Beginner's

Introduction to Computing at CARC with SLURM (Biodiversity Informatics– BIOL419/519)

Version 0.1

Goals

1) Logging in…
2) BASH (entering commands)
3) SLURM Scheduler (accessing hardware)

I hope to spend about an hour and 15 min.

Logging into Wheeler

First login to the Linux **workstation** in front
of you using the username from the sign in sheet. The initial password is “Welcome2carc”.

Use your CARC username and password.

Matthew will help you get logged in 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.

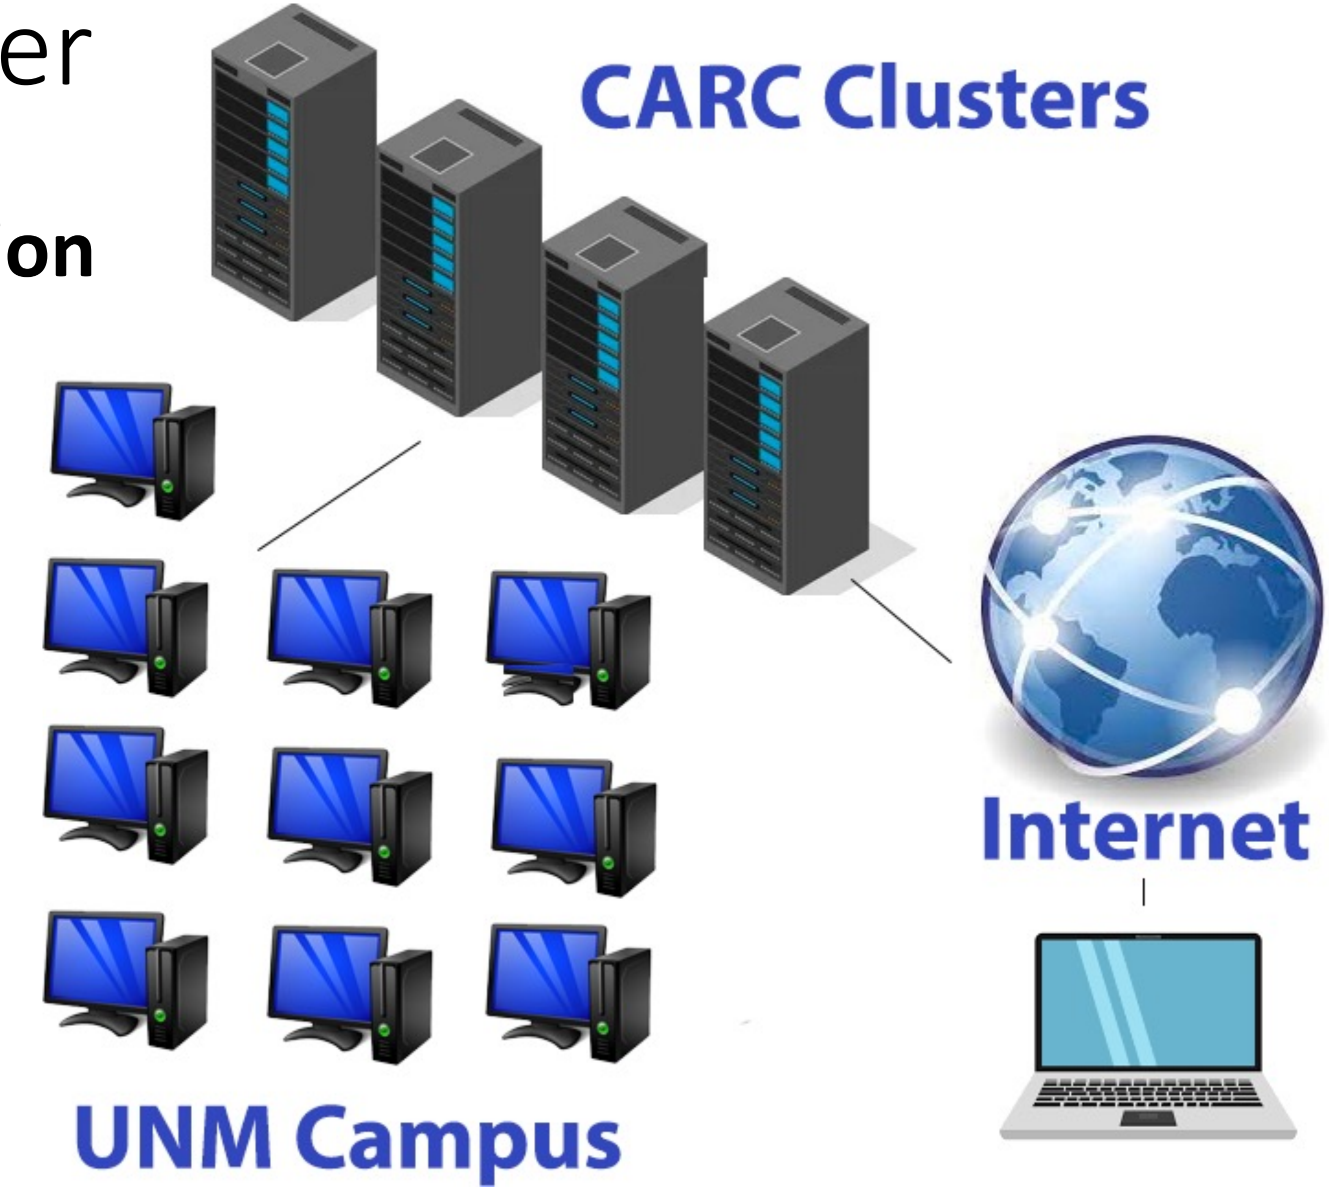

Logging into Wheeler

```
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

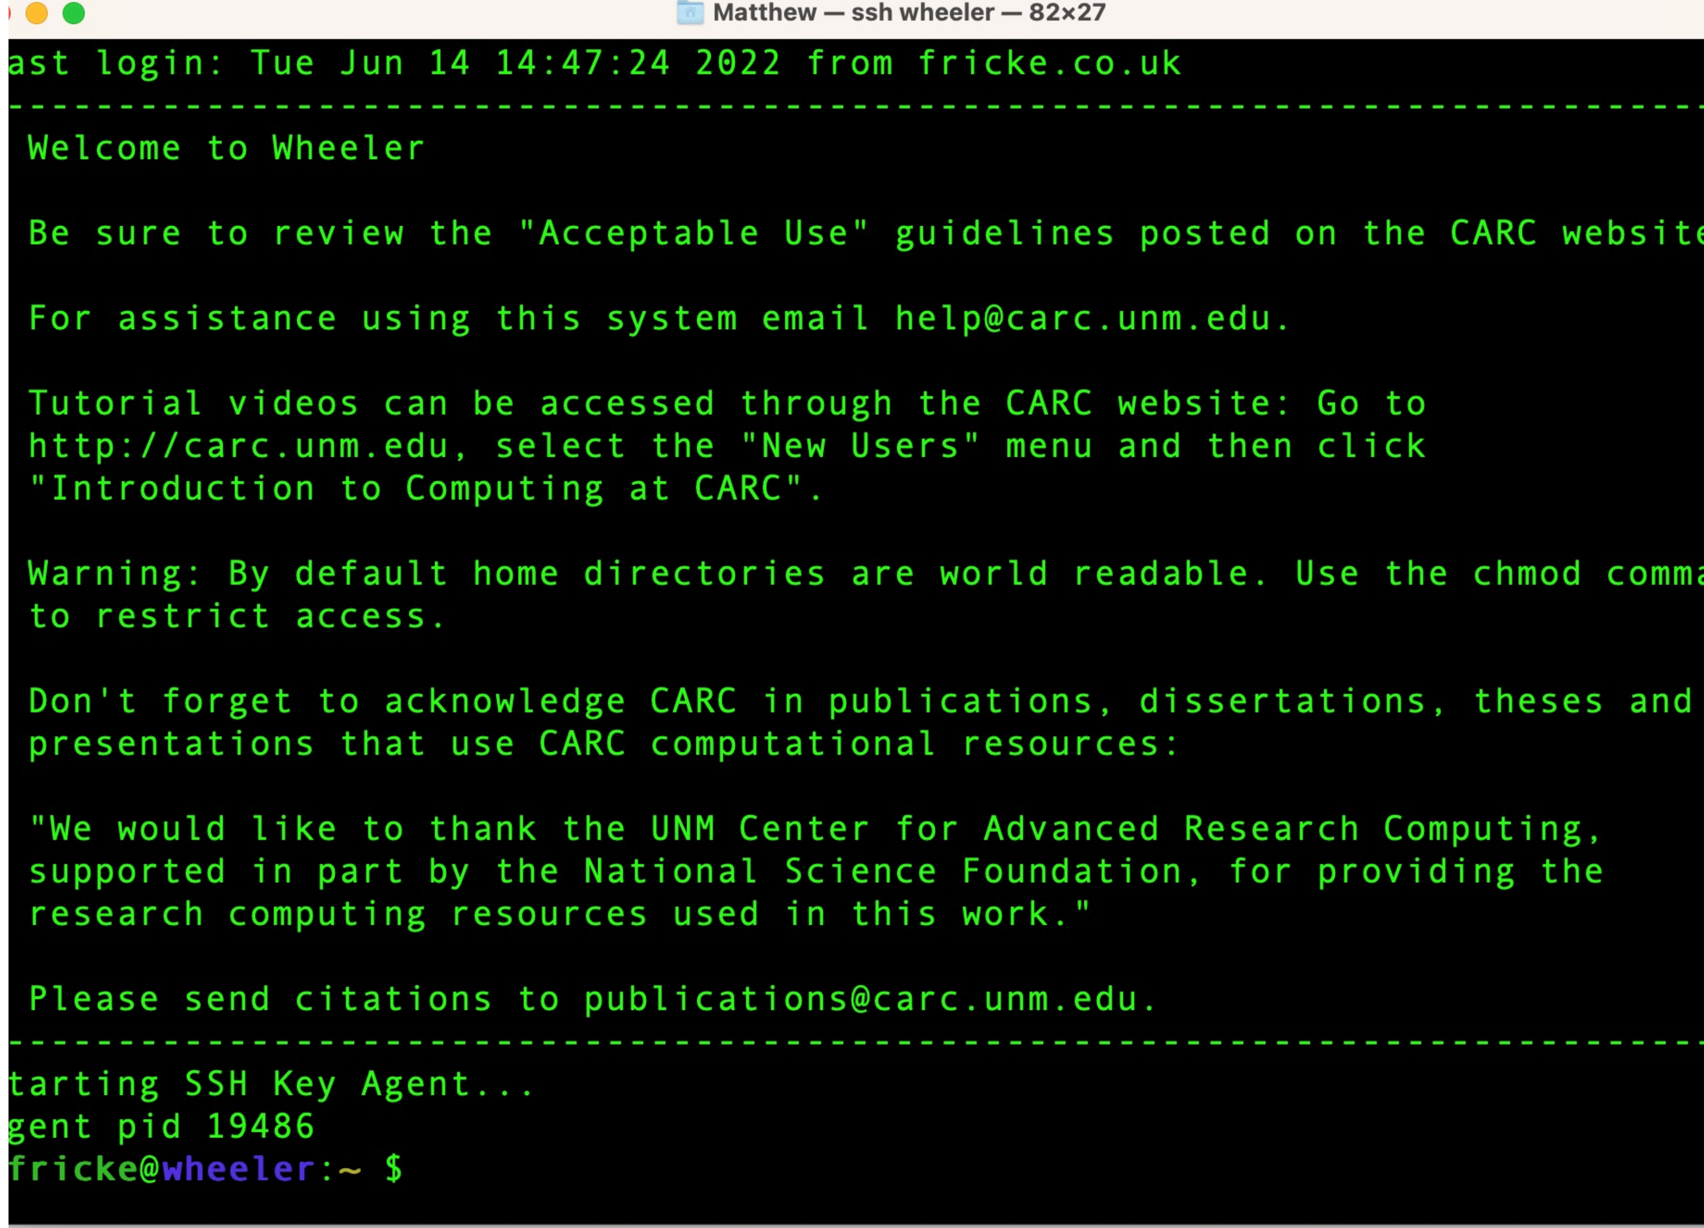

```
ast login: Tue Jun 14 14:47:24 2022 from fricke.co.uk
------------------------------------------------------------------------------
 Welcome to Wheeler

 Be sure to review the "Acceptable Use" guidelines posted on the CARC website

 For assistance using this system email help@carc.unm.edu.

 Tutorial videos can be accessed through the CARC website: Go to
 http://carc.unm.edu, select the "New Users" menu and then click
 "Introduction to Computing at CARC".

 Warning: By default home directories are world readable. Use the chmod comma
 to restrict access.

 Don't forget to acknowledge CARC in publications, dissertations, theses and
 presentations that use CARC computational resources:

 "We would like to thank the UNM Center for Advanced Research Computing,
 supported in part by the National Science Foundation, for providing the
 research computing resources used in this work."

 Please send citations to publications@carc.unm.edu.
------------------------------------------------------------------------------
tarting SSH Key Agent...
gent pid 19486
fricke@wheeler:~ $
```

Set your password

```
[vanilla@wheeler beginner]$ passwd
```

Enter “Welcome2carc”

When prompted enter a password you would like to use

Our system will check whether it is too easy to guess.

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)

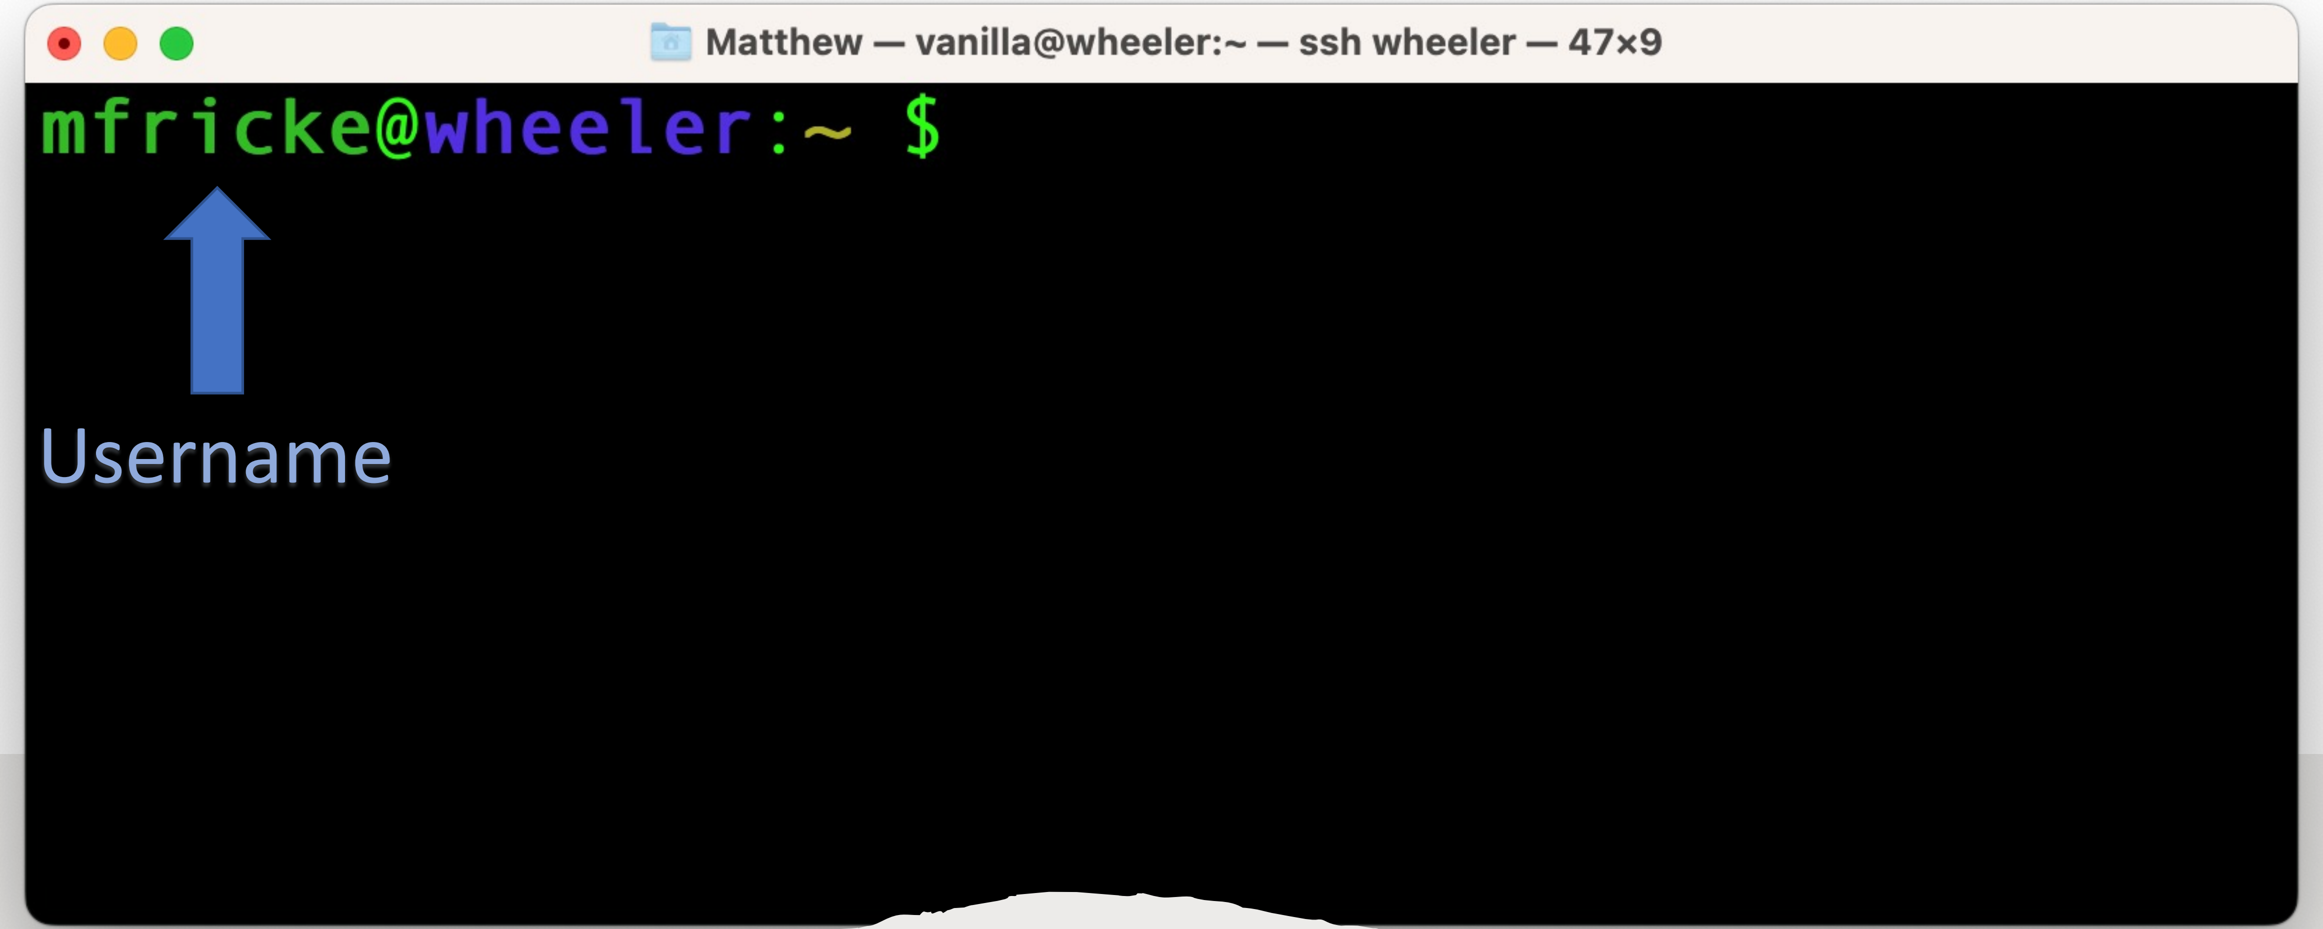

