General Purpose Cluster (GPC) and Other High Performance Computing Resources at Penn

Professor Xi Song

Department of Sociology

Methodology Working Group, Spring 2023

Why Do You Need HPC?

You Will Need HPC When You Need To...

- Have more computing power than a desktop can provide
- Run many models at the same time
- Run a model for a huge dataset that is bigger than your hard disk
- Use parallel computing to save time on a model that may take a very long time to run (e.g., Bayesian models)
- Run a model using software that is not installed in your own computer (e.g., Stata MP)

Virtual Private Network (VPN)

FortiClient

- A VPN (Virtual Private Network) Client is necessary for Remote Desktop access, when connecting to Windows Active Directory (AD) Share drives and U:'s, or other Penn resources from off-campus or on WiFi.
- You will need to install VPN in order to get access to GPC off campus
- https://computing.sas.upenn.edu/faculty_staff/VPN

Install FortiClient

- Fill out the webform to request a KITE account for VPN access https://upenn.co1.qualtrics.com/jfe/form/SV_d6lrw3Ta4gJ0Av3
- Send an email to Alex Bond (sschelp@sas.upenn.edu) to let him know about your application
- Download FortiClient from https://computing.sas.upenn.edu/faculty_staff/VPN

Connect to FortiClient

- You will be asked to use the Two-Factor Authentification sent to your Duo Mobile app.
- You will be prompted for a password, which is your typical Pennkey password.
- you will need to do a two-factor authentication method. If you have the Duo Mobile app installed and set up on your phone, simply type **1** in the password field and hit enter to get a push notification on your phone. Otherwise, enter of the other numbers for different options.

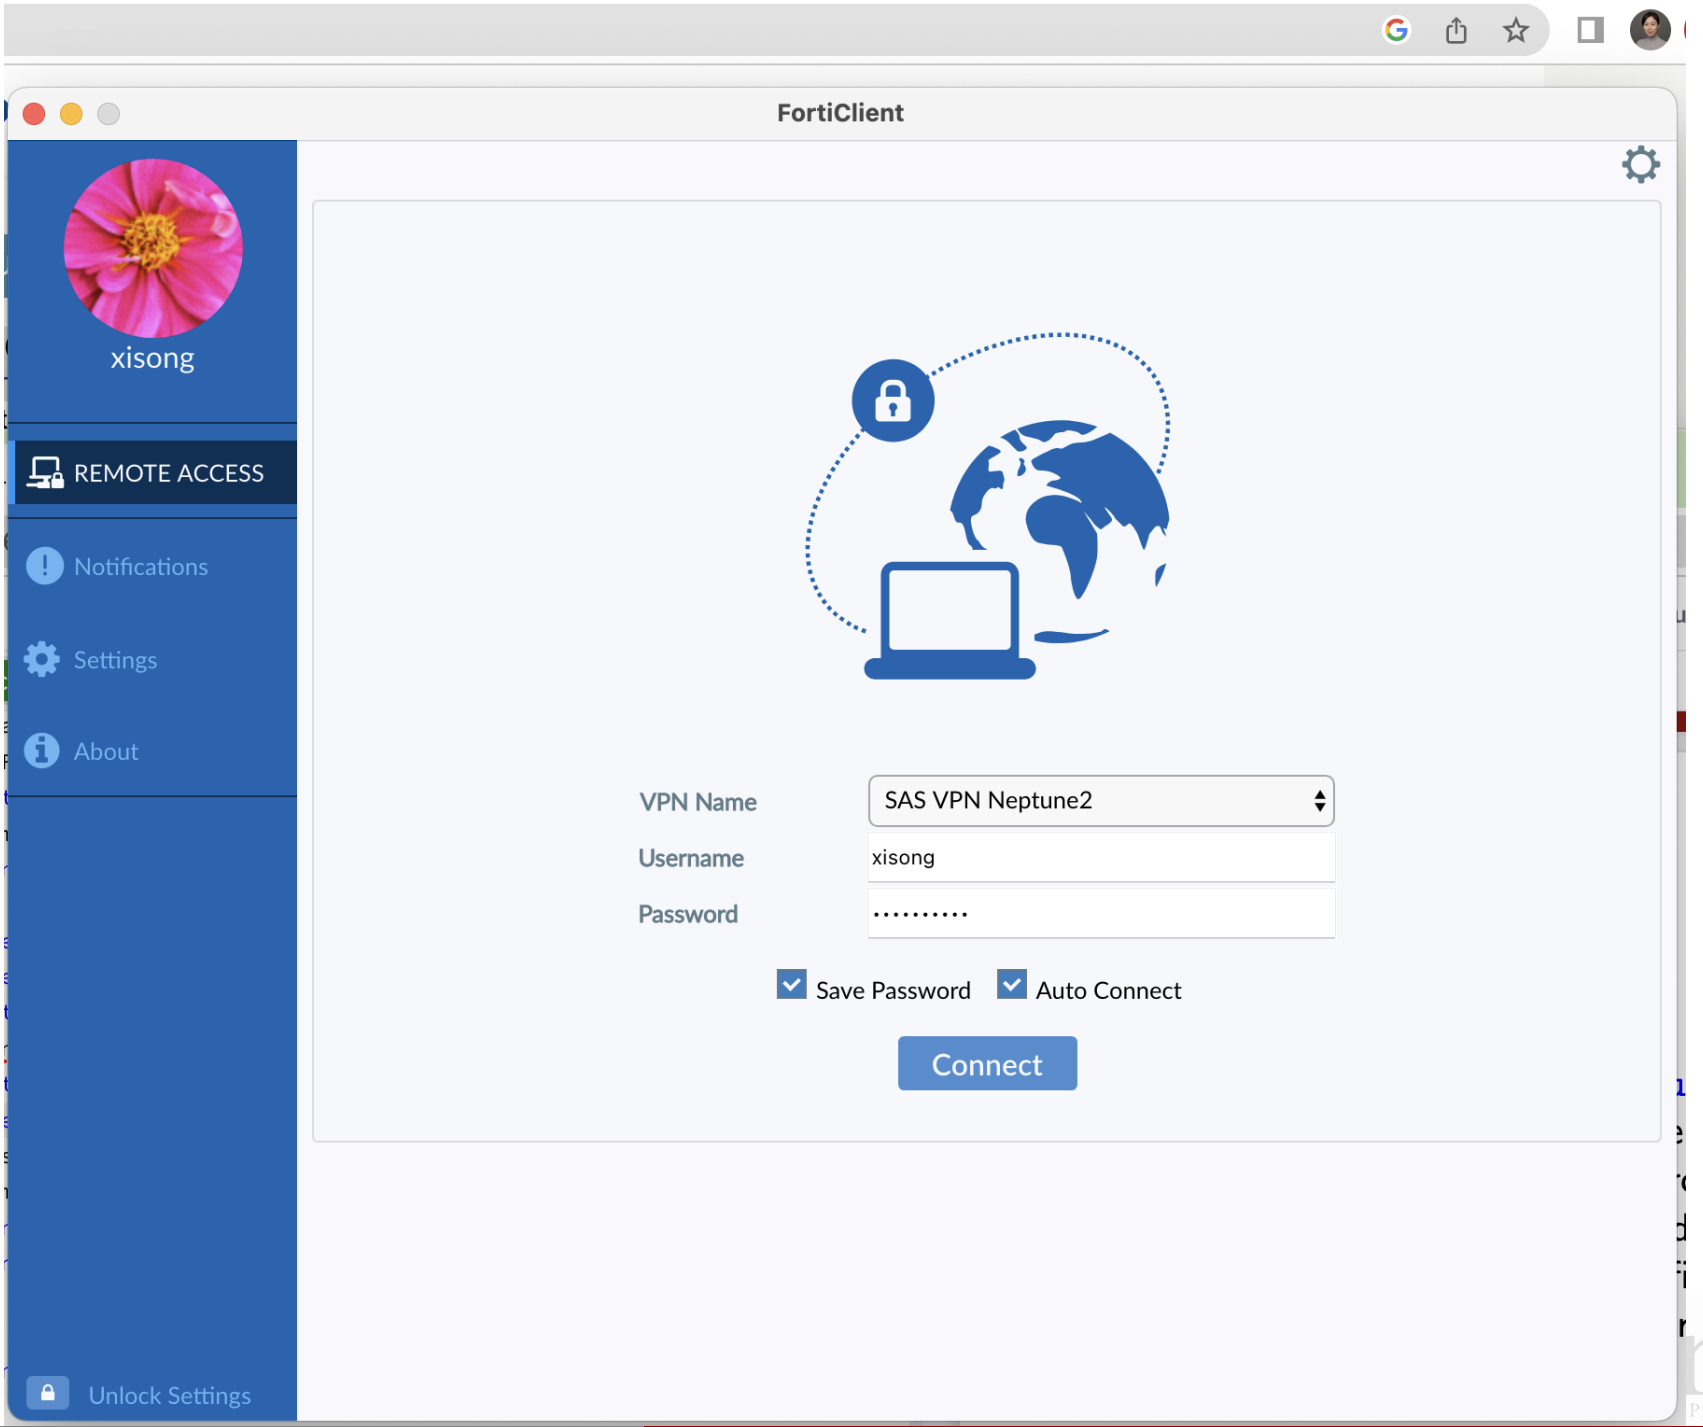
FortiClient
xisong
REMOTE ACCESS
Notifications
Settings
About
Unlock Settings
VPN Name
SAS VPN Neptune2
Username
xisong
Password
Save Password
Auto Connect
Connect

