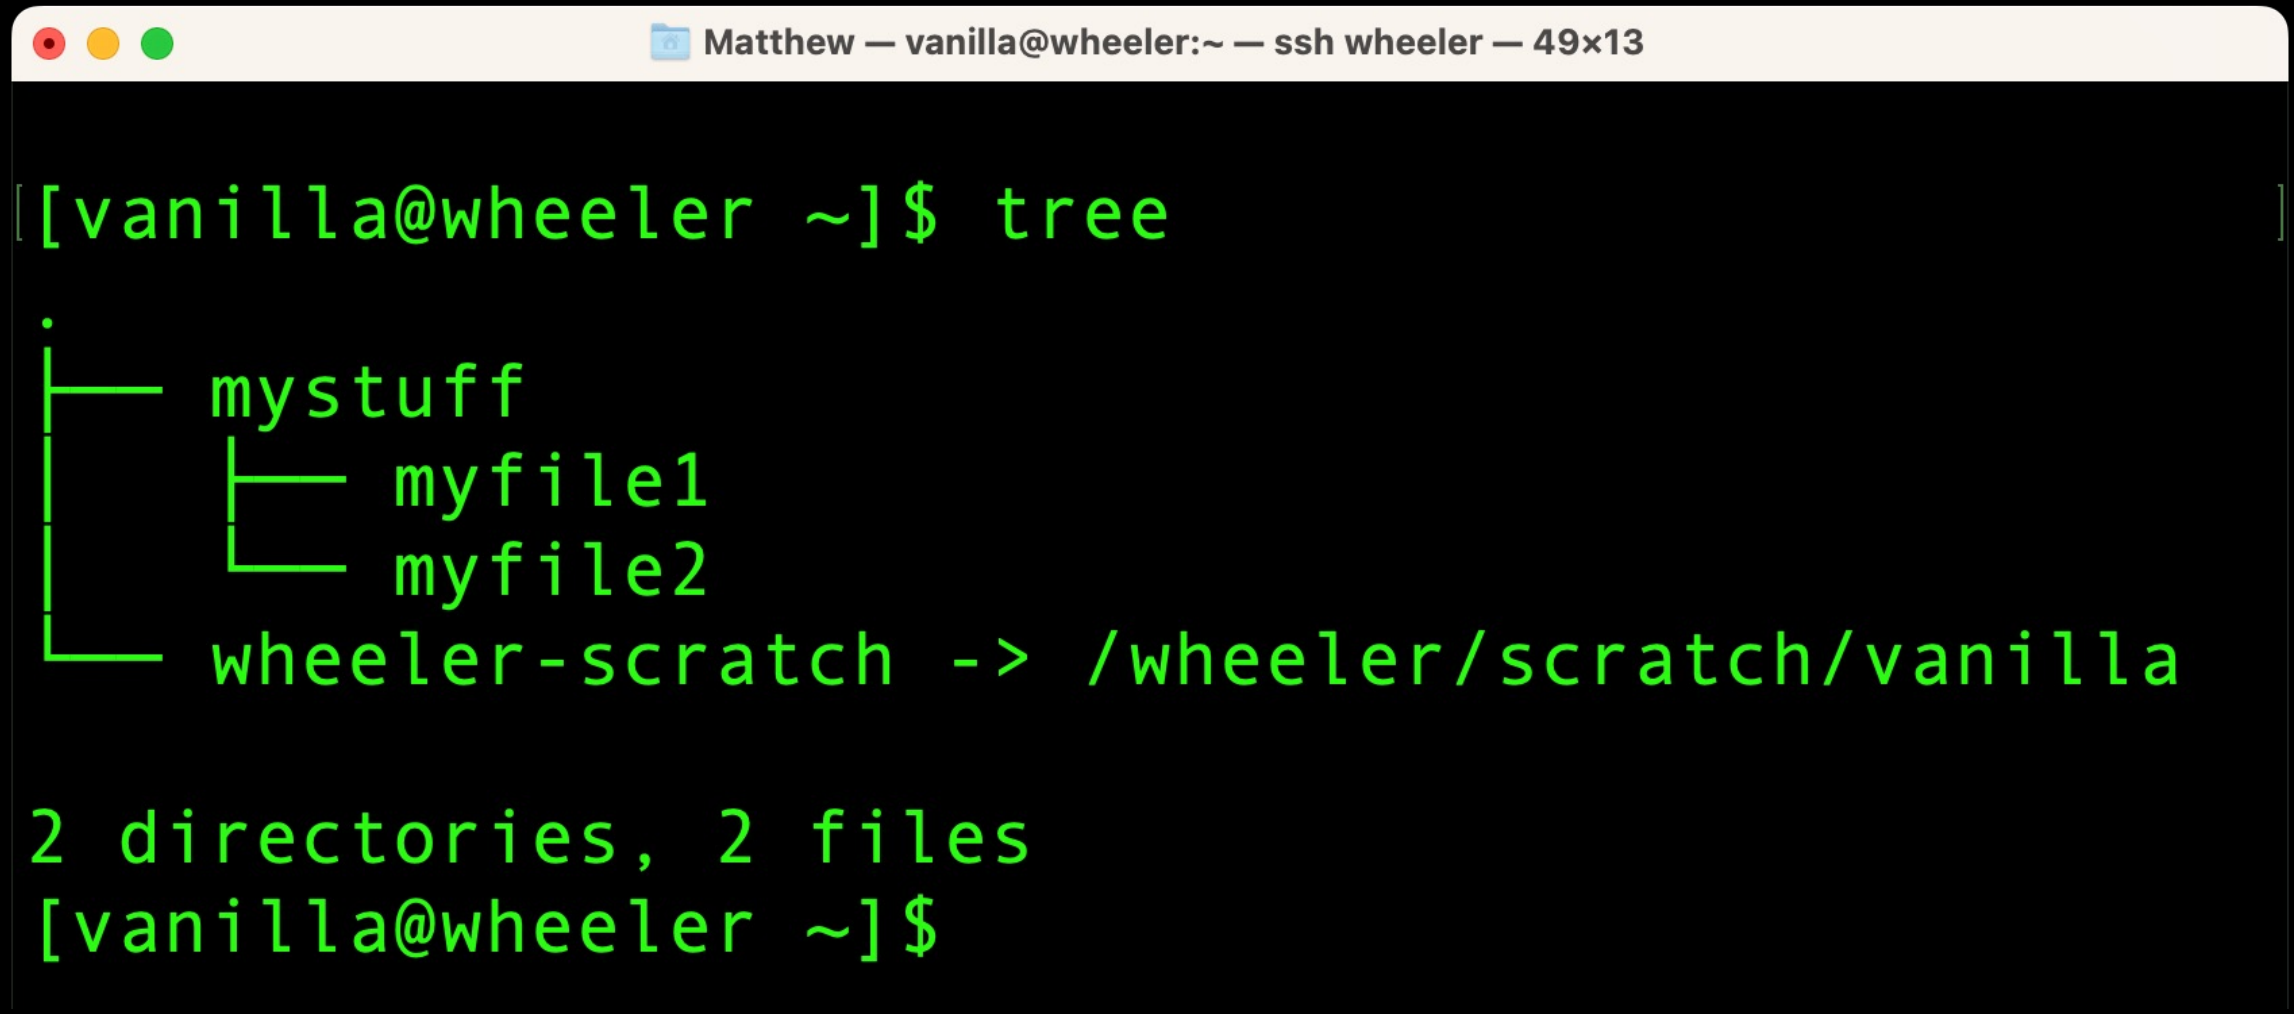

Figuring out where you are in the filesystem…

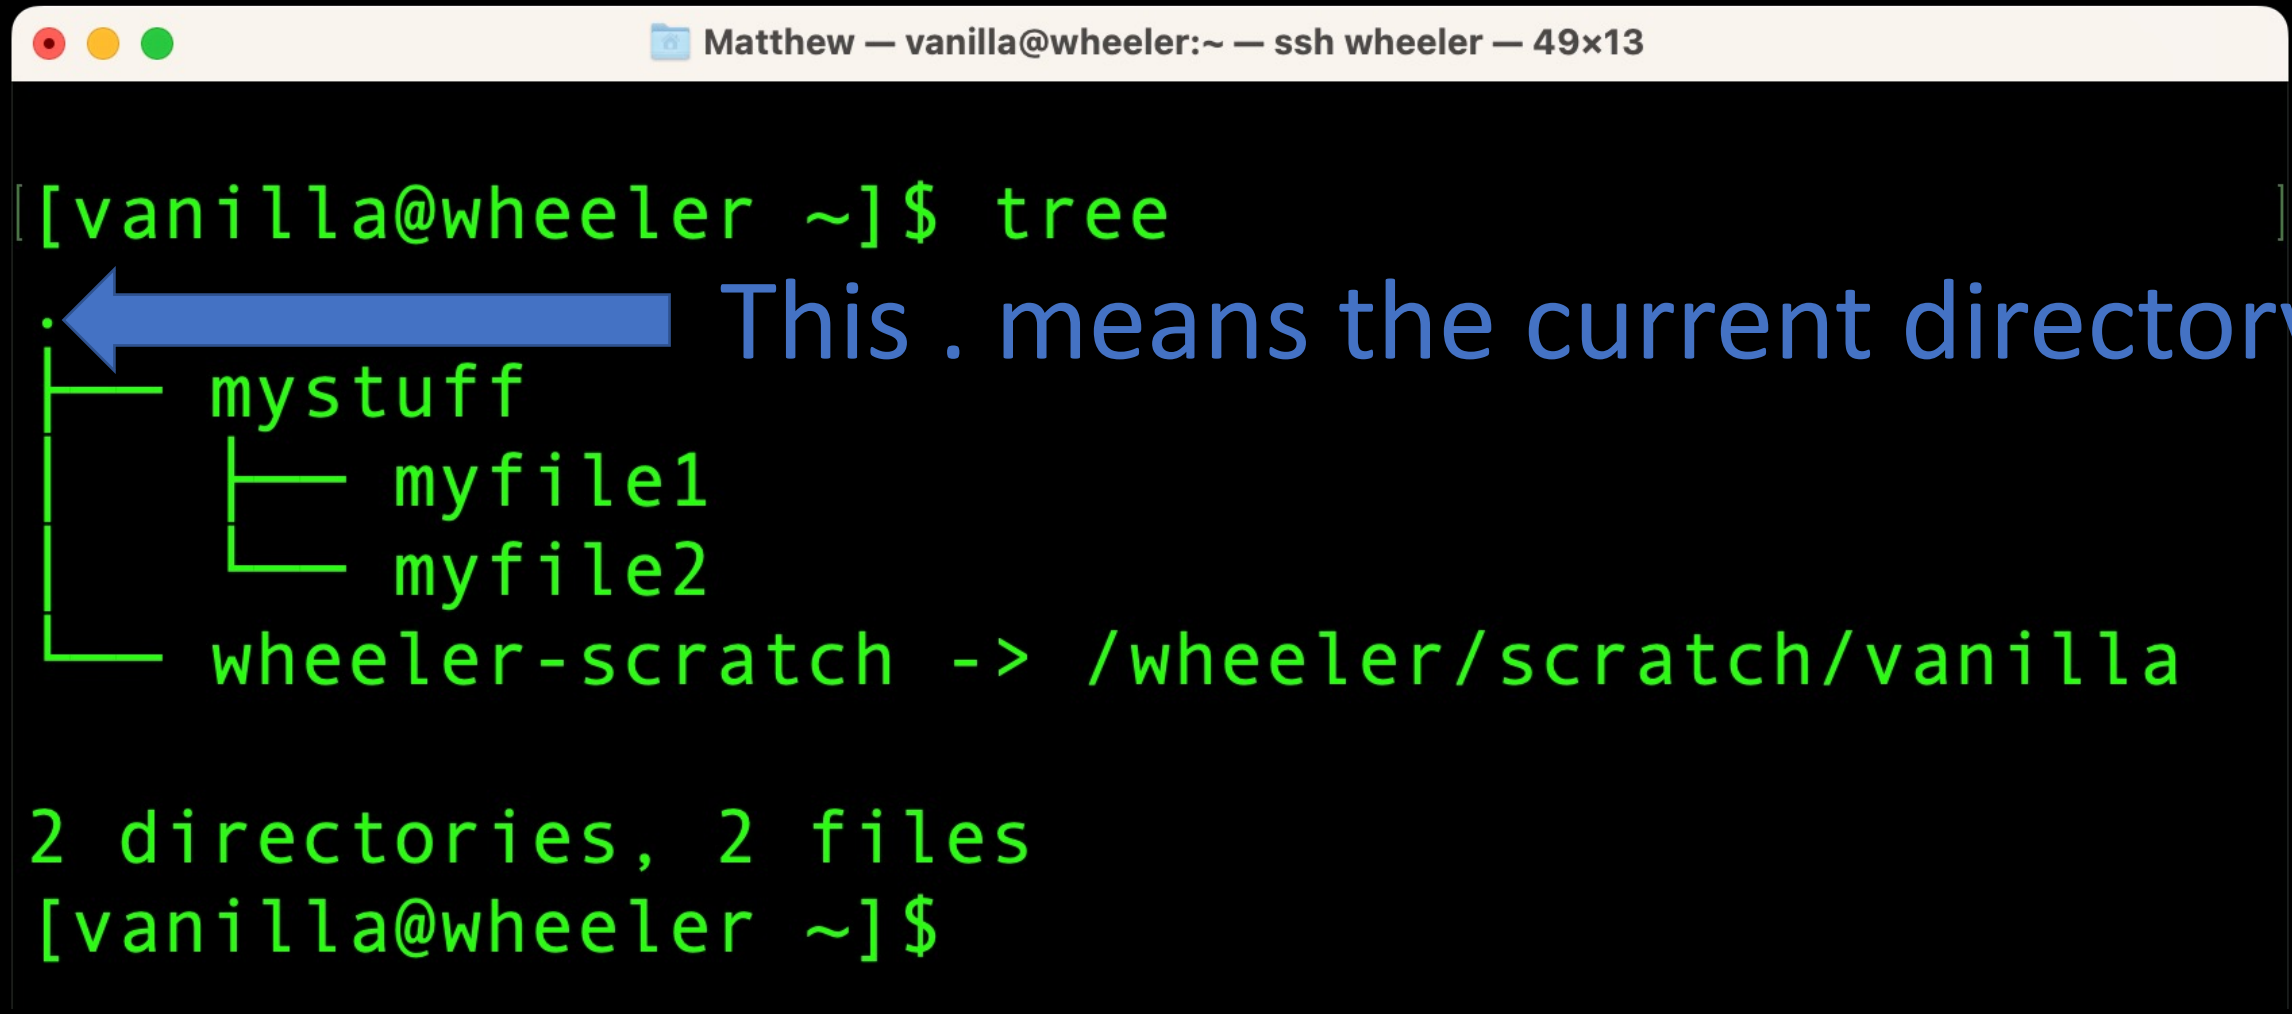

Figuring out where you are in the filesystem…

“Absolute” paths vs “relative” paths

- A path is a list of directories and/or files. It is a path through the directory tree that tells one how to get somewhere in the filesystem.

- An absolute path tells one how to get to the destination from starting from the root of the filesystem. E.g “/users/vanilla/mystuff/”

- A relative path specifies how to get there *starting from the current working directory*. E.g vanilla/mystuff/

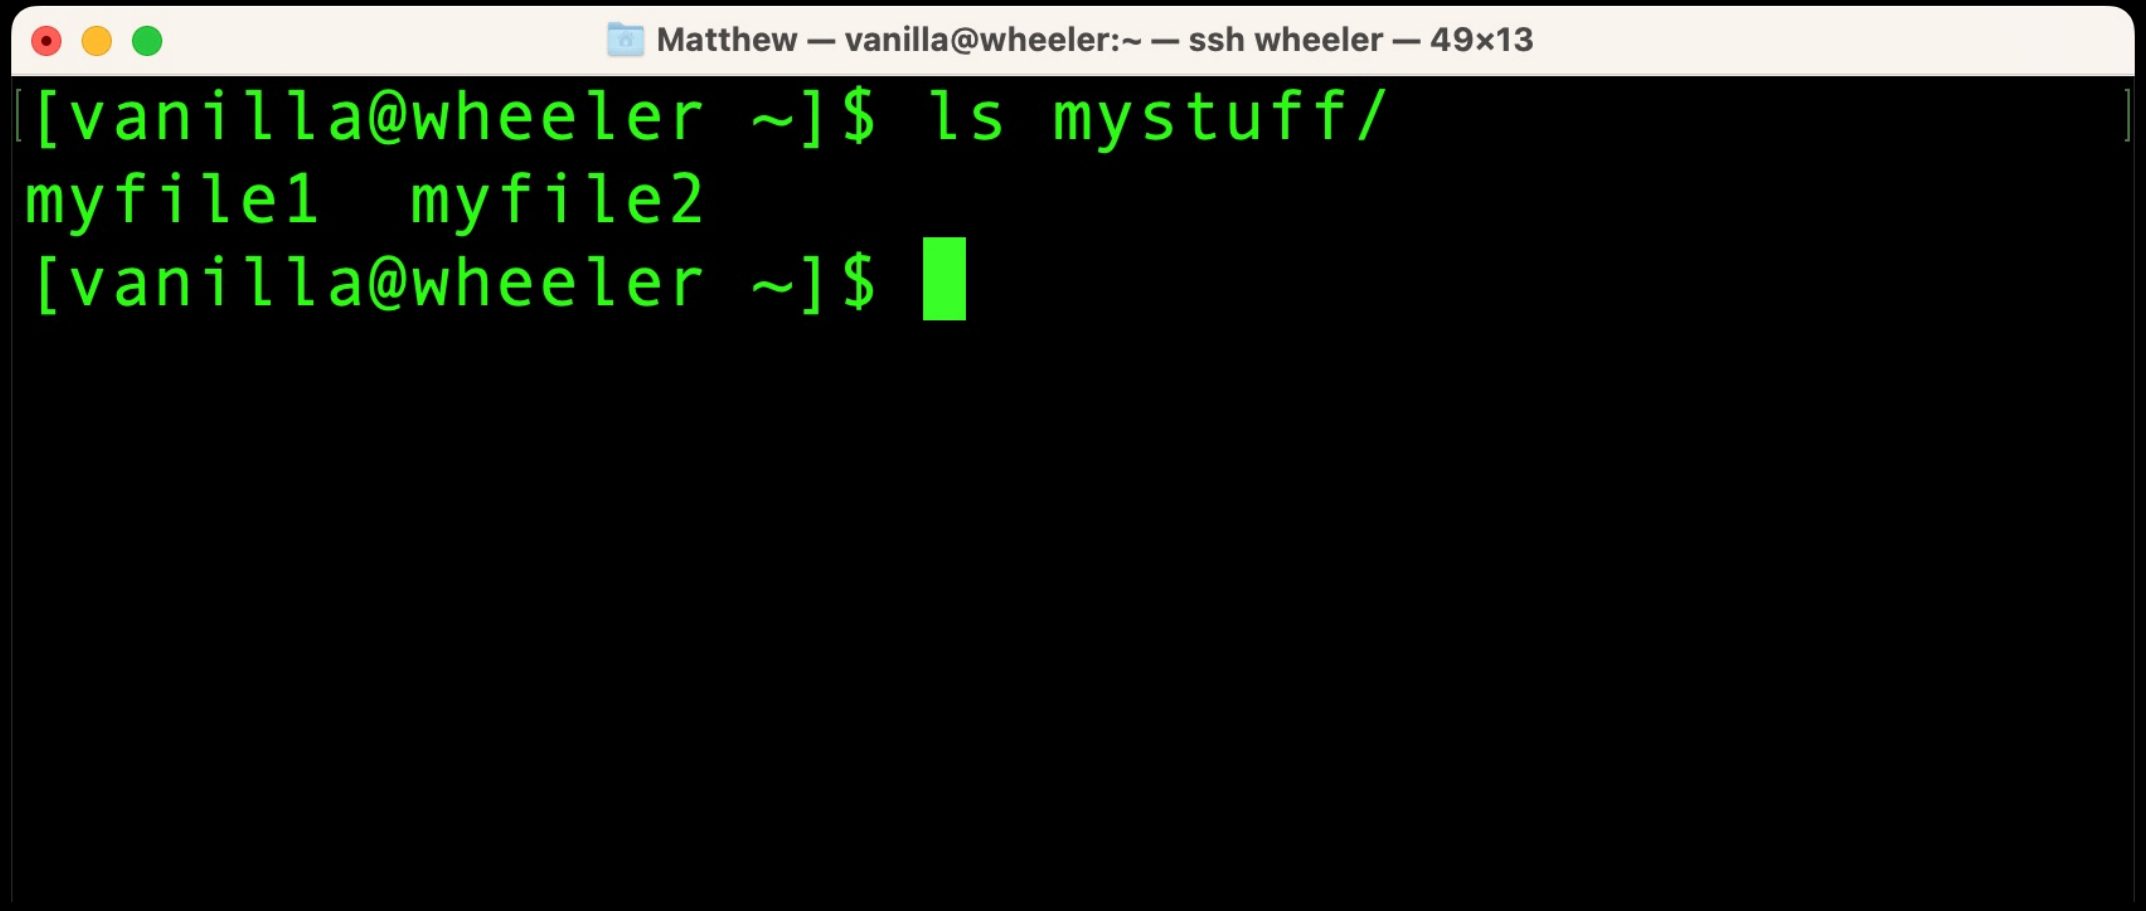

Figuring out where you going...