```
mfricke@wheeler:~ $
```

Understanding the BASH prompt…

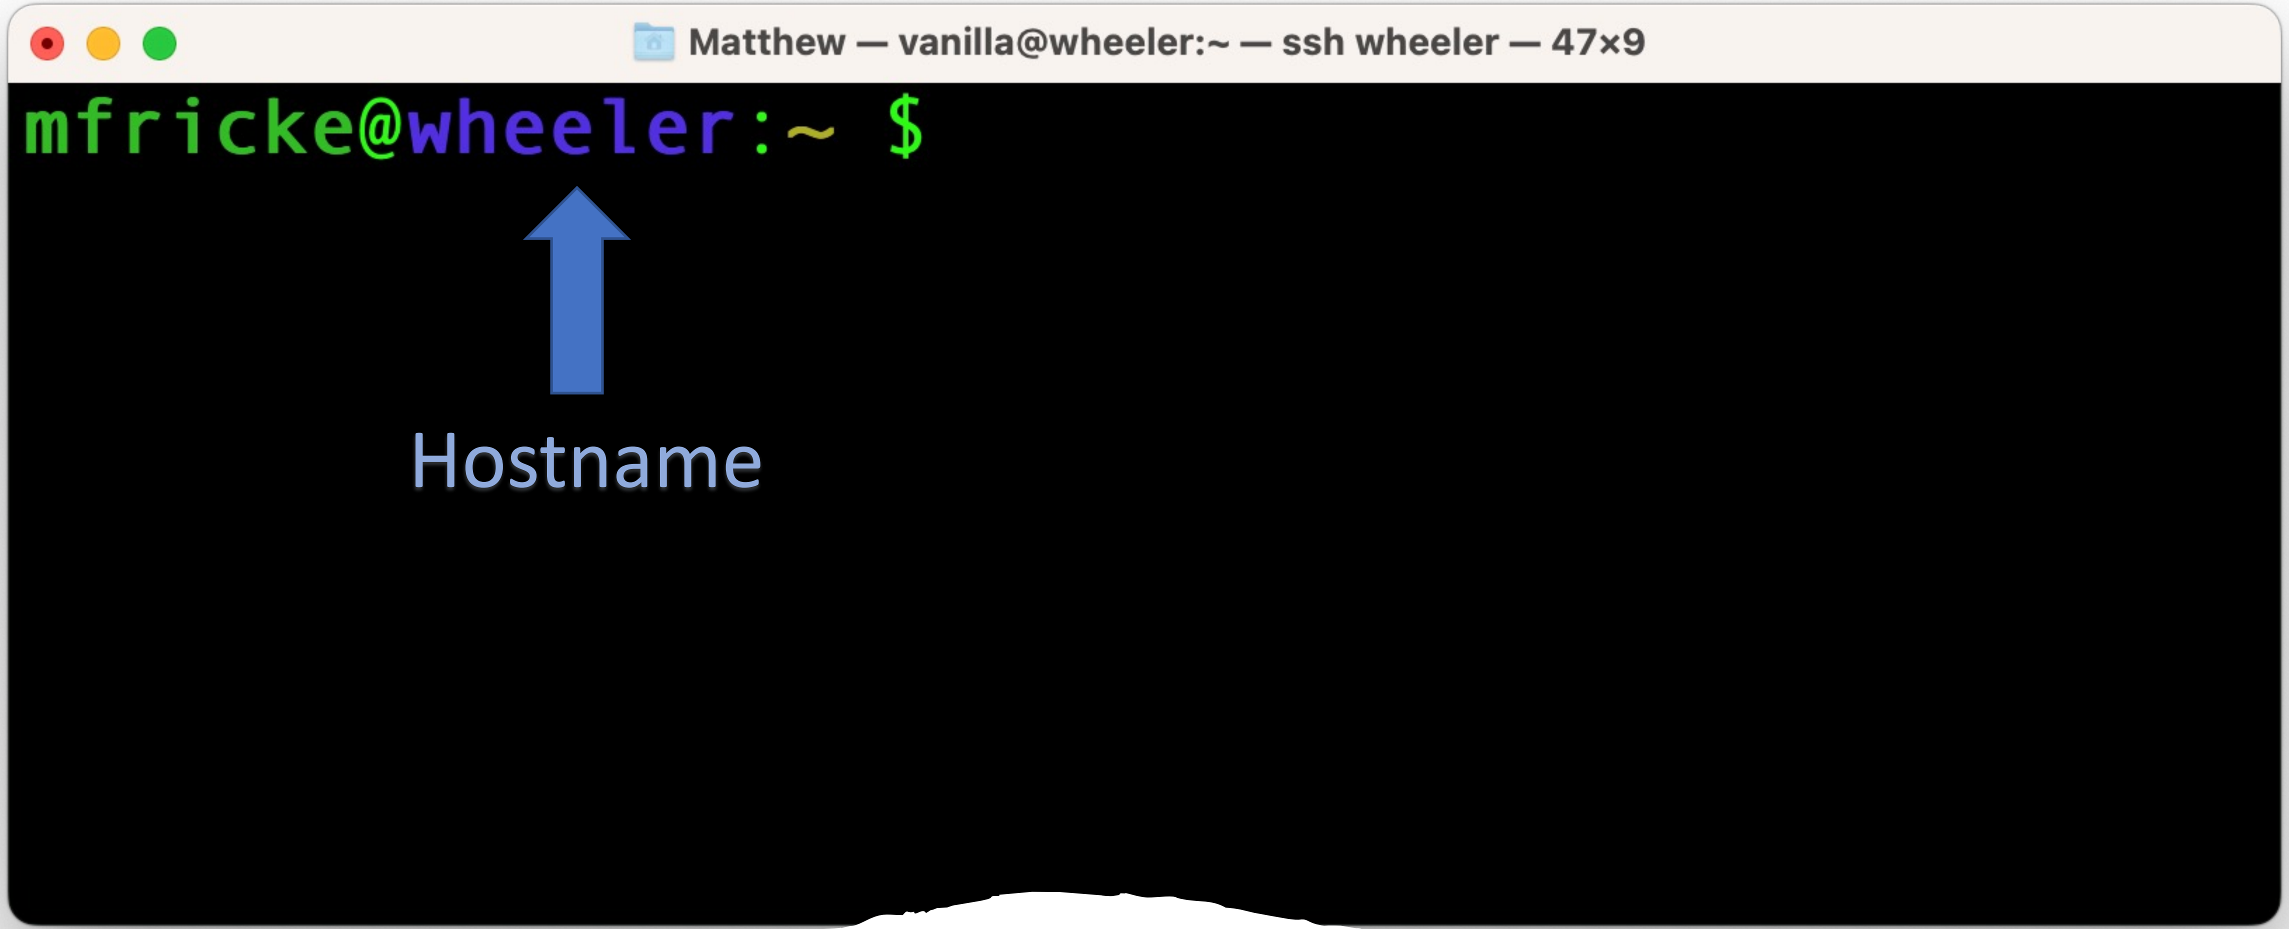

```
mfricke@wheeler:~ $
```

Hostname

Understanding the BASH prompt…

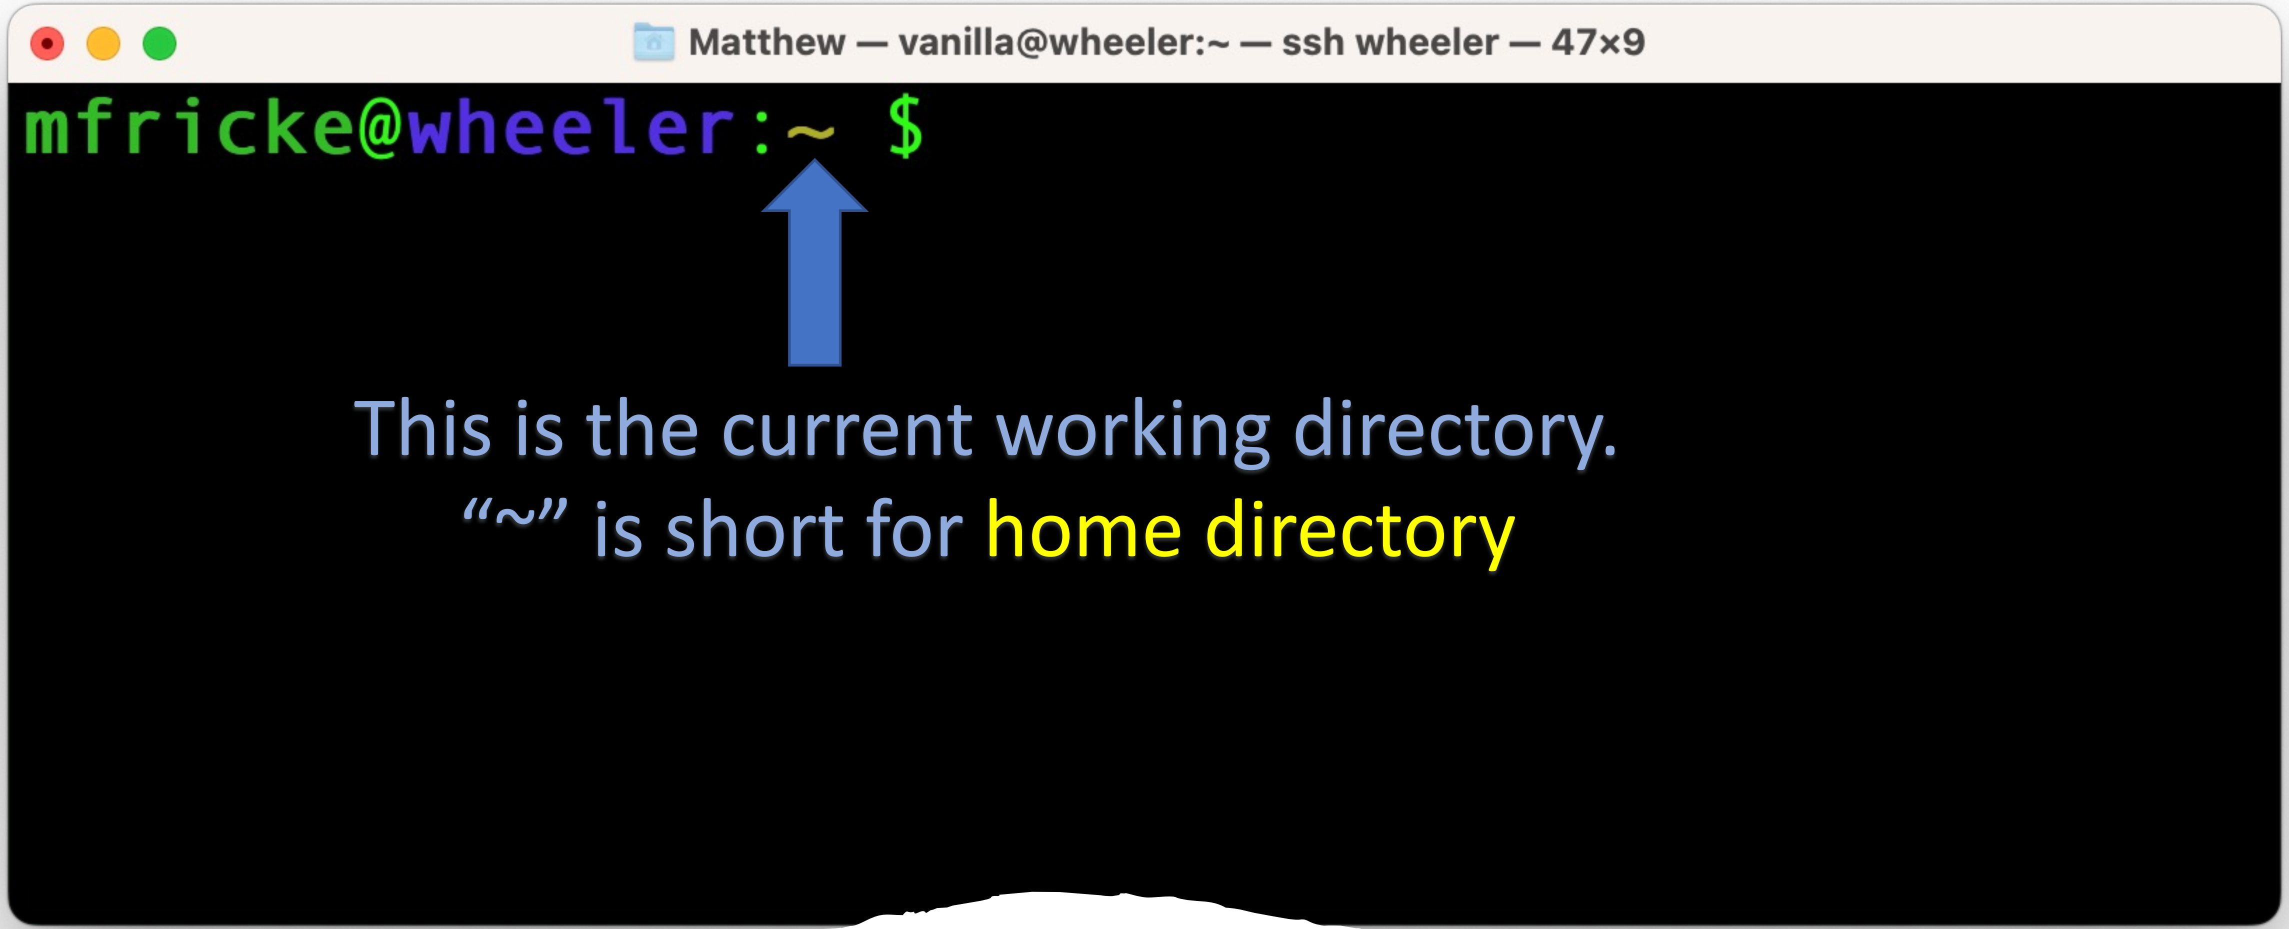

```
mfricke@wheeler:~ $
```

This is the current working directory.

“~” is short for home directory

Understanding the BASH prompt…

```
[vanilla@wheeler ~]$ pwd
/users/vanilla
[vanilla@wheeler ~]$ 
```

Figuring out where you are in the filesystem…

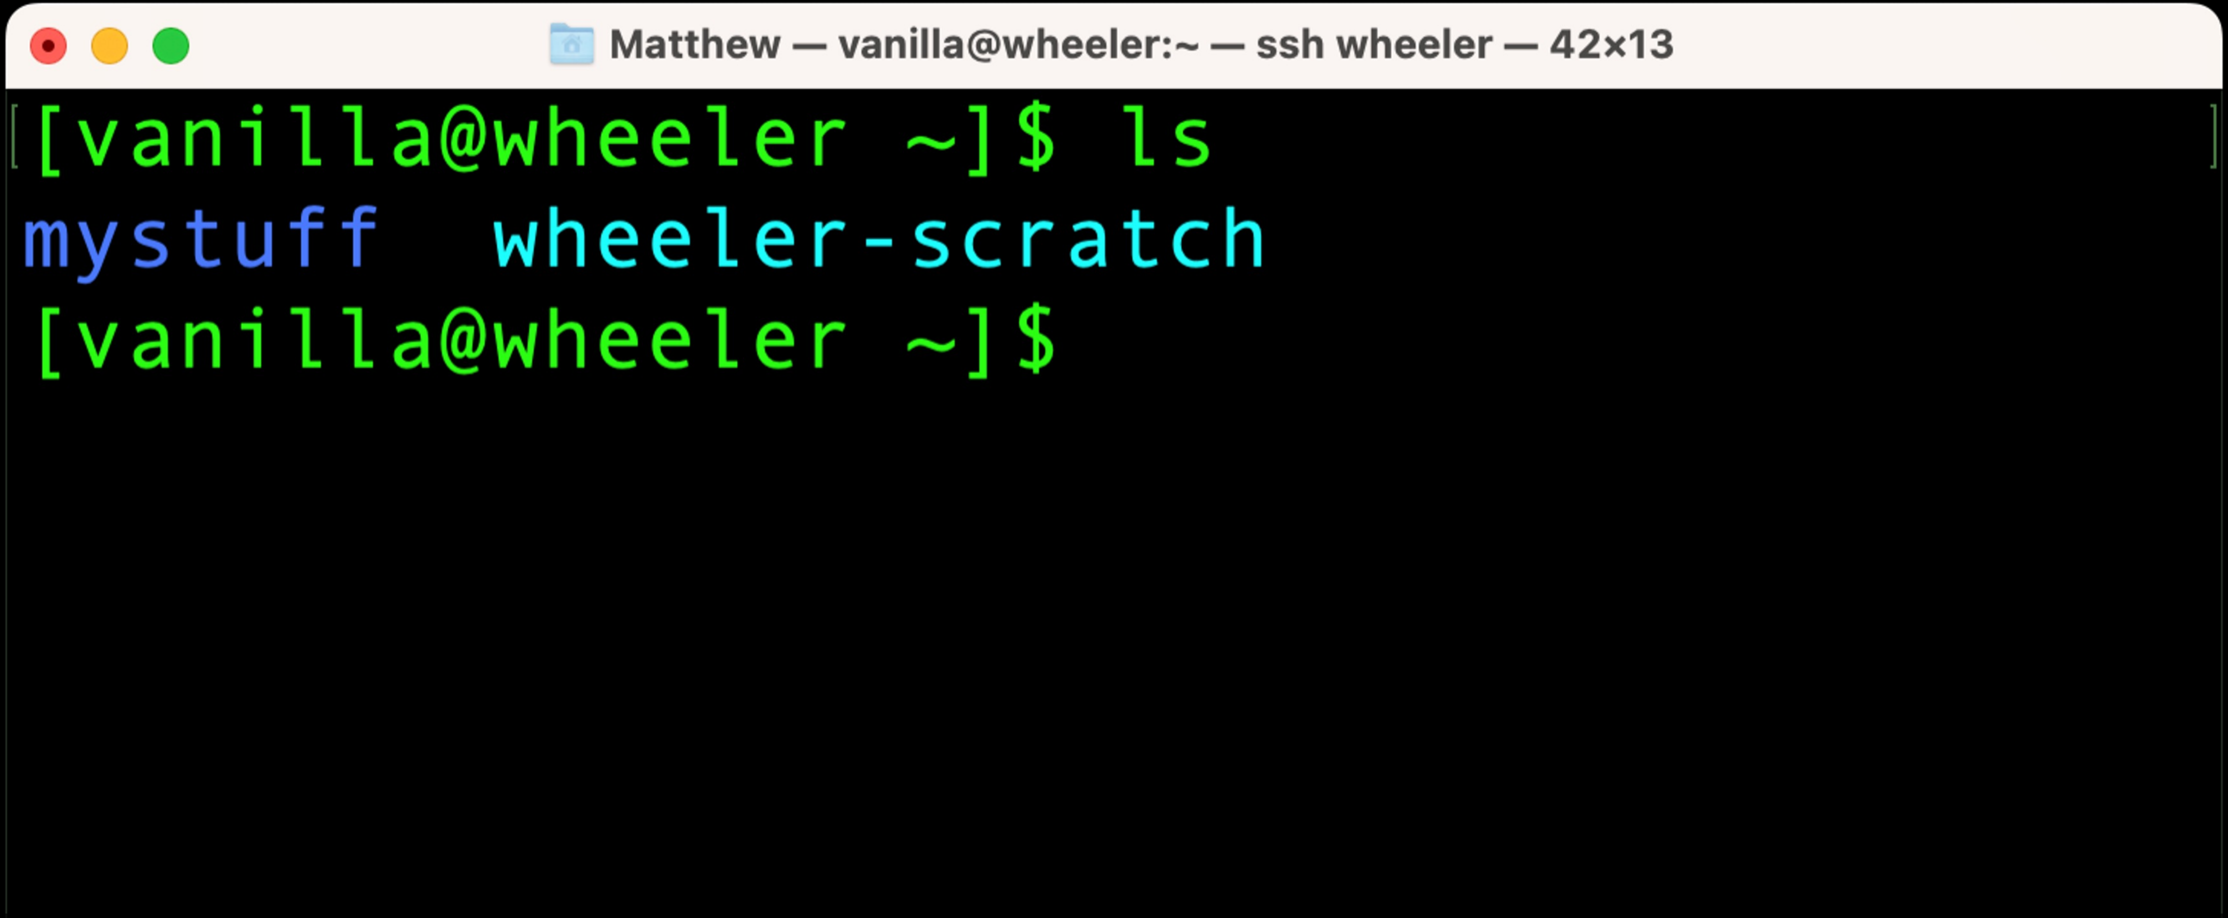

