Introduction to GACRC Teaching Cluster

Georgia Advanced Computing Resource Center (GACRC)

Enterprise Information Technology Services(EITS)

The University of Georgia

Outline

- GACRC
- Overview
- Working Environment
  - Three Folders
  - Three Computational Partitions
  - Software on Cluster
- Submit a Computational Batch Job
- GACRC Wiki and Support

GACRC

- A high-performance-computing (HPC) center at the UGA
- Provide to the UGA research and education community an advanced computing environment:
  - HPC computing and networking infrastructure located at the Boyd Data Center
  - Comprehensive collection of scientific, engineering and business applications
  - Consulting and training services

Wiki: http://wiki.gacrc.uga.edu

Support: https://wiki.gacrc.uga.edu/wiki/Getting_Help

Web Site: http://gacrc.uga.edu

Kaltura Channel: https://kaltura.uga.edu/channel/GACRC/176125031

Teaching Cluster

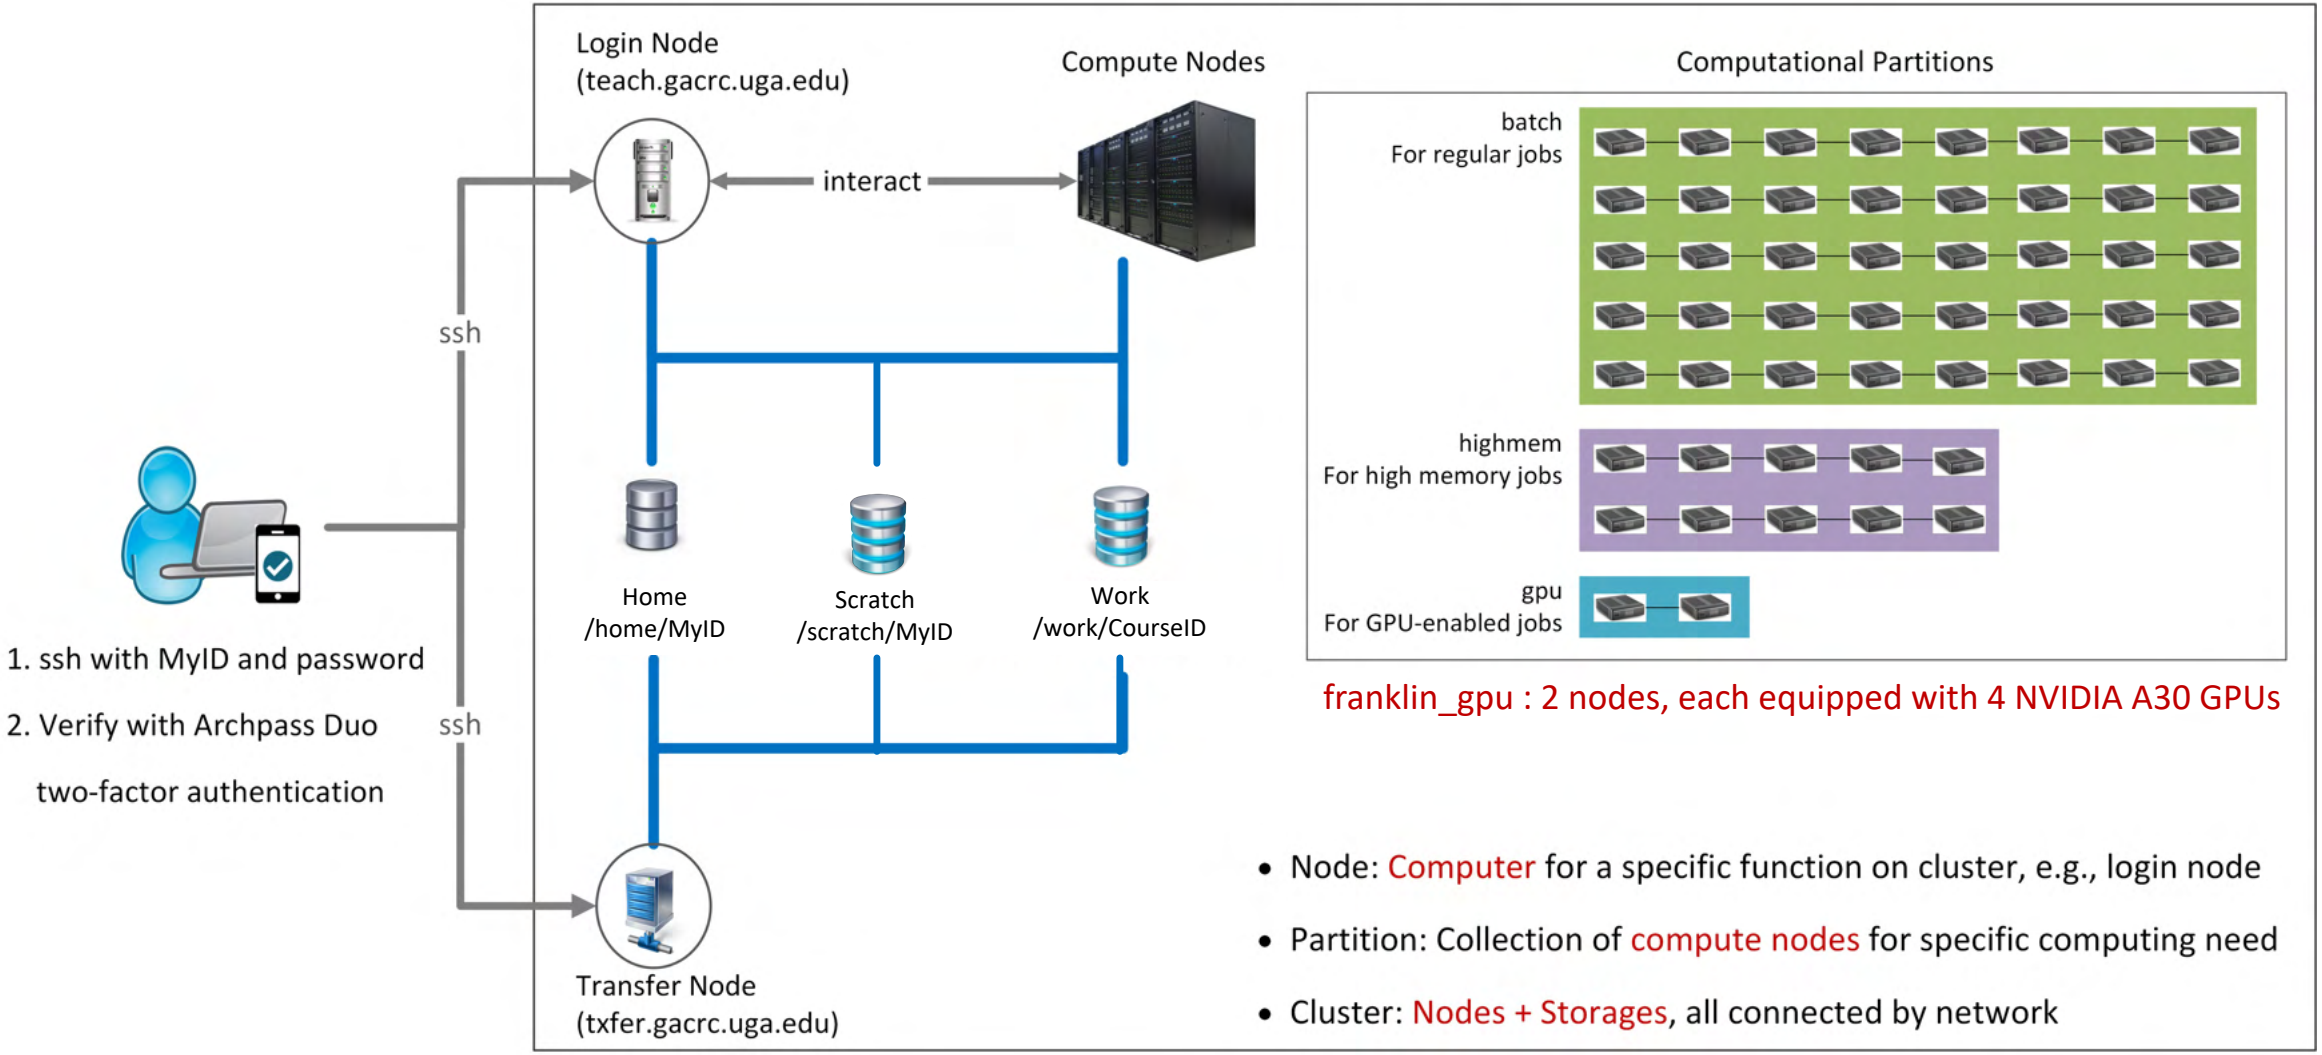

franklin_gpu : 2 nodes, each equipped with 4 NVIDIA A30 GPUs

- Node: Computer for a specific function on cluster, e.g., login node
- Partition: Collection of compute nodes for specific computing need