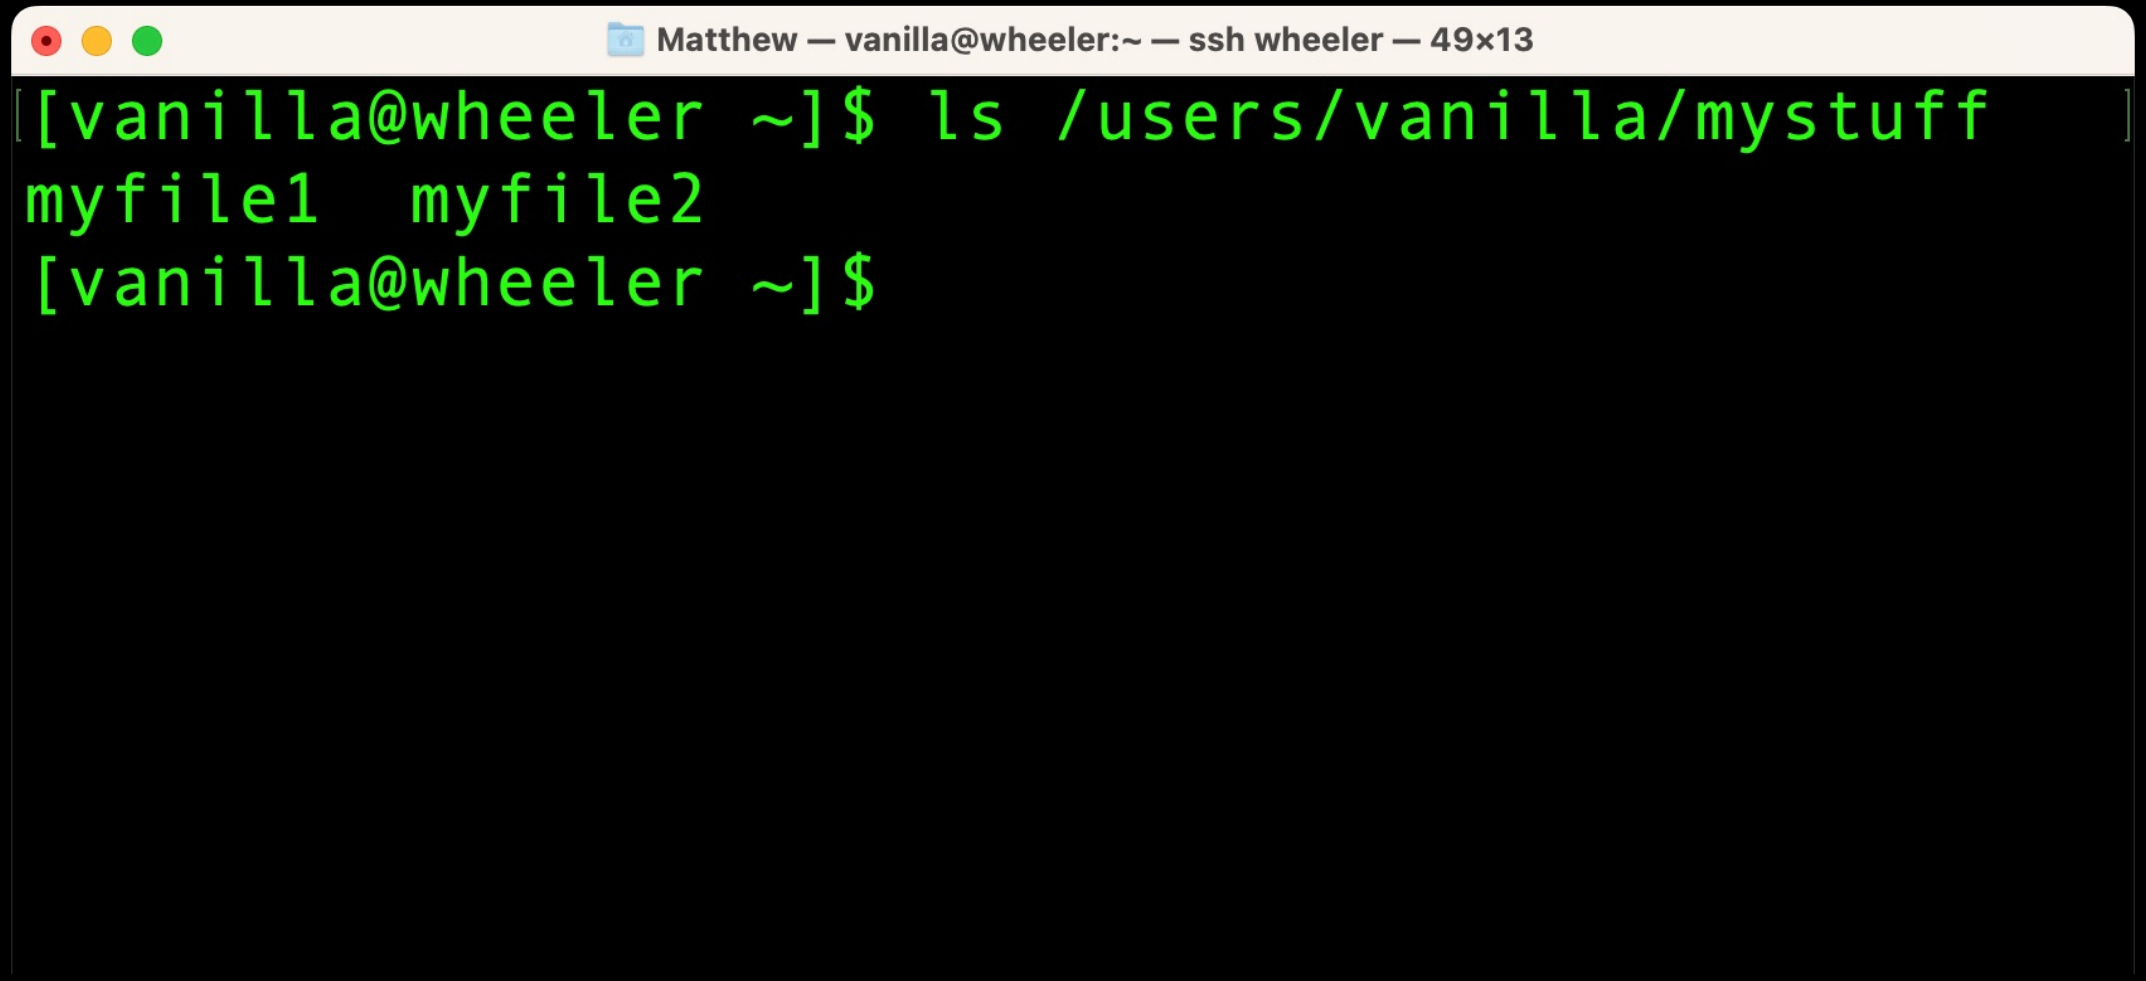

Figuring out where you going…

```
[vanilla@wheeler ~]$ ls ./mystuff/
myfile1  myfile2
[vanilla@wheeler ~]$ ls ~/mystuff/
myfile1  myfile2
[vanilla@wheeler ~]$
```

Figuring out where you going…

```
[vanilla@wheeler ~]$ ls -a
.                       .modulesbeginenv
..                      mystuff
.addressbook            .oracle_jre_usage
.addressbook.lu         .pinerc
.bashrc                 .pki
.cache                  .rhosts
.comsol                 .shosts
.config                 .spack
.flexlmrc               .ssh
```

Figuring out where you going...

```
[vanilla@wheeler ~]$ ls -l
total 4
drwxr-xr-x 2 vanilla users 4096 Jun 14 22:05 mystuff
lrwxrwxrwx 1 vanilla users   24 Jun 14 21:20 wheeler-scratch ->
/wheeler/scratch/vanilla
[vanilla@wheeler ~]$
```

Figuring out where you going…

```
[vanilla@wheeler ~]$ ls -l mystuff/
total 473704
-rw-r--r-- 1 vanilla users 483165473 Jun 14 23:20 myfile1
-rw-r--r-- 1 vanilla users         0 Jun 14 22:05 myfile2
[vanilla@wheeler ~]$
```

Figuring out where you going...

```
[vanilla@wheeler ~]$ ls -lh mystuff/
total 463M
-rw-r--r-- 1 vanilla users 461M Jun 14 23:20 myfile1
-rw-r--r-- 1 vanilla users    0 Jun 14 22:05 myfile2
[vanilla@wheeler ~]$
```

Figuring out where you going...

```
[vanilla@wheeler ~]$ du -s
499704 .
[vanilla@wheeler ~]$ du -sh
488M .
[vanilla@wheeler ~]$
```

Figuring out where you going...

```
[vanilla@wheeler ~]$ df -h
Filesystem                          Size  Used Avail Use% Mounted on
devtmpfs                             24G     0   24G   0% /dev
tmpfs                                24G   64K   24G   1% /dev/shm
tmpfs                                24G  968M   23G   5% /run
tmpfs                                24G     0   24G   0% /sys/fs/cgroup
/dev/mapper/centos-root             930G  567G  363G  61% /
/dev/sdc2                           836G   72G  764G   9% /tmp
/dev/md126p1                        2.0G  333M  1.7G  17% /boot
172.17.2.254:/mnt/wheeler-scratch    37T   28T  8.7T  77% /wheeler/scratch
172.17.2.255:/mnt/wheeler-scratch2   37T   28T  9.0T  76% /wheeler/scratch2
beegfs_nodev                        110T   51T   60T  46% /carc/scratch
chama:/home/homes                    65T   36T   30T  55% /users
chama:/home/carc_projects            65T   36T   30T  55% /projects
```

Figuring out where you going…

```
[vanilla@wheeler ~]$ quota -s
Disk quotas for user vanilla (uid 659):
     Filesystem   space   quota   limit   grace   files   quota   limit   grace
chama:/home/homes
                  488M    100G    200G             315   4295m   4295m
[vanilla@wheeler ~]$
```

Figuring out where you going...

```
[vanilla@wheeler ~]$ stat mystuff/myfile1
  File: ‘mystuff/myfile1’
  Size: 483165473 Blocks: 947408     IO Block: 65536  regular file
Device: 28h/40d Inode: 9232782834205560540  Links: 1
Access: (0644/-rw-r--r--)  Uid: (  659/ vanilla)   Gid: (  100/   users)
Access: 2022-06-14 22:05:27.503289000 -0600
Modify: 2022-06-14 23:20:26.945918000 -0600
Change: 2022-06-14 23:20:48.754917000 -0600
 Birth: -
```

Figuring out what you’ve got…

```
[vanilla@wheeler ~]$ find -name myfile2
./mystuff/myfile2

[vanilla@wheeler ~]$ find -name "myfile*"
./mystuff/myfile1
./mystuff/myfile2
./mystuff/myfile3
./mystuff/myfile0
```

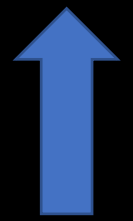

Wildcard

Figuring out what you've got…

```
[vanilla@wheeler ~]$ cd mystuff/
[vanilla@wheeler ~/mystuff]$
```

Use the tab key to autocomplete

Going somewhere new…

Now it is your turn…

- For this path:
  /projects/shared/workshops/beginner/vecadd

- What are the names of the files in that directory?
- When were they last modified?
- How large are the files?

You can find this information with the ls command.

Now it is your turn…

- For this path:
  /projects/shared/workshops/beginner/vecadd

  Now “cd” into that directory using <tab> autocomplete.

- Now you know how to find your way around filesystems using bash
- Let’s see how to modify the filesystem.

- In bash to move a file we use the mv command.
- To copy a file it is cp.
- To copy files from CARC to a personal computer use scp or rsync.

```
[vanilla@wheeler beginner]$ pwd
/projects/shared/workshops/beginner
[vanilla@wheeler beginner]$ cd ~
[vanilla@wheeler ~]$ pwd
/users/vanilla
[vanilla@wheeler ~]$
```

First return to your home directory…

```
[vanilla@wheeler ~]$ cd mystuff
[vanilla@wheeler ~/mystuff]$ mv myfile1 myfile0
[vanilla@wheeler ~/mystuff]$ ls
myfile0  myfile2  myfile3
[vanilla@wheeler ~/mystuff]
```

Modifying the filesystem…
moving a file.

```
[vanilla@wheeler ~/mystuff]$ cp myfile0 myfile1
[vanilla@wheeler ~/mystuff]$
```

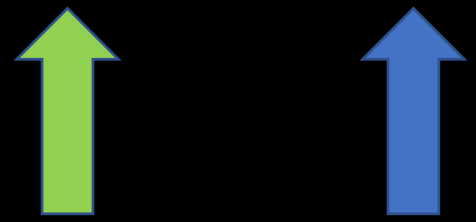

Source Destination

```
[vanilla@wheeler ~/mystuff]$ ls
myfile0  myfile1  myfile2  myfile3
[vanilla@wheeler ~/mystuff]$
```

Modifying the filesystem…
copying a file.

```
[vanilla@wheeler ~]$ cp -r mystuff mystuff2
[vanilla@wheeler ~]$
```

Source

Destination

```
[vanilla@wheeler ~]$ ls
mystuff  mystuff2  wheeler-scratch
```

Copying a whole directory tree…

```
Lycaon:~ matthew$ scp vanilla@wheeler.alliance.unm.edu:~/mystuff/myfile3 Desktop/
```

Source

Destination

```
(vanilla@wheeler.alliance.unm.edu) Password:
myfile3                       100%   40     2.0KB/s   00:00
```

Copying data to a personal computer from CARC…

```
Lycaon:~ matthew$ scp -r vanilla@wheeler.alliance.unm.edu:~/mystuff Desktop/
```

Source

Destination

```
(vanilla@wheeler.alliance.unm.edu) Password:
myfile1                       100% 1024KB   6.5MB/s   00:00
myfile2                       100% 2048KB 382.5KB/s   00:05
myfile3                       100%   40     3.2KB/s   00:00
myfile0                       100% 1024KB   8.8MB/s   00:00
```

Copying data to a personal computer from CARC…

```
Lycaon:~ matthew$ scp -r Desktop/mystuff vanilla@wheeler.alliance.unm.edu:~/
```

Source

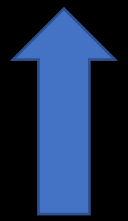

Destination

```
(vanilla@wheeler.alliance.unm.edu) Password:
myfile1                       100% 1024KB 591.5KB/s   00:01
myfile0                       100% 1024KB   2.0MB/s   00:00
myfile2                       100% 2048KB   2.1MB/s   00:00
myfile3                       100%   40     2.1KB/s   00:00
```

To copy from a personal computer to CARC…

```
ssh vanilla@wheeler.alliance.unm.edu
```

Log back into wheeler…

```
[vanilla@wheeler ~]$ file mystuff/myfile0
mystuff/myfile0: data

[vanilla@wheeler ~]$ file mystuff/myfile3
mystuff/myfile3: ASCII text
```

Figuring out file types …

```
[vanilla@wheeler ~]$ cat mystuff/myfile3
Welcome to the CARC Beginner's Workshop
```

Text files ...

```
[vanilla@wheeler ~]$ nano mystuff/myfile3
```

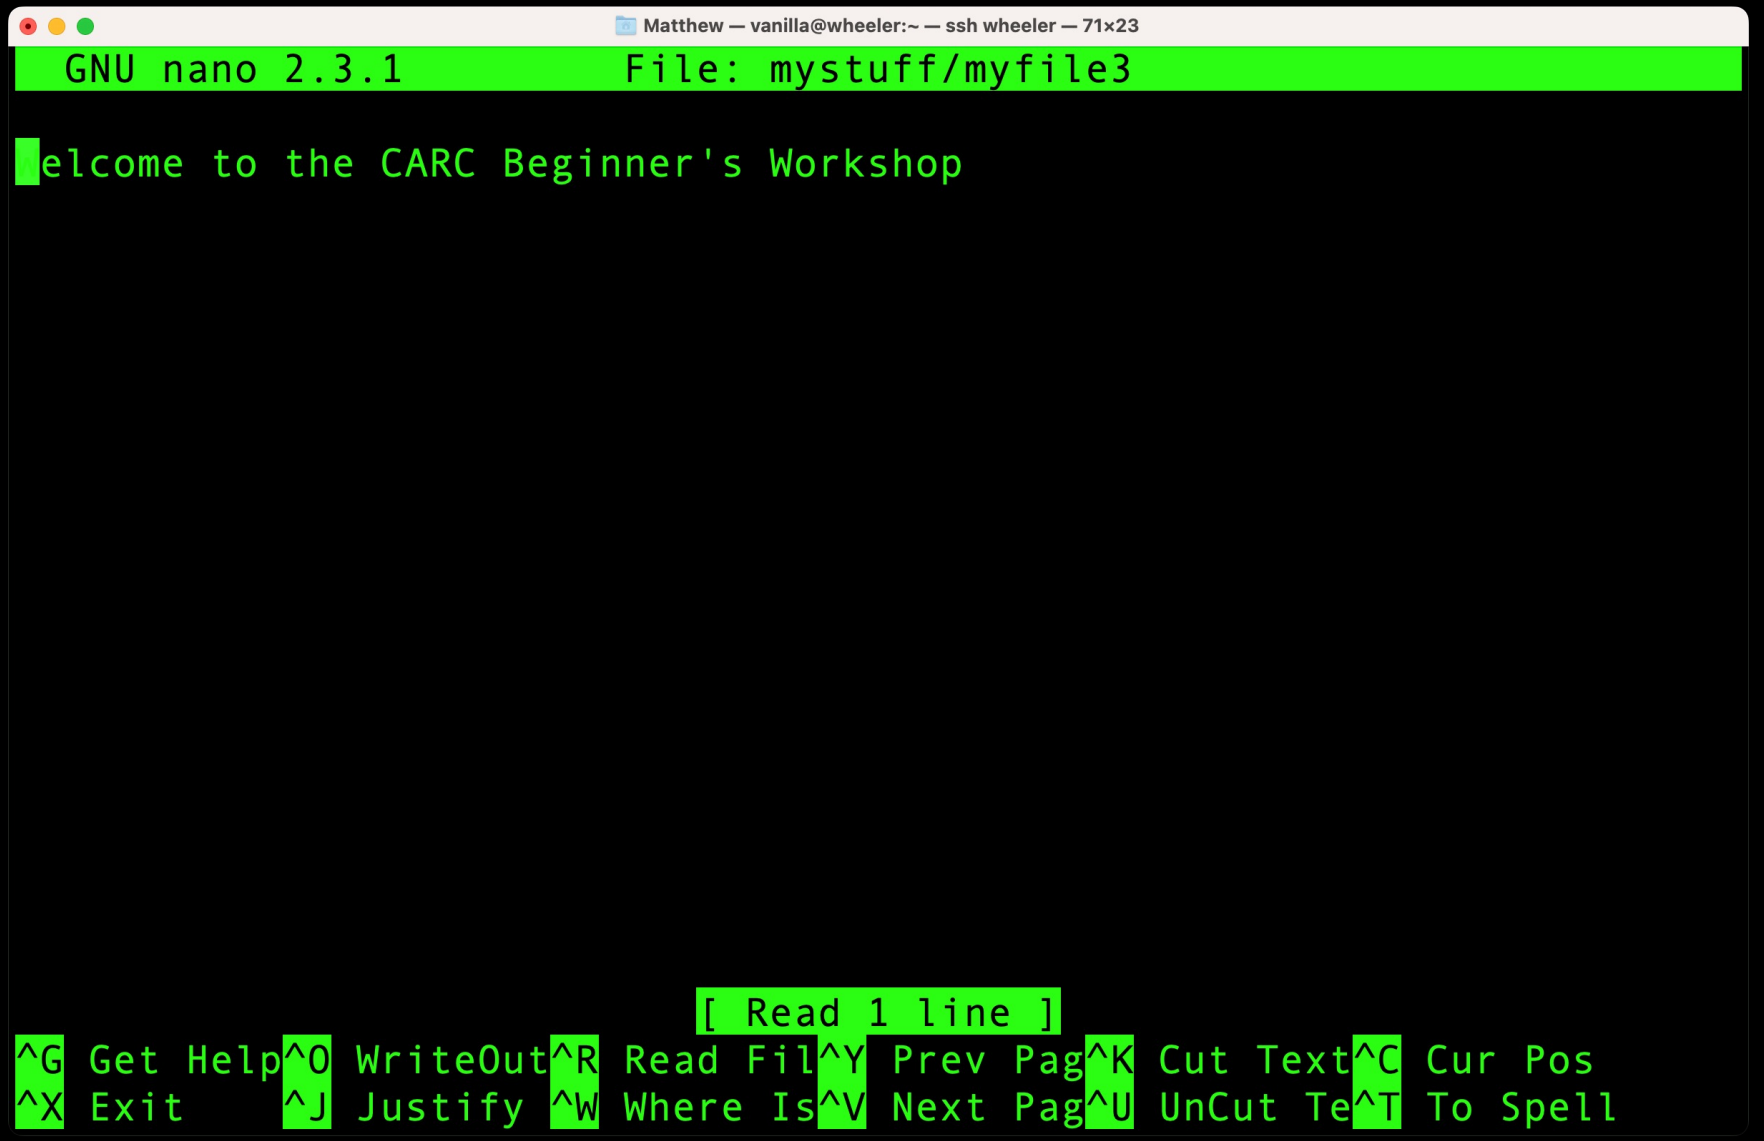

Software Access

Lmod Modules

Conda

```
[vanilla@wheeler ~]$ module spider matlab


----------------------------------------------------------------------------
-
  matlab:
----------------------------------------------------------------------------
-
     Versions:
        matlab/R2017a
        matlab/R2018b
        matlab/R2019a
        matlab/R2020a
        matlab/R2021a
```

Getting access to software…

```
[vanilla@wheeler ~]$ module load matlab/R2021a
Lmod has detected the following error: Matlab may only be run on compute
nodes. wheeler is not a compute node. Exiting...
While processing the following module(s):
    Module fullname  Module Filename
    ---------------  ---------------
    matlab/R2021a    /opt/local/modules/matlab/R2021a.lua
```

Getting access to software…

```
[vanilla@wheeler ~]$ module load matlab/R2021a
Lmod has detected the following error: Matlab may only be run on compute
nodes. wheeler is not a compute node. Exiting...
While processing the following module(s):
    Module fullname  Module Filename
    ---------------  ---------------
    matlab/R2021a    /opt/local/modules/matlab/R2021a.lua
```

What is a compute node?