```
[vanilla@wheeler ~]$ ls
mystuff  wheeler-scratch
[vanilla@wheeler ~]$
```

Figuring out where you are in the filesystem...

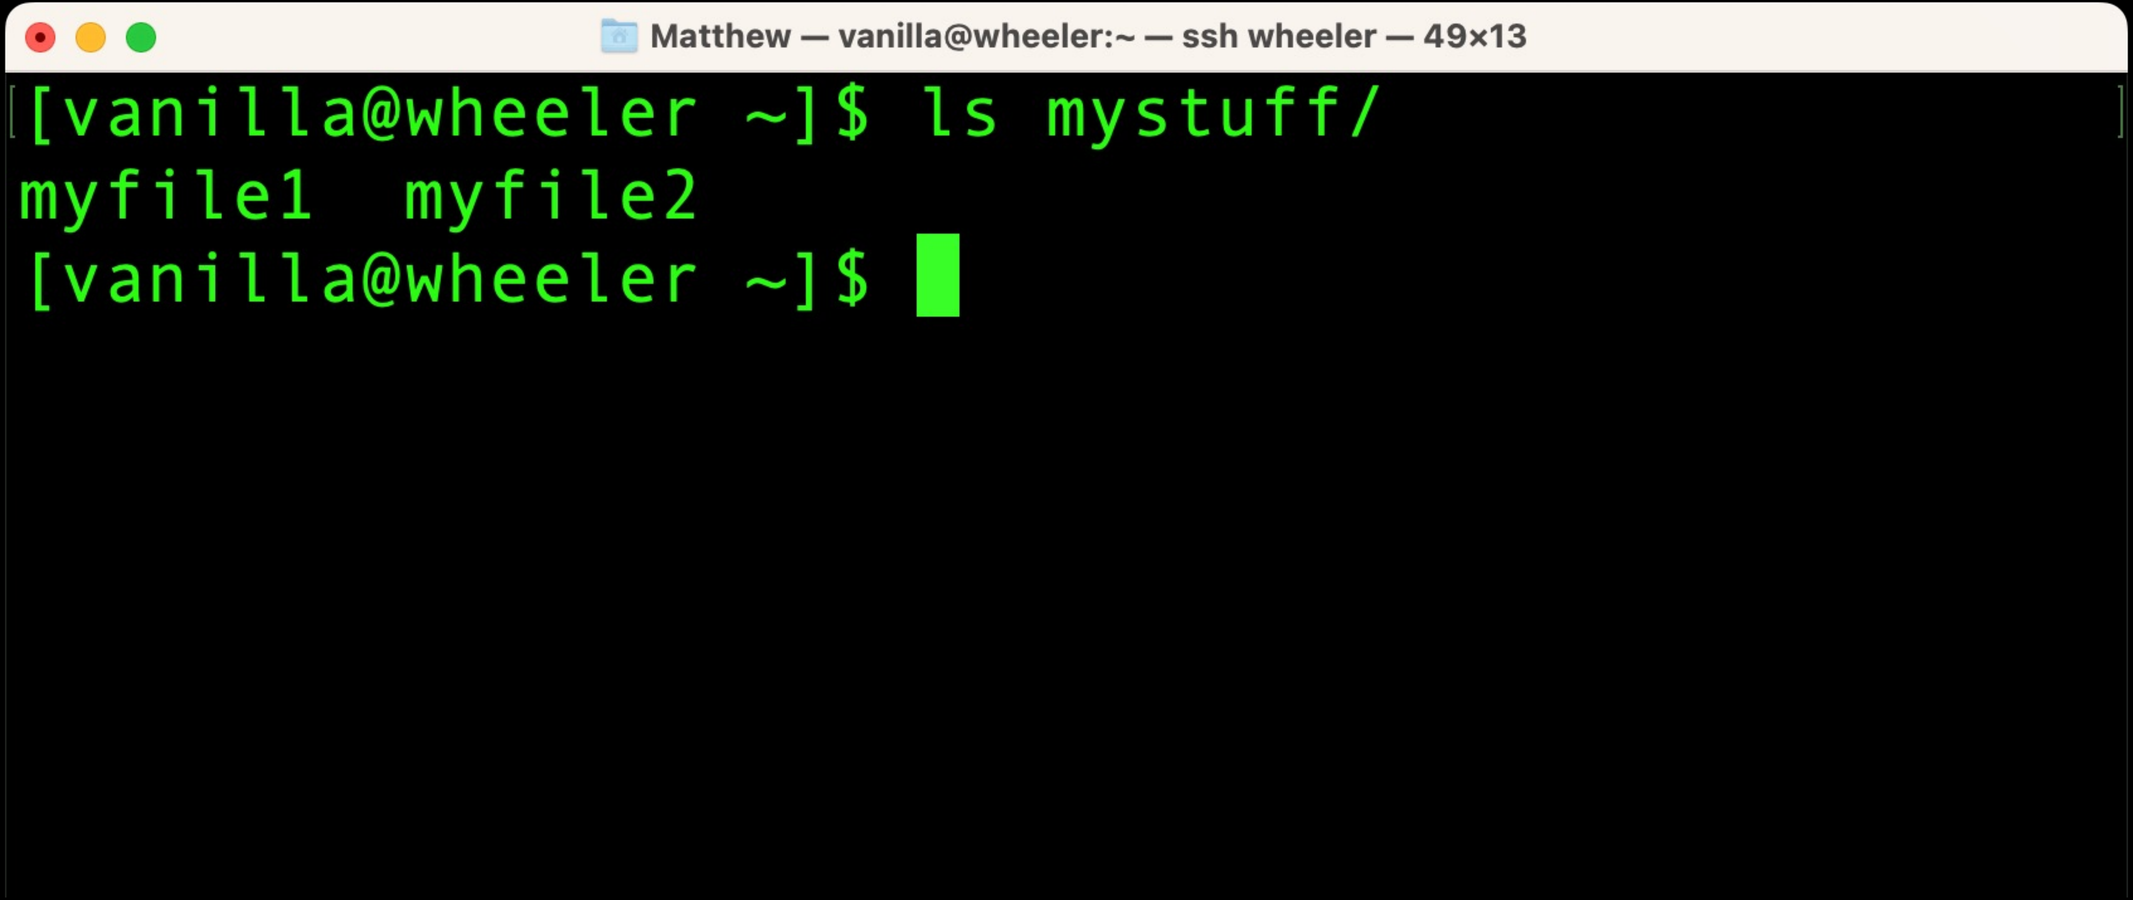

```
[vanilla@wheeler ~]$ ls mystuff/
myfile1  myfile2
[vanilla@wheeler ~]$ 
```

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

Use the tab key to autocomplete

Going somewhere new…

- 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/
```

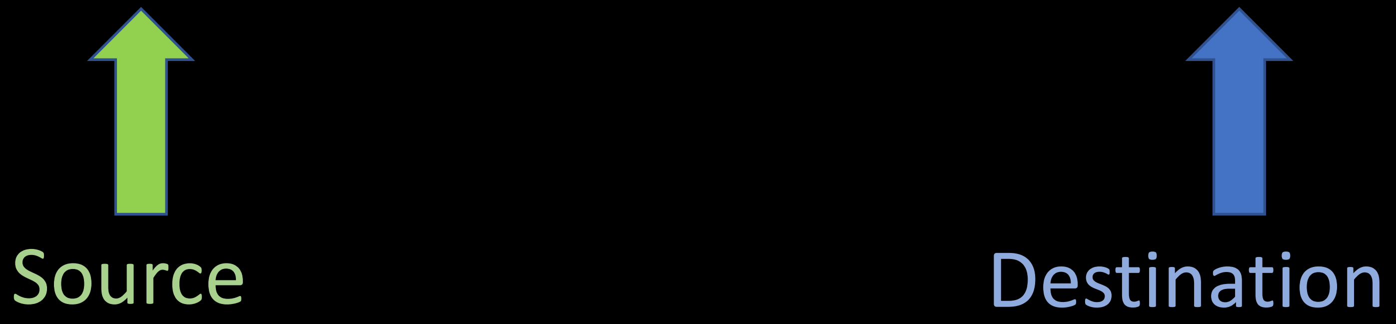

Source

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 ~]$ mkdir newdir
[vanilla@wheeler ~]$
```

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

Creating a new directory…

Software Access

Lmod Modules

Conda

```
[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

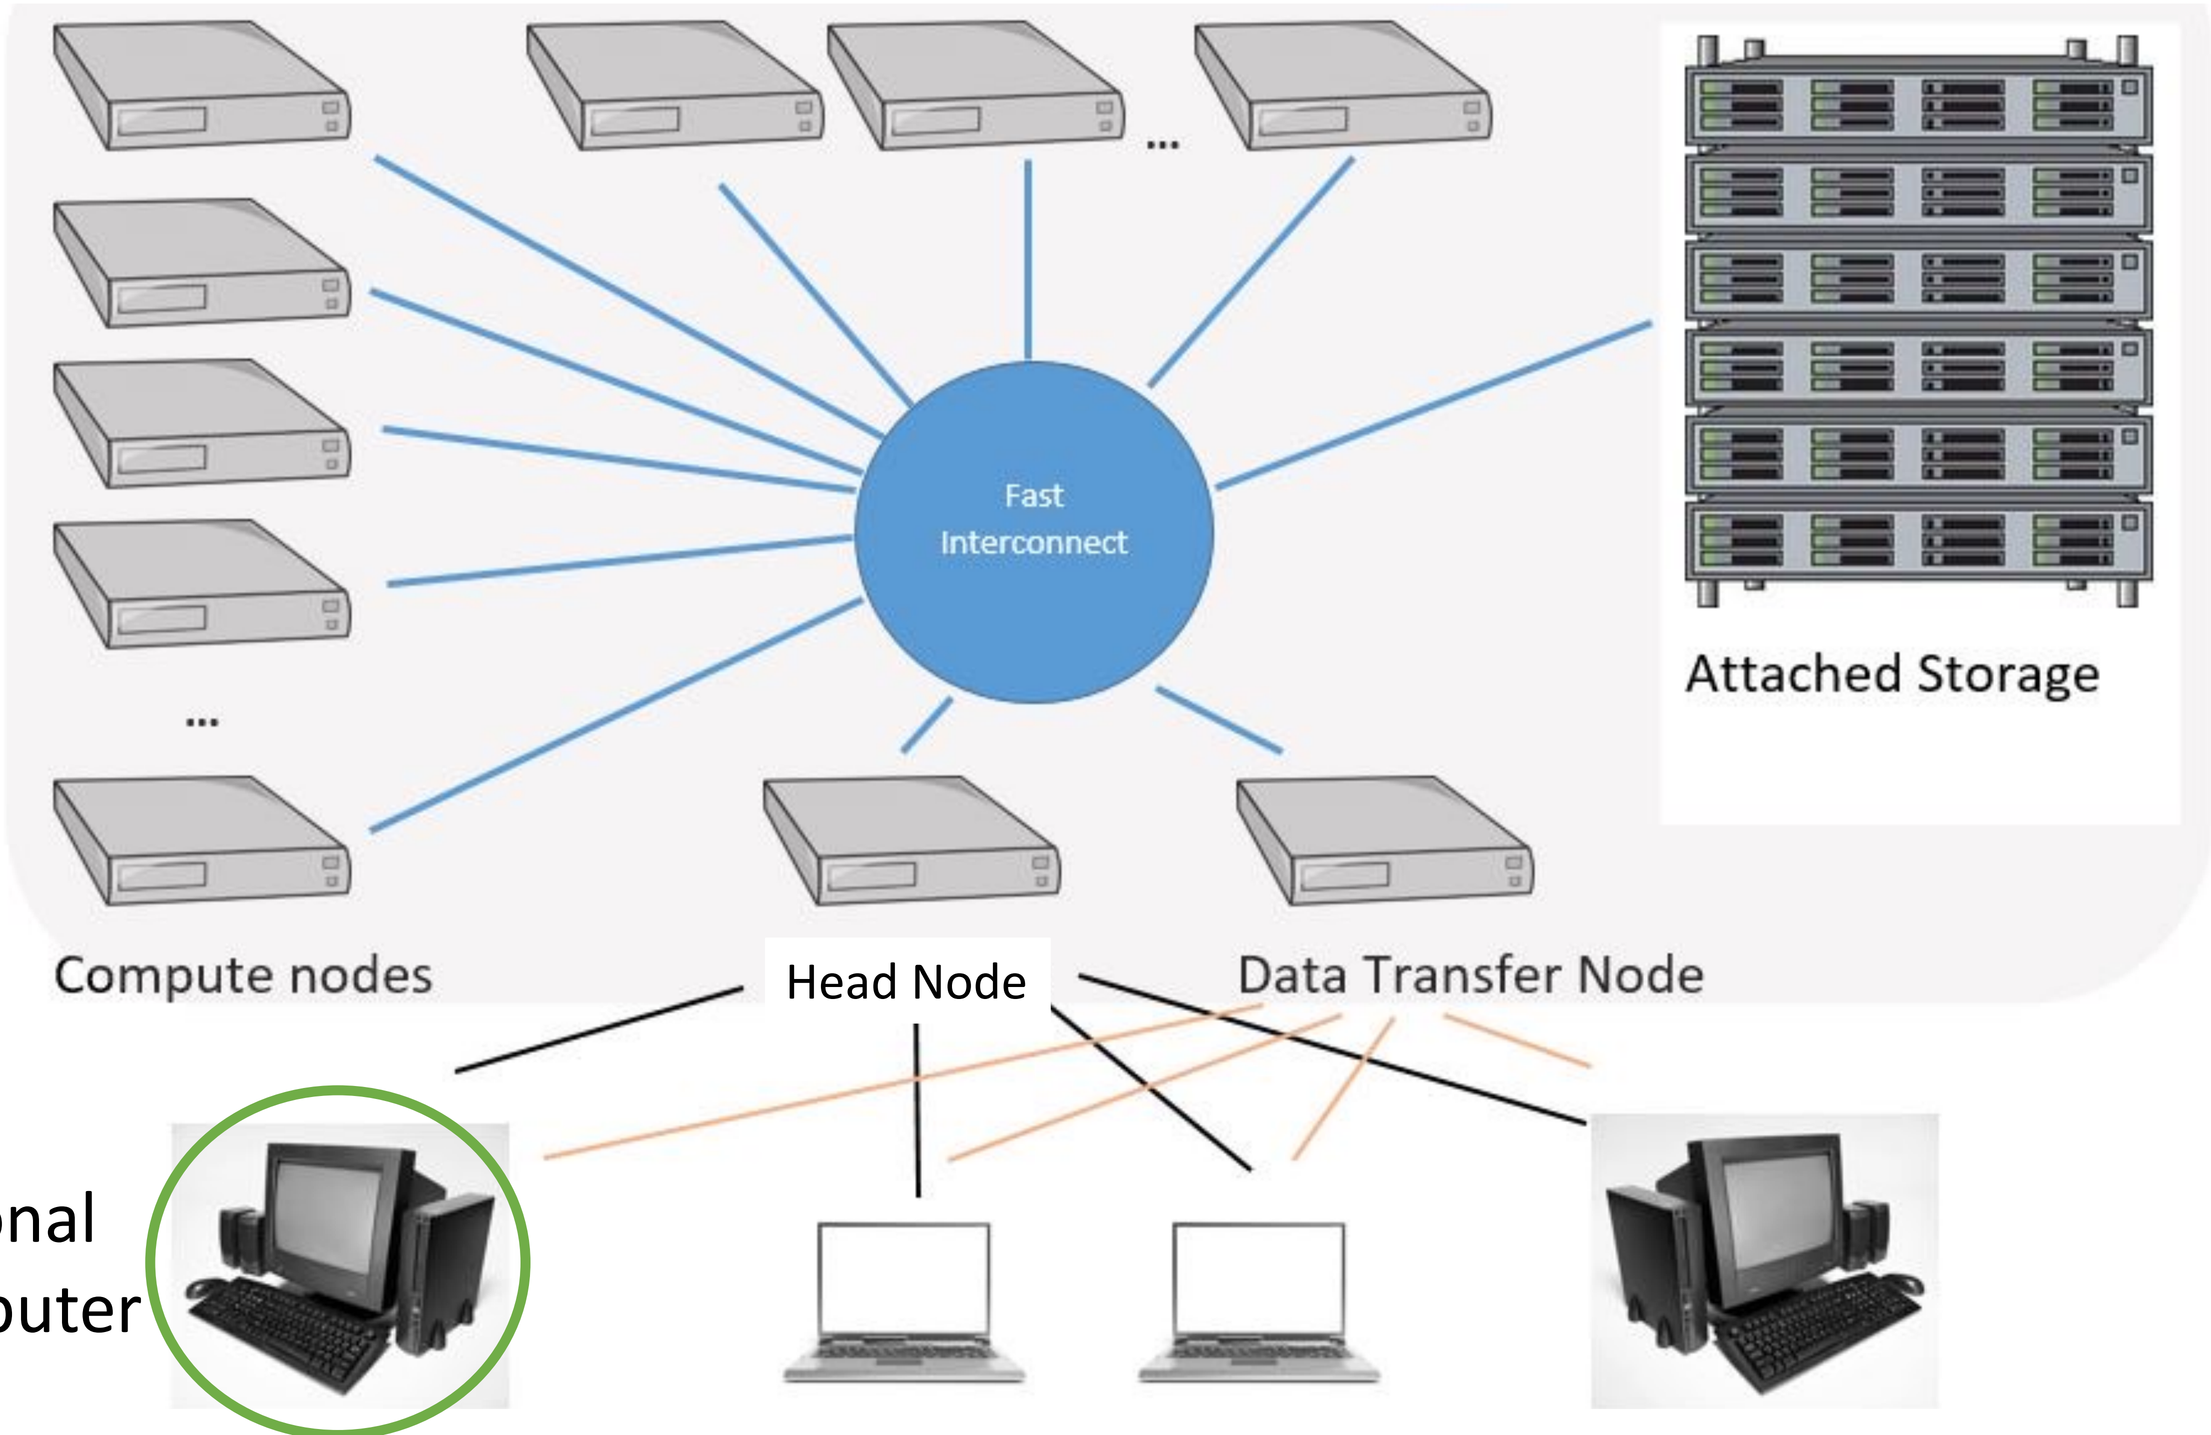

...