- Cluster: Nodes + Storages, all connected by network

Note: You need to connect to the UGA VPN at first when accessing from outside of the UGA main campus.

Working Environment

https://wiki.gacrc.uga.edu/wiki/Systems#Teaching_cluster

- Two nodes, your "username" is your MyID for both of them:
  1. For batch job workflow, the host to log into is teach.gacrc.uga.edu
  2. For file transfers, the host to log into is txfer.gacrc.uga.edu

- Three Directories:
  1. /home/MyID: directory for static data (e.g., scripts, software, etc…)
  2. /scratch/MyID: working space for running computational jobs
  3. /work/CourseID: directory for course data
     a. /work/CourseID/MyID : data storage space for individual user in a class (e.g., /work/ bcmb8330/MyID)
     b. /work/CourseID/instructor_data : data shared with class by the instructors

- Three Partitions:
  1. batch: for running regular computational jobs, e.g., MPI jobs
  2. highmem: for running high-memory jobs
  3. franklin_gpu and gpu: for running GPU jobs

Working Environment (cont.)

- Software
  1. Software names are long and have a Easybuild toolchain name associated to it
  2. Complete module name: Name/Version-toolchain, e.g., Python/3.12.3-GCCcore-13.3.0
  3. Software names are case-sensitive!
     - `module spider` *`pattern`* : Search modules using a name pattern (case-insensitive)
     - `module load/unload` *`moduleName`* : Load/remove a module
     - `module avail` : List all available modules on the cluster
     - `module list` : List modules currently loaded
     - `module purge` : Remove all modules from working environment

Submit a Batch Job

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster

1. Log on to Login node using MyID and password, and two-factor authentication with Archpass Duo:

   `ssh MyID@teach.gacrc.uga.edu`
2. Change directory to /scratch directory: `cd /scratch/MyID`
3. Create a working subdirectory for a job : `mkdir workDir`
4. Change directory to workDir : `cd workDir`
5. Transfer data from local computer to workDir : use `scp` or WinSCP to connect Transfer node

   Transfer data on cluster to workDir : log on to Transfer node and then use `cp` or `mv`
6. Make a job submission script in workDir : `nano sub_mpi.sh`
7. Submit a job from workDir : `sbatch sub_mpi.sh`
8. Check job status : `sq --me` or Cancel a job : `scancel JobID`