Getting access to software…

HPC Cluster

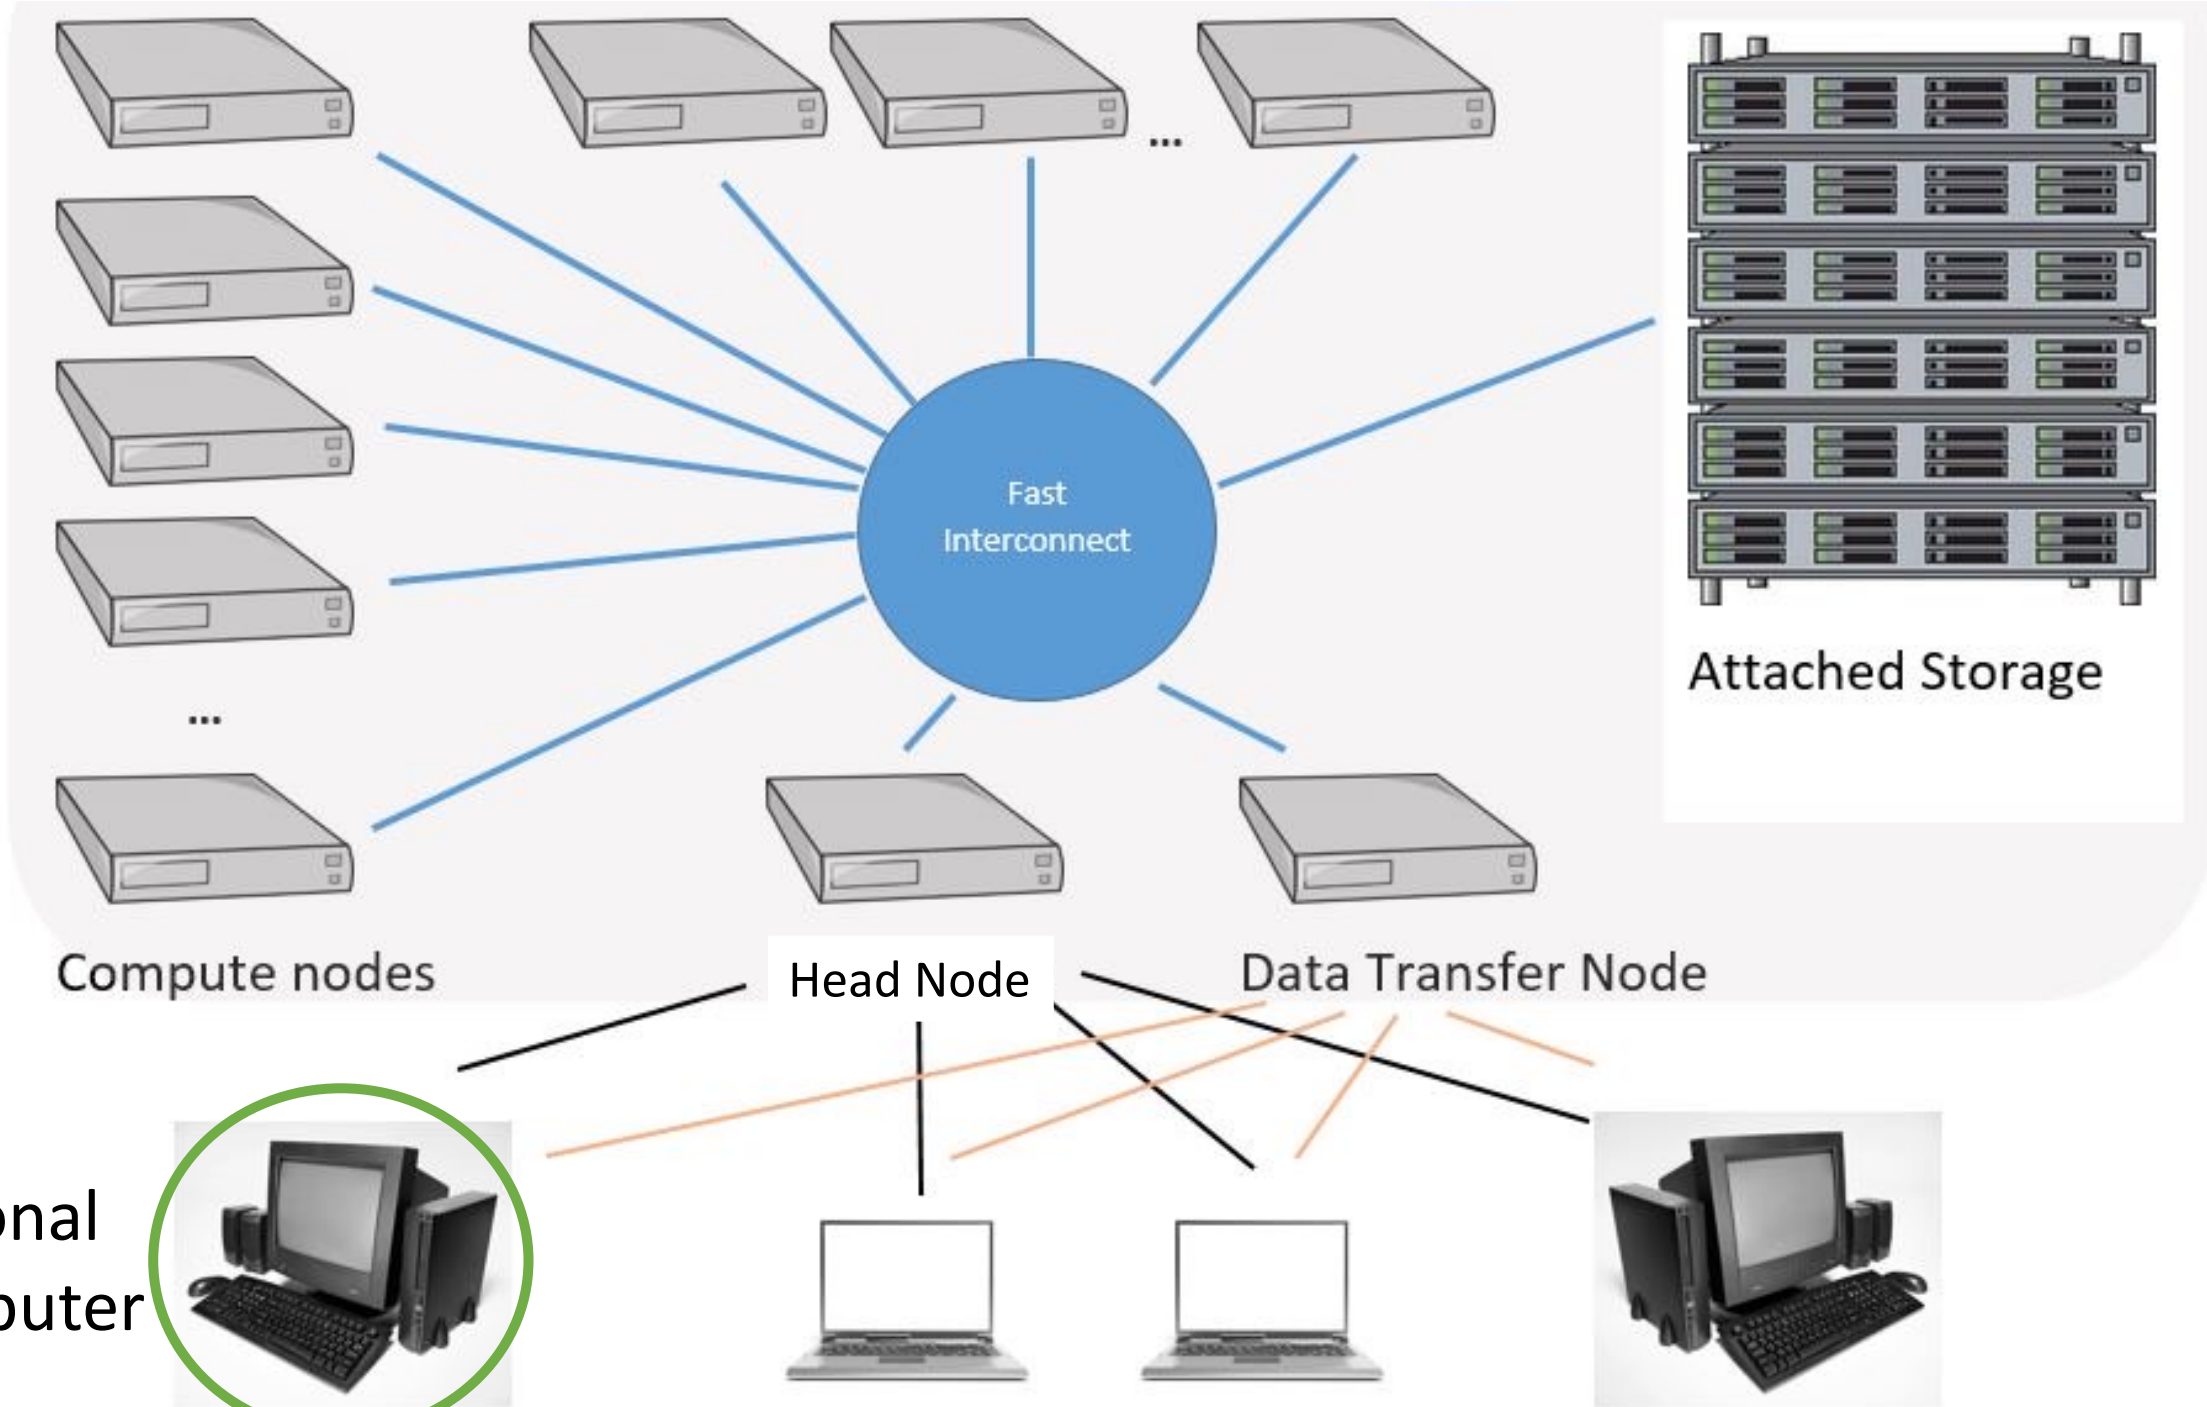
Fast
Interconnect
Attached Storage
Compute nodes
Head Node
Data Transfer Node
Personal
Computer

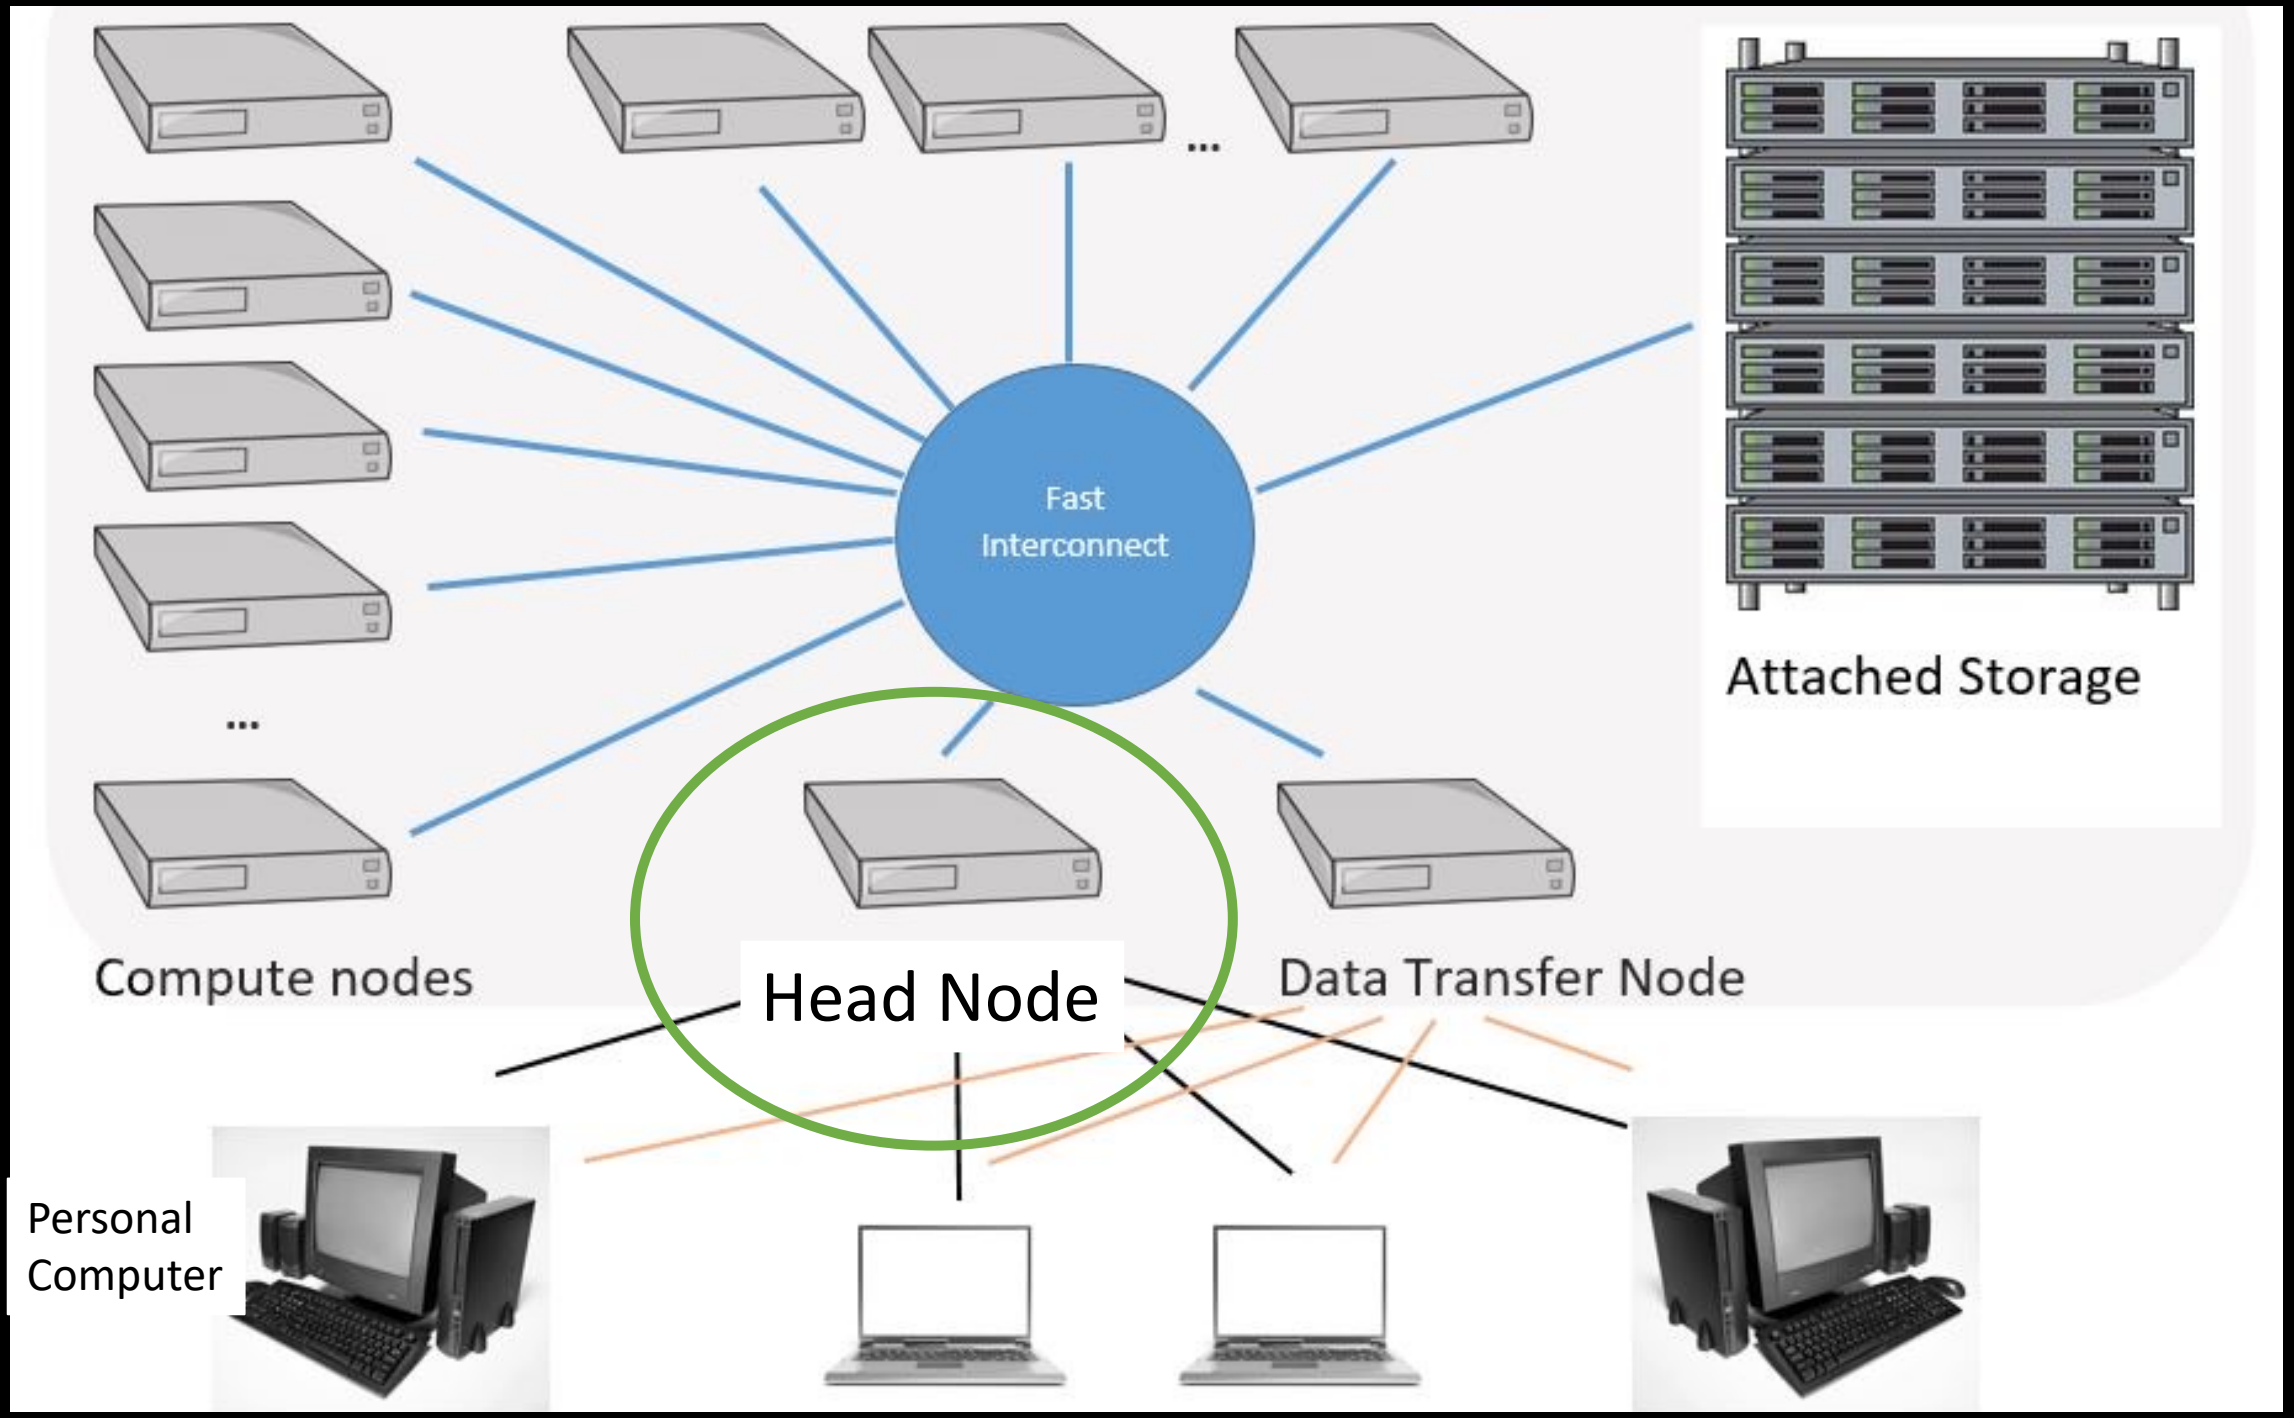
Fast
Interconnect
Attached Storage
Compute nodes
Head Node
Data Transfer Node
Personal
Computer

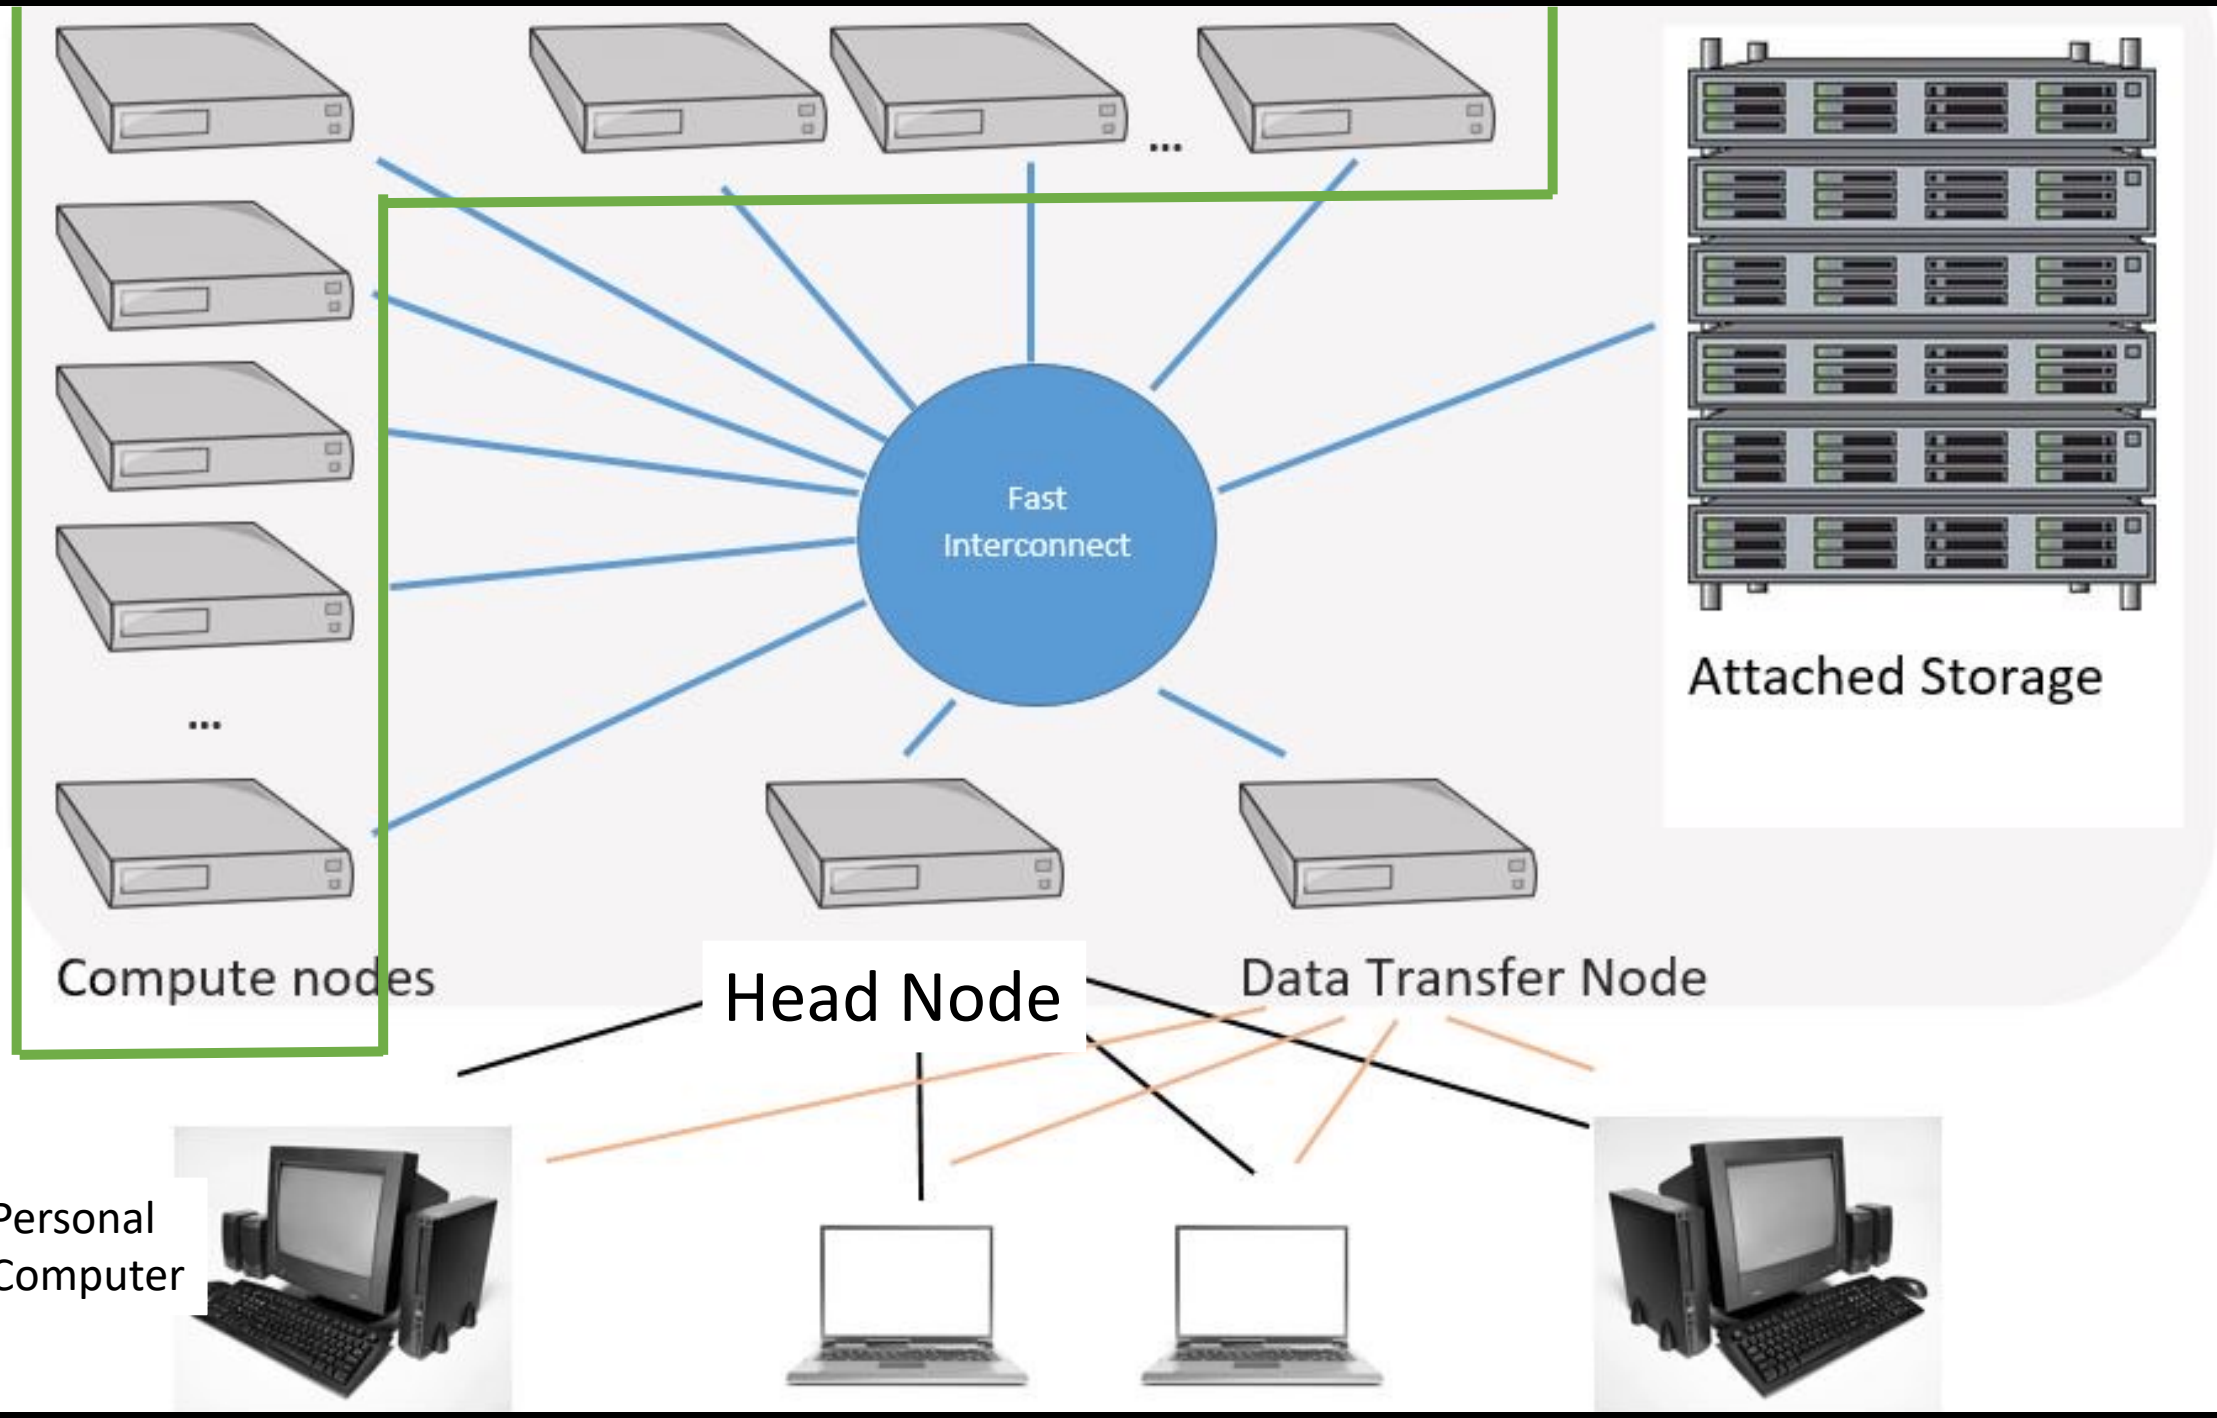
Fast
Interconnect
Attached Storage
Compute nodes
Head Node
Data Transfer Node
Personal
Computer