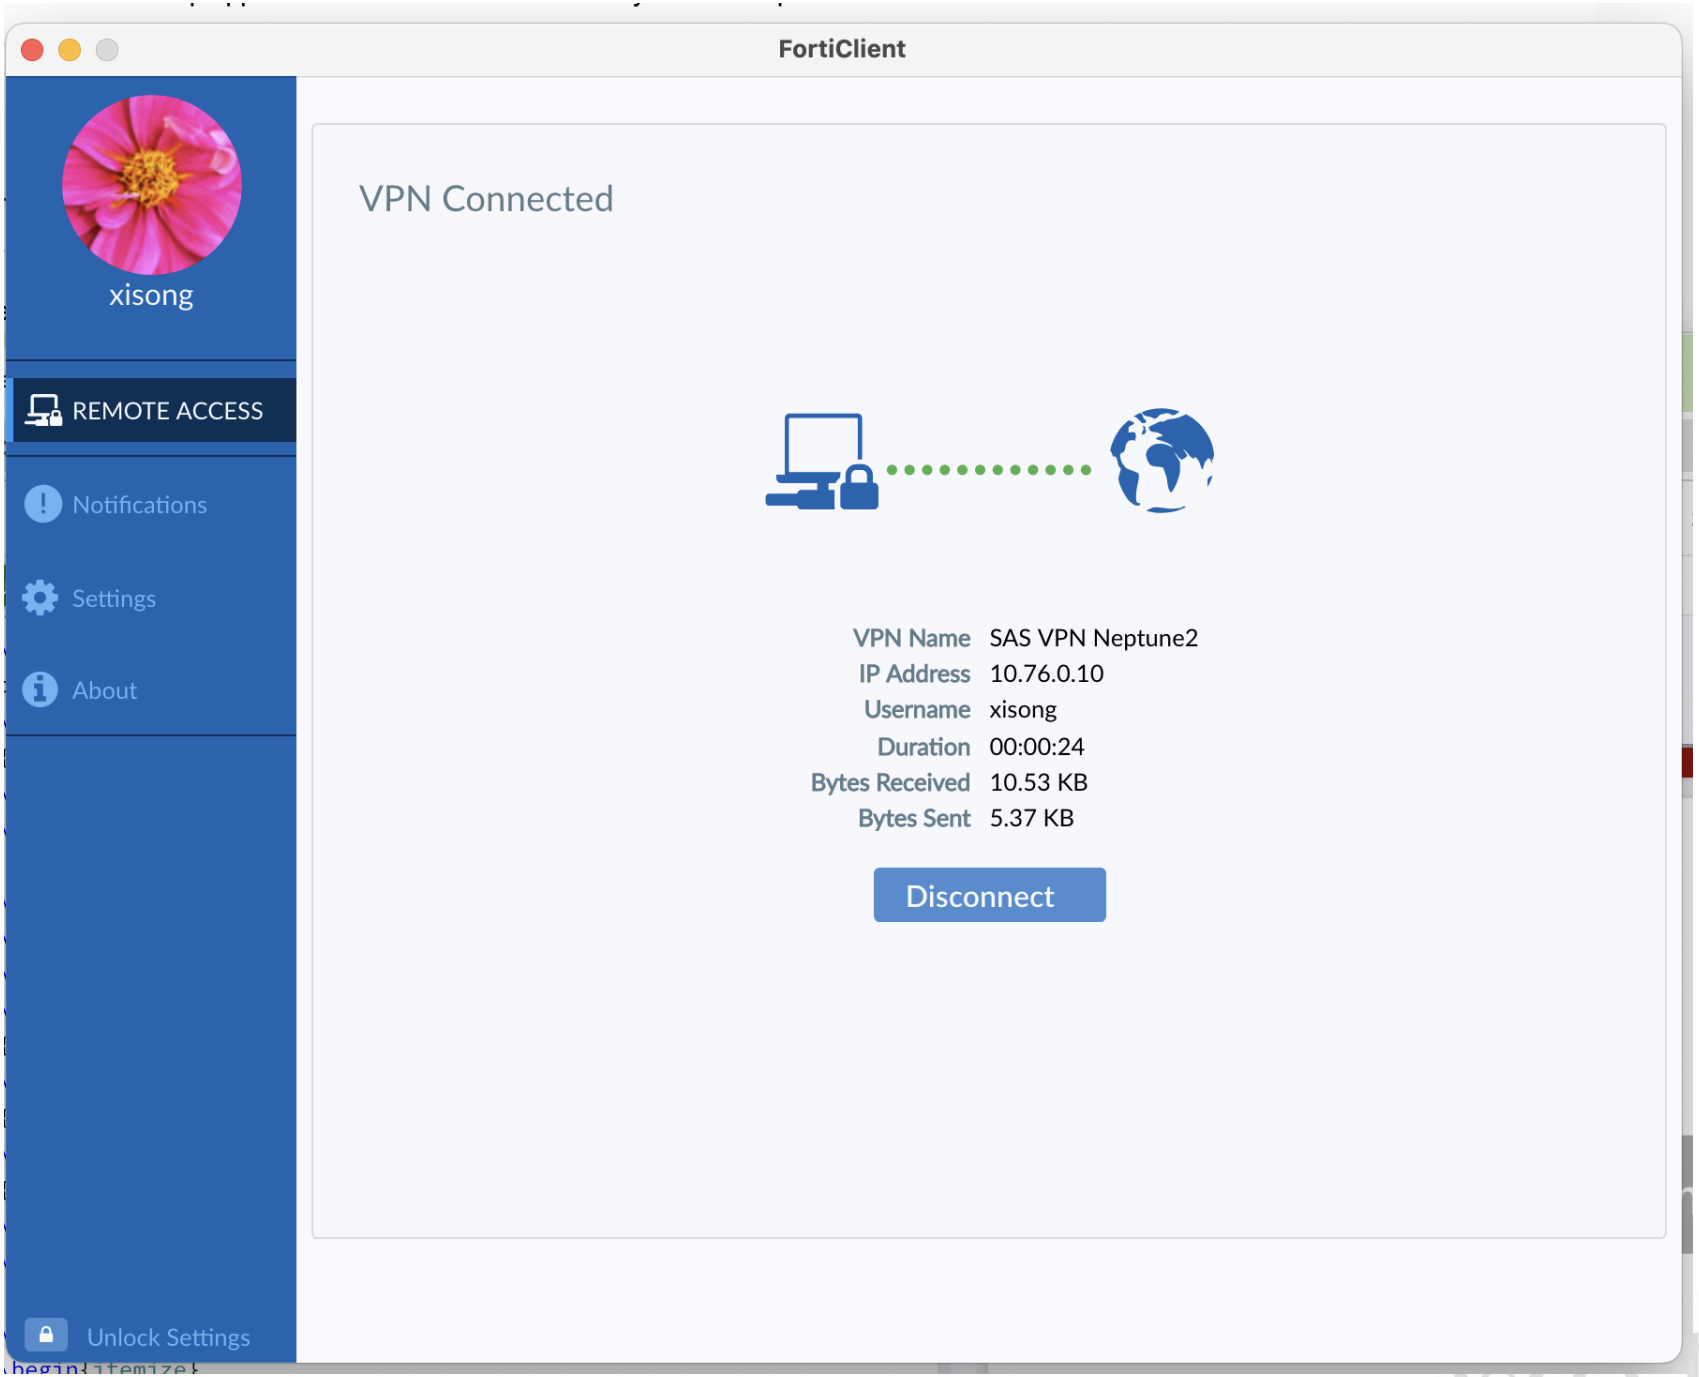
FortiClient
xisong
REMOTE ACCESS
Notifications
Settings
About
Unlock Settings
VPN Connected
VPN Name SAS VPN Neptune2
IP Address 10.76.0.10
Username xisong
Duration 00:00:24
Bytes Received 10.53 KB
Bytes Sent 5.37 KB
Disconnect

Data Transfer

How to Transfer Your Dataset to the GPC Server?

In order to use SFTP with the GPC, you will first need an SFTP client, like FileZilla: https://filezilla-project.org/

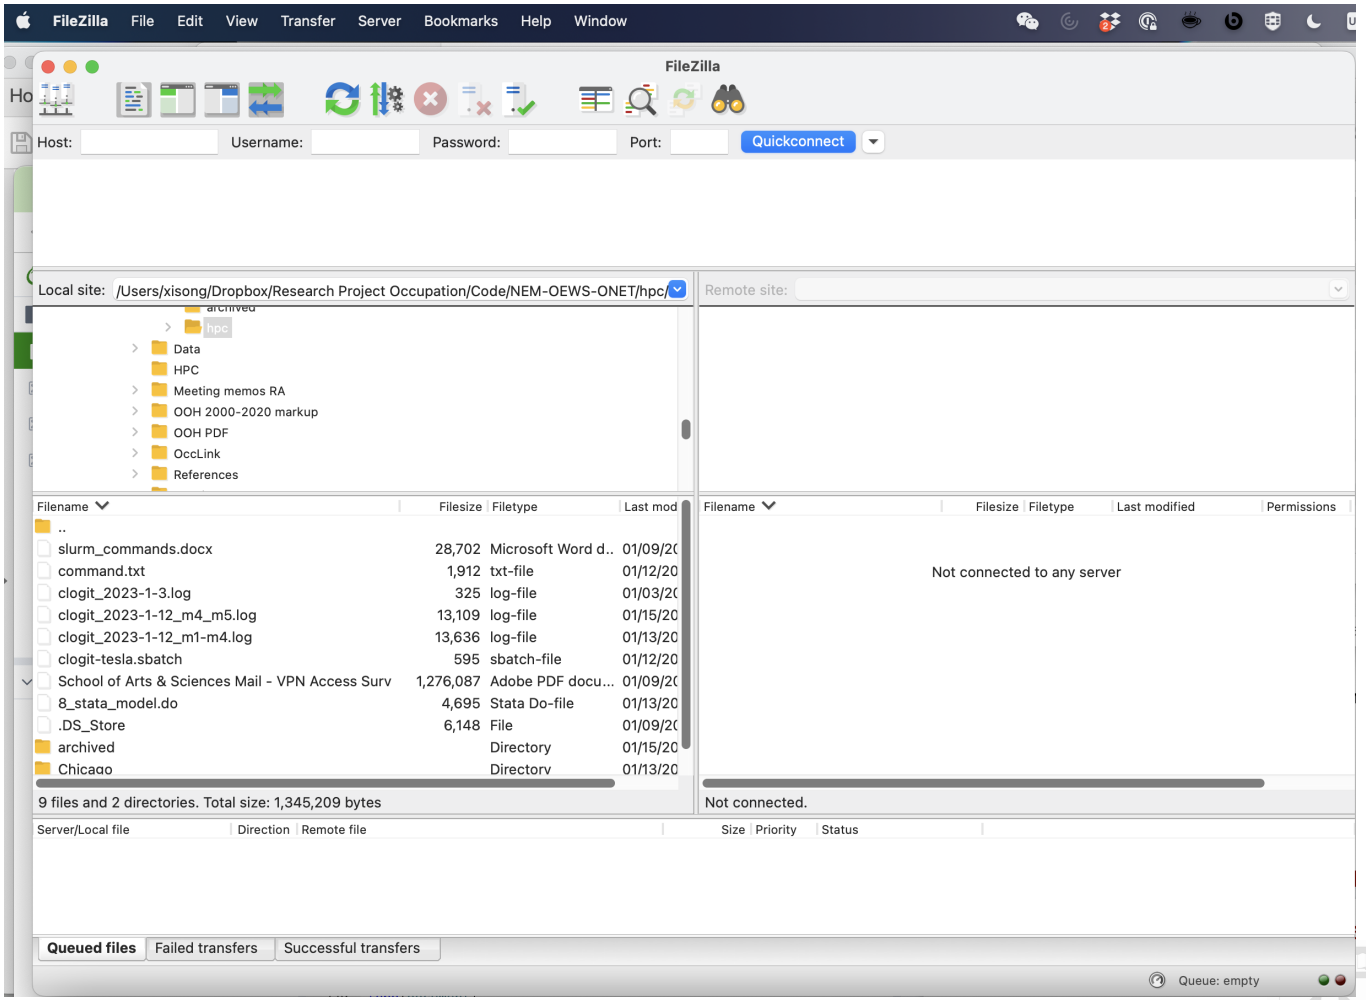

Data Transfer

Once installed, open the Site Manager under the File menu and enter the following information, where "pennkey" is your own Pennkey name:

Protocol: SFTP

Host: `gpc.sas.upenn.edu`

Port: 22

Logon Type: Interactive

User: your Pennkey name

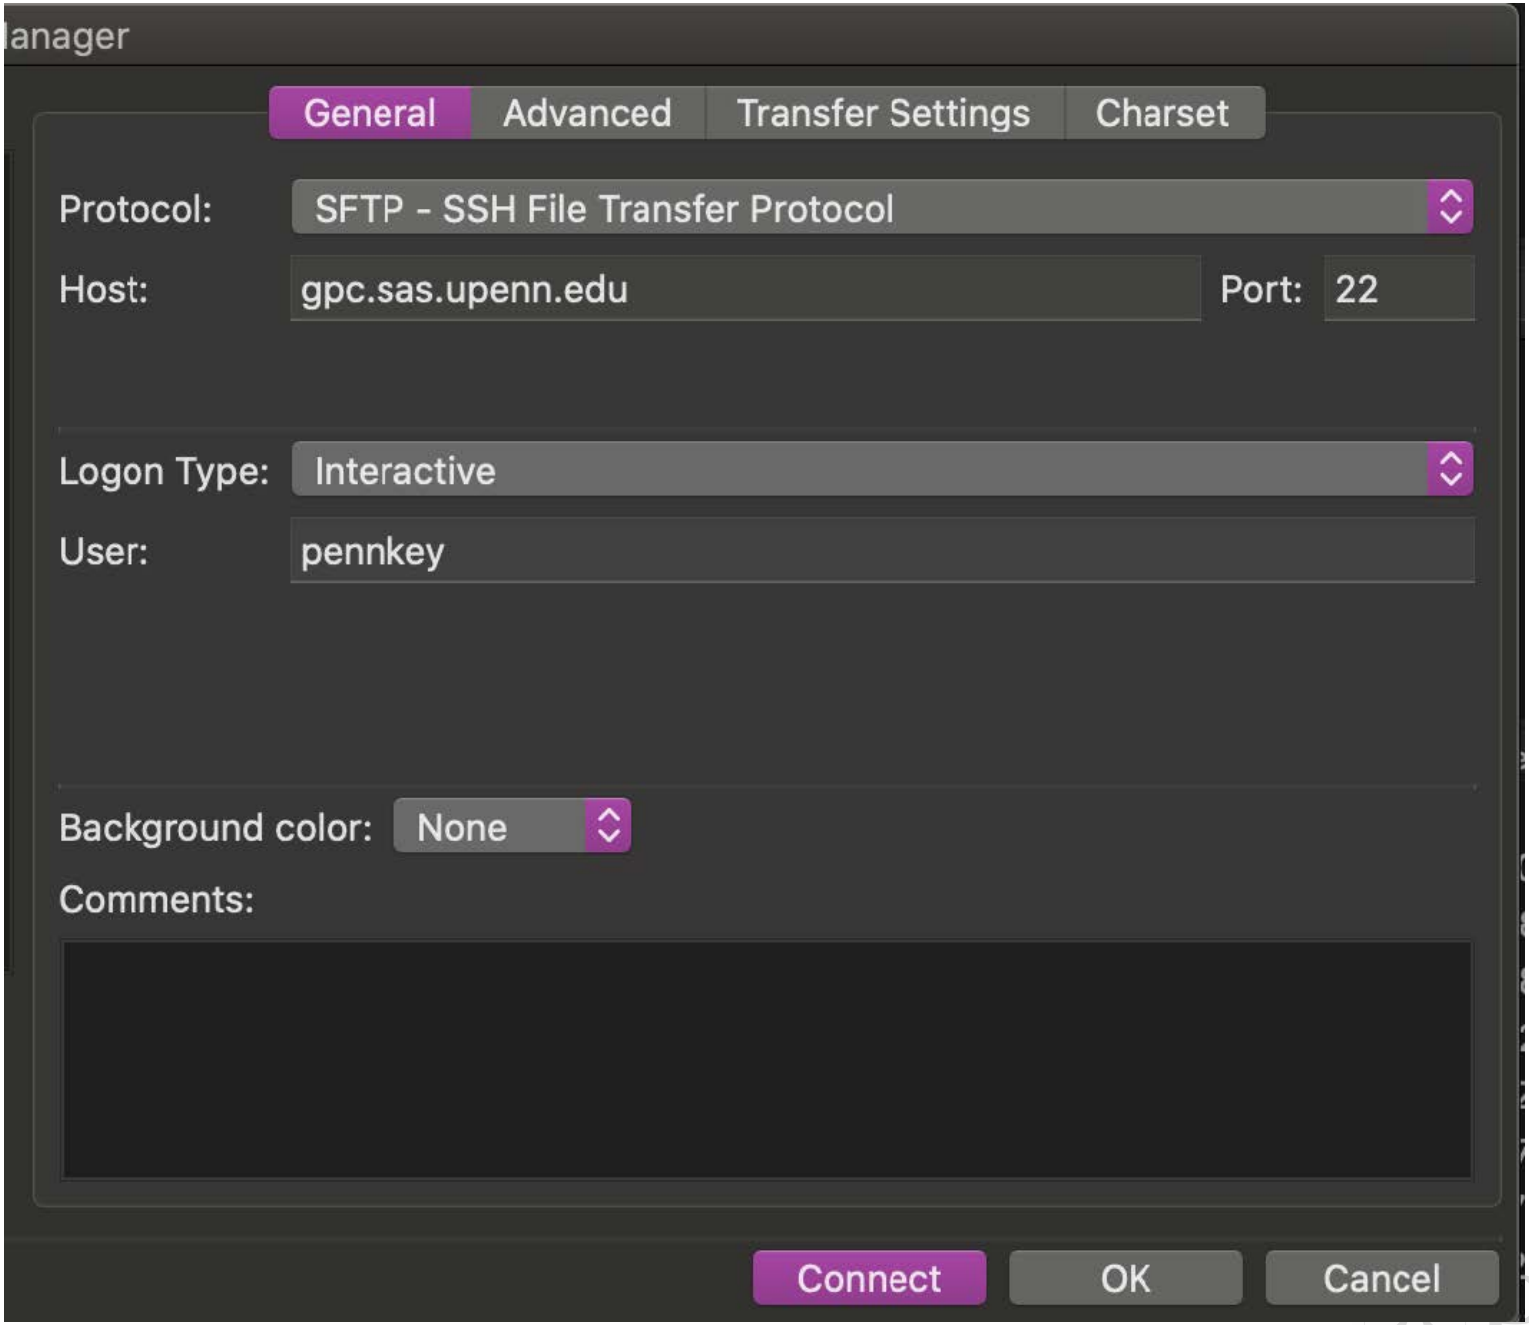
anager
General
Advanced
Transfer Settings
Charset
Protocol:
SFTP - SSH File Transfer Protocol
Host:
gpc.sas.upenn.edu
Port:
22
Logon Type:
Interactive
User:
pennkey
Background color:
None
Comments:
Connect
OK
Cancel

Once Connected to the GPC Server

- Once connected, you will see your local computer files (Local Site) on the left pane, and the GPC's files (Remote Site) on the right.
- Your own directory on the GPC will be located in /home/pennkey using your own Pennkey name.
- You can drag and drop files between panes, or do a manual transfer.