Fast Interconnect

Attached Storage

...

Compute nodes

Head Node

Data Transfer Node

Personal Computer

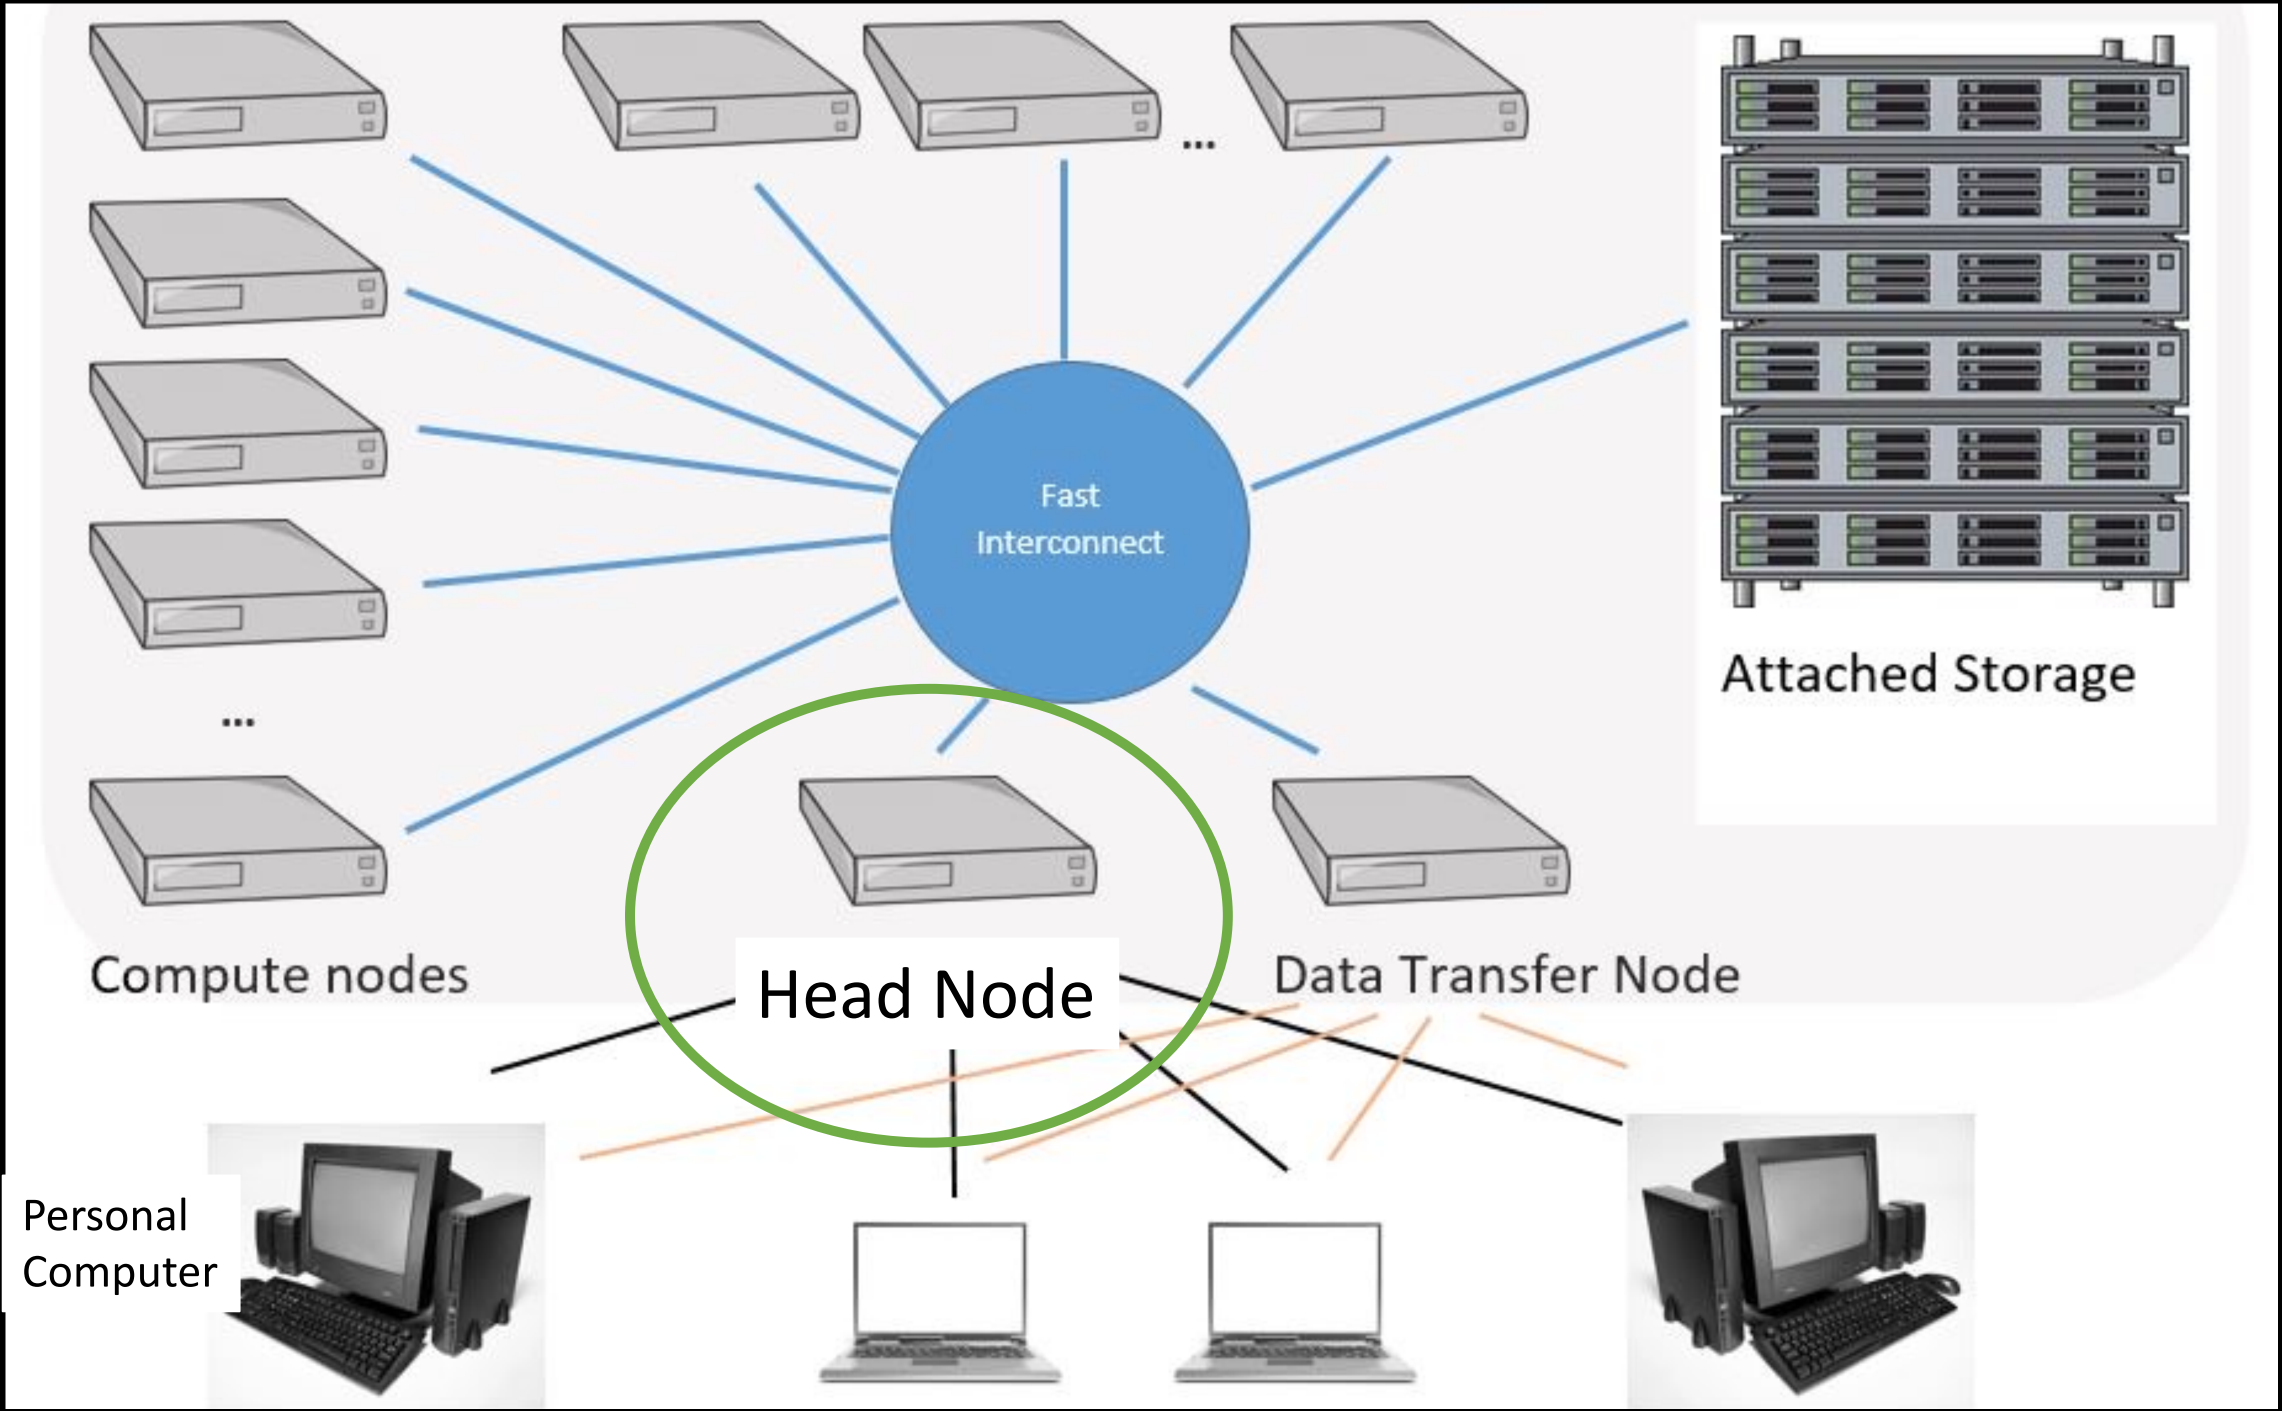

Fast Interconnect

…

…

Attached Storage

Compute nodes

Head Node

Data Transfer Node

Personal Computer

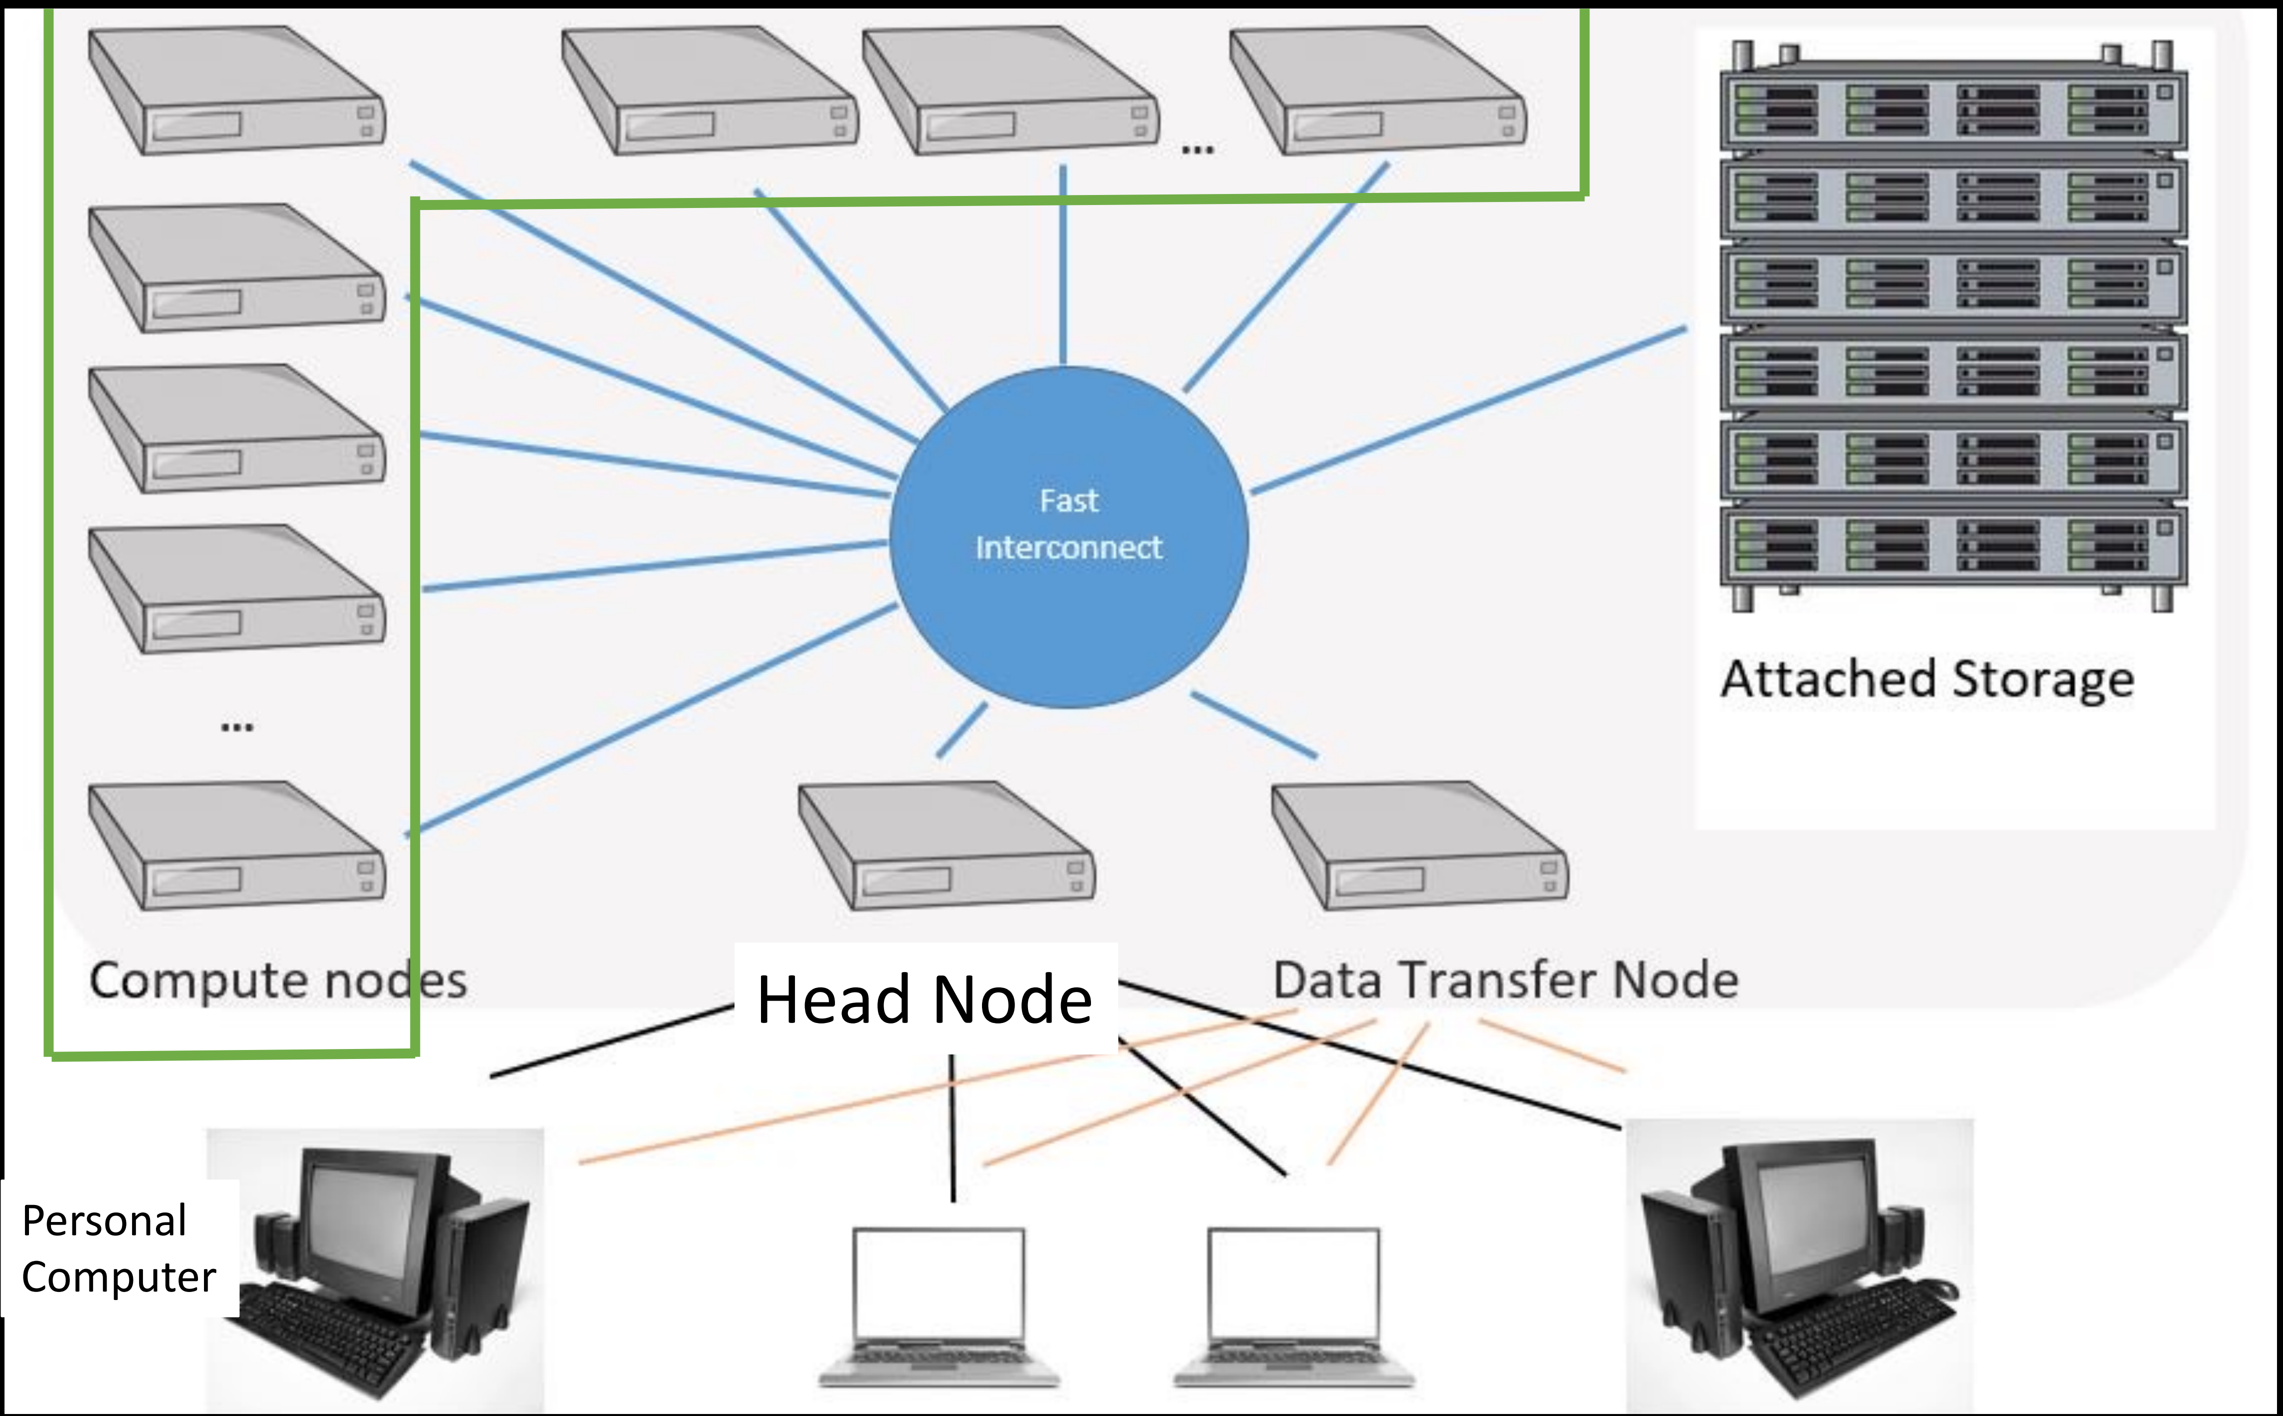
Fast Interconnect
...
...
Attached Storage
Compute nodes
Head Node
Data Transfer Node
Personal Computer