Never run computations on the head node

Always use compute nodes

```
[vanilla@wheeler ~]$ qgrok
queues    free    busy    offline   jobs    nodes   CPUs
------    -----   -----   -----     -----   -----   -----
normal    0       299     1         97      300     2400
debug     4       0       0         0       4       32
totals:   4       299     1         97      304     2432
```

Compute nodes and partitions…

```
qmfricke@hopper:~ $ qgrok
queues      free   busy   offline  jobs   nodes  CPUs
-----       -----  -----  -----    -----  -----  -----
general     12     0      0        0      12     384
condo       29     1      7        0      37     1184
bugs        2      0      0        0      2      64
pcnc        2      0      0        0      2      64
pathogen    0      1      0        1      1      32
tc          7      0      3        0      10     320
gold        2      0      0        0      2      64
fishgen     1      0      0        0      1      32
neuro-hsc   13     0      1        0      14     448
cup-ecs     2      0      0        0      2      64
tid         0      0      1        0      1      32
biocomp     0      0      1        0      1      32
chakra      0      0      1        0      1      32
totals:     70     2      14       1      86     2752
```

Compute nodes and partitions…

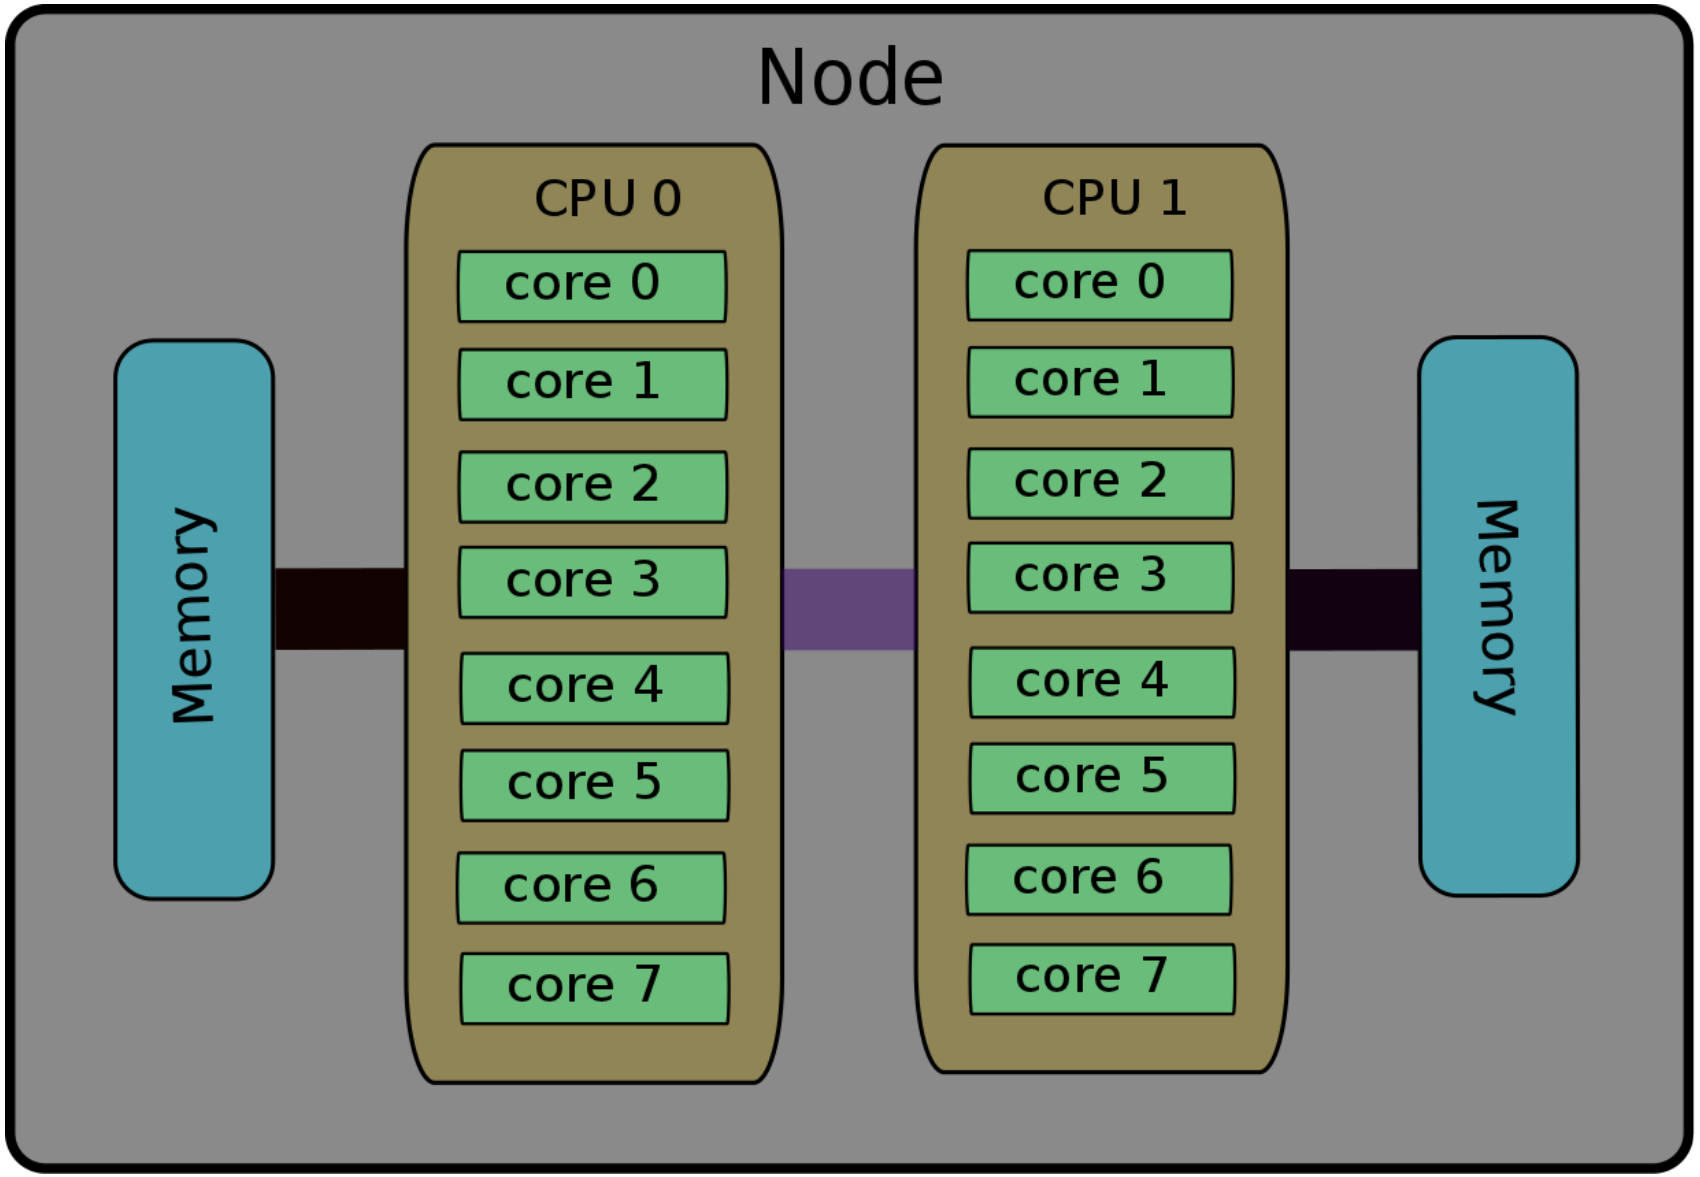
Node
CPU 0
core 0
core 1
core 2
core 3
core 4
core 5
core 6
core 7
CPU 1
core 0
core 1
core 2
core 3
core 4
core 5
core 6
core 7
Memory
Memory

Head Node (wheeler)

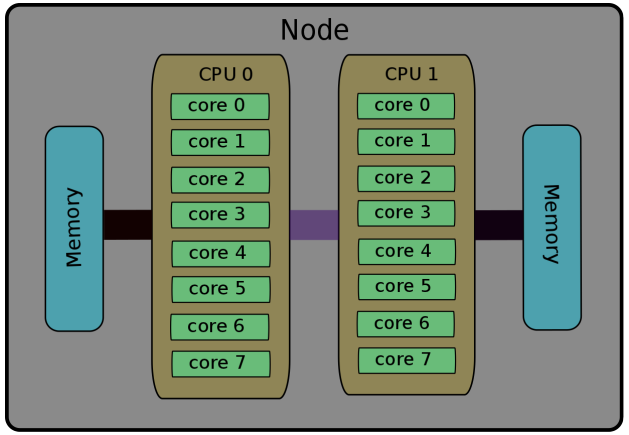

Compute Nodes

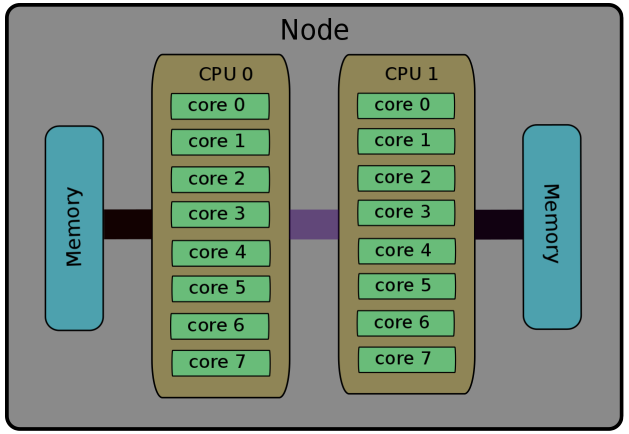

wheeler001

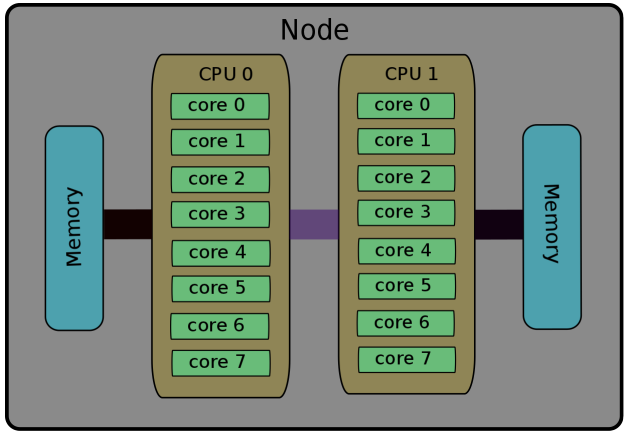

wheeler002

...

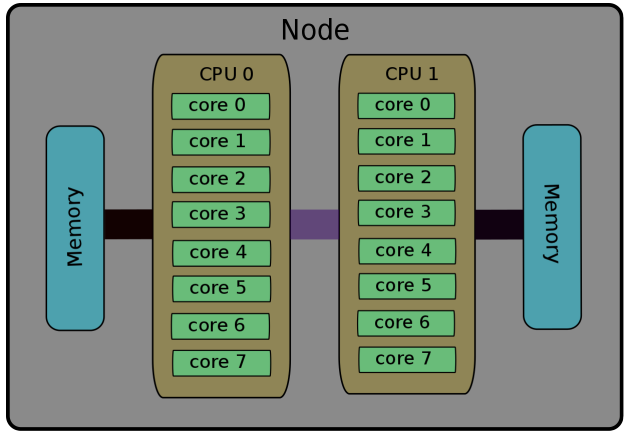

wheeler304

```
[vanilla@wheeler ~]$ lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                8
On-line CPU(s) list:   0-7
Thread(s) per core:    1
Core(s) per socket:    4
Socket(s):             2
NUMA node(s):          2
Vendor ID:             GenuineIntel
CPU family:            6
Model:                 26
Model name:            Intel(R) Xeon(R) CPU X5550  @ 2.67GHz
```

Wheeler has 8 cores per node….

```
mfricke@hopper:~ $ lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                64
On-line CPU(s) list:   0-63
Thread(s) per core:    2
Core(s) per socket:    16
Socket(s):             2
NUMA node(s):          2
Vendor ID:             GenuineIntel
CPU family:            6
Model:                 85
Model name:            Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz
```

Hopper has 32 cpus per node....

Architecture

Wheeler Cluster Grid > **Wheeler Cluster** > --Choose a Node

Overview of @ 2022-03-03 18:54

CPUs Total: **2440**
Hosts up: **305**
Hosts down: **0**

Current Load Avg (15, 5, 1m):
51%, 51%, 51%
Avg Utilization (last hour):
51%

Server Load Distribution

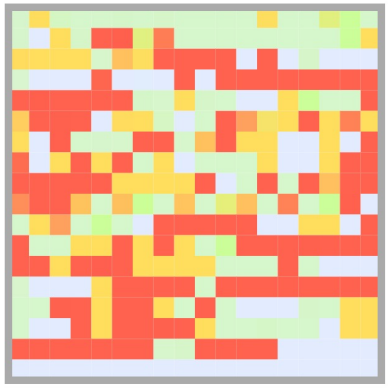

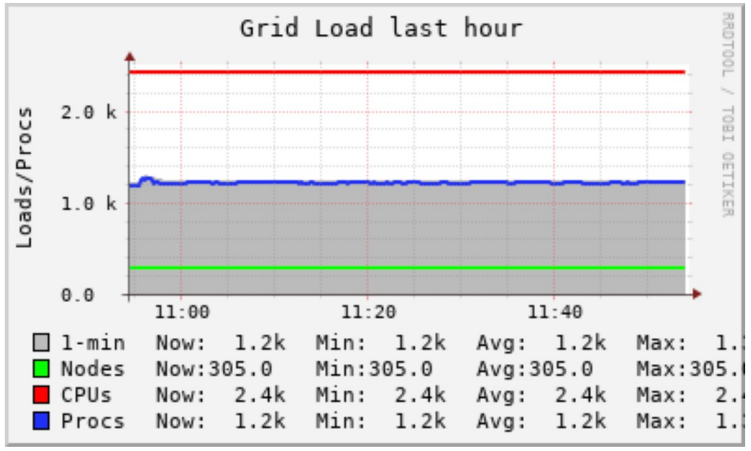

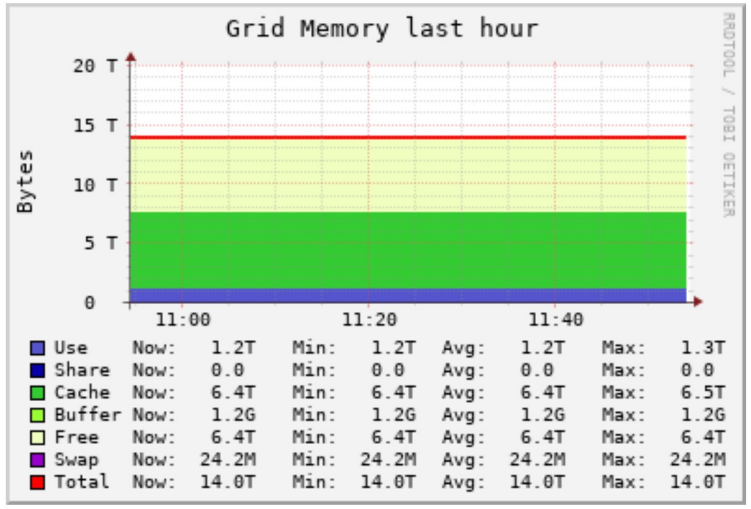

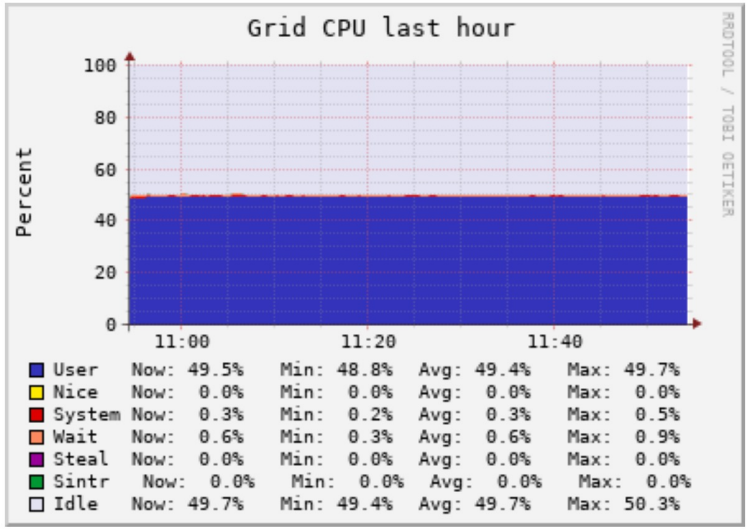

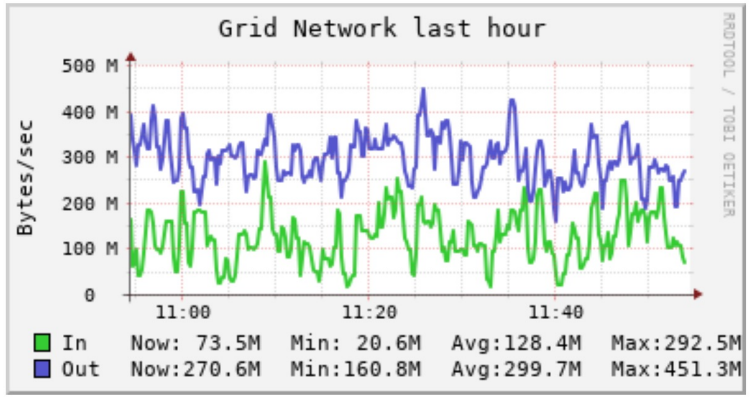

Technology, IT etc.

SLURM

means

Simple Linux Utility for Resource Management


ENJOY
Slurm
SODA
IT'S HIGHLY ADDICTIVE!
VOTED #1 SOFT DRINK OF THE 31ST CENTURY!

```
[vanilla@wheeler ~]$ squeue
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
            159914    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159915    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159916    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159917    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159918    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159919    normal co-mcpdf nsharma2 CG 2-00:00:26      1 wheeler257
            159912    normal co-mcpdf nsharma2 CG 2-00:00:28      1 wheeler257
            159913    normal co-mcpdf nsharma2 CG 2-00:00:28      1 wheeler257
166800_[21-100%10]    normal   Jannat   jannat PD       0:00      1 (JobArrayTaskLimit)
            167067    normal WINDENER rubeldas PD       0:00     36
(QOSMaxCpuPerUserLimit)
            167068    normal WINDENER rubeldas PD       0:00     24
```

Slurm….

```
[vanilla@wheeler ~]$ srun --partition debug --nodes 2 hostname
srun: Account not specified in script or
~/.default_slurm_account, using latest project
wheeler302.alliance.unm.edu
You have not been allocated GPUs. To request GPUs, use the -G
option in your submission script.
wheeler301.alliance.unm.edu
[vanilla@wheeler ~]$ srun --partition debug --nodes 2 hostname
```

The srun command…

```
[vanilla@wheeler ~]$ srun --partition debug --ntasks 8 hostname
srun: Account not specified in script or ~/.default_slurm_account, using
latest project
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
You have not been allocated GPUs. To request GPUs, use the -G option in your
submission script.
wheeler302.alliance.unm.edu
wheeler302.alliance.unm.edu
```

The srun command...