Step1: Log on to Login node

https://wiki.gacrc.uga.edu/wiki/Connecting#Connecting_to_the_teaching_cluster

1. Teaching cluster access requires verification using two-factor authentication with Archpass Duo. If you are not enrolled in Archpass Duo, please refer to https://uga.teamdynamix.com/TDClient/3190/eitsclientportal/KB/?CategoryID=23825 on how to enroll.
2. If you are connecting from off-campus, please first connect to the UGA VPN, then connect to teach.gacrc.uga.edu. Information on how to use the VPN is available at https://uga.teamdynamix.com/TDClient/3190/eitsclientportal/KB/?CategoryID=23843

Step1: Log on to Login node - Mac/Linux using ssh

1. Open Terminal utility
2. Type command line: `ssh MyID@teach.gacrc.uga.edu`
3. You will be prompted for your UGA MyID password
4. You will verify your login using Archpass Duo authentication

```
ssh zhuofei@teach.gacrc.uga.edu        ← 1. use ssh to open connection

UGA DUO authentication is required for SSH/SCP access to GACRC
systems. For additional help with UGA DUO authentication or to
report an issue please visit: https://eits.uga.edu/access_and_security...

Password:        ← 2. Enter your MyID password
                    When you enter password, no stars or dots will show
                    as you are typing. Please type password carefully!
Duo two-factor login for zhuofei

Enter a passcode or select one of the following options:

 1. Duo Push to XXX-XXX-5758
 2. Phone call to XXX-XXX-5758
 3. Phone call to XXX-XXX-1925
 4 5. SMS passcodes to XXX-XXX-5758 (next code starts with: 1)

Passcode or option (1-5): 1        ← 3. Select Duo option
Success. Logging you in...
Last login: Mon Aug  3 11:11:58 2020 from 172.18.114.119
zhuofei@teach-sub1 ~$        ← 4. Logged on!
```

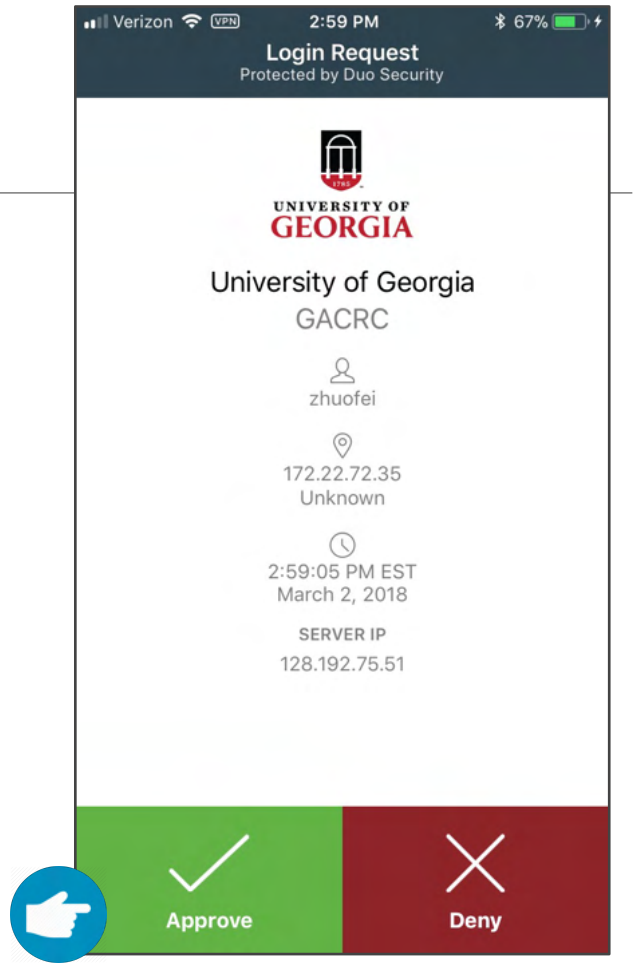

5. Verify login using Duo

Step1 (Cont.) - Windows using PuTTY

1. Download and install PuTTY:

   https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html

2. Detailed downloading and installation instructions:

   https://wiki.gacrc.uga.edu/wiki/How_to_Install_and_Configure_PuTTY

3. Detailed configuring and usage instructions:

   https://wiki.gacrc.uga.edu/wiki/How_to_Install_and_Configure_PuTTY#Configuring_PuTTY