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…

Head Node (wheeler)

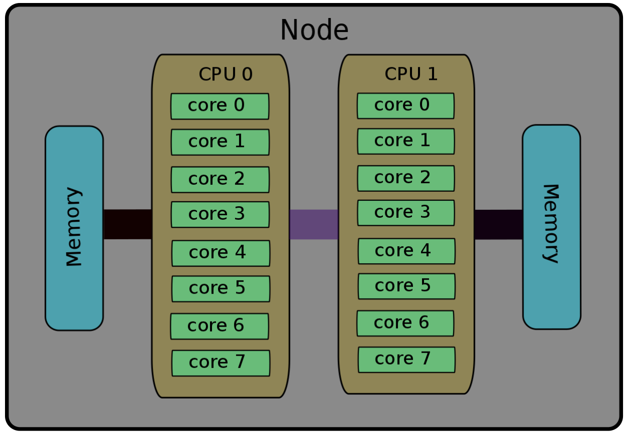

Compute Nodes

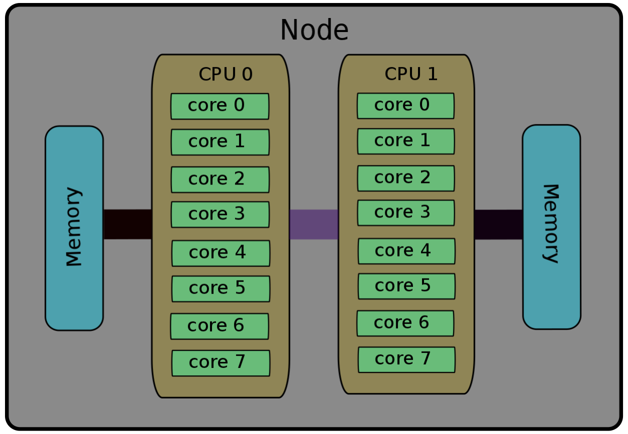

wheeler001

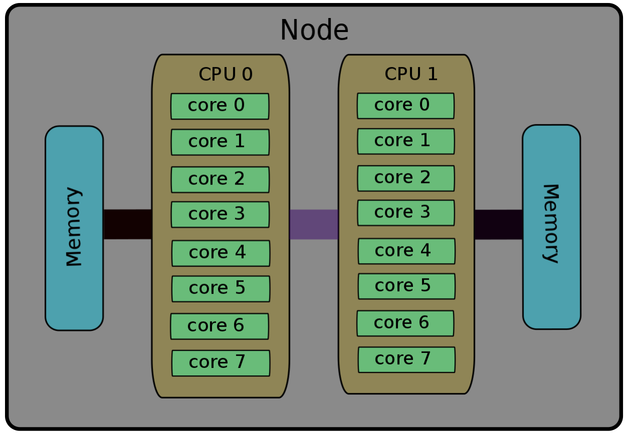

wheeler002

...

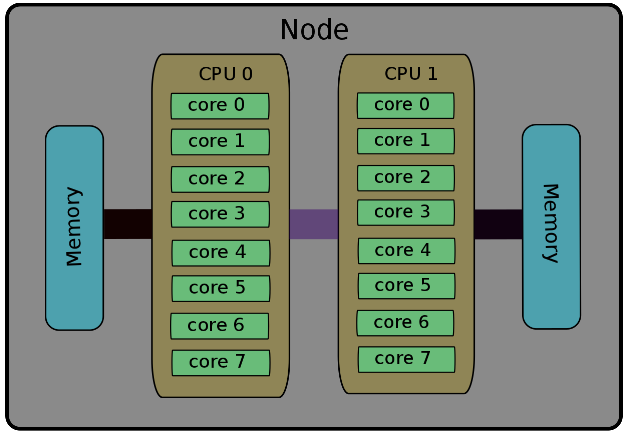

wheeler304

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!
SLURMS McKENZIE

```
[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
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…

```
[vanilla@wheeler ~]$ cp -r /projects/shared/workshops/beginner/vecadd ~
[vanilla@wheeler ~]$
```

Review, what does this command do?

```
[vanilla@wheeler ~]$ cd vecadd/
[vanilla@wheeler ~/vecadd]$ module load openmpi/4.1.2-q2zi
```

What do these commands do?

```
[vanilla@wheeler ~/vecadd]$ srun --partition debug --ntasks 4 vecaddmpi
```

Now run the program with “srun”…