JupyterHub

python™

julia

jupyter Welcome to P

File Edit View Insert Cell

jupyter

Welcome to the

This Notebook Server wa

WARNING

Don't rely on this ser

Your server is hosted tha

Run some Python

To run the code below:

1. Click on the cell to se
2. Press SHIFT+ENTER

A full tutorial for using the

```
In [ ]: %matplotlib inline

        import pandas as pd
        import numpy as np
        import matplotlib
```

Exploring the Lorenz System

In this Notebook we explore the Lorenz system of differential equations:

$$\dot{x} = \sigma(y - x)$$
$$\dot{y} = \rho x - y - xz$$
$$\dot{z} = -\beta z + xy$$

This is one of the classic systems in non-linear differential equations. It exhibits a range of complex behaviors as the parameters (σ, β, ρ) are varied, including what are known as chaotic solutions. The system was originally developed as a simplified mathematical model for atmospheric convection in 1963.

```
In [7]: interact(Lorenz, N=fixed(10), angle=(0.,360.),
                σ=(0.0,50.0),β=(0.,5), ρ=(0.0,50.0))
```

angle 308.2

max_time 12

σ 10

β 2.6

ρ 28

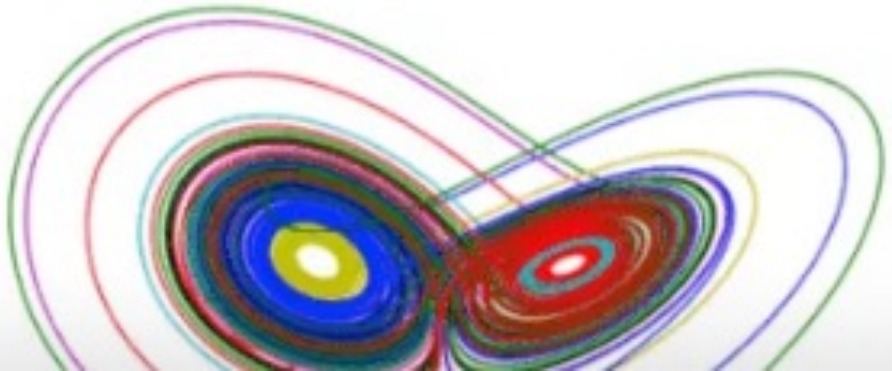

The Hopper Cluster

- Login to hopper.

- If you are already logged into wheeler this is just:

```
ssh hopper
```

If you are not then

```
ssh username@hopper.alliance.unm.edu
```

https://github.com/unm-escape/okada_jupyter

Product Team Enterprise Explore Marketplace Pricing Search / Sign in Sign up

unm-escape / okada_jupyter Public

Notifications

Code Issues Pull requests Actions Projects Wiki Security Insights

main 1 branch 0 tags

Go to file Code

ericlindsey Update README.md

AlaskaFaultsShapefiles	Initial commit of files	
Surface_Rupture_Ridgecrest_Prov...	Initial commit of files	
GPS_Alaska2021.csv	Initial commit of files	
GPS_Ridgecrest.csv	Initial commit of files	
README.md	Update README.md	
fault_model.py	Initial commit of files	
fault_plots.py	Initial commit of files	4 days ago
fit_fault_model.ipynb	Initial commit of files	4 days ago
geod_transform.py	Initial commit of files	4 days ago
moment_tensor.py	Initial commit of files	4 days ago
okada85.py	Initial commit of files	4 days ago
okada_greens.py	Initial commit of files	4 days ago

Clone

HTTPS GitHub CLI

https://github.com/unm-escape/okada_ju

Use Git or checkout with SVN using the web URL.

Open with GitHub Desktop

Download ZIP

About

Okada fault modeling to fit GPS data in Jupyter

Readme

1 star

1 watching

1 fork

Releases

No releases published

Packages

No packages published

Languages

Python 81.5%

Jupyter Notebook 18.5%

README.md

Jupyterlab example for Okada fault modeling

unm-escape / okada_jupyter

Public

Notifications

Code | Issues | Pull requests | Actions | Projects | Wiki | Security | Insights

main | 1 branch | 0 tags | Go to file | Code

ericlindsey Update README.md

AlaskaFaultsShapefiles	Initial commit of files	
Surface_Rupture_Ridgecrest_Prov...	Initial commit of files	
GPS_Alaska2021.csv	Initial commit of files	
GPS_Ridgecrest.csv	Initial commit of files	
README.md	Update README.md	
fault_model.py	Initial commit of files	
fault_plots.py	Initial commit of files	4 days ago
fit_fault_model.ipynb	Initial commit of files	4 days ago
geod_transform.py	Initial commit of files	4 days ago
moment_tensor.py	Initial commit of files	4 days ago
okada85.py	Initial commit of files	4 days ago
okada_greens.py	Initial commit of files	4 days ago

Clone

HTTPS | GitHub CLI

https://github.com/unm-escape/okada_ju

Use Git or checkout with SVN using the web URL.

Open with GitHub Desktop

Download ZIP

About

Okada fault modeling to fit GPS data in Jupyter

Readme
1 star
1 watching
1 fork

Releases

No releases published

Packages

No packages published

Languages

Python 81.5%
Jupyter Notebook 18.5%

README.md

Jupyterlab example for Okada fault modeling

```
[vanilla@hopper ~]$ git clone https://lobogit.unm.edu/CARC/workshops.git
Cloning into 'workshops'...
remote: Enumerating objects: 132, done.
remote: Counting objects: 100% (75/75), done.
remote: Compressing objects: 100% (43/43), done.
remote: Total 132 (delta 33), reused 74 (delta 32), pack-reused 57
Receiving objects: 100% (132/132), 57.58 KiB | 3.60 MiB/s, done.
Resolving deltas: 100% (51/51), done.
```

We will use a popular tool called "git" to download examples…

Dependencies

- The example code depends on these libraries:
- numpy
- matplotlib
- H5py
- Scipy
- simplekml

- Using it with jupyterhub requires ipykernel and jupyter_client

```
vanilla@hopper:~ $ module load miniconda3
vanilla@hopper:~ $ conda create --name escape numpy
scipy jupyter_client ipykernel h5py matplotlib
simplekml --channel conda-forge

Collecting package metadata (current_repodata.json): done
Solving environment: done
## Package Plan ##
  environment location: /users/mfricke/.conda/envs/fit_fault_model


Proceed ([y]/n)?
```

```
Downloading and Extracting Packages
entrypoints-0.4       | 16 KB     |
################################################################################ | 100%
bottleneck-1.3.5      | 274 KB    |
################################################################################ | 100%
[SNIP]
brotli-bin-1.0.9      | 19 KB     |
################################################################################ | 100%
Preparing transaction: done
Verifying transaction: done
Executing transaction: done
#
# To activate this environment, use
#
#     $ conda activate escape
#
# To deactivate an active environment, use
#
#     $ conda deactivate

vanilla@hopper:~ $
```

Welcome to the Center for Advanced Research Computing

The UNM Center for Advanced Research Computing is the hub of computational research at UNM and one of the largest computing centers in the State of New Mexico. It is an interdisciplinary community that uses computational resources to create new research insights. The goal is to lead and grow the computational research community at UNM.

CARC provides not just the computing resources but also the expertise and support to help the university's researchers. This service is available to faculty, staff, and student researchers free of charge through support from the UNM Office of the Vice President for Research.

Systems menu:

- CARC systems information
- CARC infrastructure
- Resource limits
- Storage and backup
- Export control
- JupyterHub cluster links
- System status and downtime
- System Usage by Principal Investigator

JupyterHub cluster links:

- Hopper
- Gibbs
- Taos
- Wheeler
- Xena

JupyterHub cluster links

Get sta

Create a help ticket

Downtime notices

Students go on-si... greenhouse gas e...

UNM Earth Science Professor to study the health of groundwater aquifers across the state

Explore or exploit: How our brains make choices

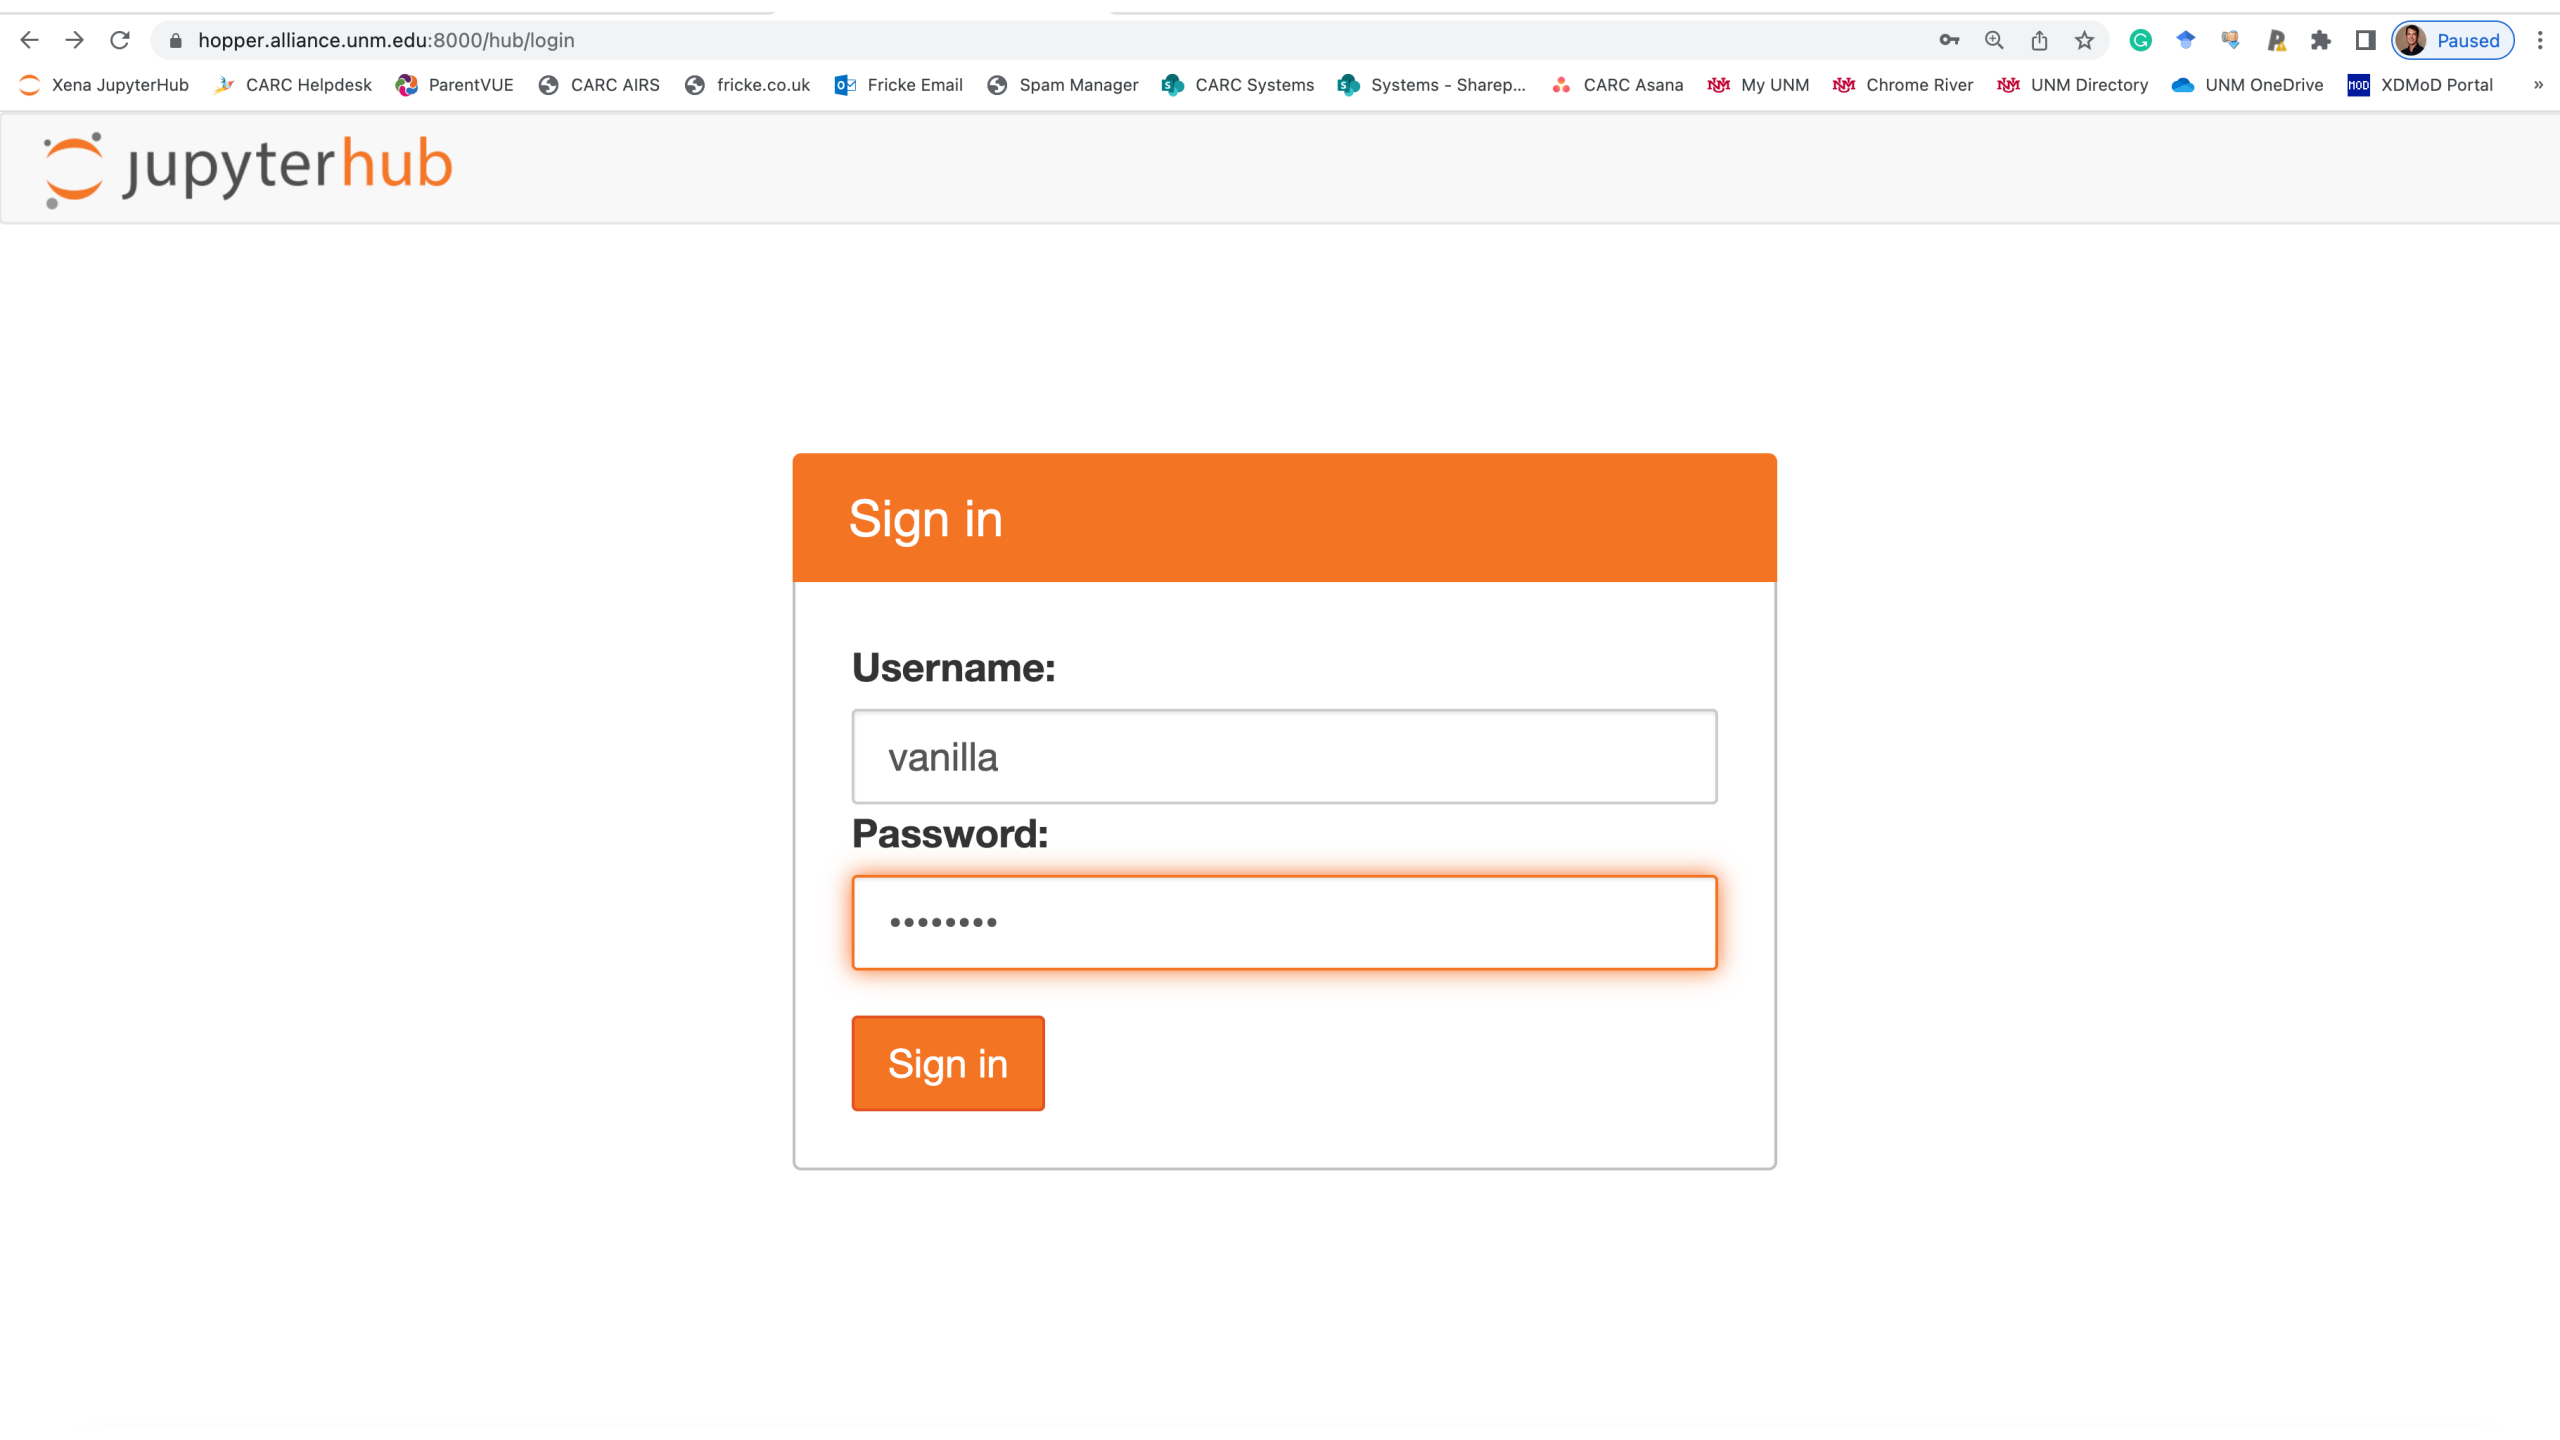
hopper.alliance.unm.edu:8000/hub/login
Paused
Xena JupyterHub
CARC Helpdesk
ParentVUE
CARC AIRS
fricke.co.uk
Fricke Email
Spam Manager
CARC Systems
Systems - Sharep...
CARC Asana
My UNM
Chrome River
UNM Directory
UNM OneDrive
XDMoD Portal
jupyterhub
Sign in
Username:
vanilla
Password:
Sign in