Once Connected to the GPC Server

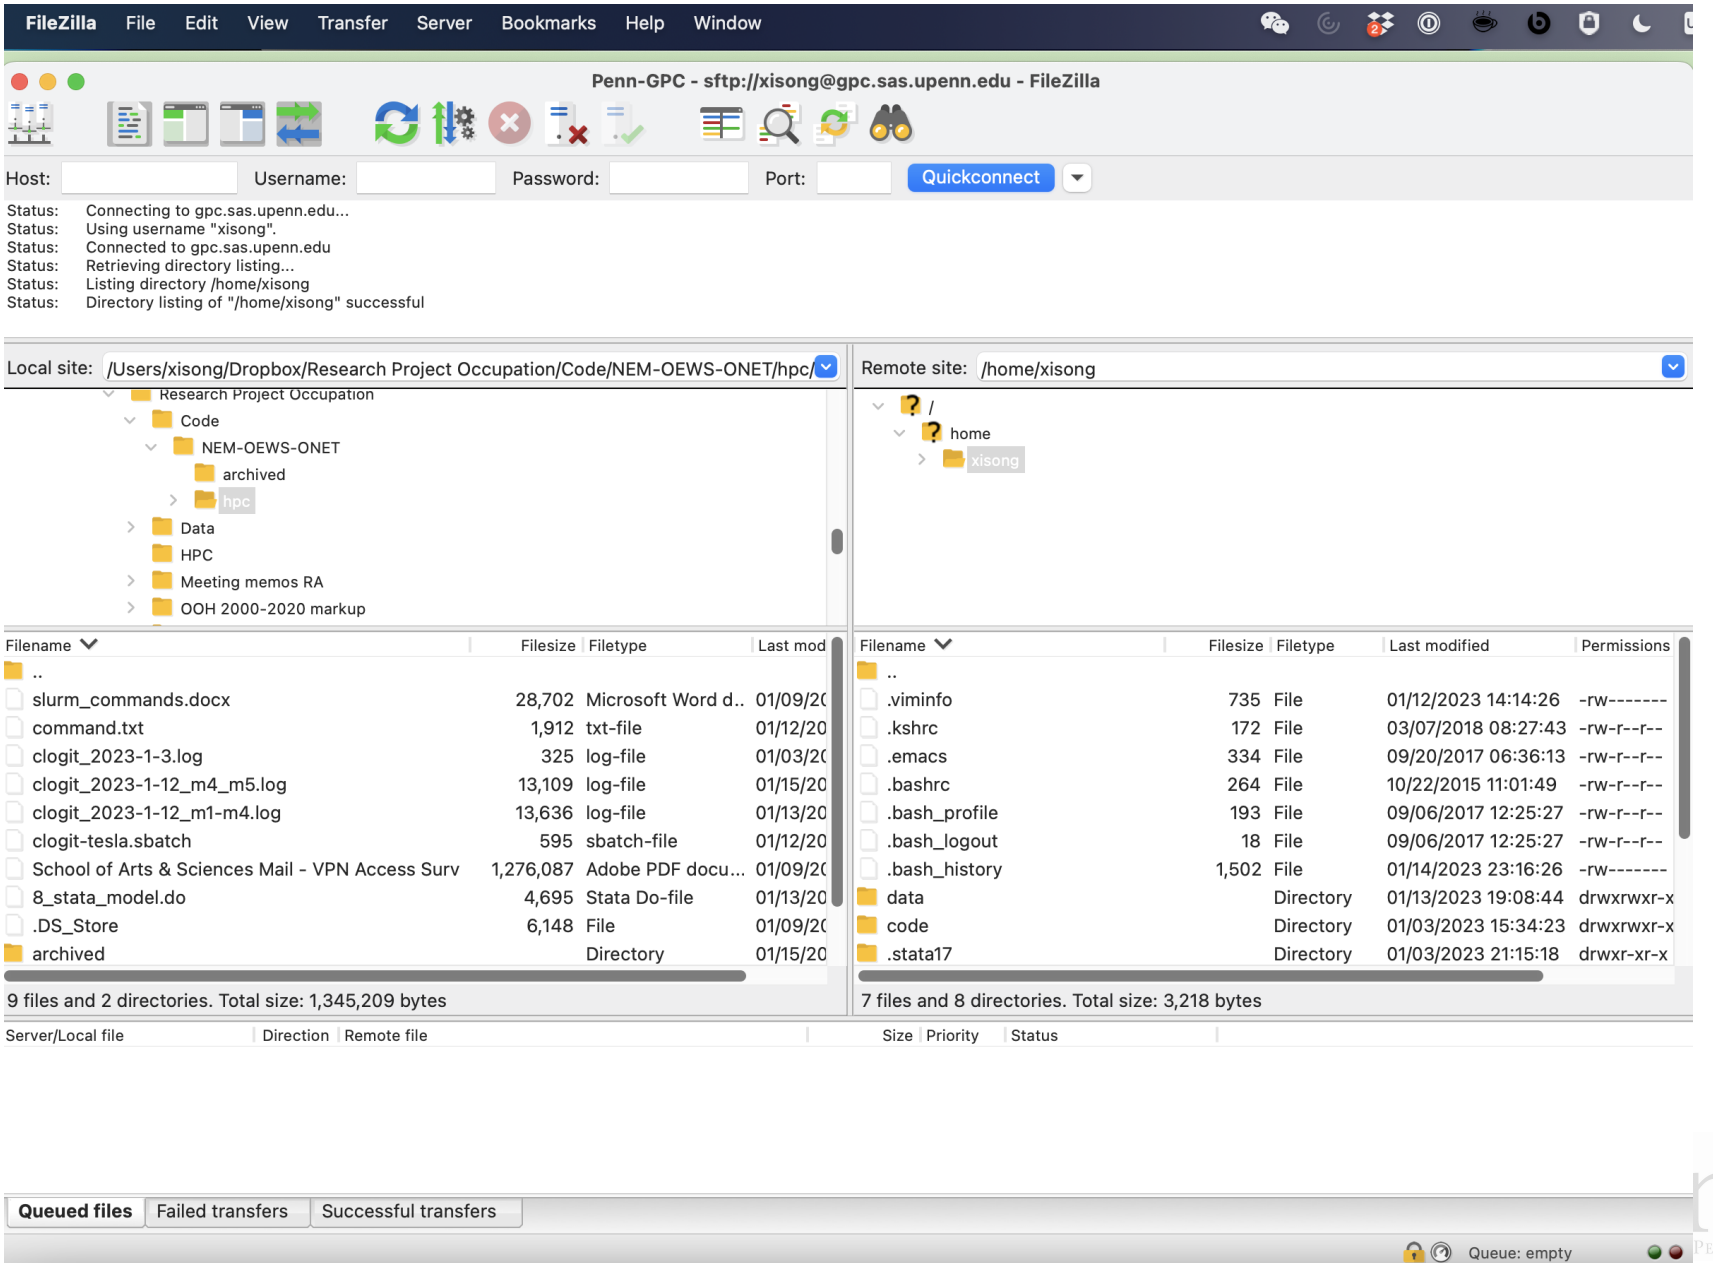

Running an Interactive Job

In Your Mac Terminal

- Windows command shell, either command dot com or cmd dot exe, or powershell is the equivalent of terminal

xisong — xisong@gpc:~ — ssh xisong@gpc.sas.upenn.edu — 142×48

```
Last login: Tue Jan 10 15:14:49 on ttys001
xisong@setups-MacBook-Pro-3 ~ % ssh-agent
SSH_AUTH_SOCK=/var/folders/cf/l7brksks0m3650fw6p61gf3m0000gr/T//ssh-CMPFaI6Bxini/agent.49450; export SSH_AUTH_SOCK;
SSH_AGENT_PID=49451; export SSH_AGENT_PID;
echo Agent pid 49451;
xisong@setups-MacBook-Pro-3 ~ % ssh-add
xisong@setups-MacBook-Pro-3 ~ % ssh xisong@gpc.sas.upenn.edu
(xisong@gpc.sas.upenn.edu) Password:
(xisong@gpc.sas.upenn.edu) Duo two-factor login for xisong

Enter a passcode or select one of the following options:

 1. Duo Push to phone_push (iOS)
 2. Phone call to XXX-XXX-0059
 3. SMS passcodes to XXX-XXX-0059

Passcode or option (1-3): 1
Success. Logging you in...
Last failed login: Sun Jan 15 06:32:08 EST 2023 from 10.76.0.3 on ssh:notty
There were 3 failed login attempts since the last successful login.
Last login: Sat Jan 14 10:48:03 2023 from 10.76.0.5
Welcome to the

        GGGGGGGGGGGG   PPPPPPPPPPPPPPPP             CCCCCCCCCCCC
     GGG::::::::::::G   P::::::::::::::::P         CCC::::::::::::C
   GG::::::::::::::G   P::::::PPPPPP:::::P       CC:::::::::::::::C
  G::::::GGGGGGGG::::G   PP:::::P     P:::::P     C:::::CCCCCCCC::::C
 G::::::G       GGGGGG     P::::P     P:::::P    C:::::C       CCCCCC
G:::::G                   P::::P     P:::::P   C:::::C
G:::::G                   P::::PPPPPP:::::P    C:::::C
G:::::G    GGGGGGGGGG     P::::::::::::::PP     C:::::C
G:::::G    G::::::::G     P::::PPPPPPPPP       C:::::C
G:::::G    GGGGG::::G     P::::P               C:::::C
G:::::G        G::::G     P::::P               C:::::C
 G::::::G       G::::G     P::::P                C:::::C       CCCCCC
  G::::::GGGGGGGG::::G   PP::::::PP               C:::::CCCCCCCC::::C
   GG::::::::::::::G   P::::::::P                 CC:::::::::::::::C
     GGG::::::GGG:::G   P::::::::P                   CCC::::::::::::C
        GGGGGG   GGGG   PPPPPPPPPP                      CCCCCCCCCCCCC

                  General Purpose Cluster

Use the following commands to adjust your environment:

'module avail'            - show available modules
'module load <module>'    - adds a module to your environment for this session
'module unload <module>'  - removes a module from your environment for this session
'module switch <m1> <m2>' - switch module version 1 with version 2
```

Add your identification key and use the key to log into other server

Prompt the user for a private key password and add it to the list maintained by ssh-agent

use yourpennkey@gpc.sas.upenn.edu

First enter your Pennkey password and then use the two-factor authentication

Available Partitions

- **normal:** General jobs, 192GB memory available per node (many nodes available)
- **highmem:** High memory jobs; 256GB memory available per node (only three nodes available)
- **tesla:** Faster, memory unknown (only 1 node available). For students who only need a single-node server to run jobs in R or Stata, SAS provides the Tesla server, which can be run in GPC. In addition, Tesla contains one Tesla M2075 graphics processing unit (GPU) card available to faculty and students whose work lends itself to GPU processing. The server supports development is C and Fortran (through the PGI compilers) in CUDA 2.0.
- However, my jobs were killed even when I was requesting 128GB in normal or highmem nodes. If such cases happen, you'll need to email the gpc manager to increase the memory for you or help debug errors

Available Partitions

- **normal:** General jobs, 192GB memory available per node (many nodes available)
- **highmem:** High memory jobs; 256GB memory available per node (only three nodes available)
- **tesla:** Faster, memory unknown (only 1 node available). For students who only need a single-node server to run jobs in R or Stata, SAS provides the Tesla server, which can be run in GPC. In addition, Tesla contains one Tesla M2075 graphics processing unit (GPU) card available to faculty and students whose work lends itself to GPU processing. The server supports development is C and Fortran (through the PGI compilers) in CUDA 2.0.
- However, my jobs were killed even when I was requesting 128GB in normal or highmem nodes. If such cases happen, you'll need to email the gpc manager to increase the memory for you or help debug errors

Available Partitions

- **normal:** General jobs, 192GB memory available per node (many nodes available)
- **highmem:** High memory jobs; 256GB memory available per node (only three nodes available)
- **tesla:** Faster, memory unknown (only 1 node available). For students who only need a single-node server to run jobs in R or Stata, SAS provides the Tesla server, which can be run in GPC. In addition, Tesla contains one Tesla M2075 graphics processing unit (GPU) card available to faculty and students whose work lends itself to GPU processing. The server supports development is C and Fortran (through the PGI compilers) in CUDA 2.0.
- However, my jobs were killed even when I was requesting 128GB in normal or highmem nodes. If such cases happen, you'll need to email the gpc manager to increase the memory for you or help debug errors