Step1 (Cont.) - Windows using PuTTY

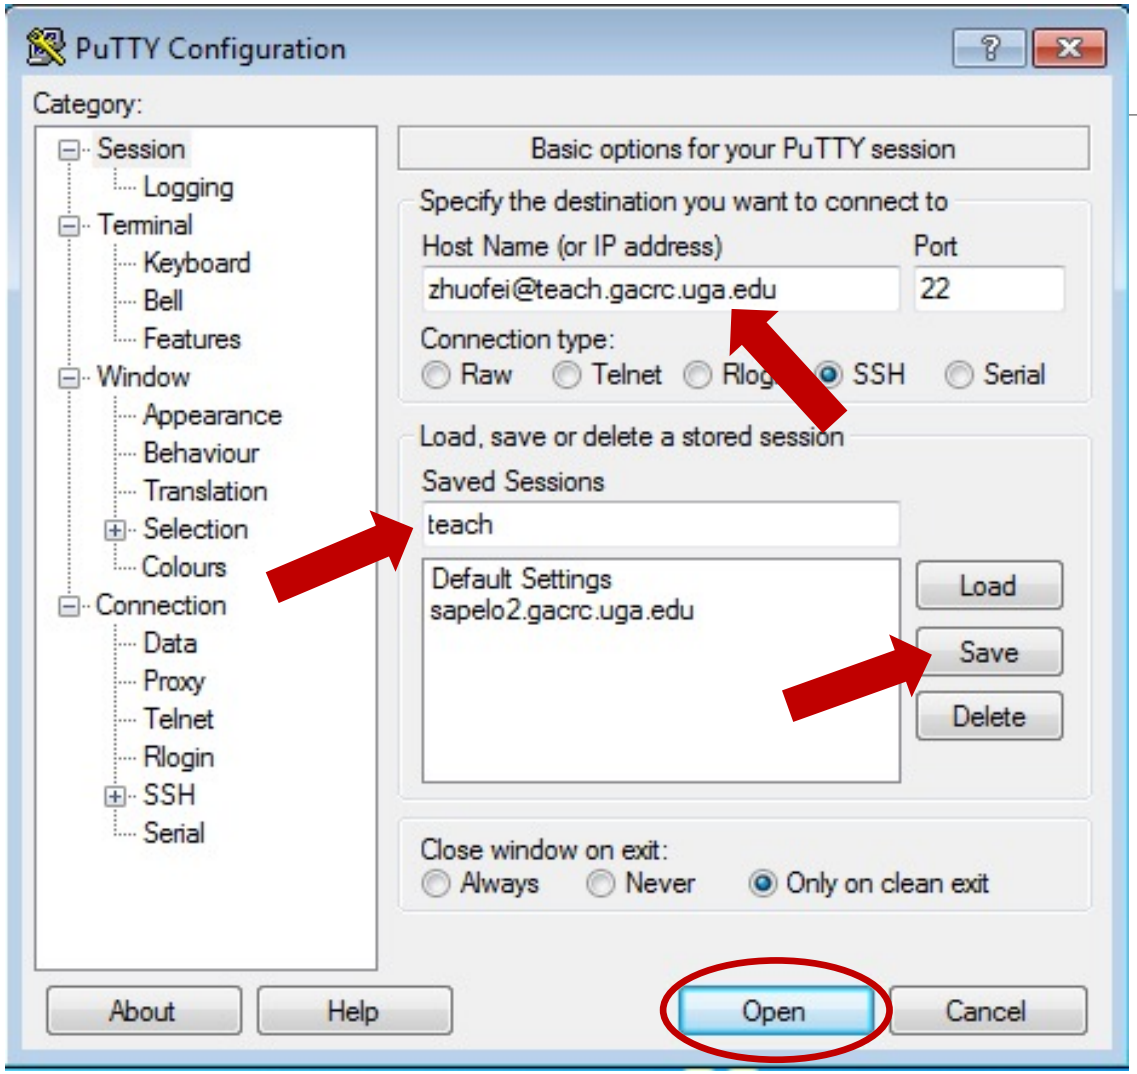

The first time you connect to login node, PuTTY will give you this security alert window. Please click "Yes"