```
[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
```

srun is good but HPC centers are busy!

Workflow

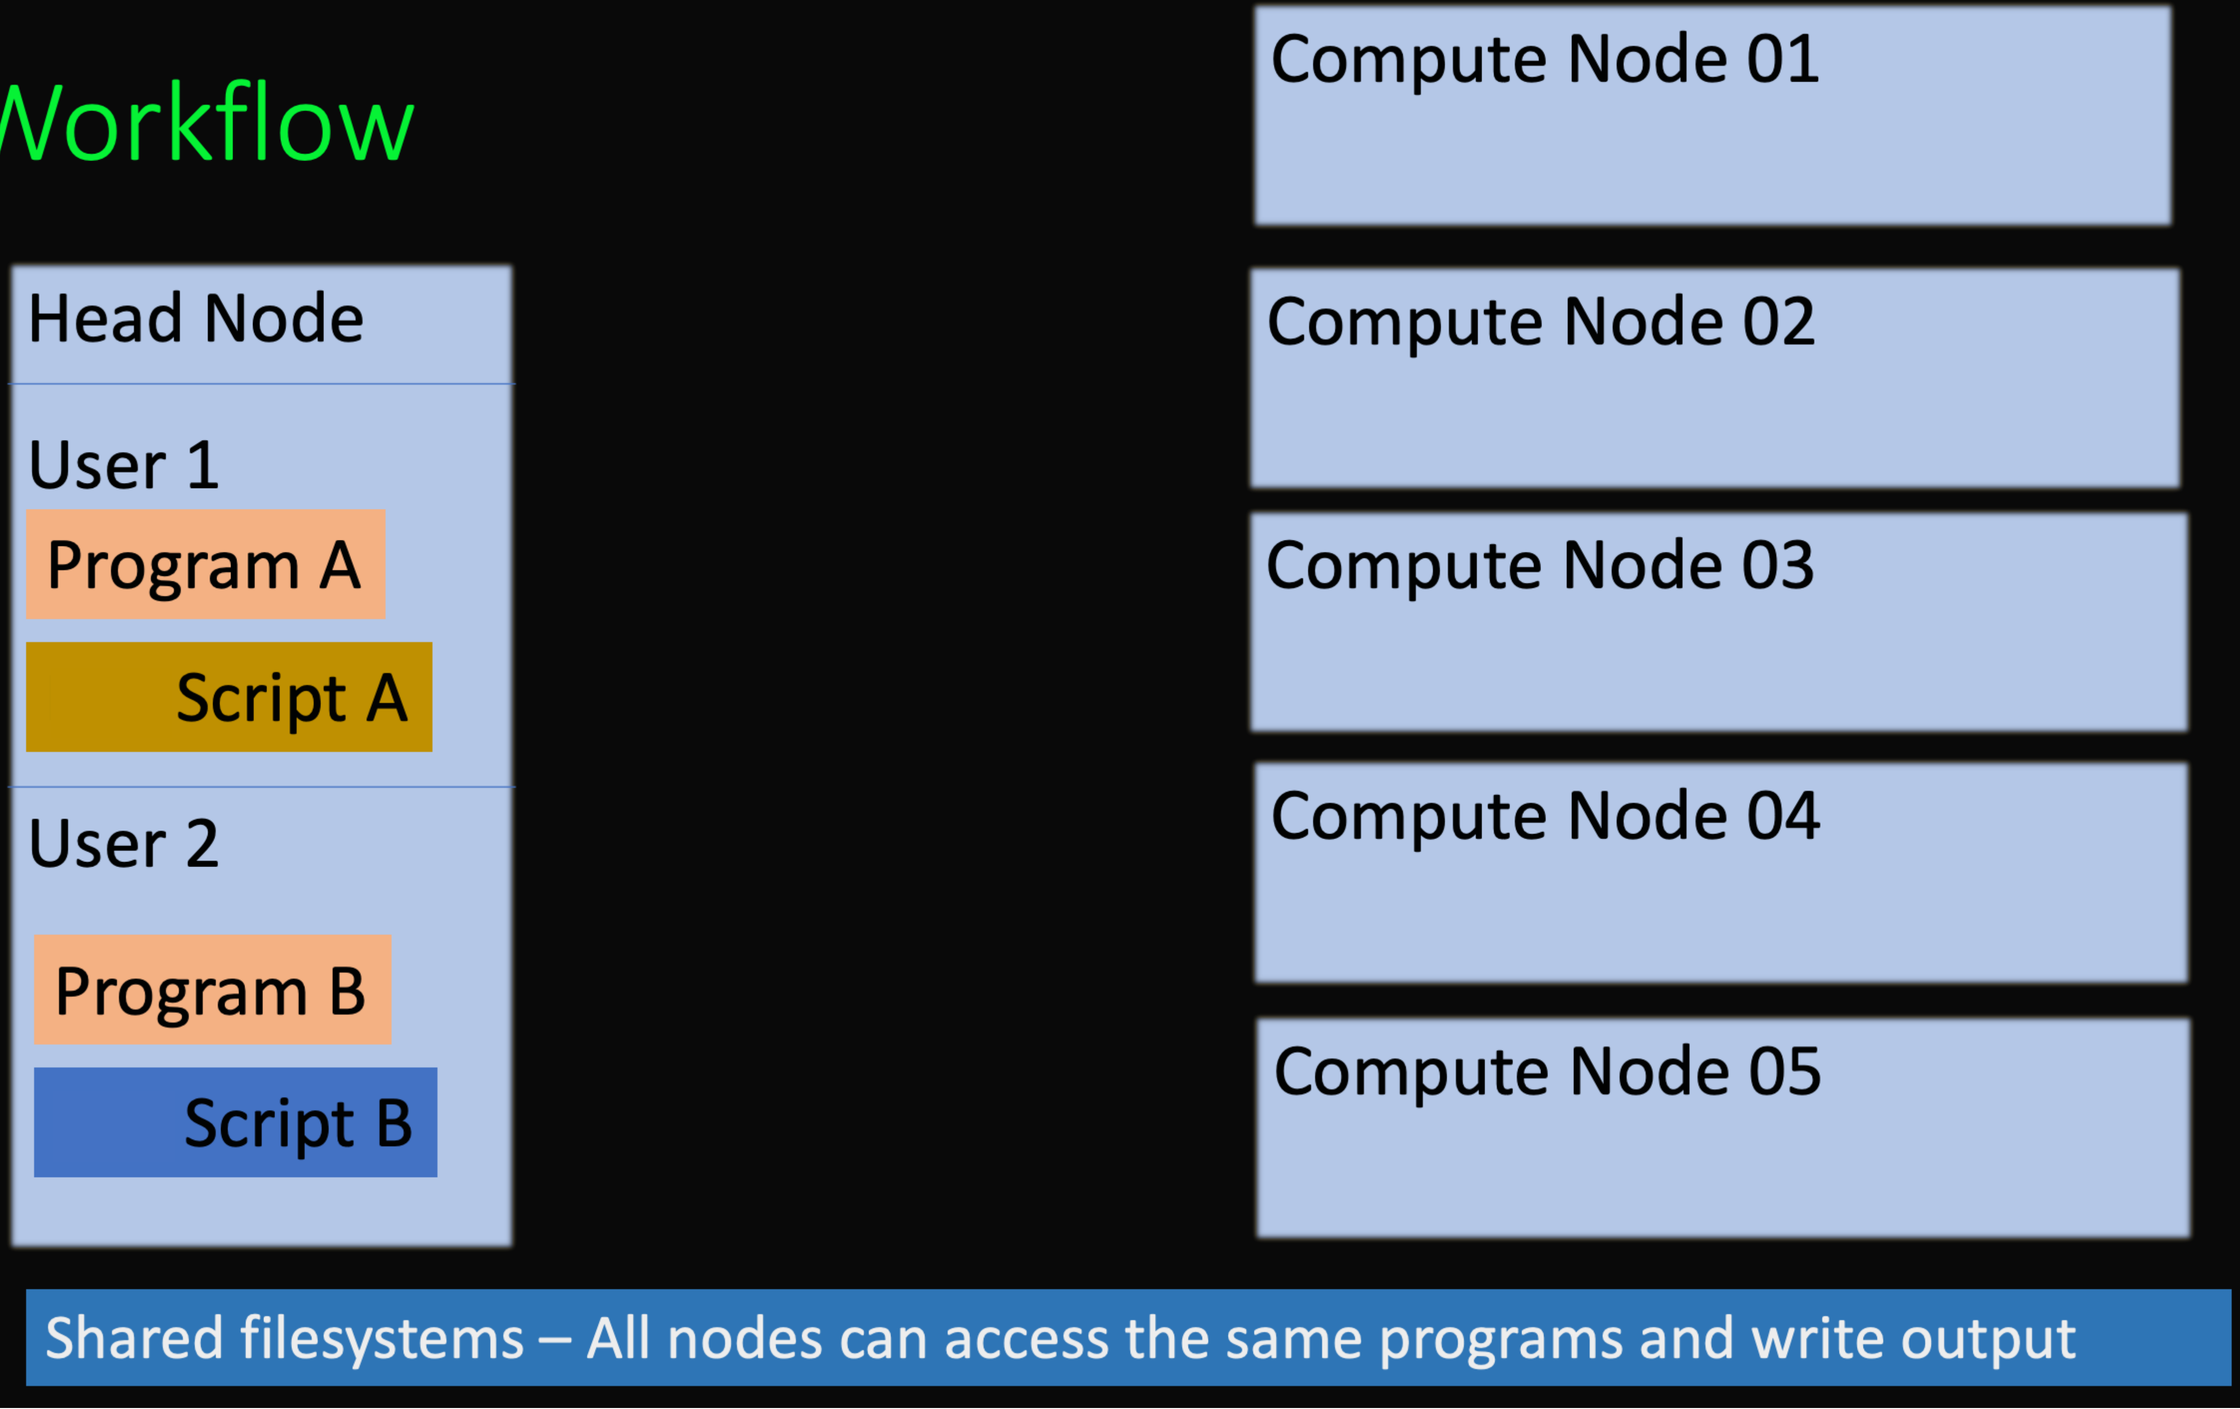

Head Node

User 1

Program A

Script A

User 2

Program B

Script B

Compute Node 01

Compute Node 02

Compute Node 03

Compute Node 04

Compute Node 05

Shared filesystems – All nodes can access the same programs and write output

Workflow

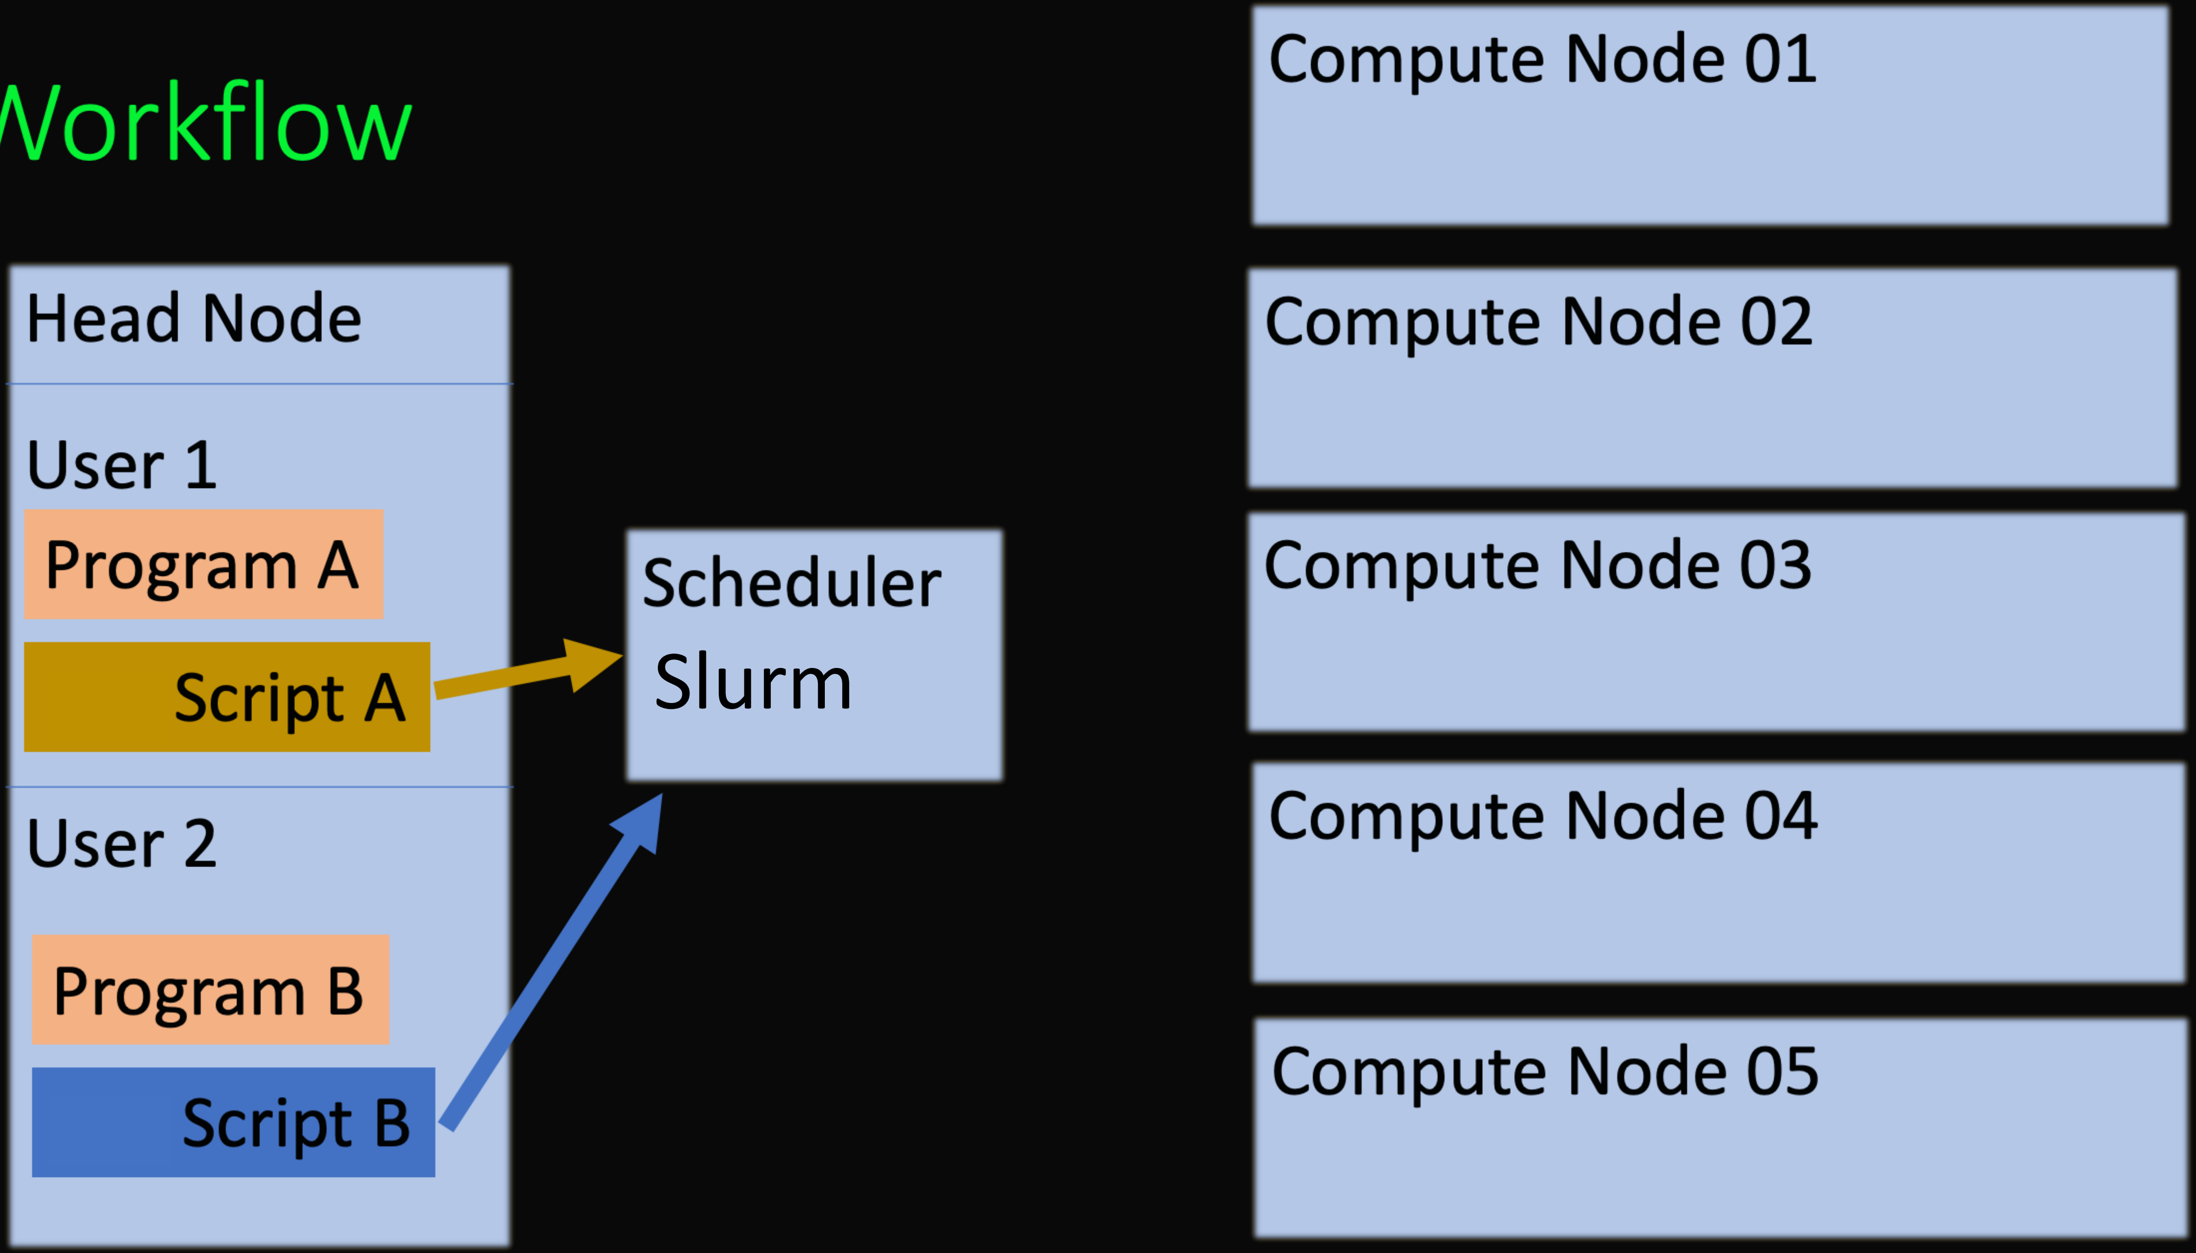

Head Node

User 1

Program A

Script A

User 2

Program B

Script B

Scheduler

Slurm

Compute Node 01

Compute Node 02

Compute Node 03

Compute Node 04

Compute Node 05

Shared filesystems – All nodes can access the same programs and write output

Workflow

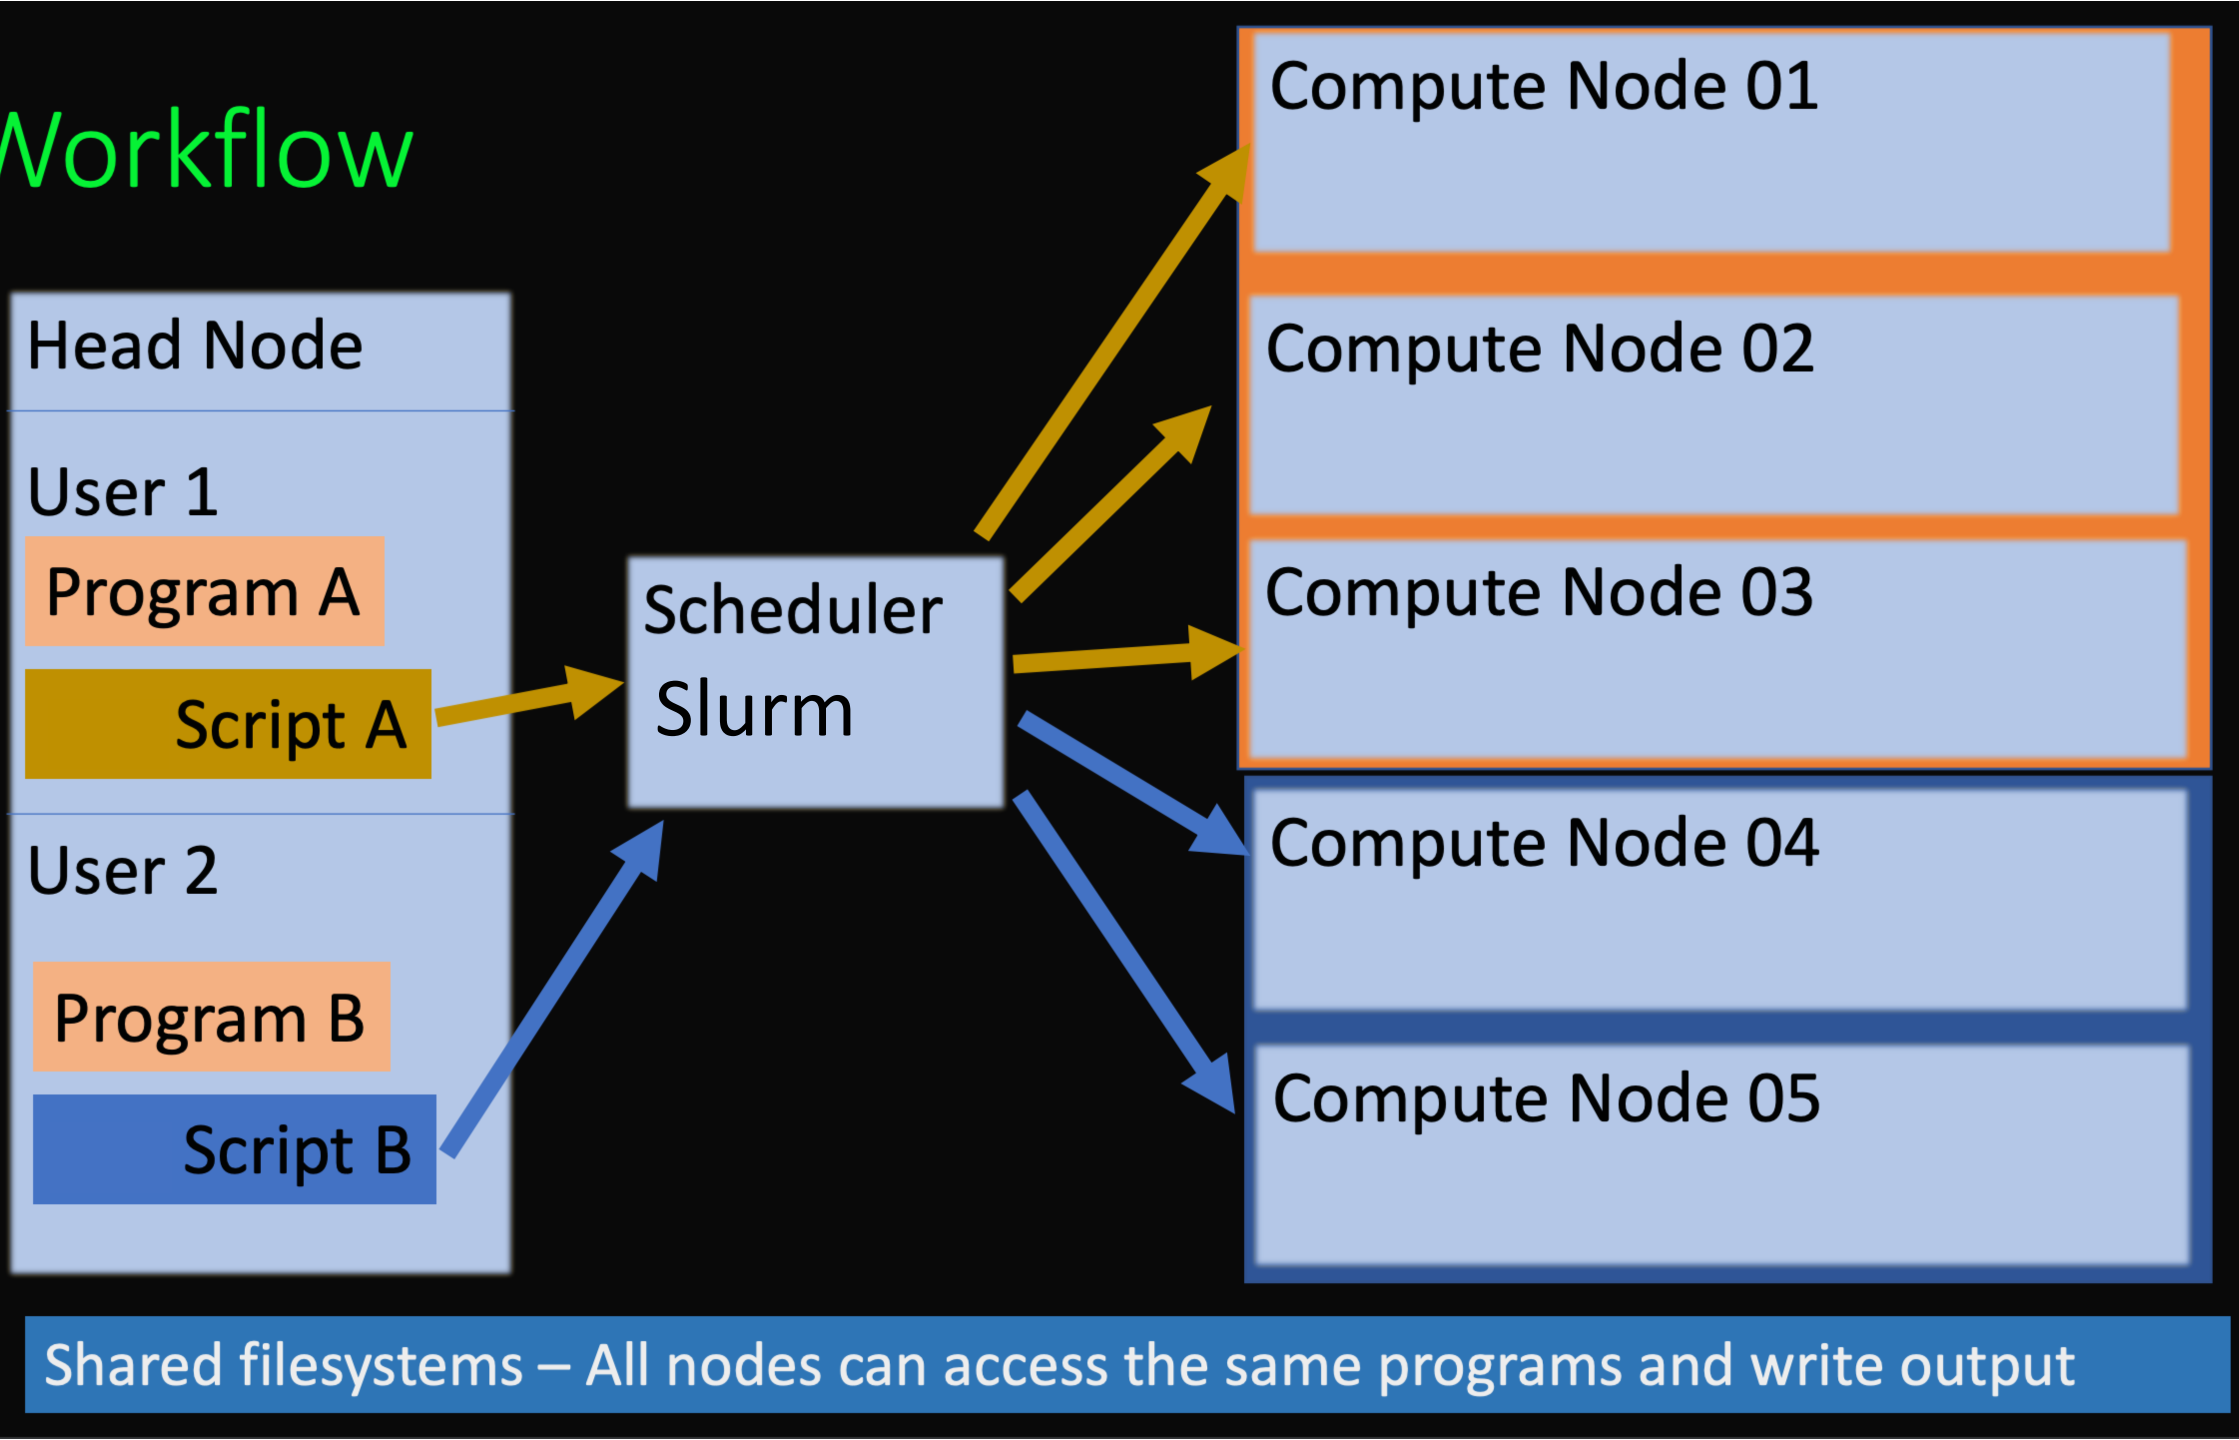

Head Node