Available Partitions

- **normal:** General jobs, 192GB memory available per node (many nodes available)
- **highmem:** High memory jobs; 256GB memory available per node (only three nodes available)
- **tesla:** Faster, memory unknown (only 1 node available). For students who only need a single-node server to run jobs in R or Stata, SAS provides the Tesla server, which can be run in GPC. In addition, Tesla contains one Tesla M2075 graphics processing unit (GPU) card available to faculty and students whose work lends itself to GPU processing. The server supports development is C and Fortran (through the PGI compilers) in CUDA 2.0.
- However, my jobs were killed even when I was requesting 128GB in normal or highmem nodes. If such cases happen, you'll need to email the gpc manager to increase the memory for you or help debug errors

Slurm Commands

- `sinfo`: Check the status of computing nodes
  - Idle means not in use. You can then request the node.
- `squeue`: View the status of all running jobs
  - Check the statuses of your jobs: `squeue | grep yourpennkey`
- `srun`: Request a partition or a node (what resources you need to run your job)
  - If you don't request anything, GPC will randomly assign you to a normal node
  - `srun --pty bash`
  - `srun --nodelist=node03 --pty bash`
  - `srun --pty bash --partition=tesla`
  - `srun -p highmem --mem=160000 --pty bash`
  - Type `exit` to exit the node and switch to a new node
- `sbatch`: Submit your job in a batch file
- `scancel`: Cancel your jobs using `scancel yourjobid`
  - You can find your job id using `squeue`

Slurm Commands

- `sinfo`: Check the status of computing nodes
  - Idle means not in use. You can then request the node.
- `squeue`: View the status of all running jobs
  - Check the statuses of your jobs: `squeue | grep yourpennkey`
- `srun`: Request a partition or a node (what resources you need to run your job)
  - If you don't request anything, GPC will randomly assign you to a normal node
  - `srun --pty bash`
  - `srun --nodelist=node03 --pty bash`
  - `srun --pty bash --partition=tesla`
  - `srun -p highmem --mem=160000 --pty bash`
  - Type `exit` to exit the node and switch to a new node
- `sbatch`: Submit your job in a batch file
- `scancel`: Cancel your jobs using `scancel yourjobid`
  - You can find your job id using `squeue`

Slurm Commands

- `sinfo`: Check the status of computing nodes
  - Idle means not in use. You can then request the node.
- `squeue`: View the status of all running jobs
  - Check the statuses of your jobs: `squeue | grep yourpennkey`
- `srun`: Request a partition or a node (what resources you need to run your job)
  - If you don't request anything, GPC will randomly assign you to a normal node
  - `srun --pty bash`
  - `srun --nodelist=node03 --pty bash`
  - `srun --pty bash --partition=tesla`
  - `srun -p highmem --mem=160000 --pty bash`
  - Type `exit` to exit the node and switch to a new node
- `sbatch`: Submit your job in a batch file
- `scancel`: Cancel your jobs using `scancel yourjobid`
  - You can find your job id using `squeue`

Slurm Commands

- `sinfo`: Check the status of computing nodes
  - Idle means not in use. You can then request the node.
- `squeue`: View the status of all running jobs
  - Check the statuses of your jobs: `squeue | grep yourpennkey`
- `srun`: Request a partition or a node (what resources you need to run your job)
  - If you don't request anything, GPC will randomly assign you to a normal node
  - `srun --pty bash`
  - `srun --nodelist=node03 --pty bash`
  - `srun --pty bash --partition=tesla`
  - `srun -p highmem --mem=160000 --pty bash`
  - Type `exit` to exit the node and switch to a new node
- `sbatch`: Submit your job in a batch file
- `scancel`: Cancel your jobs using `scancel yourjobid`
  - You can find your job id using `squeue`

Slurm Commands

- `sinfo`: Check the status of computing nodes
  - Idle means not in use. You can then request the node.
- `squeue`: View the status of all running jobs
  - Check the statuses of your jobs: `squeue | grep yourpennkey`
- `srun`: Request a partition or a node (what resources you need to run your job)
  - If you don't request anything, GPC will randomly assign you to a normal node
  - `srun --pty bash`
  - `srun --nodelist=node03 --pty bash`
  - `srun --pty bash --partition=tesla`
  - `srun -p highmem --mem=160000 --pty bash`
  - Type `exit` to exit the node and switch to a new node
- `sbatch`: Submit your job in a batch file
- `scancel`: Cancel your jobs using `scancel yourjobid`
  - You can find your job id using `squeue`

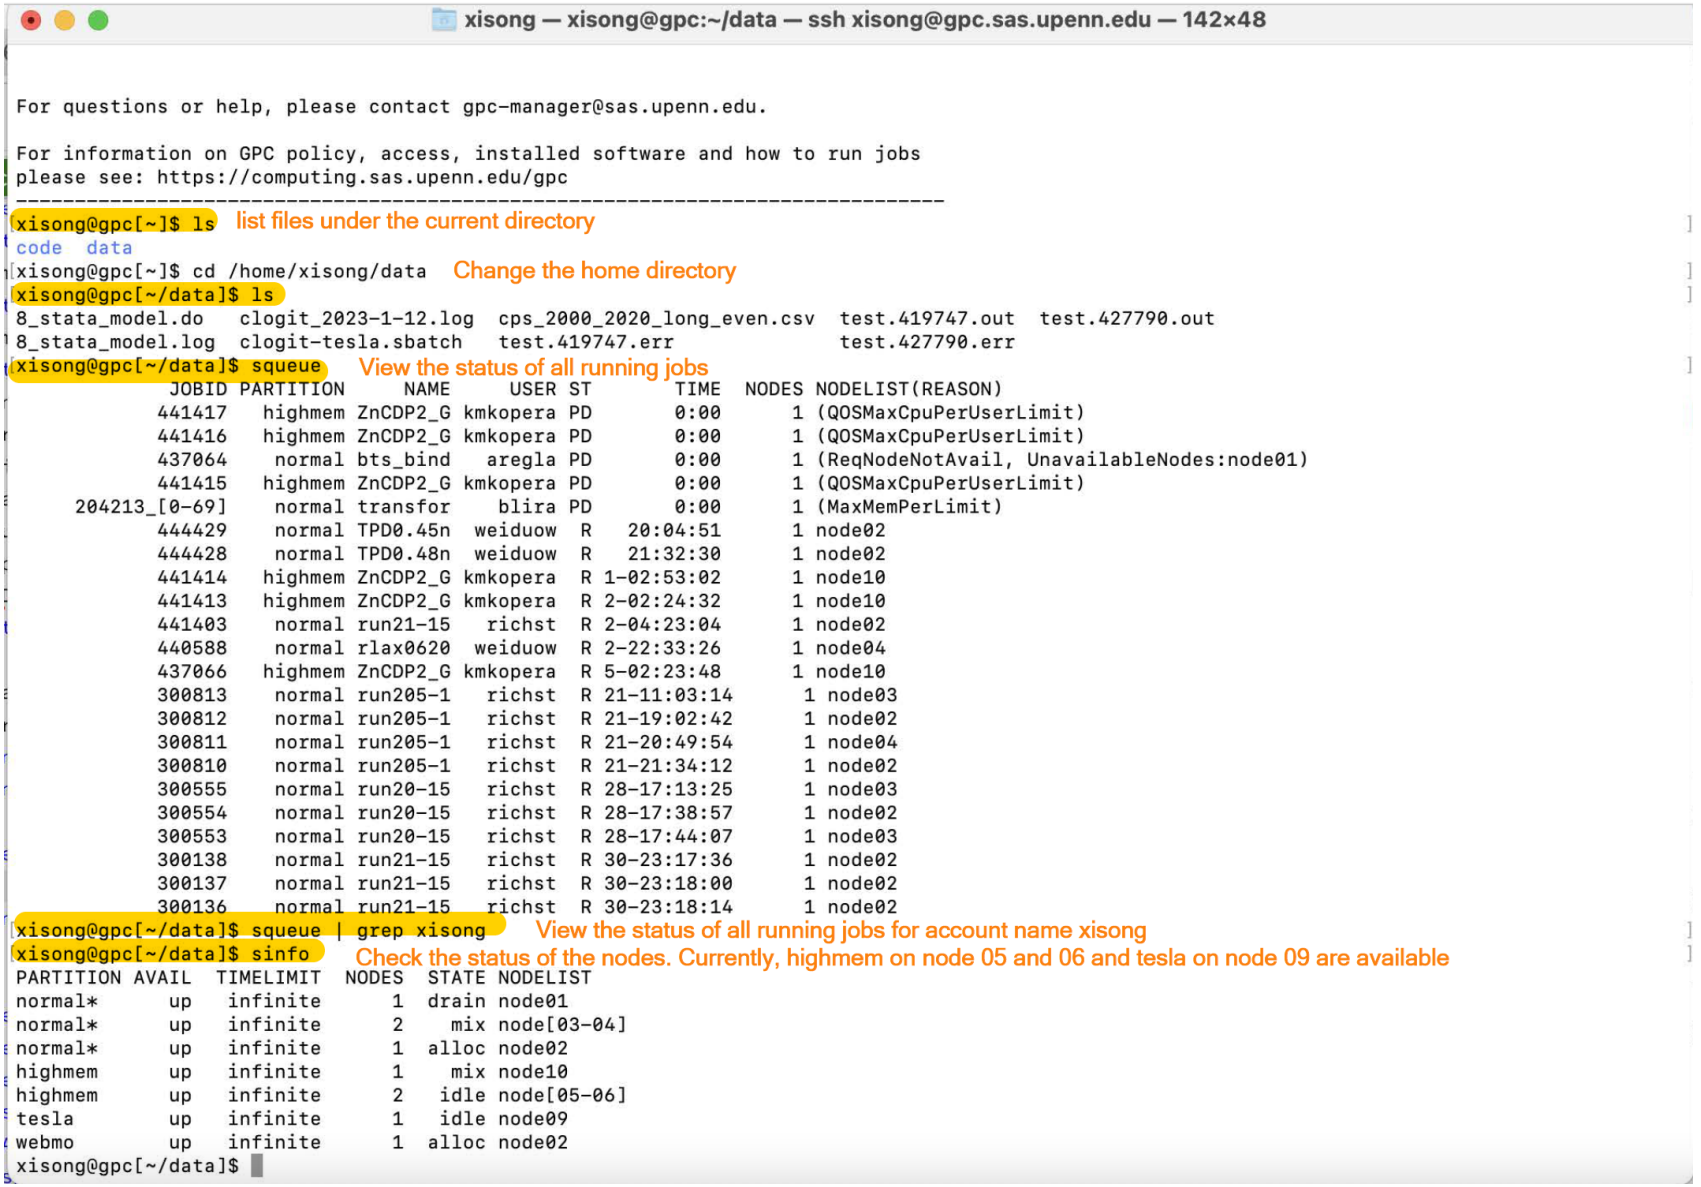
xisong — xisong@gpc:~/data — ssh xisong@gpc.sas.upenn.edu — 142×48
For questions or help, please contact gpc-manager@sas.upenn.edu.
For information on GPC policy, access, installed software and how to run jobs
please see: https://computing.sas.upenn.edu/gpc
xisong@gpc[~]$ ls
list files under the current directory
code data
xisong@gpc[~]$ cd /home/xisong/data
Change the home directory
xisong@gpc[~/data]$ ls
8_stata_model.do clogit_2023-1-12.log cps_2000_2020_long_even.csv test.419747.out test.427790.out
8_stata_model.log clogit-tesla.sbatch test.419747.err test.427790.err
xisong@gpc[~/data]$ squeue
View the status of all running jobs
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
441417 highmem ZnCDP2_G kmkopera PD 0:00 1 (QOSMaxCpuPerUserLimit)
441416 highmem ZnCDP2_G kmkopera PD 0:00 1 (QOSMaxCpuPerUserLimit)
437064 normal bts_bind aregla PD 0:00 1 (ReqNodeNotAvail, UnavailableNodes:node01)
441415 highmem ZnCDP2_G kmkopera PD 0:00 1 (QOSMaxCpuPerUserLimit)
204213_[0-69] normal transfor blira PD 0:00 1 (MaxMemPerLimit)
444429 normal TPD0.45n weiduow R 20:04:51 1 node02
444428 normal TPD0.48n weiduow R 21:32:30 1 node02
441414 highmem ZnCDP2_G kmkopera R 1-02:53:02 1 node10
441413 highmem ZnCDP2_G kmkopera R 2-02:24:32 1 node10
441403 normal run21-15 richst R 2-04:23:04 1 node02
440588 normal rlax0620 weiduow R 2-22:33:26 1 node04
437066 highmem ZnCDP2_G kmkopera R 5-02:23:48 1 node10
300813 normal run205-1 richst R 21-11:03:14 1 node03
300812 normal run205-1 richst R 21-19:02:42 1 node02
300811 normal run205-1 richst R 21-20:49:54 1 node04
300810 normal run205-1 richst R 21-21:34:12 1 node02
300555 normal run20-15 richst R 28-17:13:25 1 node03
300554 normal run20-15 richst R 28-17:38:57 1 node02
300553 normal run20-15 richst R 28-17:44:07 1 node03
300138 normal run21-15 richst R 30-23:17:36 1 node02
300137 normal run21-15 richst R 30-23:18:00 1 node02
300136 normal run21-15 richst R 30-23:18:14 1 node02
xisong@gpc[~/data]$ squeue | grep xisong
View the status of all running jobs for account name xisong
xisong@gpc[~/data]$ sinfo
Check the status of the nodes. Currently, highmem on node 05 and 06 and tesla on node 09 are available
PARTITION AVAIL TIMELIMIT NODES STATE NODELIST
normal* up infinite 1 drain node01
normal* up infinite 2 mix node[03-04]
normal* up infinite 1 alloc node02
highmem up infinite 1 mix node10
highmem up infinite 2 idle node[05-06]
tesla up infinite 1 idle node09
webmo up infinite 1 alloc node02
xisong@gpc[~/data]$