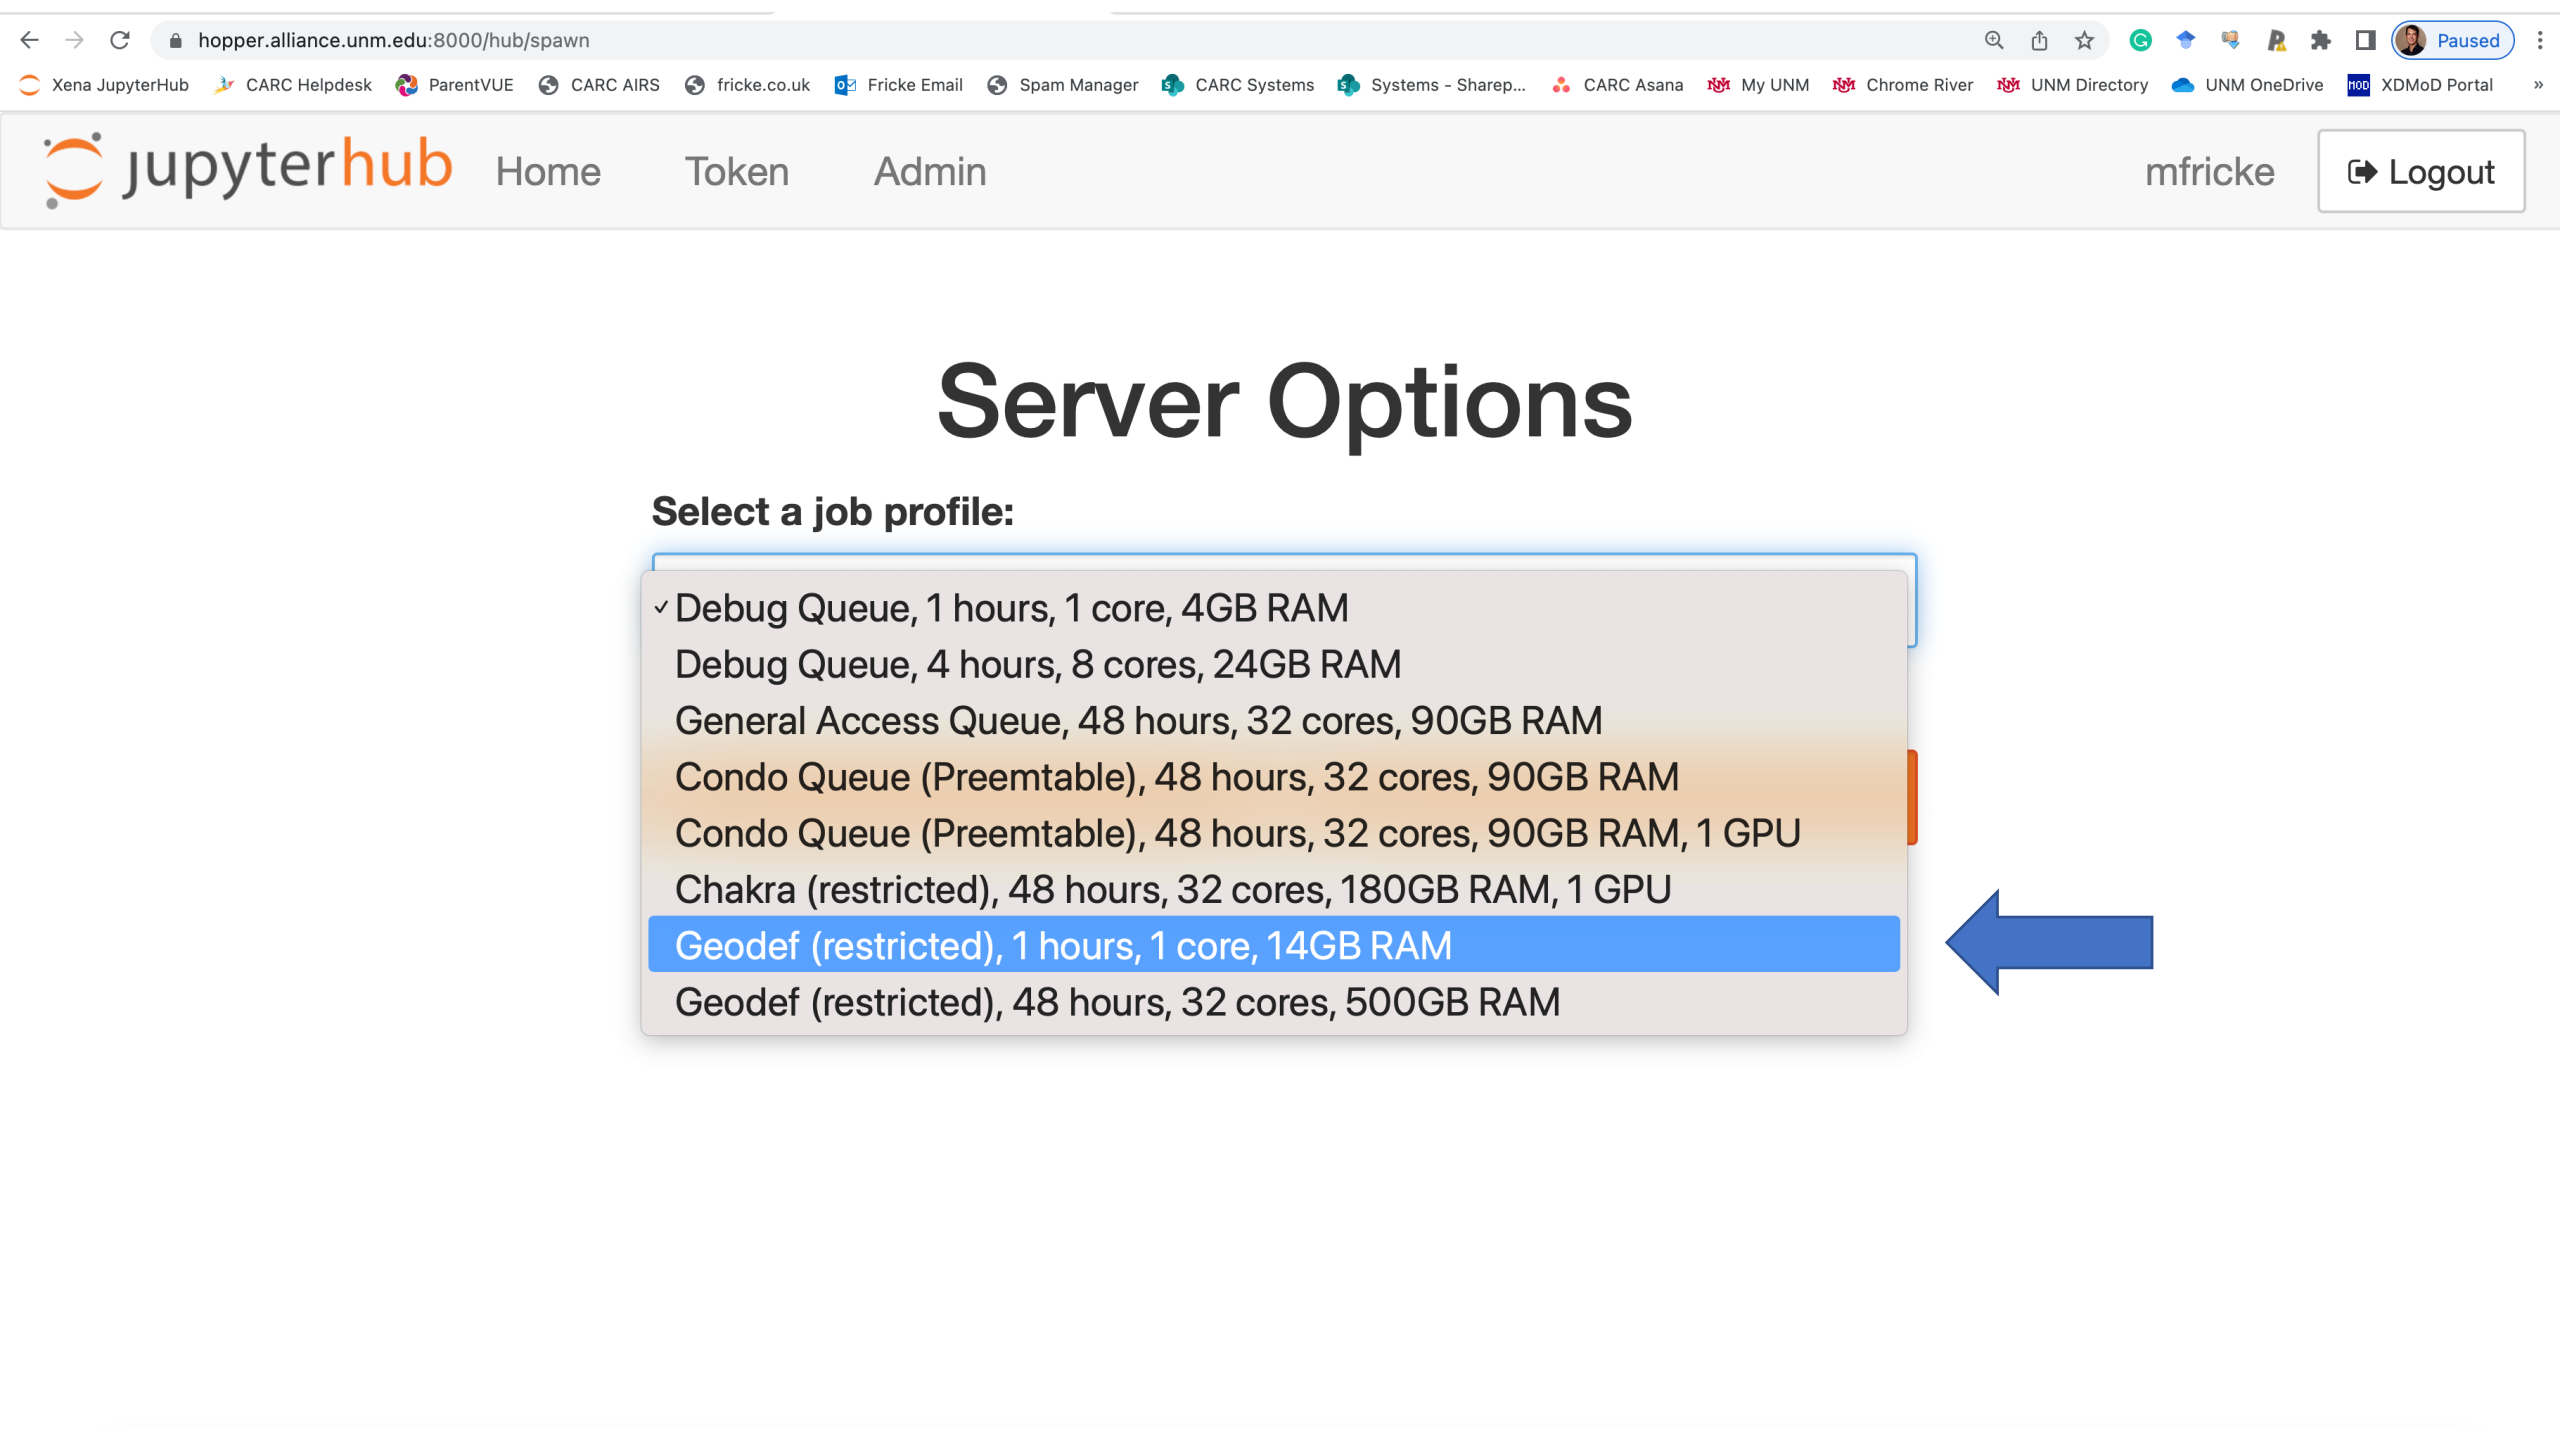
hopper.alliance.unm.edu:8000/hub/spawn
Paused
Xena JupyterHub
CARC Helpdesk
ParentVUE
CARC AIRS
fricke.co.uk
Fricke Email
Spam Manager
CARC Systems
Systems - Sharep...
CARC Asana
My UNM
Chrome River
UNM Directory
UNM OneDrive
XDMoD Portal
jupyterhub
Home
Token
Admin
mfricke
Logout
Server Options
Select a job profile:
Debug Queue, 1 hours, 1 core, 4GB RAM
Debug Queue, 4 hours, 8 cores, 24GB RAM
General Access Queue, 48 hours, 32 cores, 90GB RAM
Condo Queue (Preemtable), 48 hours, 32 cores, 90GB RAM
Condo Queue (Preemtable), 48 hours, 32 cores, 90GB RAM, 1 GPU
Chakra (restricted), 48 hours, 32 cores, 180GB RAM, 1 GPU
Geodef (restricted), 1 hours, 1 core, 14GB RAM
Geodef (restricted), 48 hours, 32 cores, 500GB RAM

Your server is starting up.

You will be redirected automatically when it's ready for you.

Cluster job running... waiting to connect

Event log

Logout | Control Panel

Files | Running | Clusters

Select items to perform actions on them.

Upload | New | 

Notebook:
- Python 3
- Python [conda env:.conda-escape]

Other:
- Text File
- Folder
- Terminal

0 / ze

carc-scratch		
mystuff		
simcov		
vecadd		
wheeler-scratch	a year ago	
workshops	11 minutes ago	
hopper-scratch	a year ago	23 B
jupyterhub_wheeler_batchspawner_552957.log	7 minutes ago	85 B