User 1

Program A

Script A

User 2

Program B

Script B

Scheduler

Slurm

Compute Node 01

Compute Node 02

Compute Node 03

Compute Node 04

Compute Node 05

Shared filesystems – All nodes can access the same programs and write output

Workflow

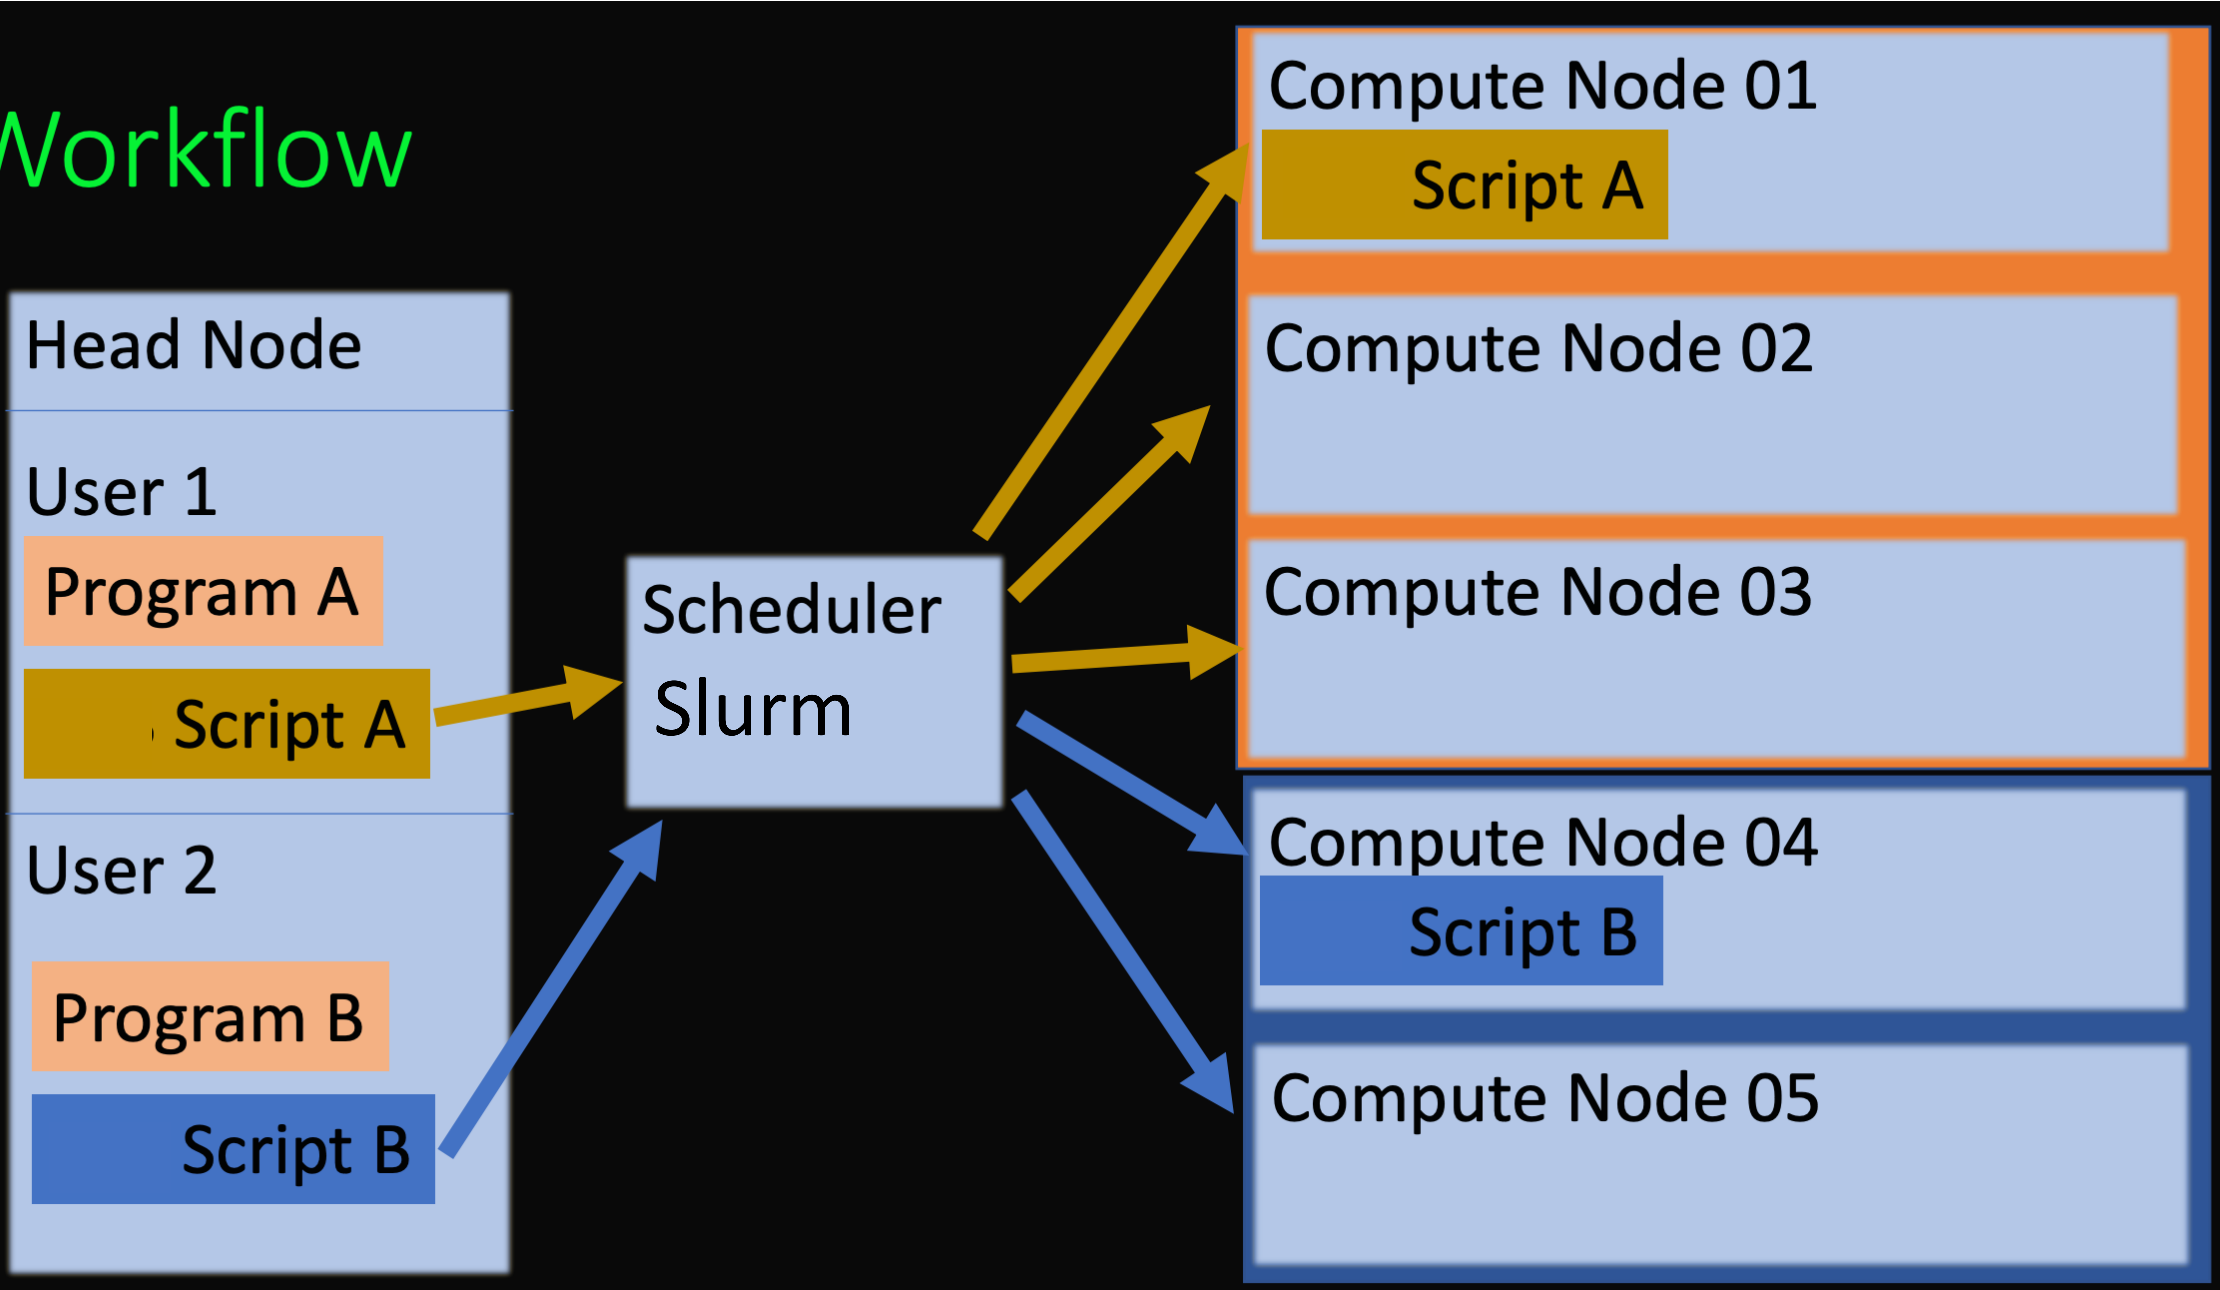

Shared filesystems – All nodes can access the same programs and write output

Workflow

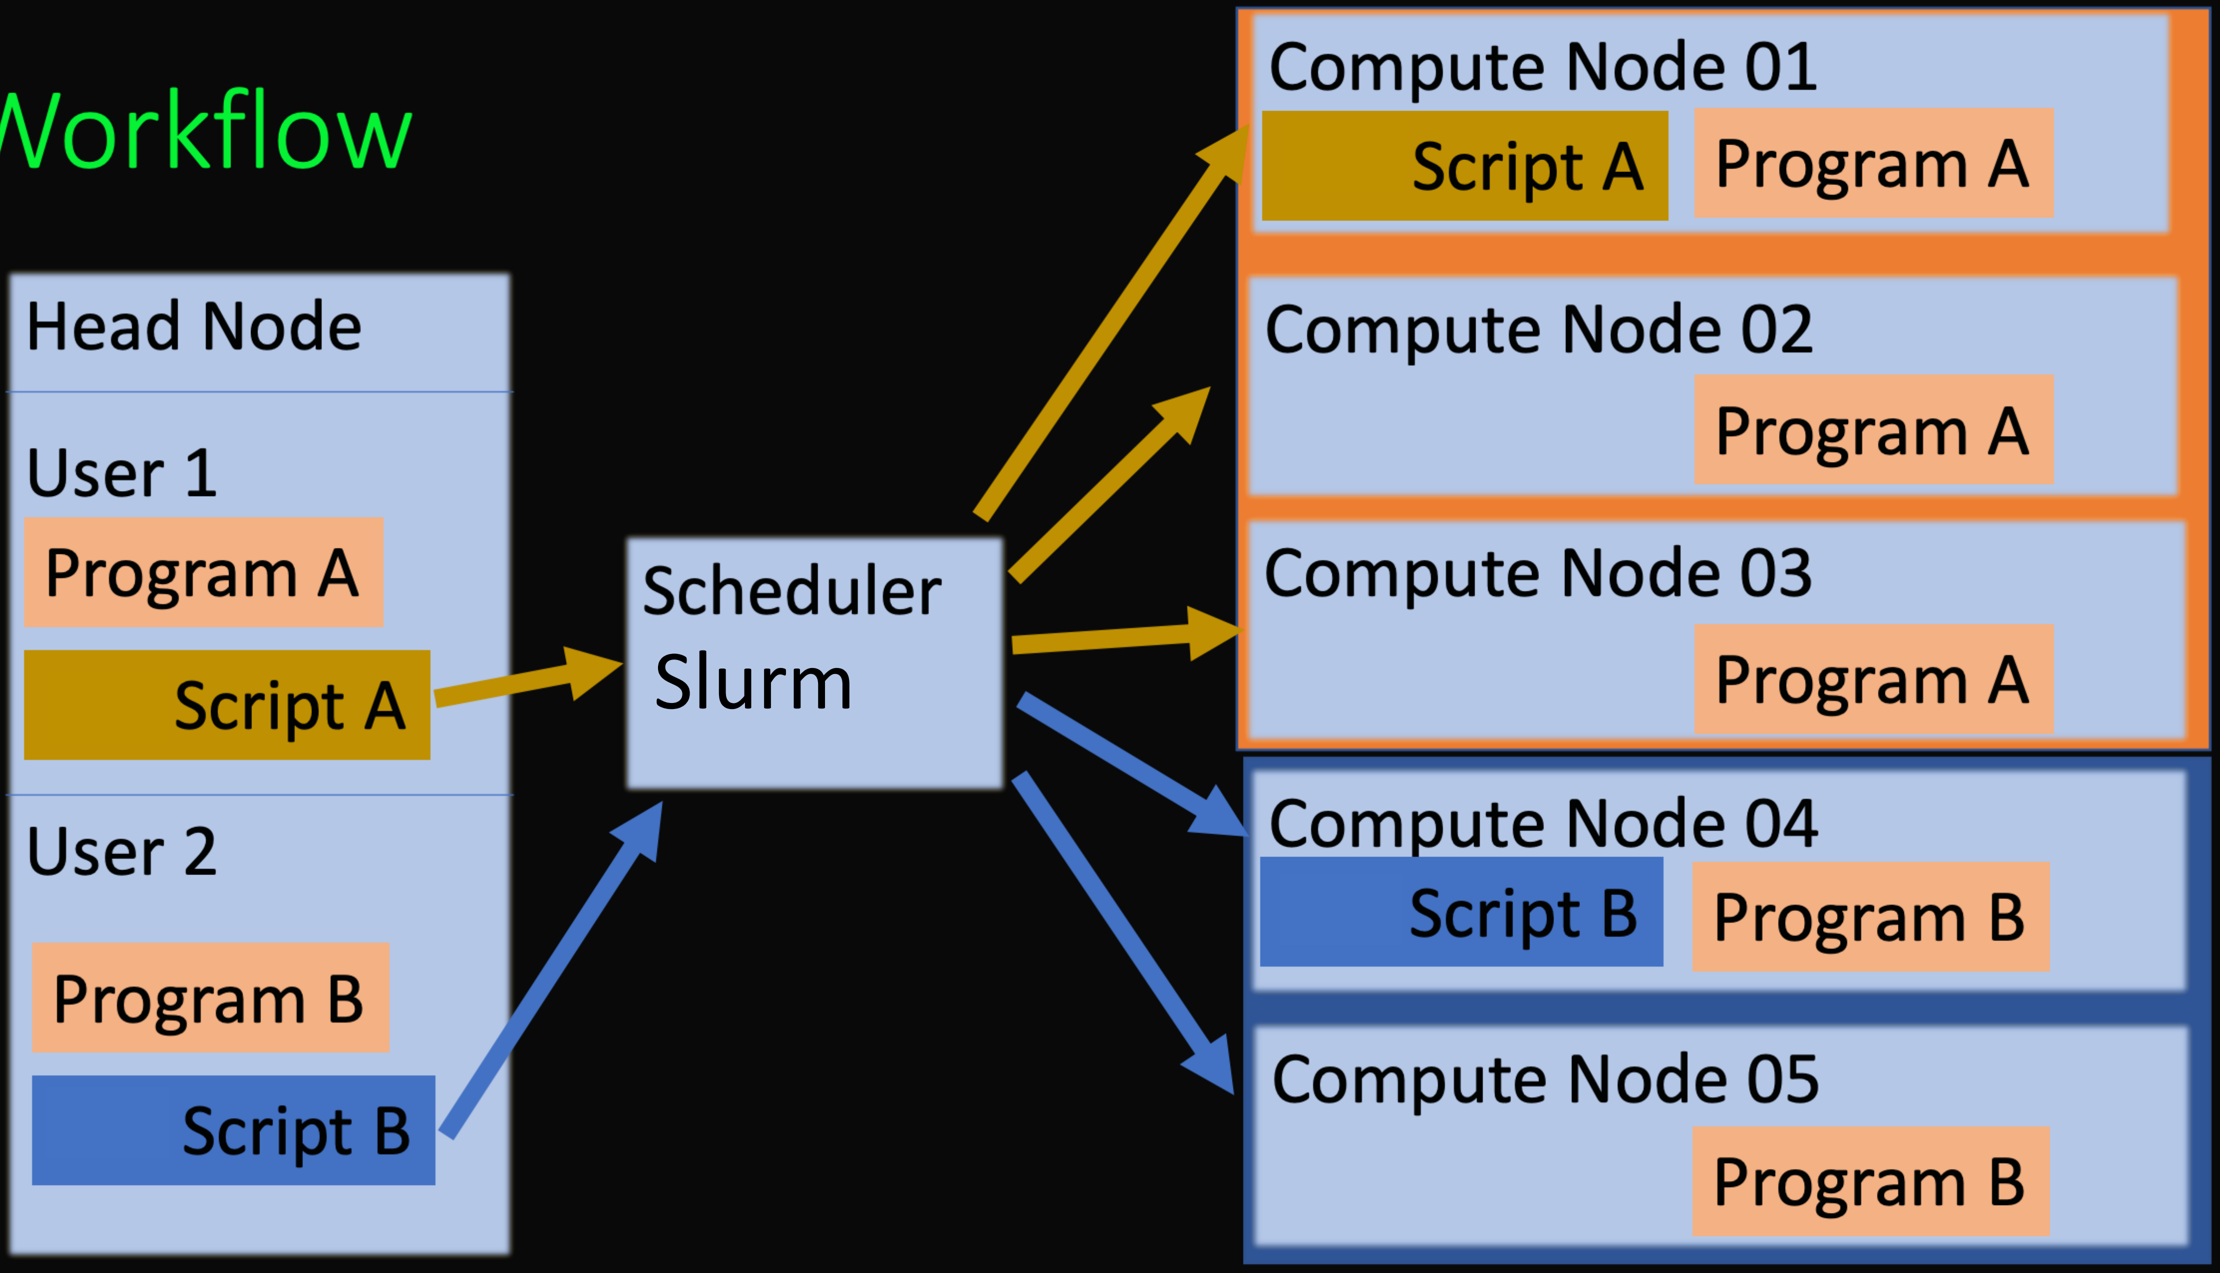

Head Node

User 1

Program A

Script A

User 2

Program B

Script B

Scheduler

Slurm

Compute Node 01

Script A

Program A

Compute Node 02

Program A

Compute Node 03

Program A

Compute Node 04

Script B

Program B

Compute Node 05

Program B

Shared filesystems – All nodes can access the same programs and write output

Scheduler