Check Available Software on the Server

Use `module avail`

```
[xisong@gpc[~/data]$ srun --nodelist=node03 --pty bash
[[xisong@node03 data]$ module avail

------------------------------------------------------------ /cm/local/modulefiles ------------------------------------------------------------
cluster-tools/7.2            docker/engine/1.9.1          ipmitool/1.8.15              openmpi/mlnx/gcc/64/3.0.0rc6
cluster-tools-dell/7.2       dot                          module-git                   shared
cm-cloud-copy/7.2            flannel/0.5.4                module-info                  use.own
cmd                          freeipmi/1.4.11              null                         version
cmsub                        gcc/5.2.0                    openldap

------------------------------------------------------------ /cm/shared/modulefiles ------------------------------------------------------------
acml/gcc/64/5.3.1                      intel/compiler/64/15.0/2015.5.223      nbo/7.32
acml/gcc/fma4/5.3.1                    intel/compiler/64/2017/17.0.6          ncbi-blast/2.13.0
acml/gcc/mp/64/5.3.1                   intel/mkl/64/2017/6.256                netcdf/4.8.1
acml/gcc/mp/fma4/5.3.1                 intel/mkl/mic/2017/6.256               netcdf/gcc/64/4.3.3.1
```

...

```
git/2.7.4                              mess/2021.05                           SHOREmap/3.8
GLFW/3.0.4                             mpc/1.1.0                              slurm/15.08.13
globalarrays/openmpi/gcc/64/5.4        mpfr/4.0.1                             slurm/20.02.4
globalarrays/openmpi/open64/64/5.4     mpich/ge/gcc/64/3.2                    smalt/0.7.6
gmp/6.1.2                              mpich/ge/open64/64/3.2                 sratoolkit/2.11.2
golang/1.13.6                          mpiexec/0.84_432                       STAR/2.7.10a
gromacs/5.1.2                          mvapich/gcc/64/1.2rc1                  stata/13
gsl/2.4                                mvapich/open64/64/1.2rc1               stata/14
gurobi/8.0.0                           mvapich2/gcc/64/2.2b                   stata/17
hdf5/1.12.1                            mvapich2/open64/64/2.2b                tmux/3.1
hdf5_18/1.8.16                         namd/mpi/2.11                          torque/6.0.0.1
hpl/2.1                                namd/multicore/2.11                    trimmomatic/0.39
hwloc/1.10.1                           nbo/6                                  vcftools/0.1.15
intel/compiler/32/2017/17.0.6          nbo/7                                  vowpal_wabbit/current
[xisong@node03 data]$
```

```
[xisong@node02 ~]$ module load stata/17
[xisong@node02 ~]$ stata-mp
```

Load the module you want to use

Open Stata MP version; now you can use it interactively

```
  ___  ____  ____  ____  ____ ©
 /__    /   ____/   /   ____/      17.0
___/   /   /___/   /   /___/       MP—Parallel Edition

 Statistics and Data Science       Copyright 1985-2021 StataCorp LLC
                                   StataCorp
                                   4905 Lakeway Drive
                                   College Station, Texas 77845 USA
                                   800-STATA-PC        https://www.stata.com
                                   979-696-4600        stata@stata.com

Stata license: 2-user 8-core network perpetual
Serial number: 501706320672
  Licensed to: Research Data Services
               University of Pennsylvania

Notes:
      1. Unicode is supported; see help unicode_advice.
      2. More than 2 billion observations are allowed; see help obs_advice.
      3. Maximum number of variables is set to 5,000; see help set_maxvar.

. set obs 10000
Number of observations (_N) was 0, now 10,000.

. gen x = rnormal(0,1)

. sum x

    Variable |        Obs        Mean    Std. dev.       Min        Max
-------------+---------------------------------------------------------
           x |     10,000    .0003138    .9946919  -4.468302   4.327502

. exit
no; dataset in memory has changed since last saved
    Save the data or specify option clear to exit anyway.
r(4);
```

Running a Batch Job (recommend!)

Run Your Models with Slurm Job Scheduler

If you want to run your jobs by means of wrapper scripts rather than by entering the commands line-by-line, you will need to use Slurm

- Slurm documentation
  https://computing.sas.upenn.edu/gpc/job/slurm
- More information about Slurm
  - Slurm Quickstart guide
    https://slurm.schedmd.com/quickstart.html
  - Slurm command summary
    https://slurm.schedmd.com/pdfs/summary.pdf
  - Slurm manual pages
    https://slurm.schedmd.com/man_index.html

A Slurm Example: clogit.batch

Save the following commands in a text file with the format .batch

clogit-tesla.sbatch — hpc

```
#!/bin/bash
#SBATCH -p highmem
#SBATCH -J clogit-2023-1-12
#SBATCH --mail-type=END,FAIL
#SBATCH --mail-user=xisong@upenn.edu
#SBATCH --nodes=1
#SBATCH --ntasks=1
#SBATCH --mem=100G
#SBATCH --time=36:00:00
#SBATCH -o test."%j".out
#SBATCH -e test."%j".err
#SBATCH --chdir=/home/xisong/data

module load stata/17
stata-mp -b do "8_stata_model.do"

```

A Slurm Example: clogit.batch

See other slurm examples:
https://computing.sas.upenn.edu/gpc/job/slurm

clogit-tesla.sbatch — hpc

```
1  #!/bin/bash
2  #SBATCH -p highmem      Request the high memory partition
3  #SBATCH -J clogit-2023-1-12   Create the following error and output files using this name
4  #SBATCH --mail-type=END,FAIL   Send emails to the following address when the job ends or fails
5  #SBATCH --mail-user=xisong@upenn.edu
6  #SBATCH --nodes=1      Request 1 node
7  #SBATCH --ntasks=1     Request 1 task
8  #SBATCH --mem=100G     Request 100GB memory
9  #SBATCH --time=36:00:00   Run the model for 36 hours
10 #SBATCH -o test."%j".out
11 #SBATCH -e test."%j".err
12 #SBATCH --chdir=/home/xisong/data   Save the files in this folder
13
14 module load stata/17          Load Stata 17
15 stata-mp -b do "8_stata_model.do"      Run the do file using the multiple processor version of Stata
16
```

Submit the sbatch file

- Submit your sbatch file by typing `clogit.sbatch` in the terminal (make sure that the file is in the current directory; otherwise, use cd to change the directory)

Other Useful Commands

- use `squeue` to see the status of all running jobs on the server
- use `squeue | grep yourpennkey` to see your running jobs

- `tail -f filename` to see the tail of a file. `-f` means follows
- `less filename` to see the whole file
- e.g., I used `tail -f clogit.log` to view my log file from Stata. You can also use `tail -f clogit.err` to see if there is any error report.