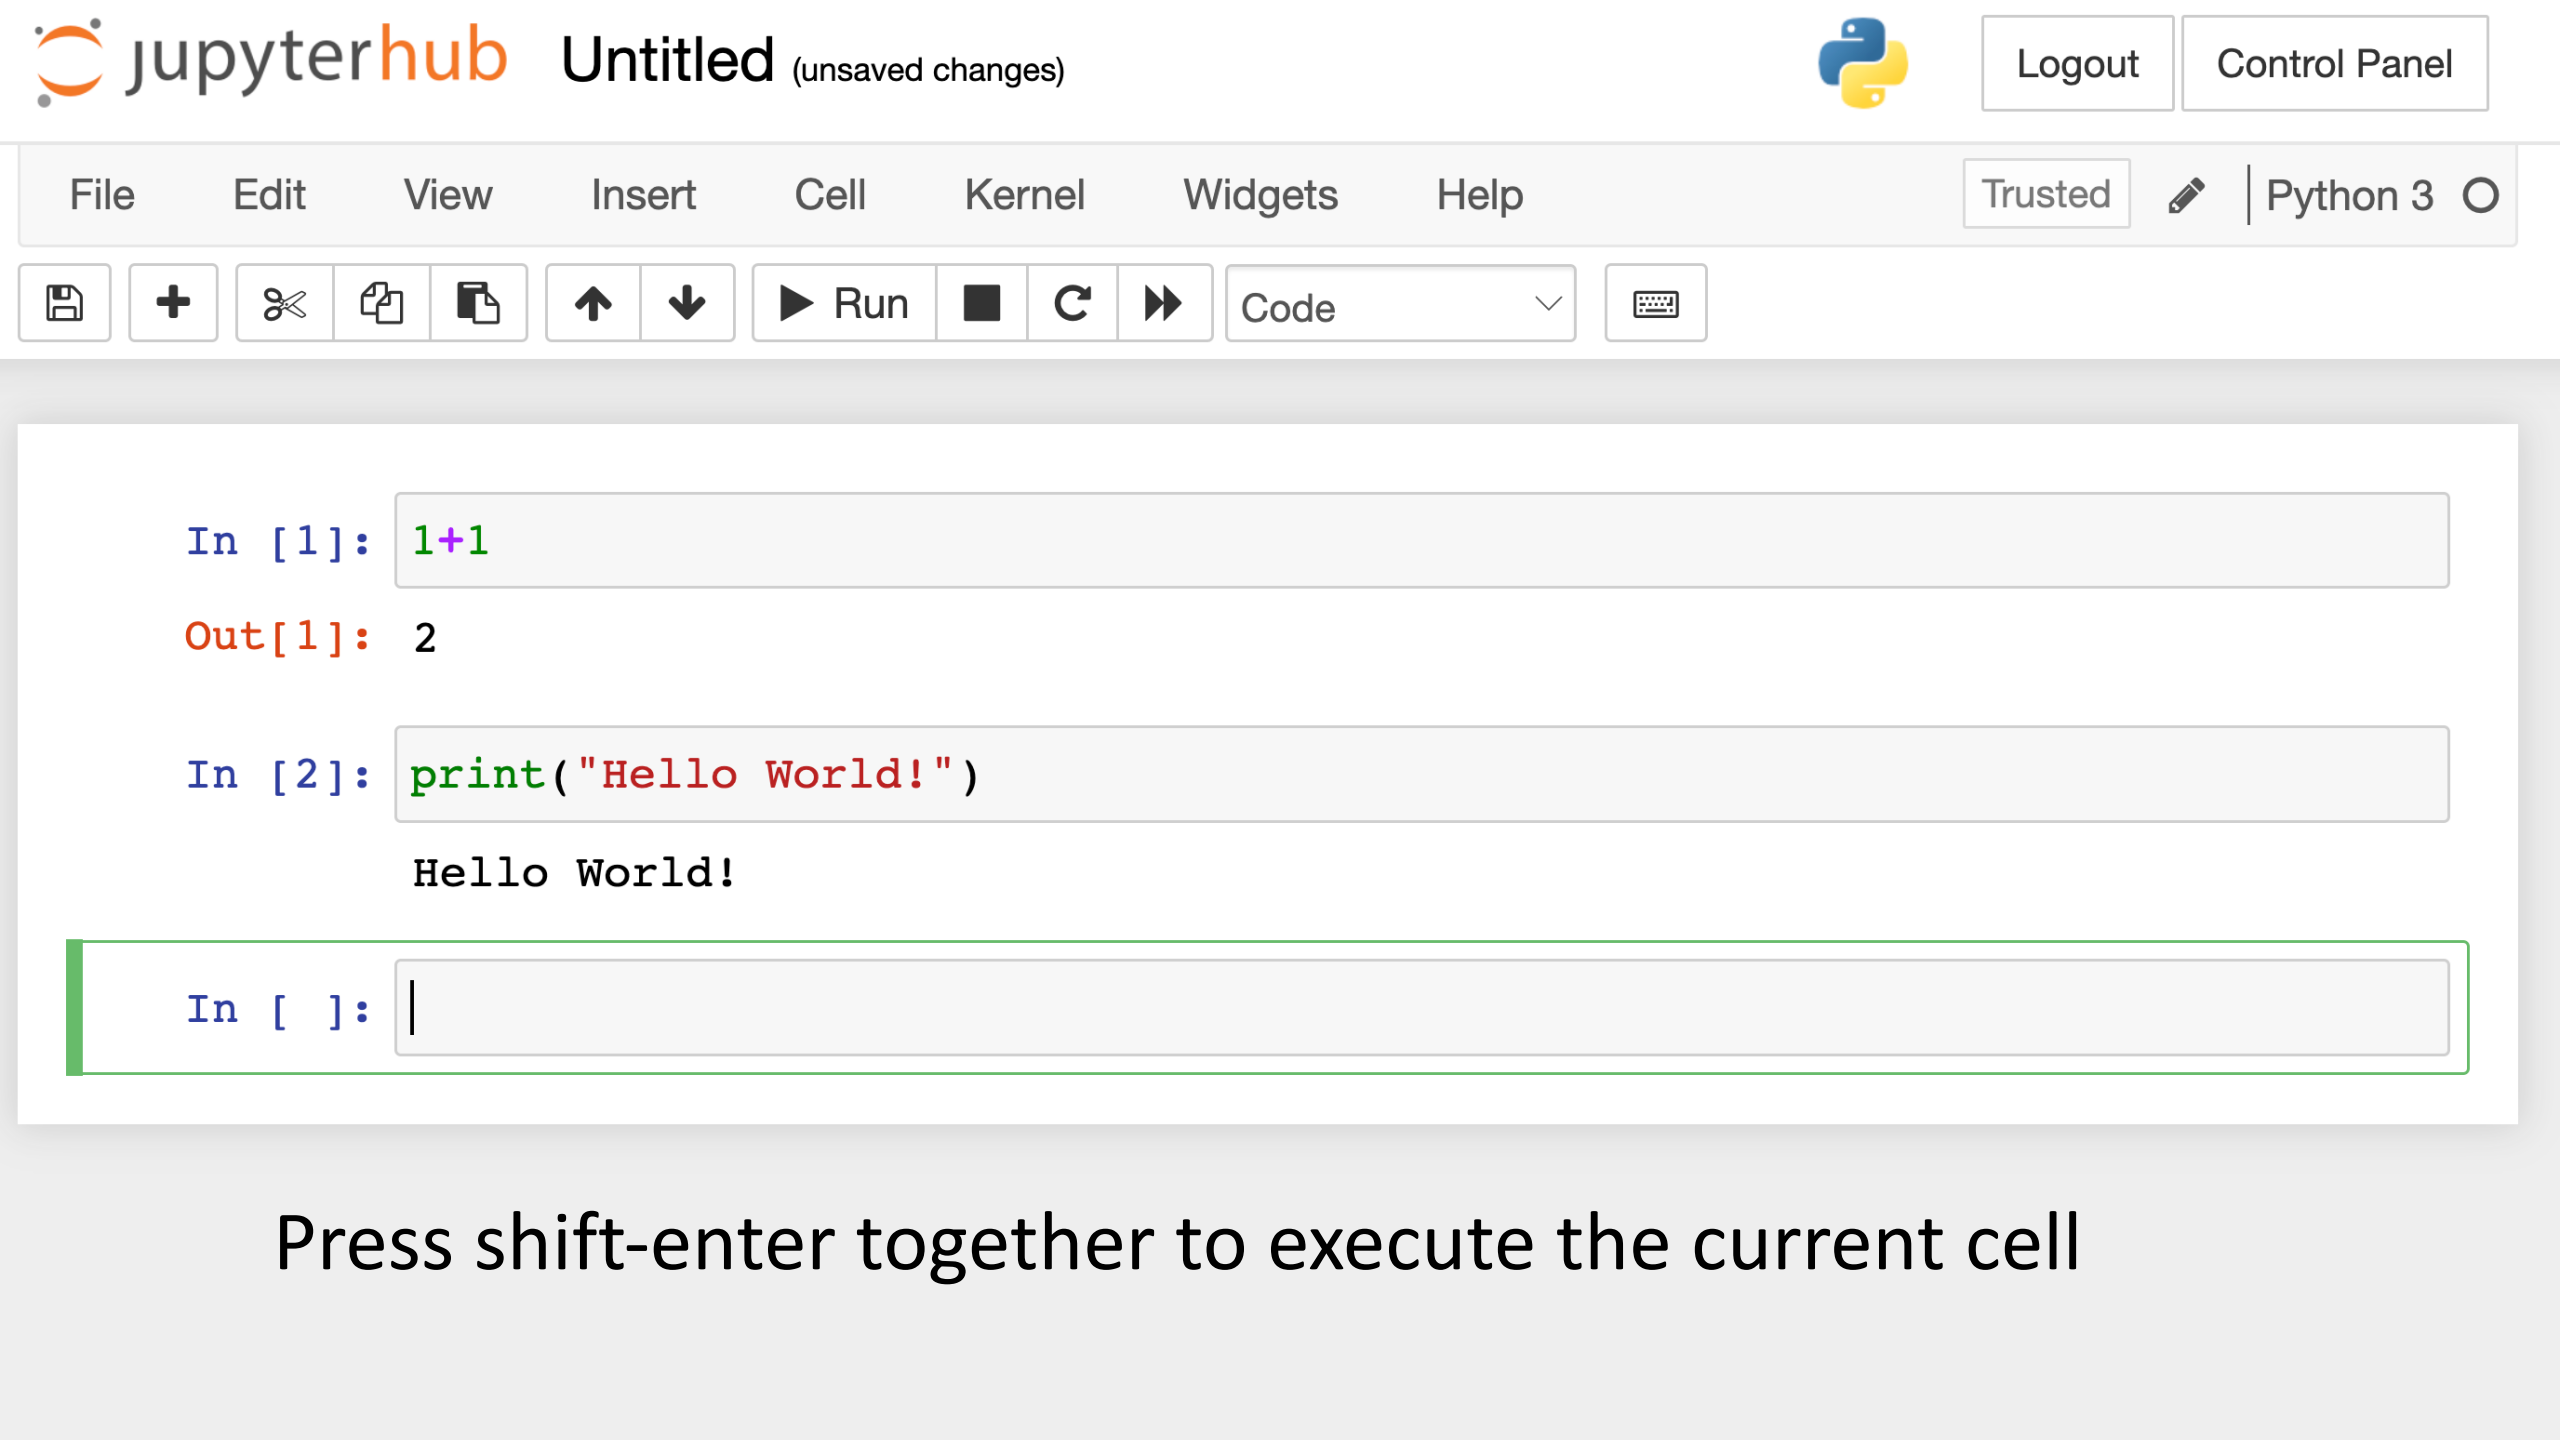

Press shift-enter together to execute the current cell

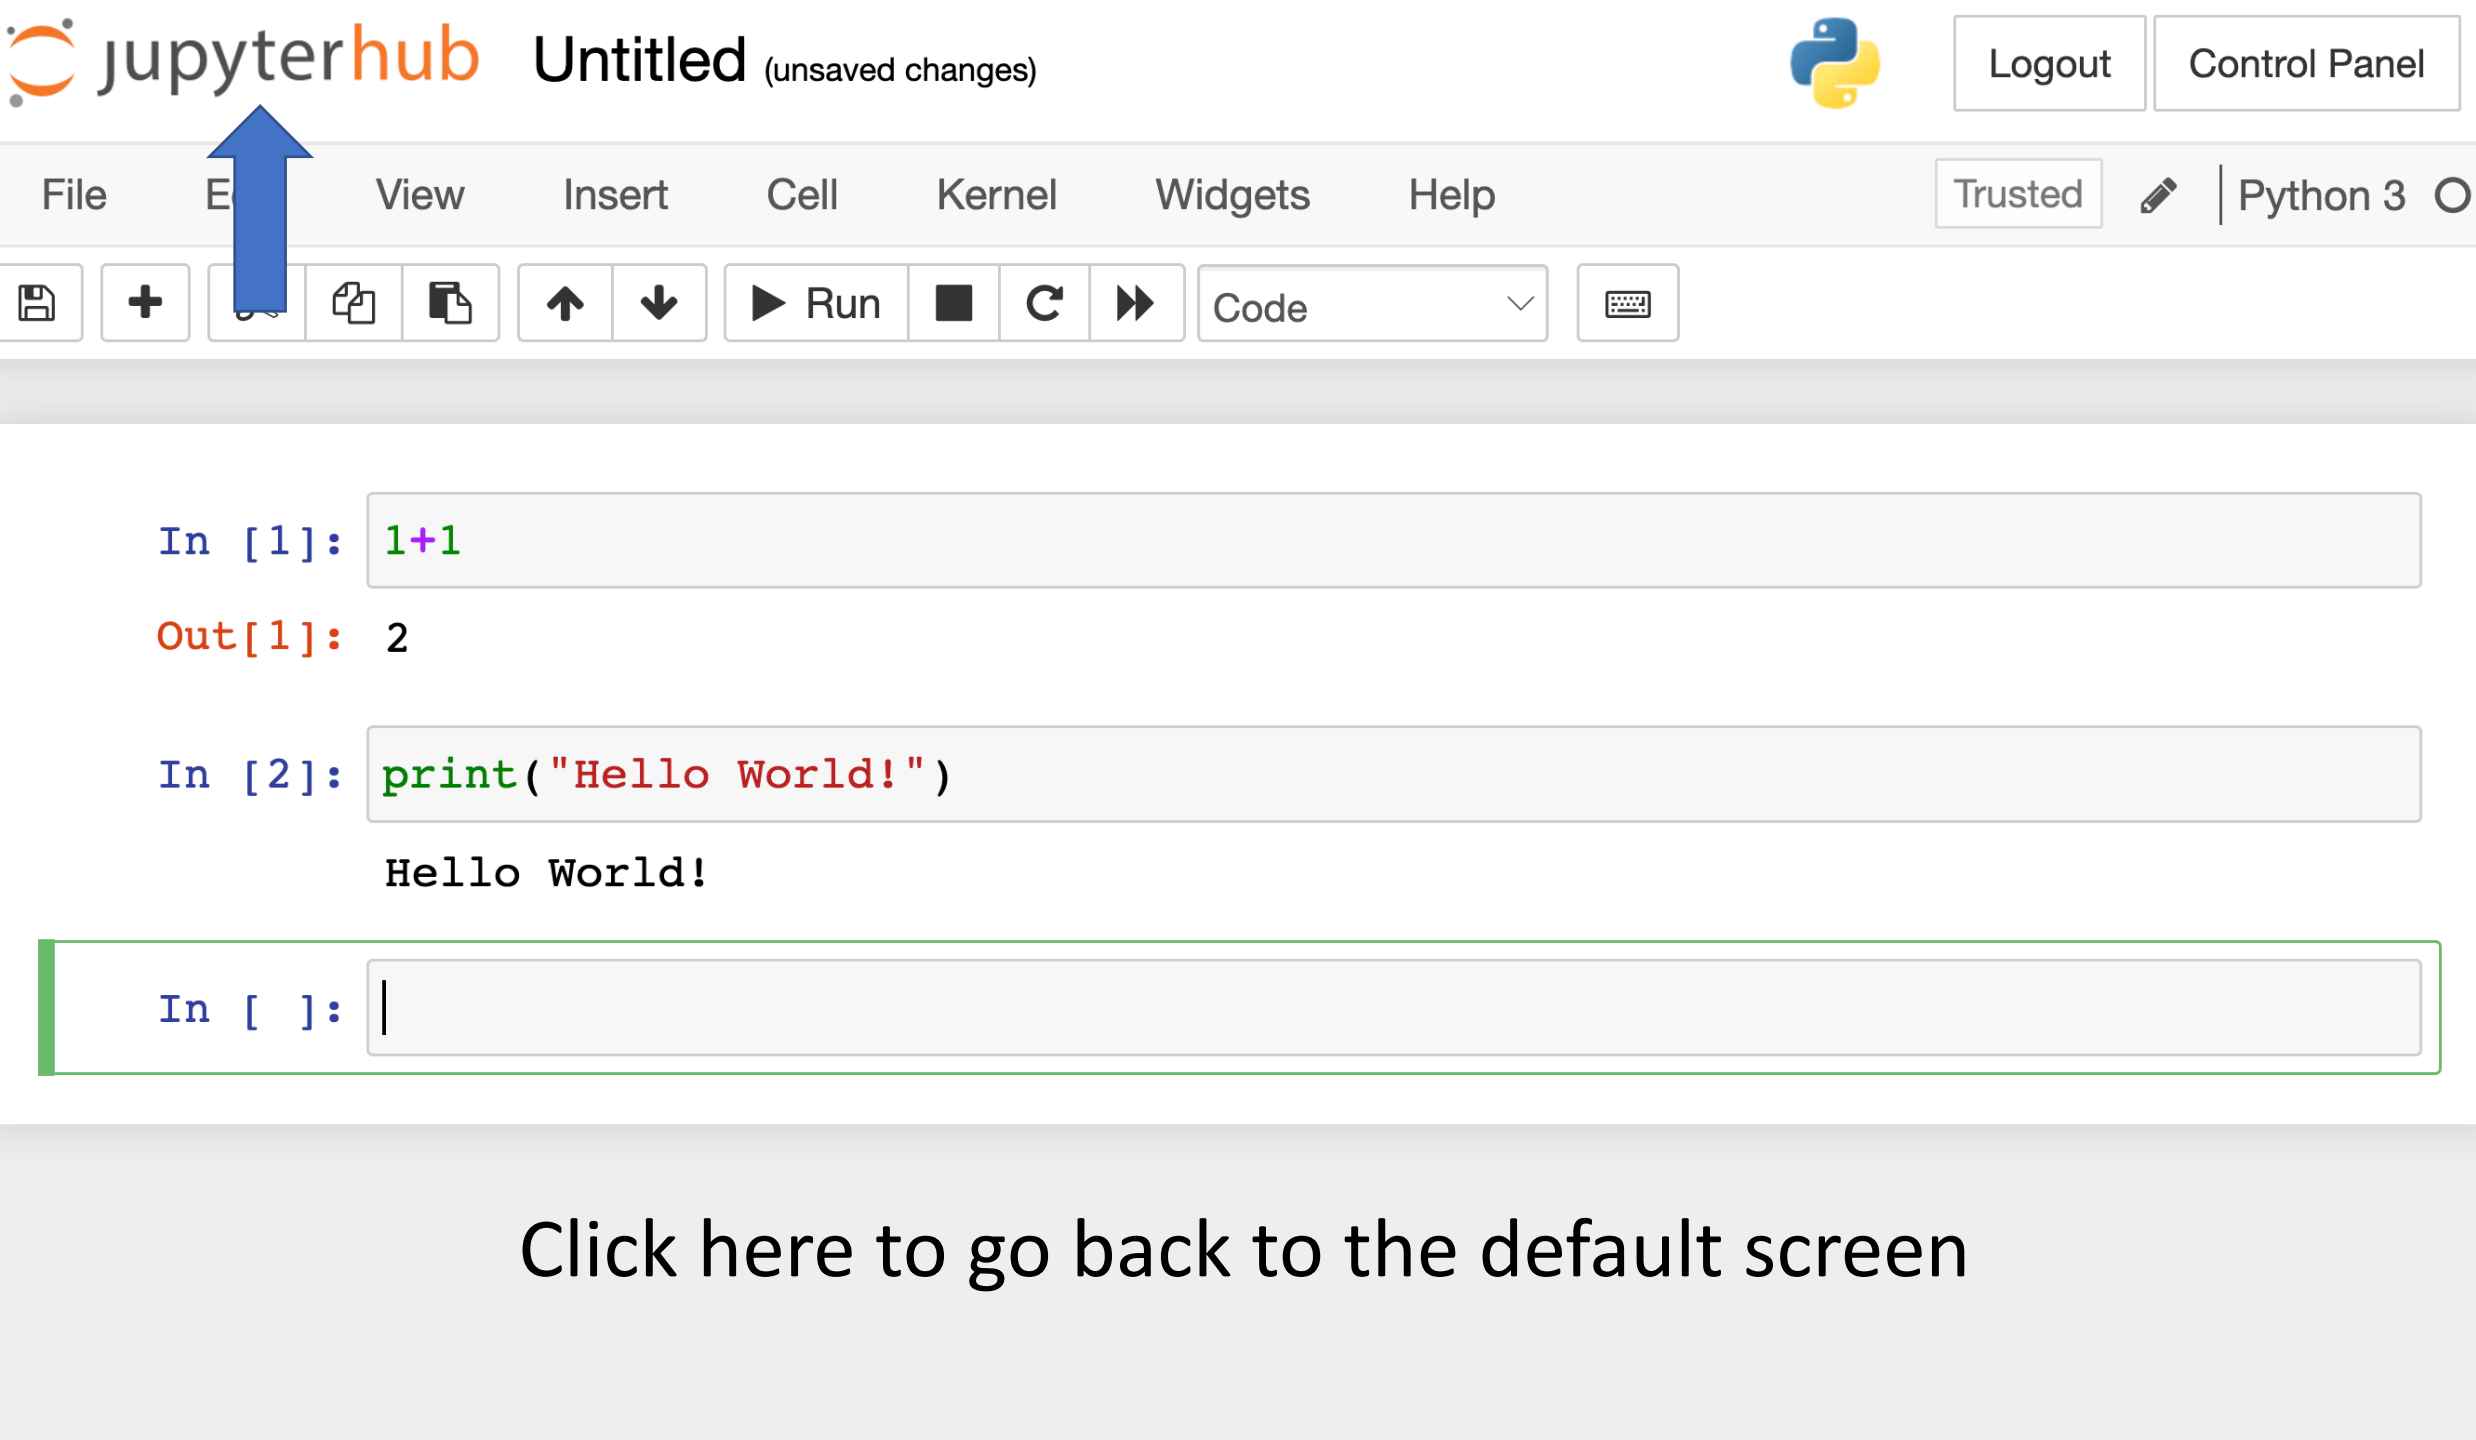
jupyterhub
Untitled (unsaved changes)
Logout
Control Panel
File
E
View
Insert
Cell
Kernel
Widgets
Help
Trusted
Python 3
Run
Code
In [1]: 1+1
Out[1]: 2
In [2]: print("Hello World!")
Hello World!
In [ ]:
Click here to go back to the default screen

jupyterhub

Logout | Control Panel

Files | Running | Clusters

Select items to perform actions on them.

Upload | New | 

0 /	Name	Last Modified	File size
carc-scratch		a year ago	
mystuff		a year ago	
simcov		7 months ago	
vecadd		a year ago	
wheeler-scratch		a year ago	
workshops		seconds ago	
hopper-scratch		a year ago	23 B
jupyterhub_wheeler_batchspawner_552957.log		22 minutes ago	0 B
jupyterhub_wheeler_slurm_cmd.err		in a few seconds	5.57 kB
jupyterhub_wheeler_slurm_cmd.out		22 minutes ago	0 B

jupyterhub

Logout | Control Panel

Files | Running | Clusters

Select items to perform actions on them.

Upload | New | 

0 / workshops

	Name	Last Modified	File size
..		seconds ago	
escape		a minute ago	
intro_workshop		a minute ago	
namd		a minute ago	
radio_astronomy		a minute ago	
README.md		a minute ago	47 B

jupyterhub

Logout | Control Panel

Files | Running | Clusters

Select items to perform actions on them.

Upload | New | ↻

0 / workshops / escape

	Name	Last Modified	File size
..		seconds ago	
fit_insar_data.ipynb		2 minutes ago	14 MB
kml_functions.py		2 minutes ago	4.41 kB

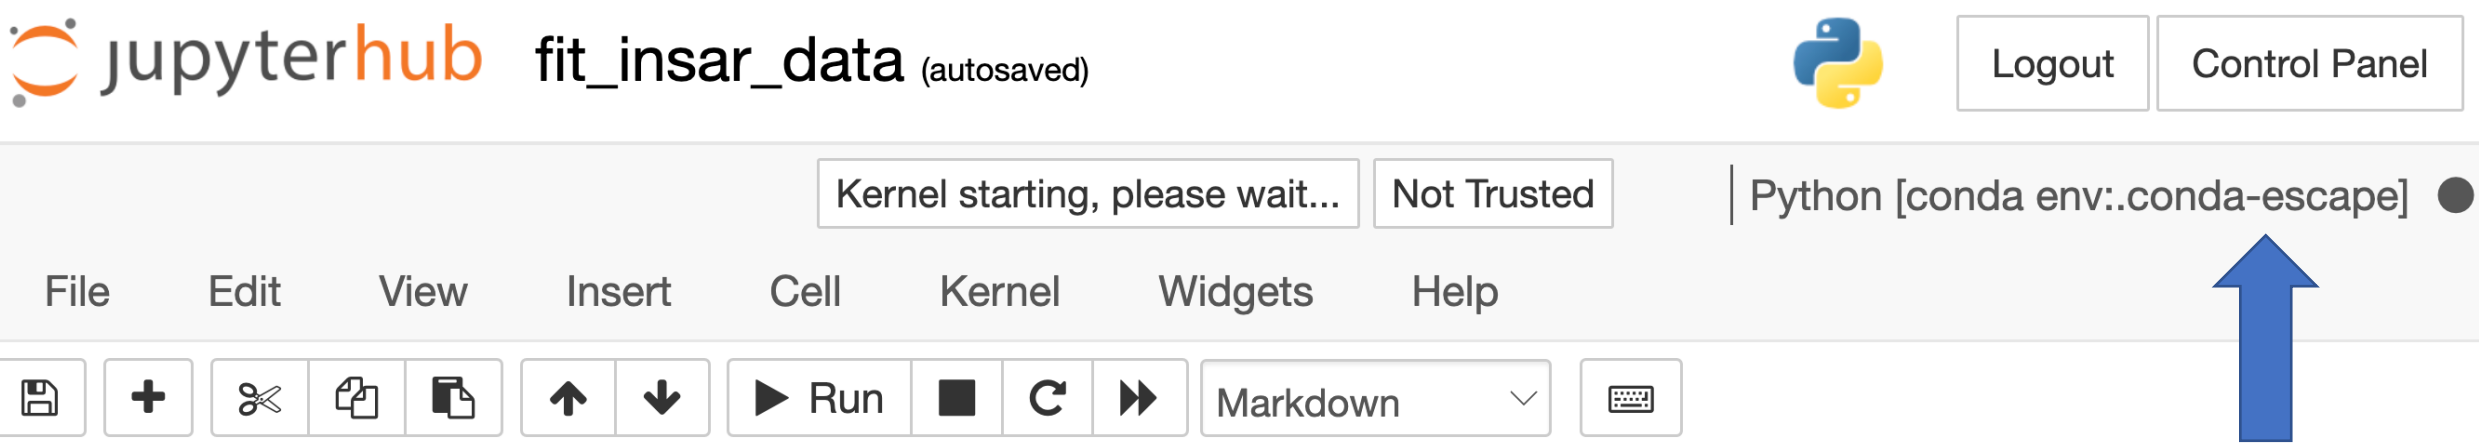

Viewing and interacting with a large InSAR dataset on CARC

The InSAR dataset "geo_timeseries_ERA5_ramp_demErr_msk.h5" contains the complete time-series of deformation between 2014 and 2023 over the city of Jakarta, Indonesia. The file "geo_velocity_msk.h5" contains the average rate of deformation at each pixel. Together, these datasets are about 17 GB in size, so it's inconvenient to download and work with them on our laptop. Thankfully, we can use CARC to make life easier!

In [1]:
```
import h5py
import matplotlib.pyplot as plt
import numpy as np
```

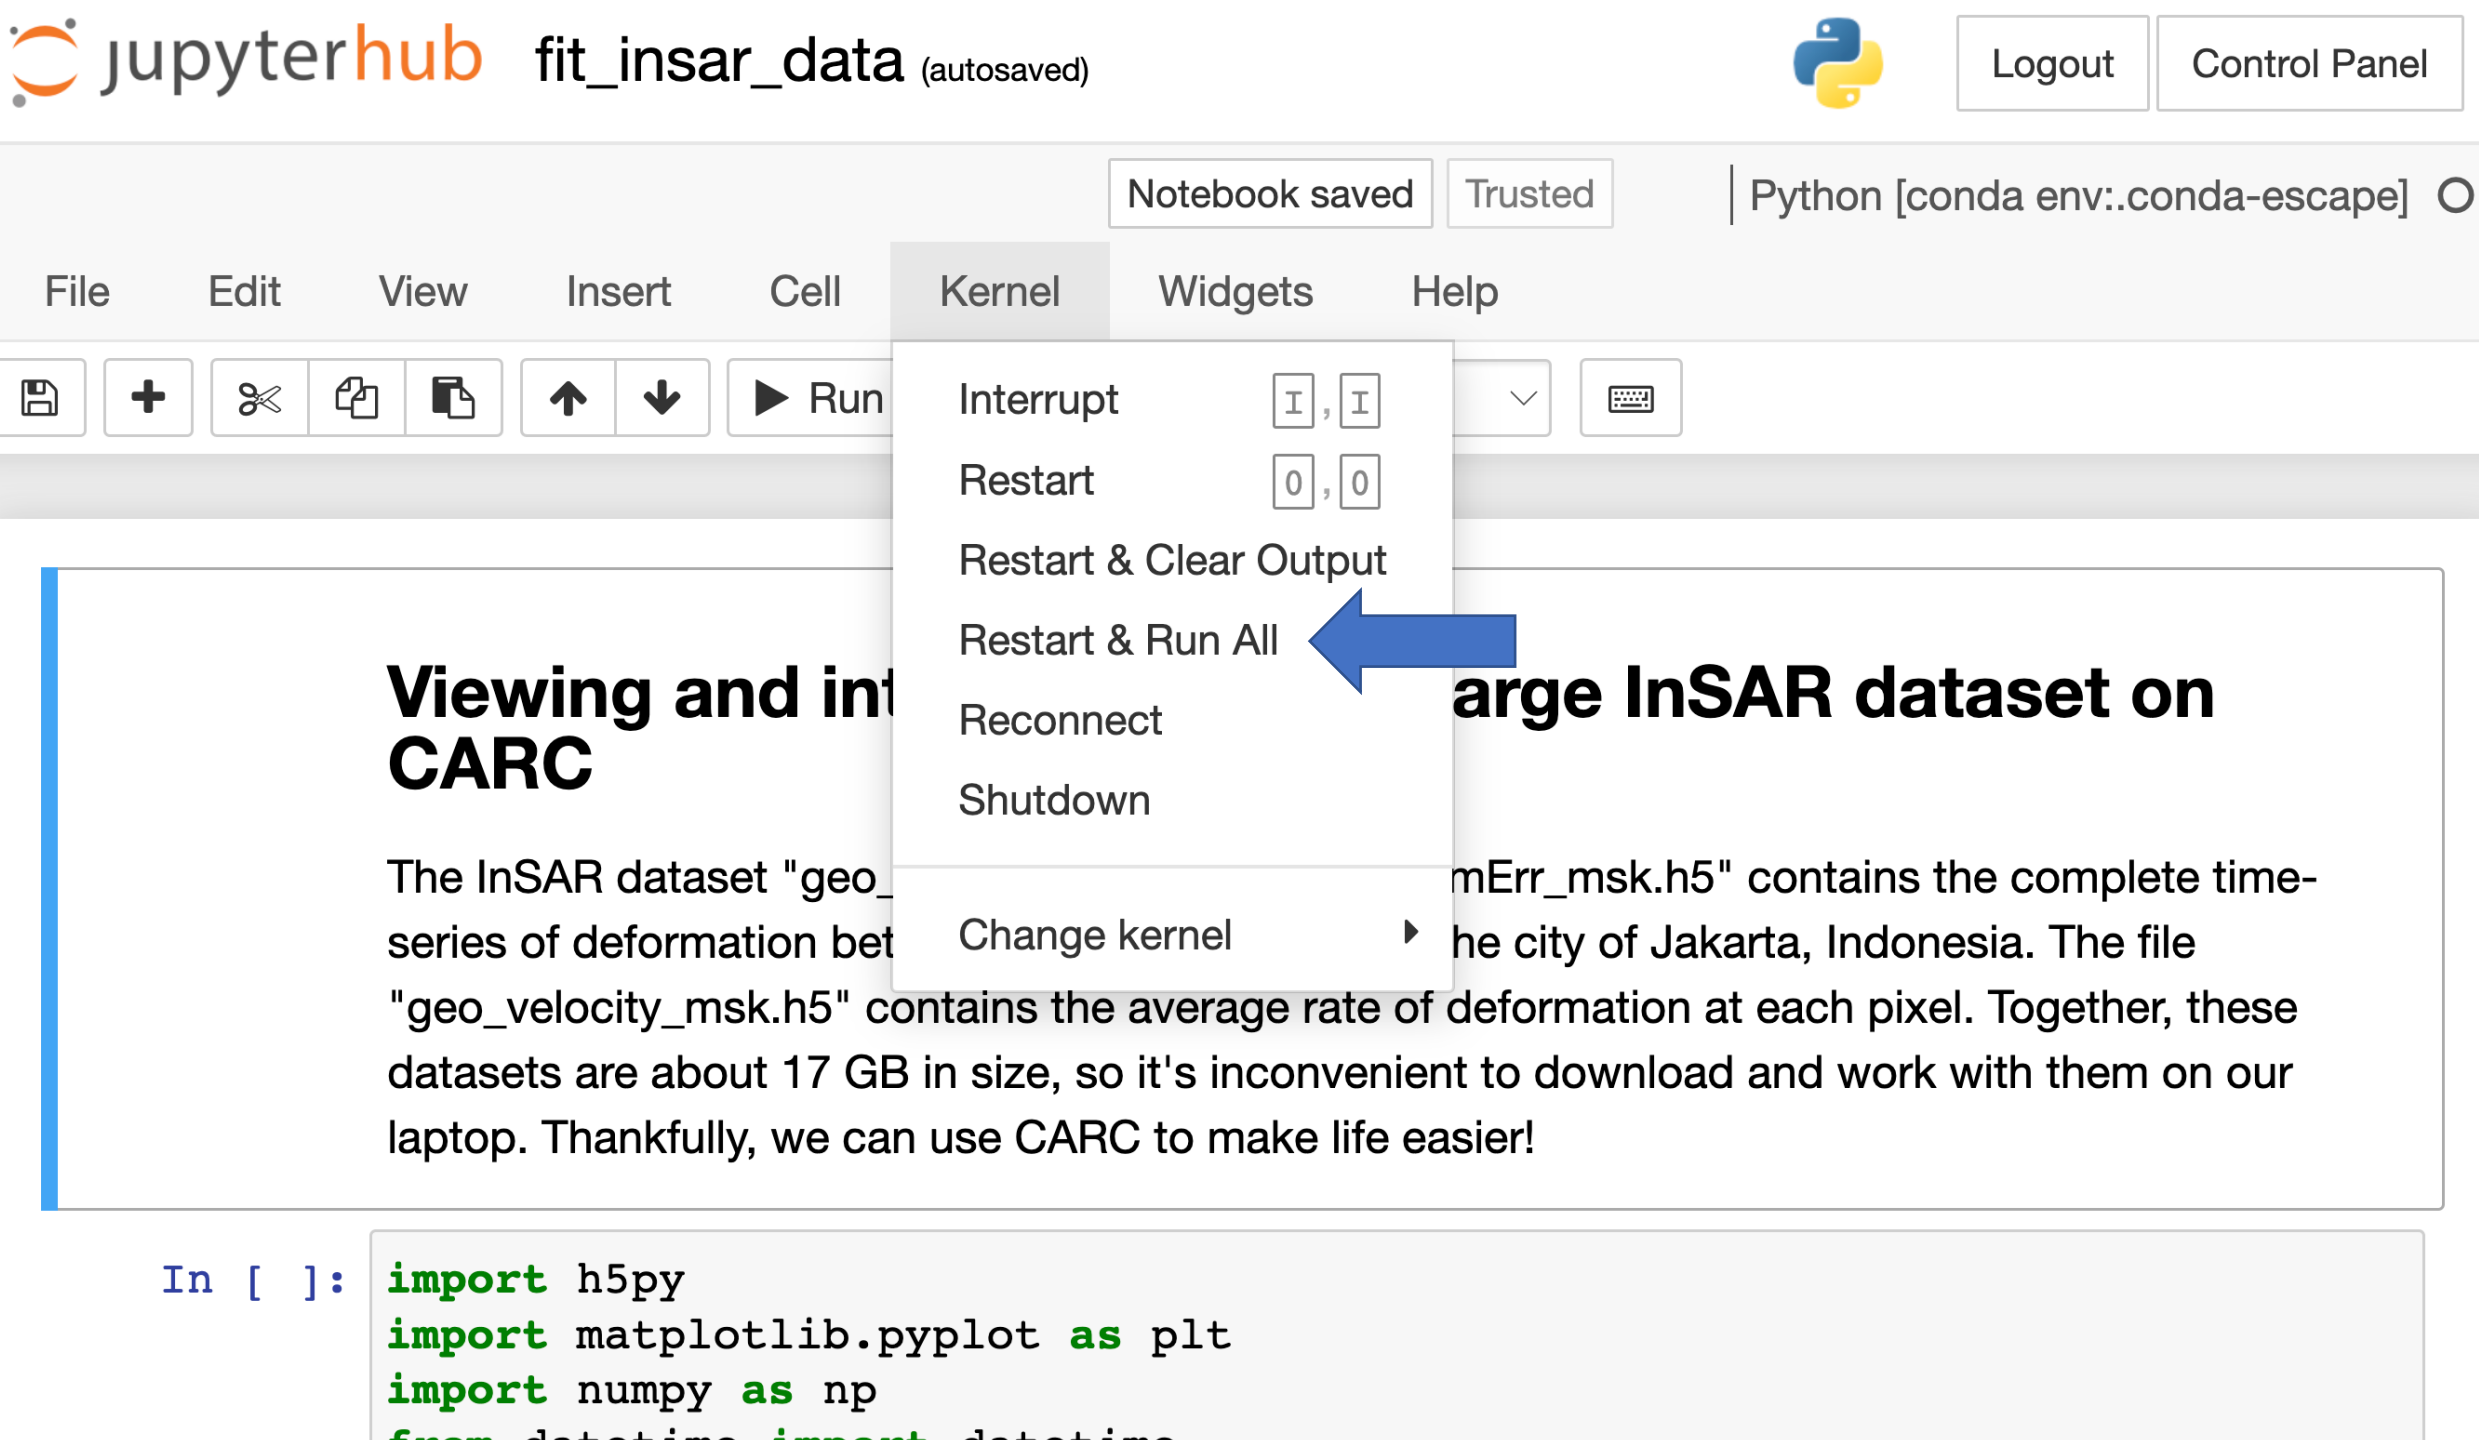
jupyterhub
fit_insar_data (autosaved)
Logout
Control Panel
Notebook saved
Trusted
Python [conda env:.conda-escape]
File
Edit
View
Insert
Cell
Kernel
Widgets
Help
Run
Interrupt
I, I
Restart
0, 0
Restart & Clear Output
Restart & Run All
Reconnect
Shutdown
Change kernel
Viewing and int
arge InSAR dataset on CARC
The InSAR dataset "geo_
mErr_msk.h5" contains the complete time-series of deformation bet
he city of Jakarta, Indonesia. The file "geo_velocity_msk.h5" contains the average rate of deformation at each pixel. Together, these datasets are about 17 GB in size, so it's inconvenient to download and work with them on our laptop. Thankfully, we can use CARC to make life easier!
In [ ]:
import h5py
import matplotlib.pyplot as plt
import numpy as np