```
qgrok
```

```
sbatch
```

```
squeue –u USERNAME
```

Running Programs on Compute Nodes

- qgrok
- Intro to the Slurm Scheduler
- The srun command
- sbatch
- sinfo
- squeue
- squeue -u username --start
- https://www.cism.ucl.ac.be/Services/Formations/slurm/2016/slurm.pdf
- https://www.nrel.gov/hpc/assets/pdfs/slurm-advanced-topics.pdf

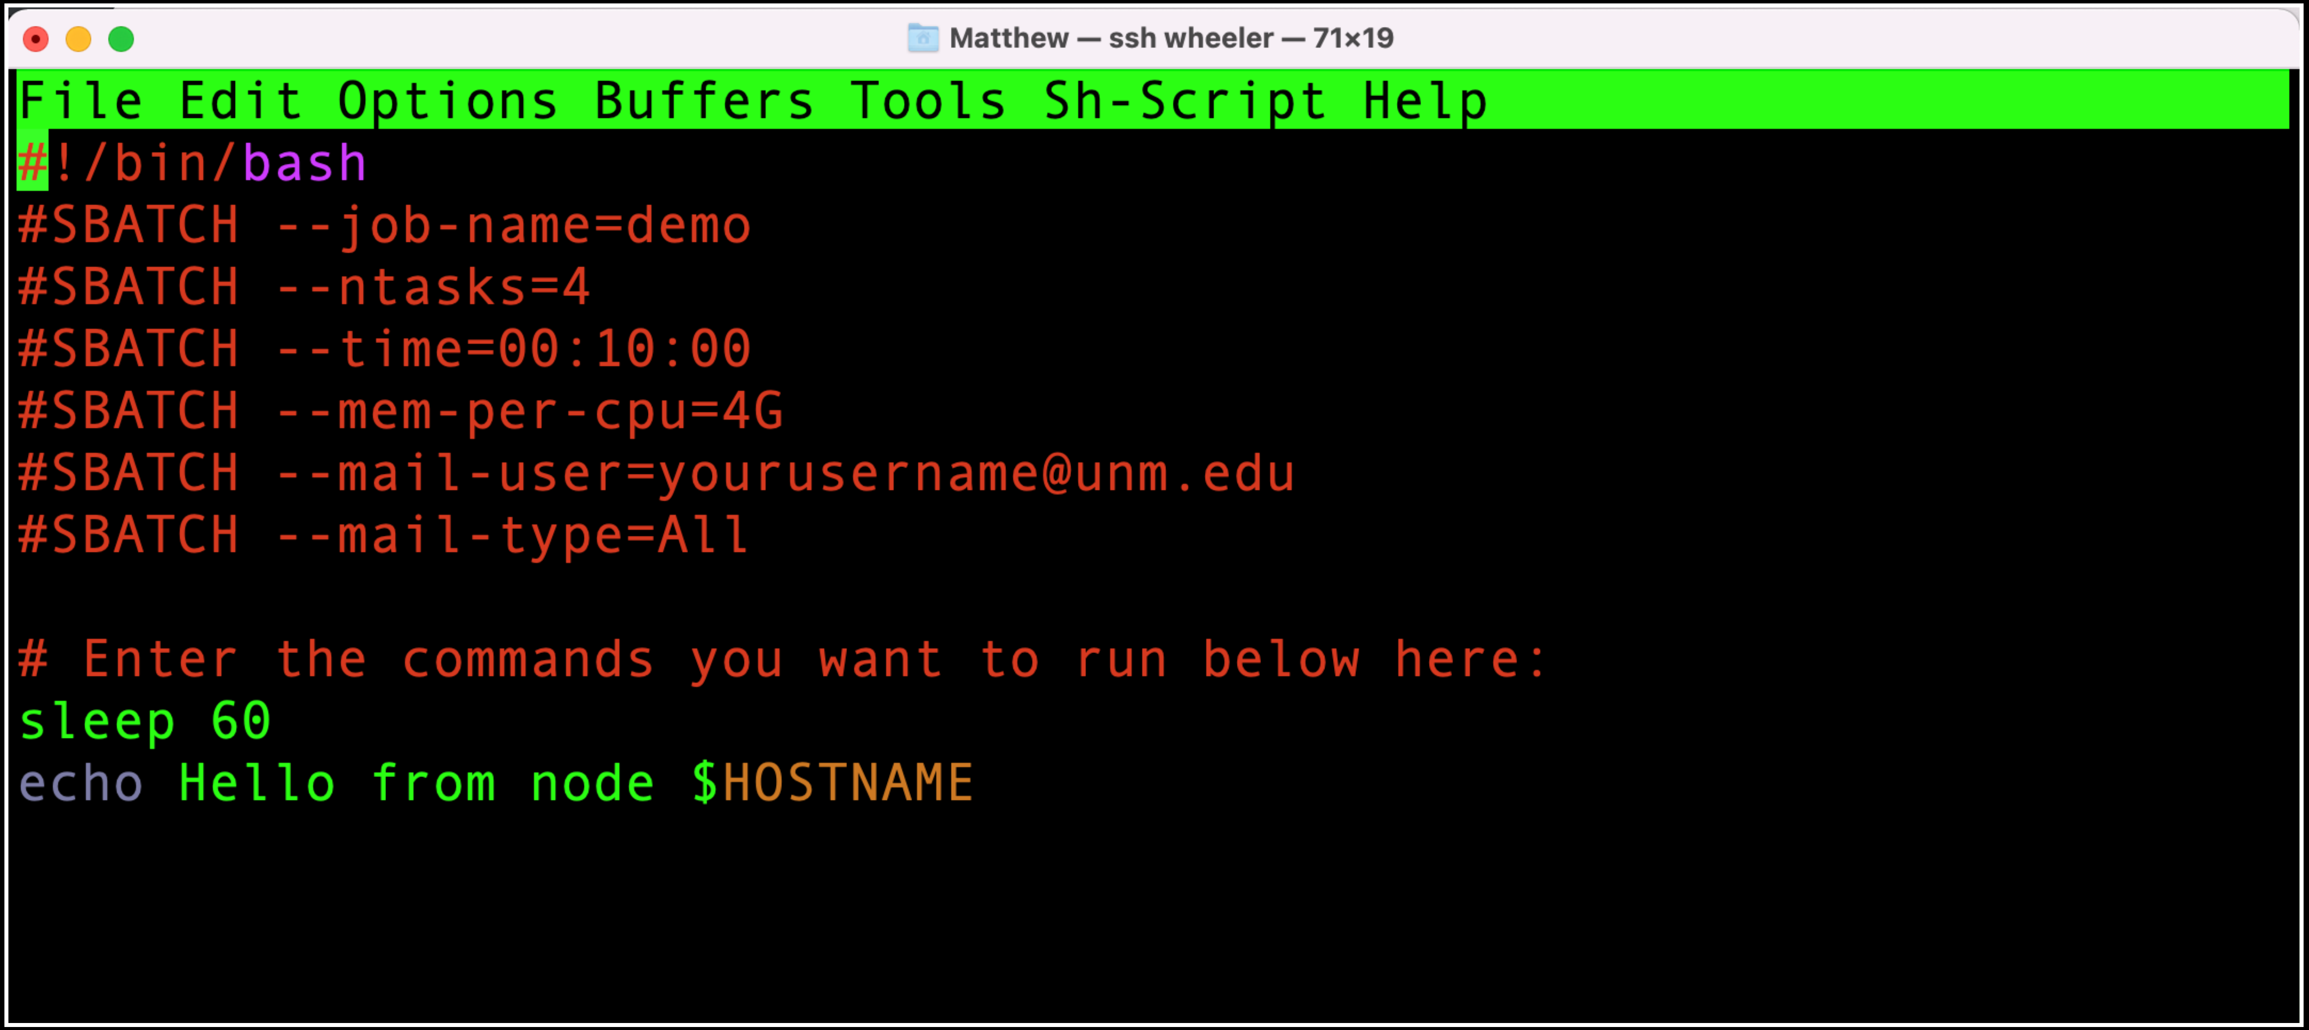

```
#!/bin/bash
#SBATCH --job-name=demo
#SBATCH --ntasks=4
#SBATCH --time=00:10:00
#SBATCH --mem-per-cpu=4G
#SBATCH --mail-user=yourusername@unm.edu
#SBATCH --mail-type=All

# Enter the commands you want to run below here:
sleep 60
echo Hello from node $HOSTNAME
```

Slurm Script

```
[vanilla@wheeler ~/vecadd]$ cat vecaddmpi.sh
#!/bin/bash
#SBATCH --job-name=vecaddmpi
#SBATCH --ntasks=4
#SBATCH --time=00:10:00
#SBATCH --mem-per-cpu=4G
#SBATCH --mail-user=yourusername@unm.edu
#SBATCH --mail-type=All
#SBATCH --output=vecaddmpi.out


module load openmpi/4.1.2-q2zi
srun --mpi=pmi2 ./vecaddmpi
```

Slurm Script

```
vanilla@wheeler:~/vecadd $ sbatch vecaddmpi.sh
sbatch: Using account 2016199 from ~/.default_slurm_account
Submitted batch job 167571


vanilla@wheeler:~/vecadd $ squeue --me
             JOBID PARTITION     NAME     USER ST       TIME  NODES
NODELIST(REASON)
            167571    normal vecaddmp  vanilla  R       0:07      1 wheeler145
```

Slurm Script

```
vanilla@wheeler:~/vecadd $ tail -f vecaddmpi.out
You have not been allocated GPUs. To request GPUs, use the -G option in your
submission script.
Assigning compute node to rank 1.
ComputeNode: Starting with rank 1.
ComputeNode (1): Waiting for vectors from dataserver with rank 3...
Assigning compute node to rank 2.
ComputeNode: Starting with rank 2.
ComputeNode (2): Waiting for vectors from dataserver with rank 3...
Will try to allocate a vector of size 1 GB.
```

Slurm Script

```
vanilla@wheeler:~/vecadd $ tail -f vecaddmpi.out
You have not been allocated GPUs. To request GPUs, use the -G option in your
submission script.
Assigning compute node to rank 1.
ComputeNode: Starting with rank 1.
ComputeNode (1): Waiting for vectors from dataserver with rank 3...
Assigning compute node to rank 2.
ComputeNode: Starting with rank 2.
ComputeNode (2): Waiting for vectors from dataserver with rank 3...
Will try to allocate a vector of size 1 GB.
```

Slurm Script

Useful Slurm Commands

squeue --me --long	shows information about jobs you submitted
squeue --me --start	shows when slurm expects your job to start
scancel jobid	cancels a job
sacct	shows your job history

```
vanilla@wheeler:~/vecadd $ seff 167573
Job ID: 167573
Cluster: wheeler
User/Group: mfricke/users
State: COMPLETED (exit code 0)
Nodes: 1
Cores per node: 4
CPU Utilized: 00:01:03
CPU Efficiency: 78.75% of 00:01:20 core-walltime
Job Wall-clock time: 00:00:20
Memory Utilized: 39.55 MB (estimated maximum)
Memory Efficiency: 0.24% of 16.00 GB (4.00 GB/core)
```

Slurm Script

CARC Resources

- Tutorial Videos
- Written Tutorials

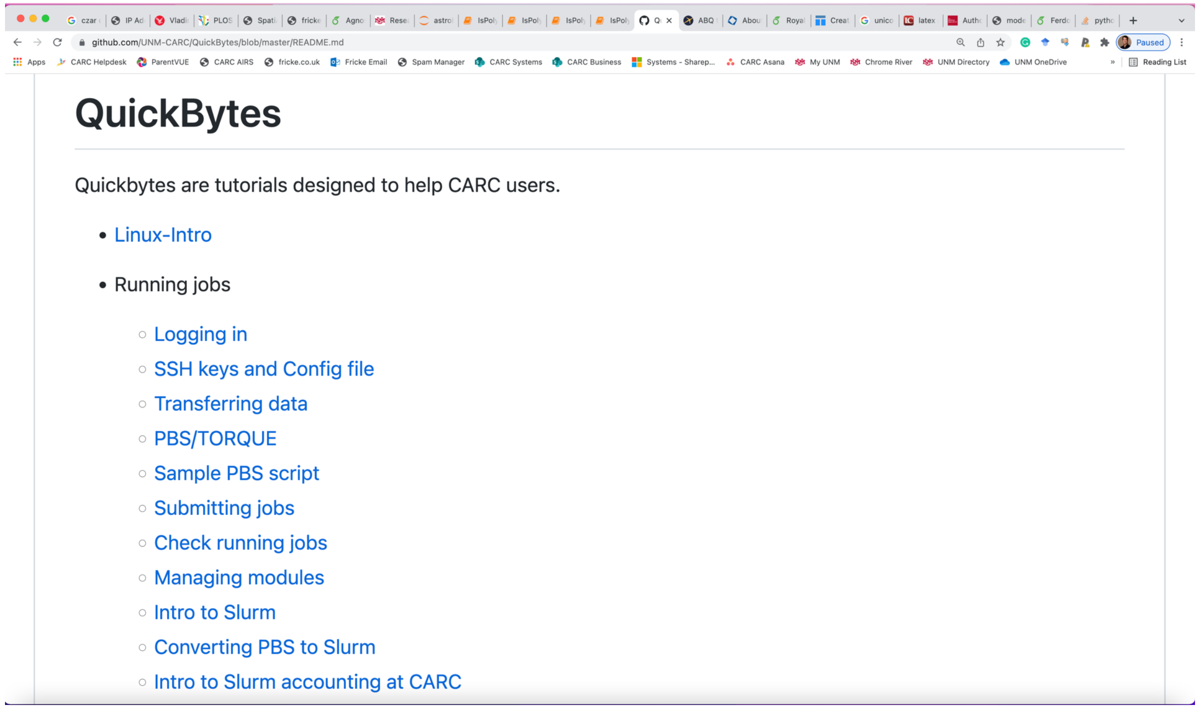

QuickBytes

Quickbytes are tutorials designed to help CARC users.

- Linux-Intro
- Running jobs
  - Logging in
  - SSH keys and Config file
  - Transferring data
  - PBS/TORQUE
  - Sample PBS script
  - Submitting jobs
  - Check running jobs
  - Managing modules
  - Intro to Slurm
  - Converting PBS to Slurm
  - Intro to Slurm accounting at CARC

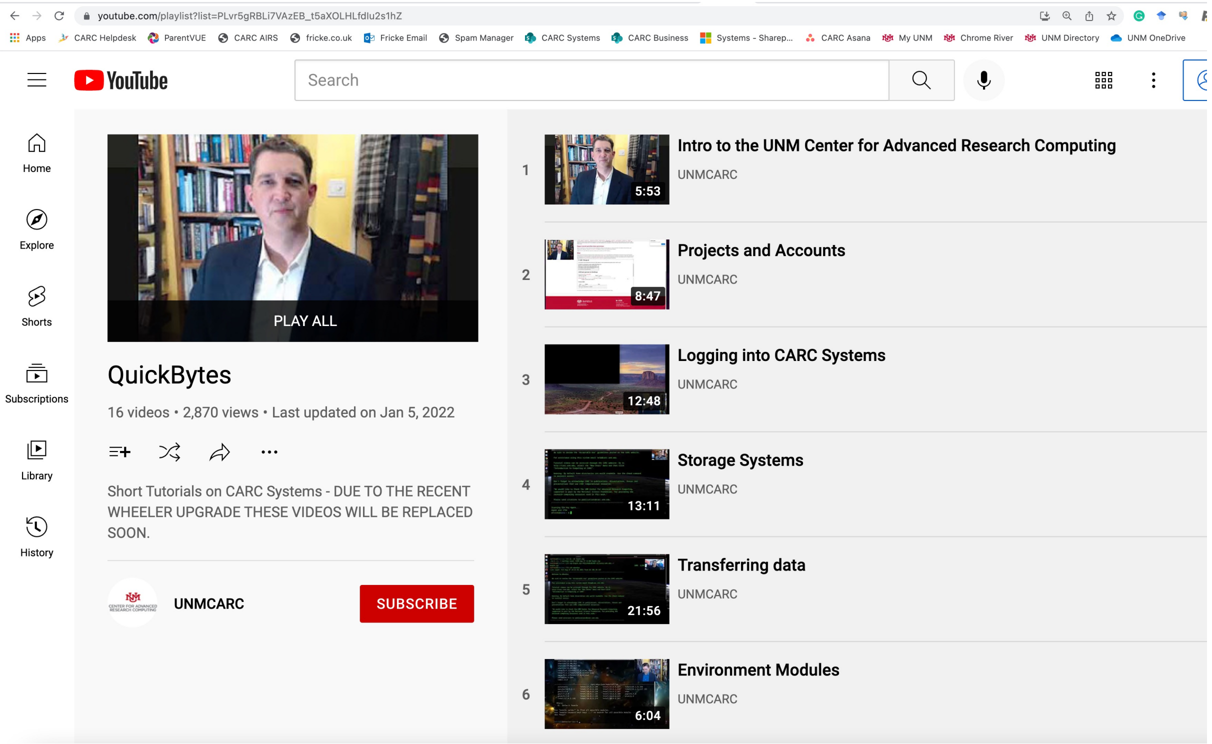

QuickBytes

16 videos • 2,870 views • Last updated on Jan 5, 2022

Short Tutorials on CARC Systems - DUE TO THE RECENT WHEELER UPGRADE THESE VIDEOS WILL BE REPLACED SOON.

UNMCARC

	Video	Length	Channel
1	Intro to the UNM Center for Advanced Research Computing	5:53	UNMCARC
2	Projects and Accounts	8:47	UNMCARC
3	Logging into CARC Systems	12:48	UNMCARC
4	Storage Systems	13:11	UNMCARC
5	Transferring data	21:56	UNMCARC
6	Environment Modules	6:04	UNMCARC

Getting Help

help@carc.unm.edu

Office hours