- use `nano` to quickly edit your files on the server rather than upload a new file to the server using Filezilla.
- e.g., `nano clogit.sbatch`. Then use control+C to quit and Y to save.

More Linux and Slurm Commands

- Basic Overview of Linux:

  https://www.digitalocean.com/community/tutorials/an-introduction-to-linux-basics

- Basic File Management Commands: (e.g., ls, cat, less, tail commands)

  https://www.linuxtrainingacademy.com/linux-commands-cheat-sheet/#5_8211_FILE_AND_DIRECTORY_COMMANDS

- Nano Text Editor:

  https://linuxize.com/post/how-to-use-nano-text-editor/

- Basic Slurm Commands and SBATCH configuration Options (understanding the SBATCH -J and other options, etc):

  https://hpc.nmsu.edu/discovery/slurm/slurm-commands/

- Official Slurm Documentation Cheat Sheet:

  https://slurm.schedmd.com/pdfs/summary.pdf

More Information

- GPC website: https://computing.sas.upenn.edu/gpc/getstarted
- Contact GPC manager: gpc-manager@sas.upenn.edu
- Be patient and keep sending emails if you don't receive a reply!

Other HPC Resources on Campus

Wharton HPC

- https://research-it.wharton.upenn.edu/documentation/
  You will need to be a Wharton affiliate

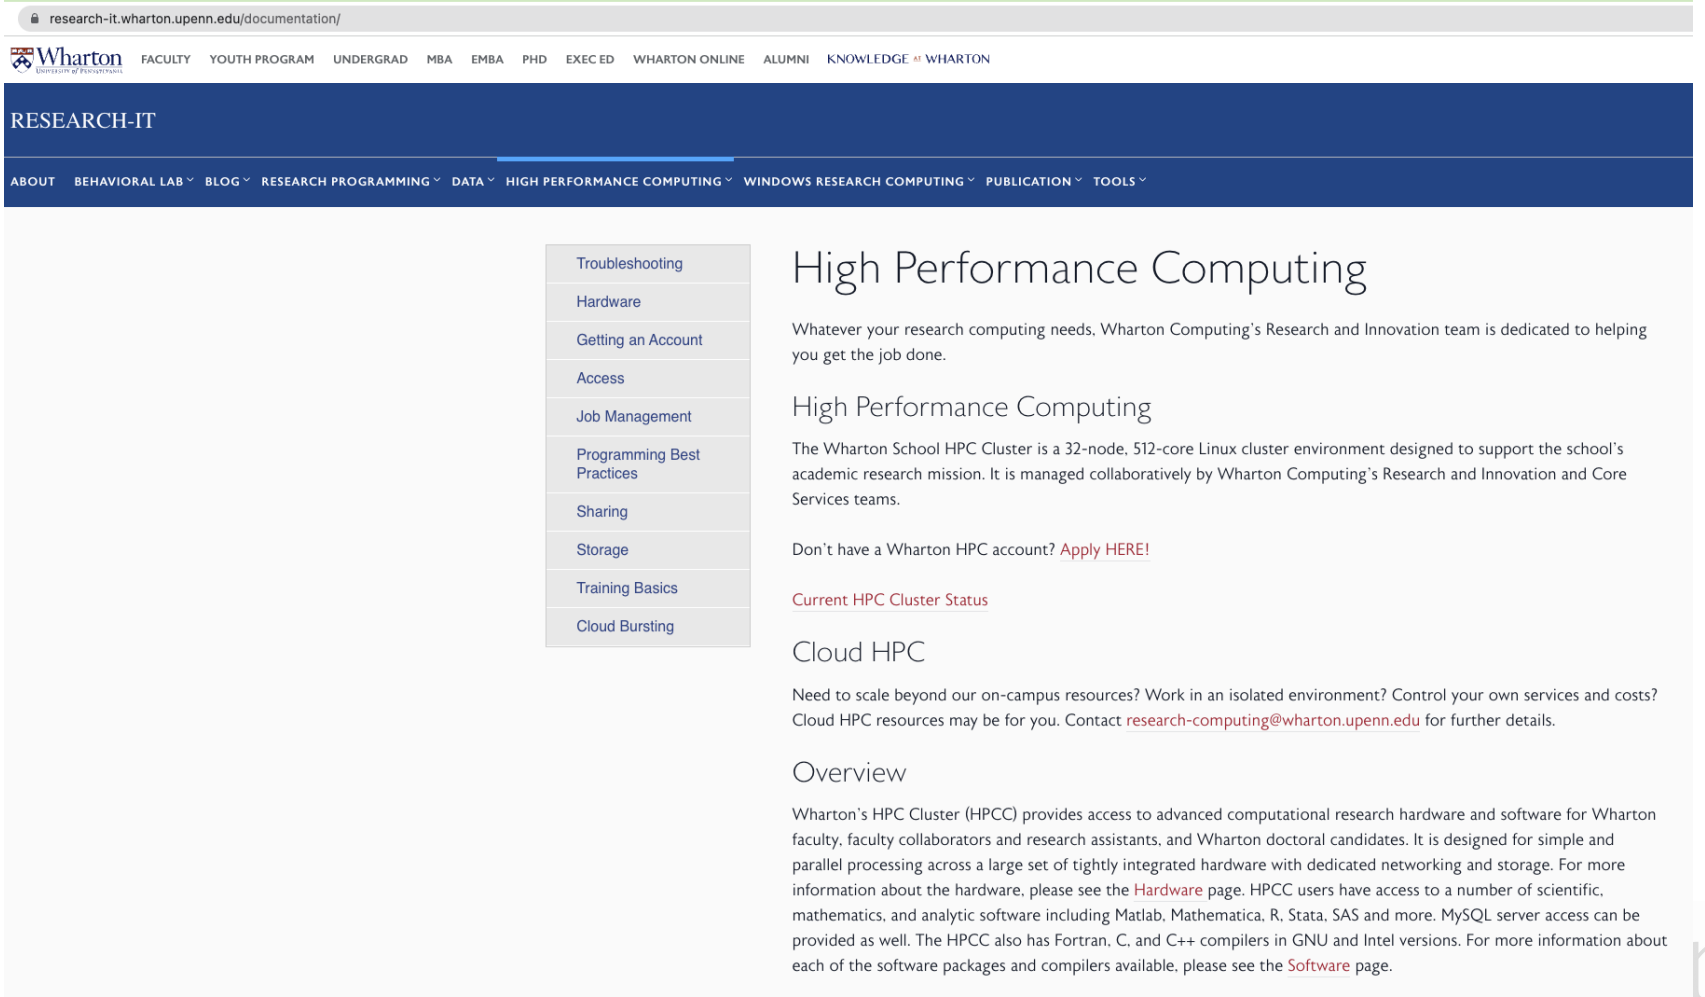

School of Medicine HPC

- Although PSOM HPC not free, it is accessible to all Penn researchers even if they are not affiliated with the medical school. https://www.med.upenn.edu/hpc/billing-pricing-service-descriptions.html
- Given that HPC is a commercial software, it is not available on Penn Medicine's HPC. Alternatively, you can use Penn Medicine Limited Performance Computing (LPC) if you are a PSOM faculty, staff or affiliate
- Contact `psom-pmacshpc@pennmedicine.upenn.edu`