jupyterhub fit_insar_data (unsaved changes)

Logout | Control Panel

Trusted | Python [conda env:.conda-escape]

File Edit View Insert Cell Kernel Widgets Help

Run | Code

In [4]:

[N] Means finished

```
# first, read the map coordinates

fname = data_path+'geo_timeseries_ERA5_ramp_demErr_msk.h5'
with h5py.File(fname,'r') as f:
    # read the metadata attributes:
    # to print all the attributes in the file (a lot!): print(f.attrs.keys(
    x0=float(f.attrs['X_FIRST'])
    dx=float(f.attrs['X_STEP'])
    nx=int(f.attrs['WIDTH'])
    y0=float(f.attrs['Y_FIRST'])
    dy=float(f.attrs['Y_STEP'])
    ny=int(f.attrs['LENGTH'])
    rlon=float(f.attrs['REF_LON'])
    rlat=float(f.attrs['REF_LAT'])
    # so, we have the x0 and dx, but need to turn them into vectors (x0,x1,
    lonvec = [x0 + i*dx for i in range(nx)]
    latvec = [y0 + i*dy for i in range(ny)]
```

jupyterhub fit_insar_data (unsaved changes)

Logout | Control Panel

Trusted | Python [conda env:.conda-escape]

File Edit View Insert Cell Kernel Widgets Help

Run | Code

In [*]:

* Means running

Hopper took 100 s
Wheeler 250 s.

```
# now, read some of the timeseries images and plot them in a grid (3x3)

fname = data_path+'geo_timeseries_ERA5_ramp_demErr_msk.h5'
with h5py.File(fname, 'r') as f:
    # Read the dates from the file
    st = time.time()
    dates = [date.decode('utf-8') for date in f['date'][:]]
    timestep0 = f['timeseries'][0, :, :]
    print('Time to read dates:', round(time.time() - st, 3), 's')

    # Calculate nearly-evenly spaced timesteps
    st = time.time()
    total_timesteps = len(dates)
    spacing = total_timesteps // 9  # 9 is 3x3
    selected_t_indices = np.arange(0, total_timesteps, spacing)[:9]  # Get
    print('Time to calculate nearly-evenly spaced timesteps:', round(time.t
```

```
/users/vanilla/workshops/escape/kml_functions.py:25: RuntimeWarning: Mean of empty slice
  downsampled[i, j] = np.nanmean(block)
```

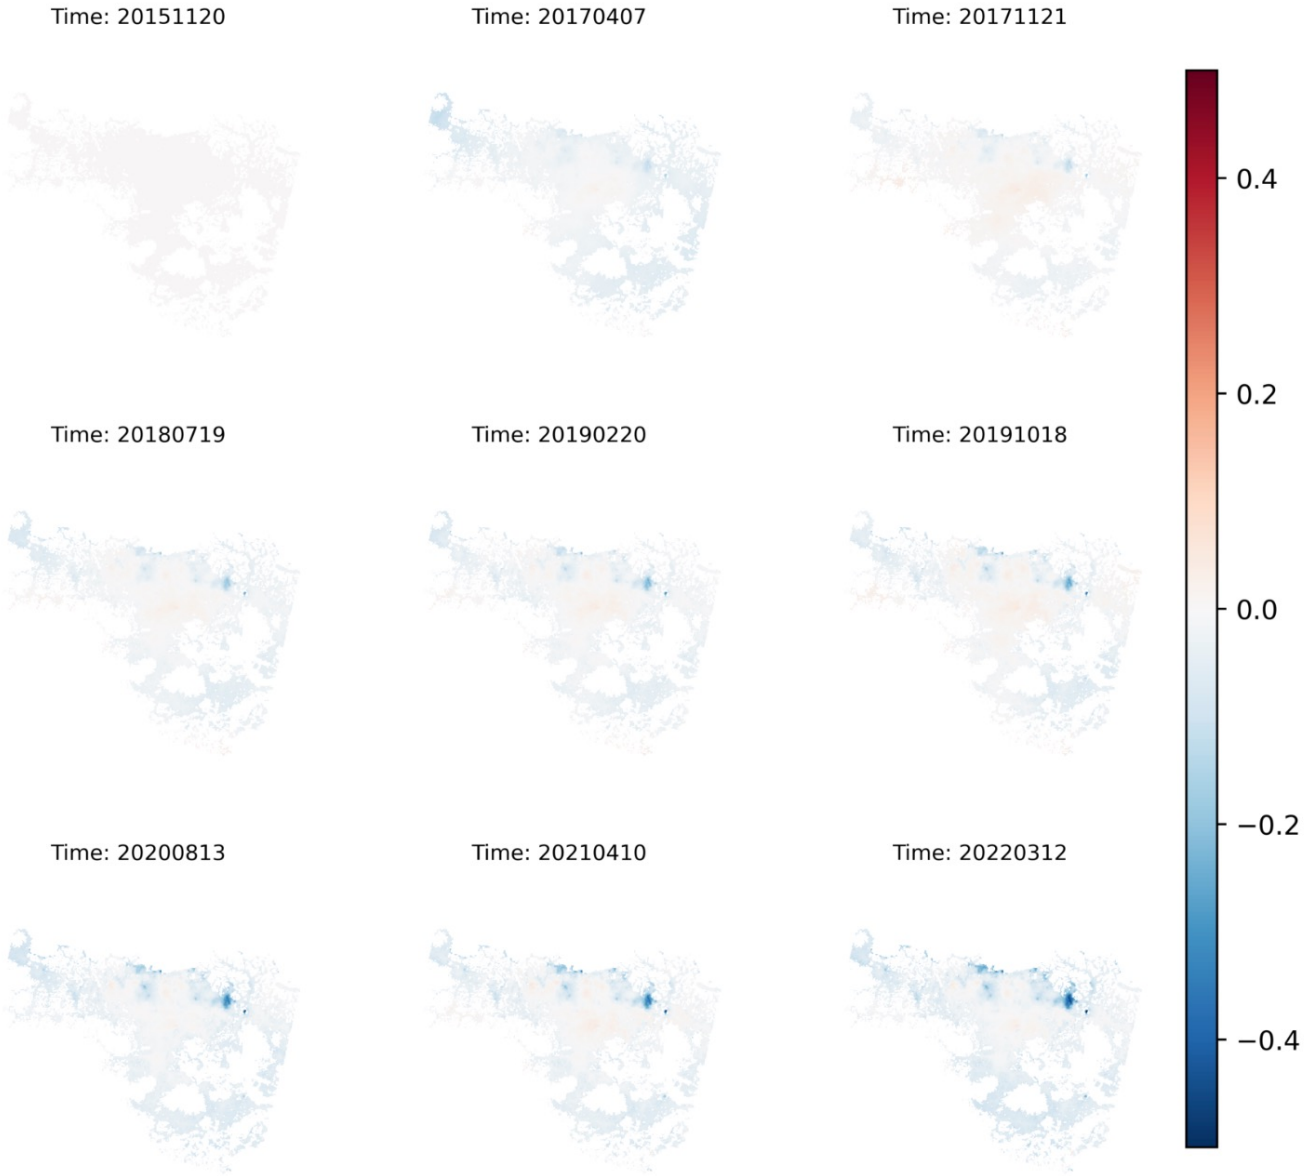

Shut down your jupyterhub server

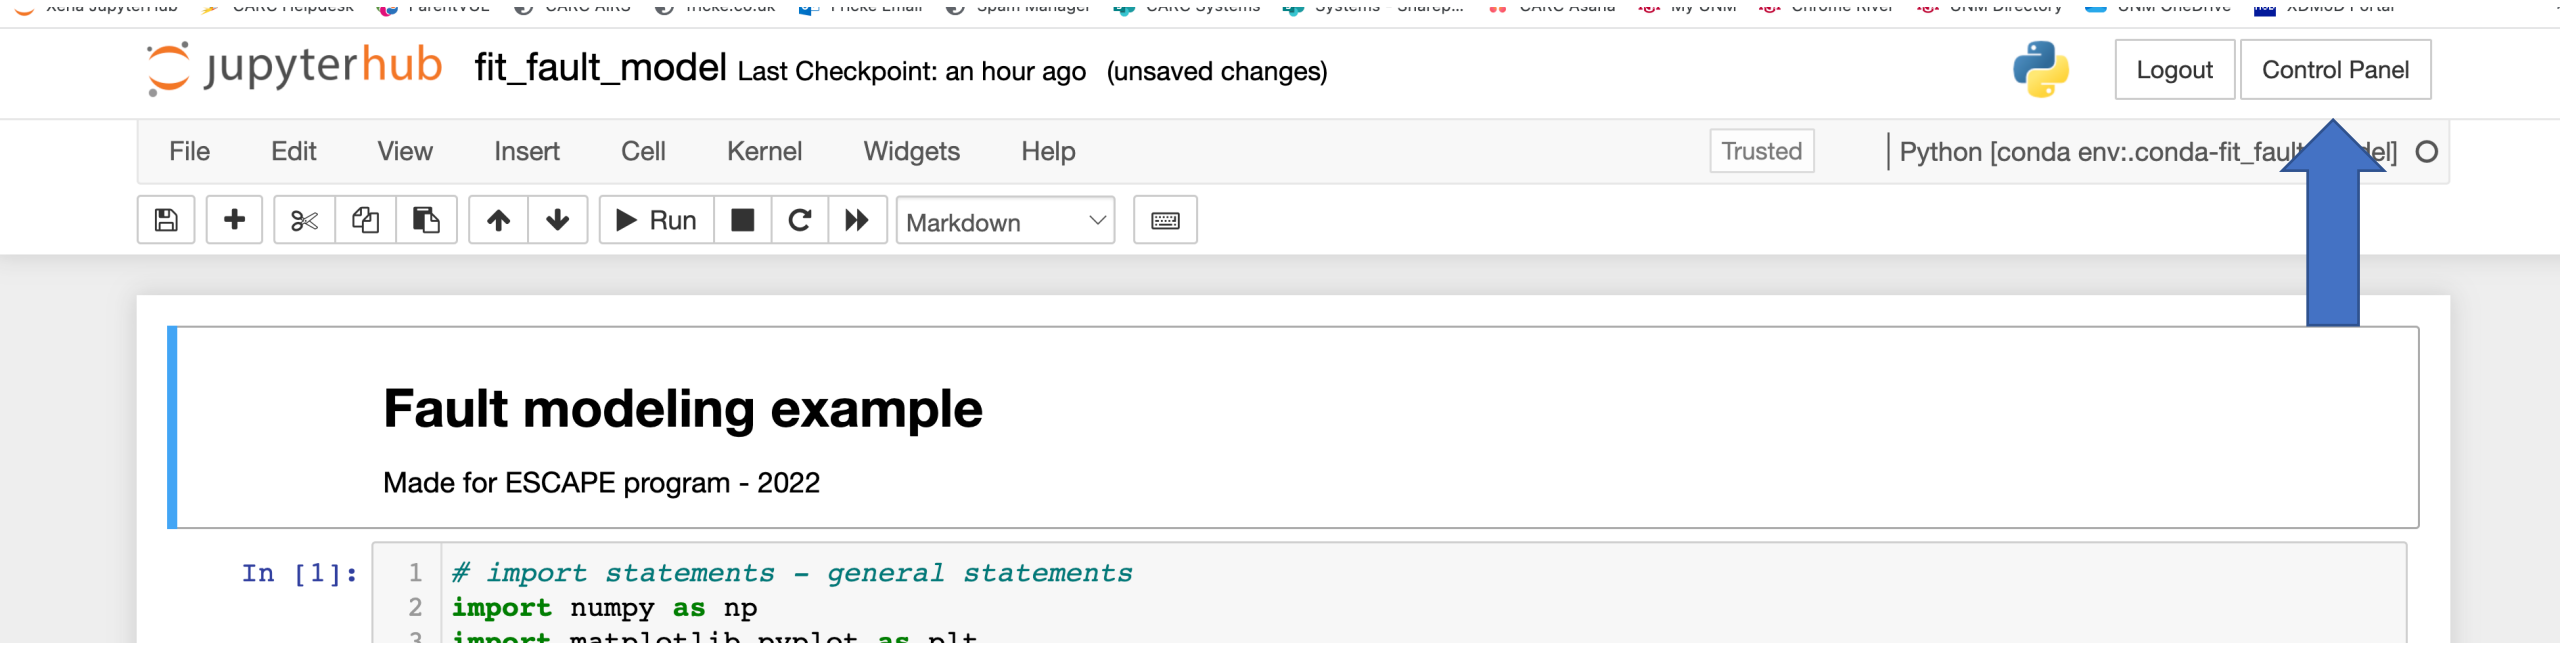

CARC Resources

- Tutorial Videos
- Written Tutorials

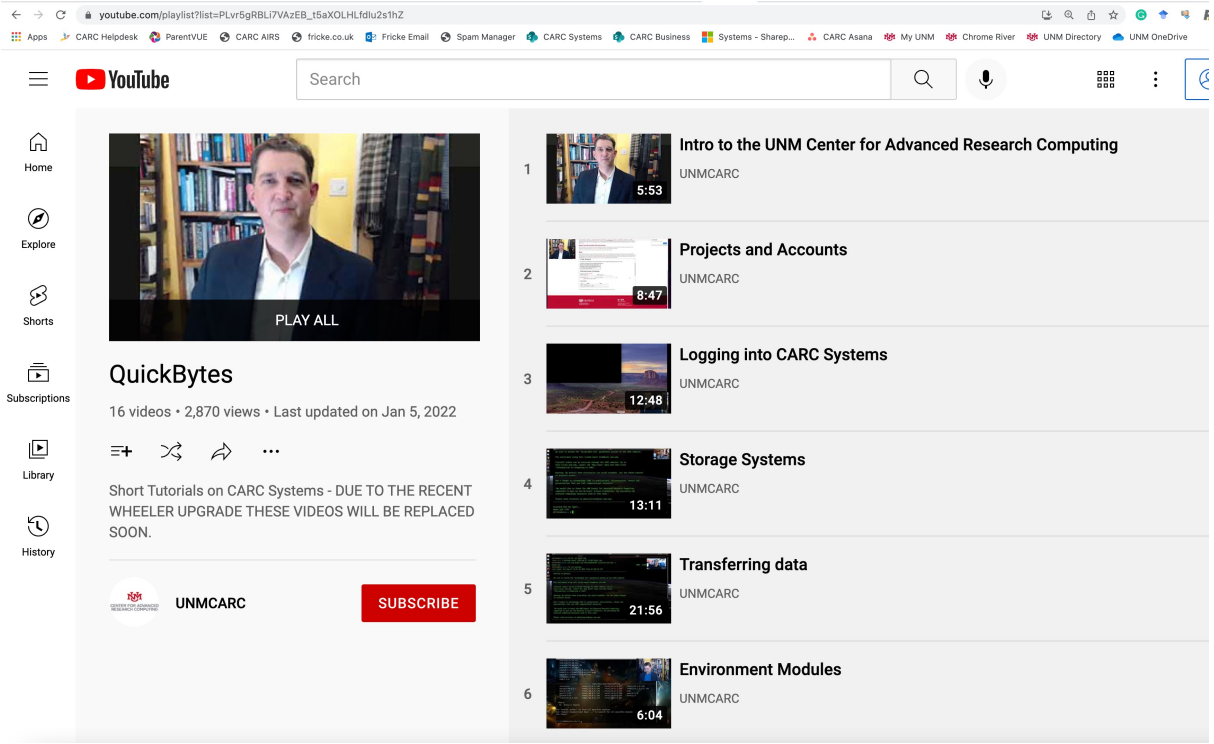

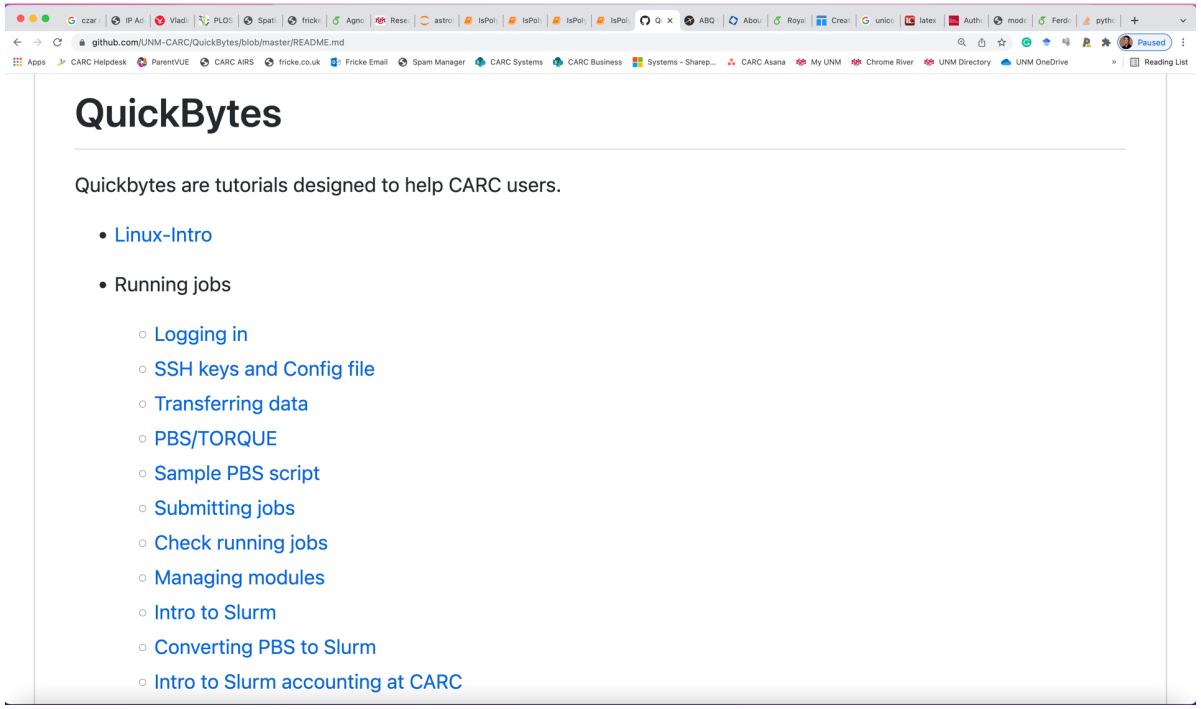

Getting Help

help@carc.unm.edu

Office hours