School of Medicine HPC

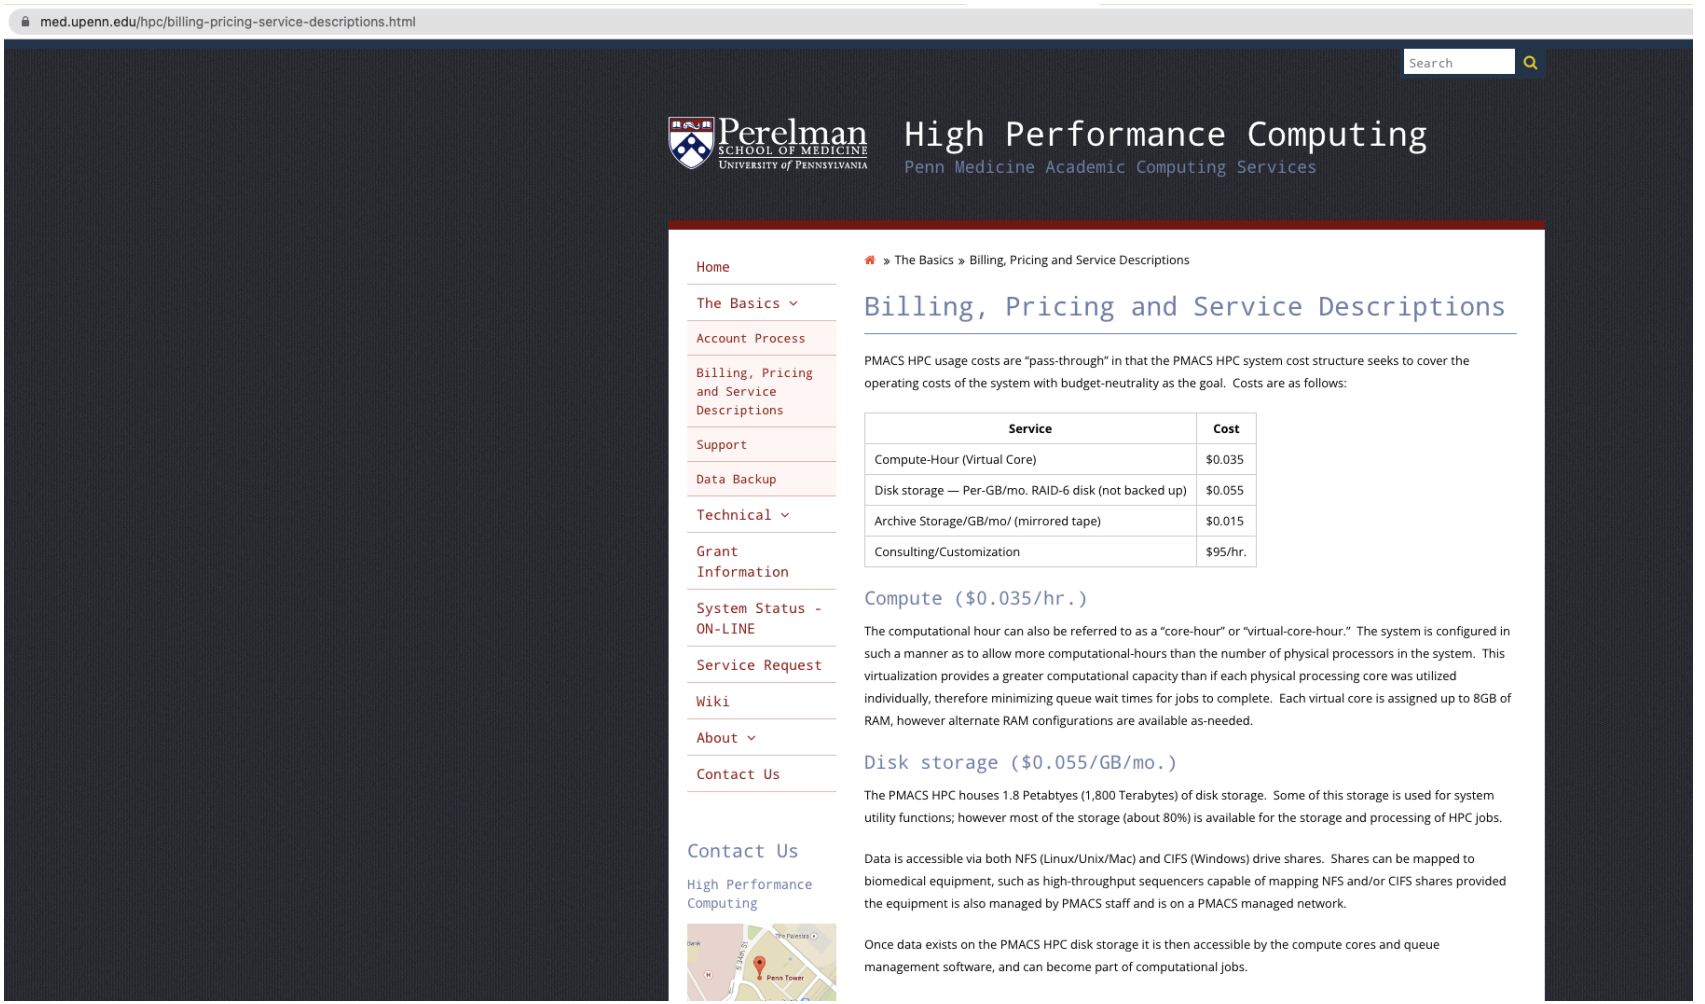

Department of Economics Linux Hawk Cluster

- The cluster contains 144 compute nodes. Students can take advantage of parallel processing, which allows a program to run on more than one compute node at a time.
- Students can access up to 36 nodes at a time. The hawk cluster runs the Linux operating system and has available both Fortran (Intel and GNU) and C compilers.
- The cluster also has available a parallel optimizer called HOPSPACK.
- To get an account on hawk, contact Professor Petra Todd.

Amazon Web Services (AWS)

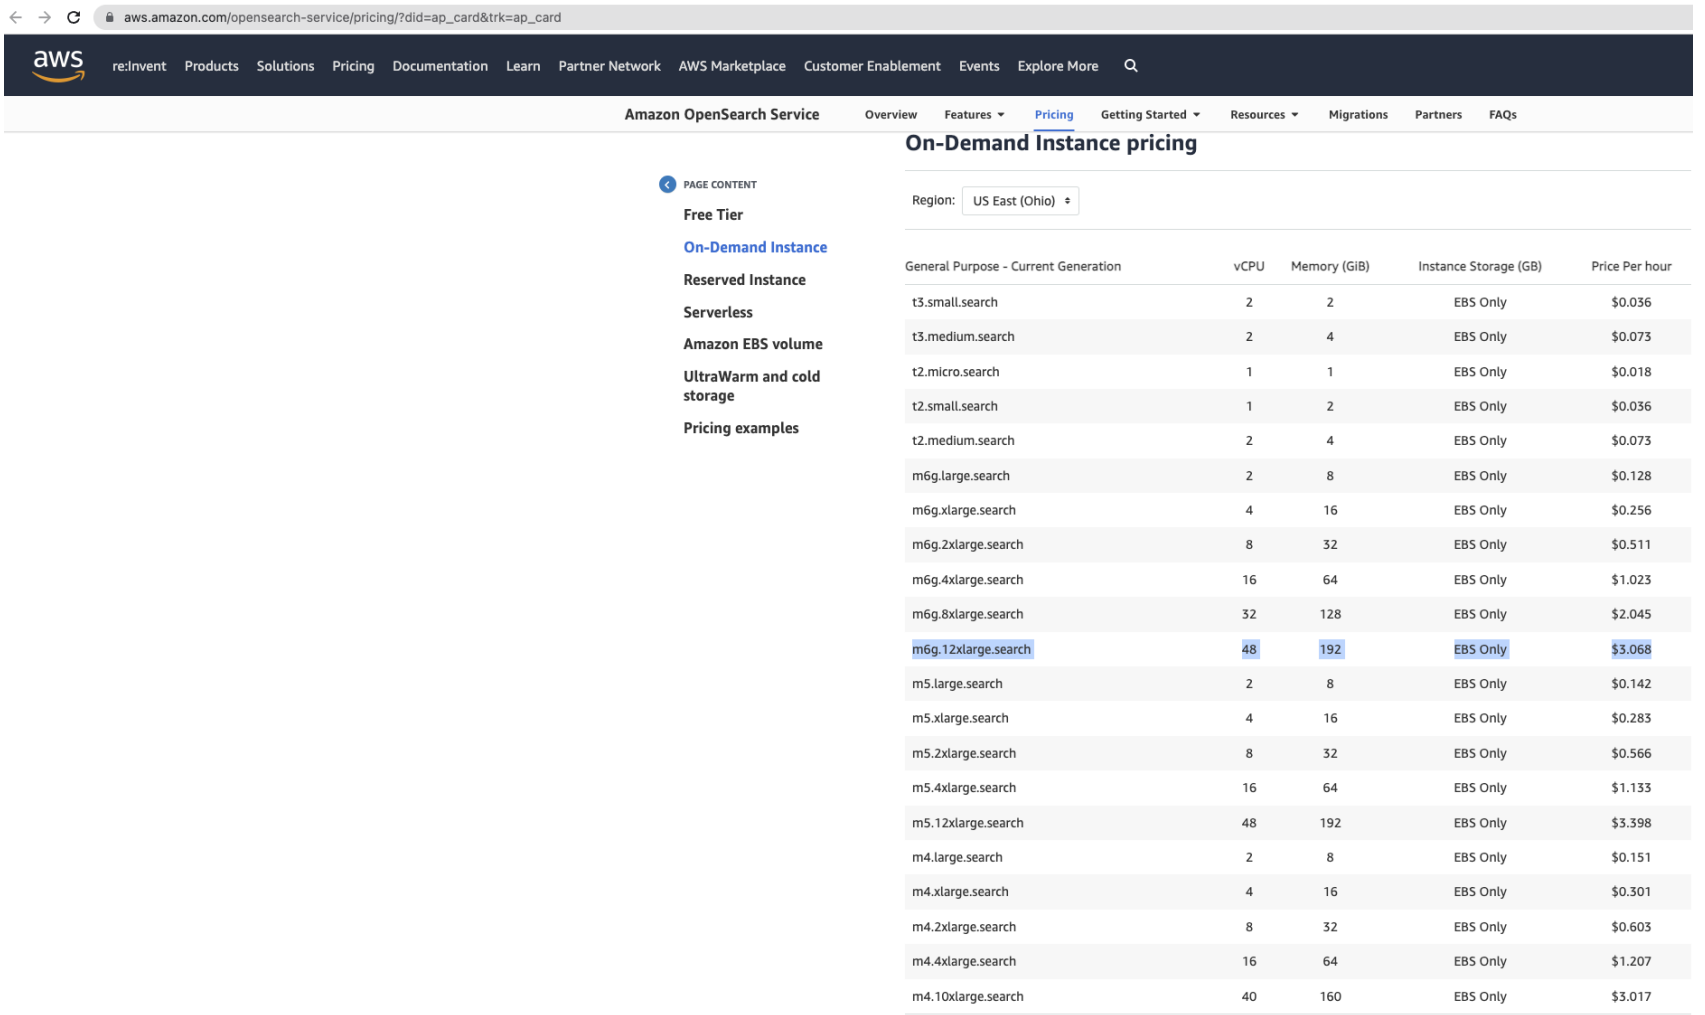

General Purpose - Current Generation	vCPU	Memory (GiB)	Instance Storage (GB)	Price Per hour
t3.small.search	2	2	EBS Only	$0.036
t3.medium.search	2	4	EBS Only	$0.073
t2.micro.search	1	1	EBS Only	$0.018
t2.small.search	1	2	EBS Only	$0.036
t2.medium.search	2	4	EBS Only	$0.073
m6g.large.search	2	8	EBS Only	$0.128
m6g.xlarge.search	4	16	EBS Only	$0.256
m6g.2xlarge.search	8	32	EBS Only	$0.511
m6g.4xlarge.search	16	64	EBS Only	$1.023
m6g.8xlarge.search	32	128	EBS Only	$2.045
m6g.12xlarge.search	48	192	EBS Only	$3.068
m5.large.search	2	8	EBS Only	$0.142
m5.xlarge.search	4	16	EBS Only	$0.283
m5.2xlarge.search	8	32	EBS Only	$0.566
m5.4xlarge.search	16	64	EBS Only	$1.133
m5.12xlarge.search	48	192	EBS Only	$3.398
m4.large.search	2	8	EBS Only	$0.151
m4.xlarge.search	4	16	EBS Only	$0.301
m4.2xlarge.search	8	32	EBS Only	$0.603
m4.4xlarge.search	16	64	EBS Only	$1.207
m4.10xlarge.search	40	160	EBS Only	$3.017

Acknowledgement

I would like to thank Alex Bond, Hanming Fang, Hadiya Gaiters, Michael Lachanski, Emilio Parrado, Petra Todd, Kyle Jermaine Small, Zihan Sun, and Xiuqi Yang for helpful discussions that made my work and this presentation possible.