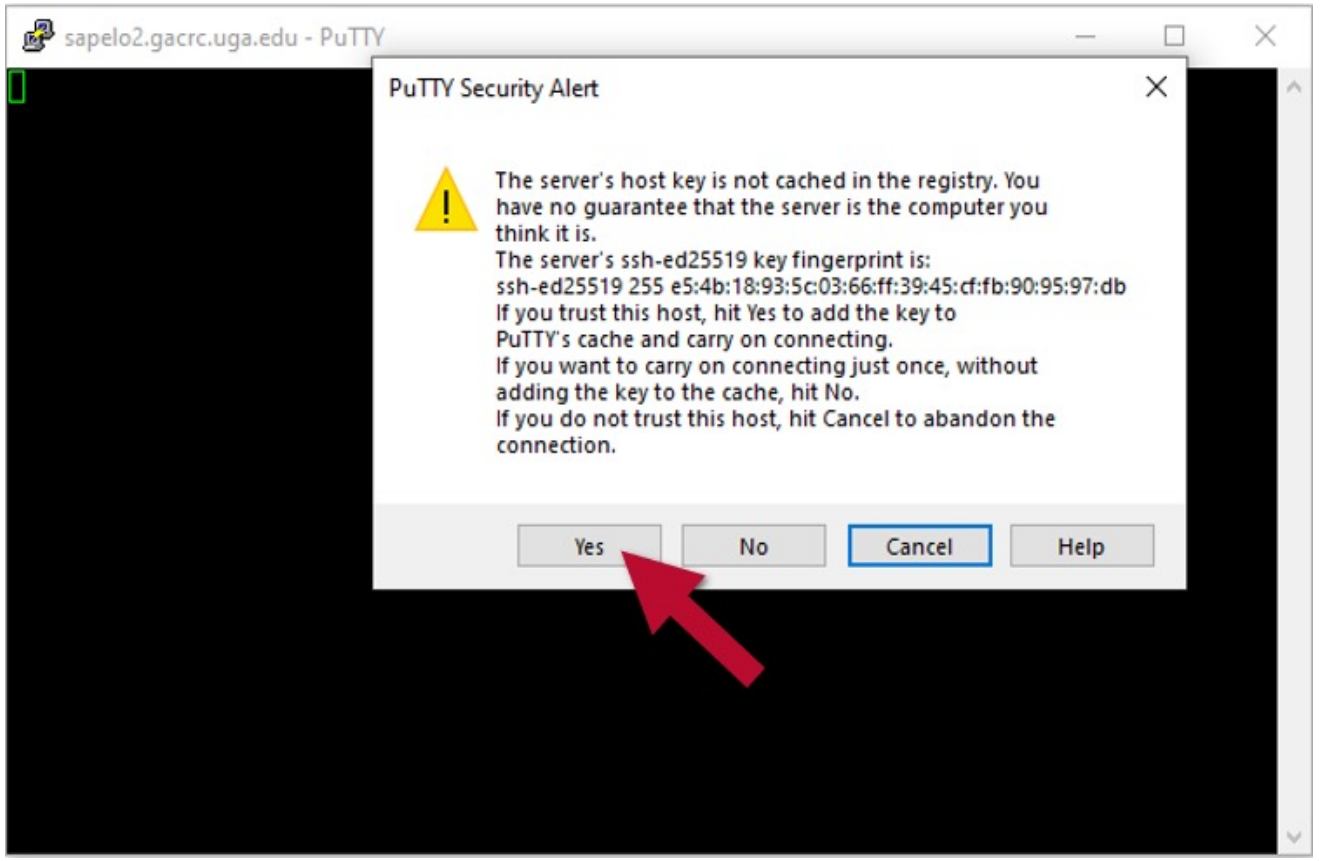

Step1 (Cont.) - Windows using PuTTY

Next you will enter your UGA MyID password and initiate DUO authentication procedure:

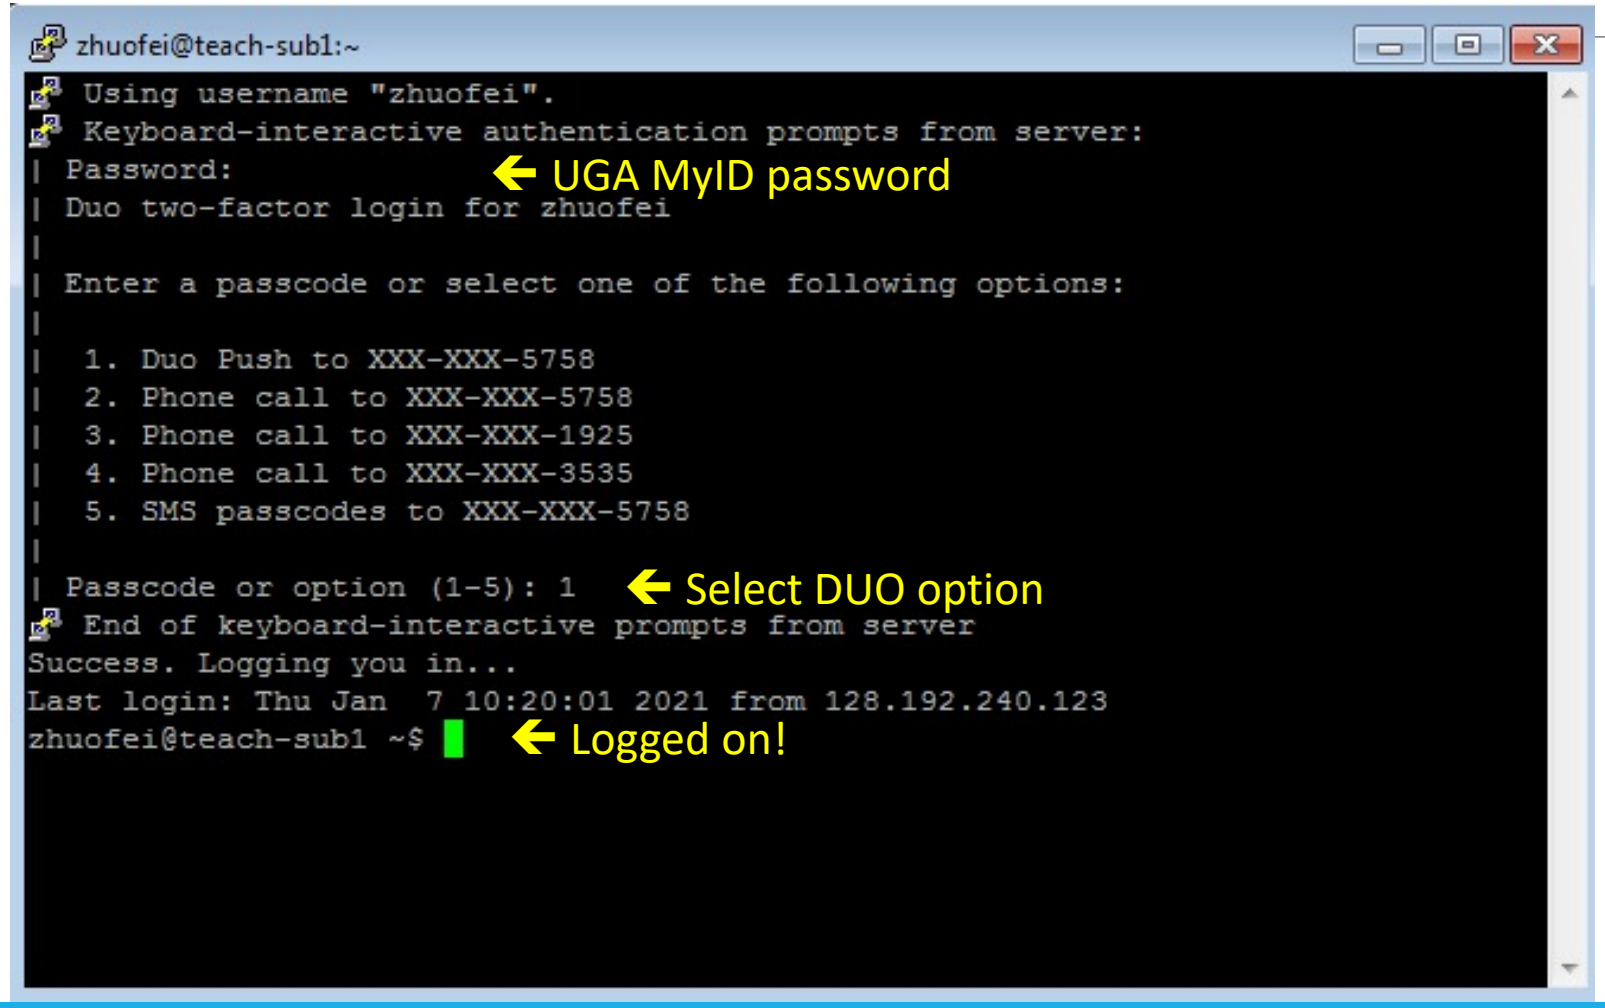

Step2 - 4: cd to /scratch dir, make and cd into workDir

```
zhuofei@teach-sub1 ~$ cd /scratch/MyID          ← cd command to change directory
zhuofei@teach-sub1 zhuofei$ mkdir workDir       ← mkdir command to create a subdirectory
zhuofei@teach-sub1 zhuofei$ cd workDir/         ← cd command to change directory
zhuofei@teach-sub1 workDir$ ls                  ← ls command to list contents of directory
zhuofei@teach-sub1 workDir$                     ← it is empty in workDir!
```

Step 5: Transfer data from local computer to workDir - Mac/Linux

https://wiki.gacrc.uga.edu/wiki/Transferring_Files#Using_scp_2

1. Connect to Transfer node (txfer.gacrc.uga.edu) in Terminal from your local computer
2. Use scp command: scp (-r) [Source] [Target]
3. Enter your MyID password, then select Duo option to verify connection

E.g. 1: use scp on local computer, from Local ➔ workDir on cluster

```
scp ./file zhuofei@txfer.gacrc.uga.edu:/home/zhuofei/workDir
```

```
scp -r ./folder/ zhuofei@txfer.gacrc.uga.edu:/home/zhuofei/workDir
```

E.g. 2: use scp on local computer, from workDir on cluster ➔ Local

```
scp zhuofei@txfer.gacrc.uga.edu:/home/zhuofei/workDir/file .
```

```
scp -r zhuofei@txfer.gacrc.uga.edu:/home/zhuofei/workDir/folder/ .
```

Step 5 (Cont.) - Windows using WinSCP

https://wiki.gacrc.uga.edu/wiki/Transferring_Files#Using_WinSCP_2

1. You need to connect to cluster's Transfer node (txfer.gacrc.uga.edu)
2. Use WinSCP on local computer
   - WinSCP can be downloaded from https://winscp.net/eng/index.php
   - Default installation procedure is simple
3. Alternative FileZilla https://wiki.gacrc.uga.edu/wiki/Transferring_Files#Using_FileZilla_2

Step 5 (Cont.) - Windows using WinSCP

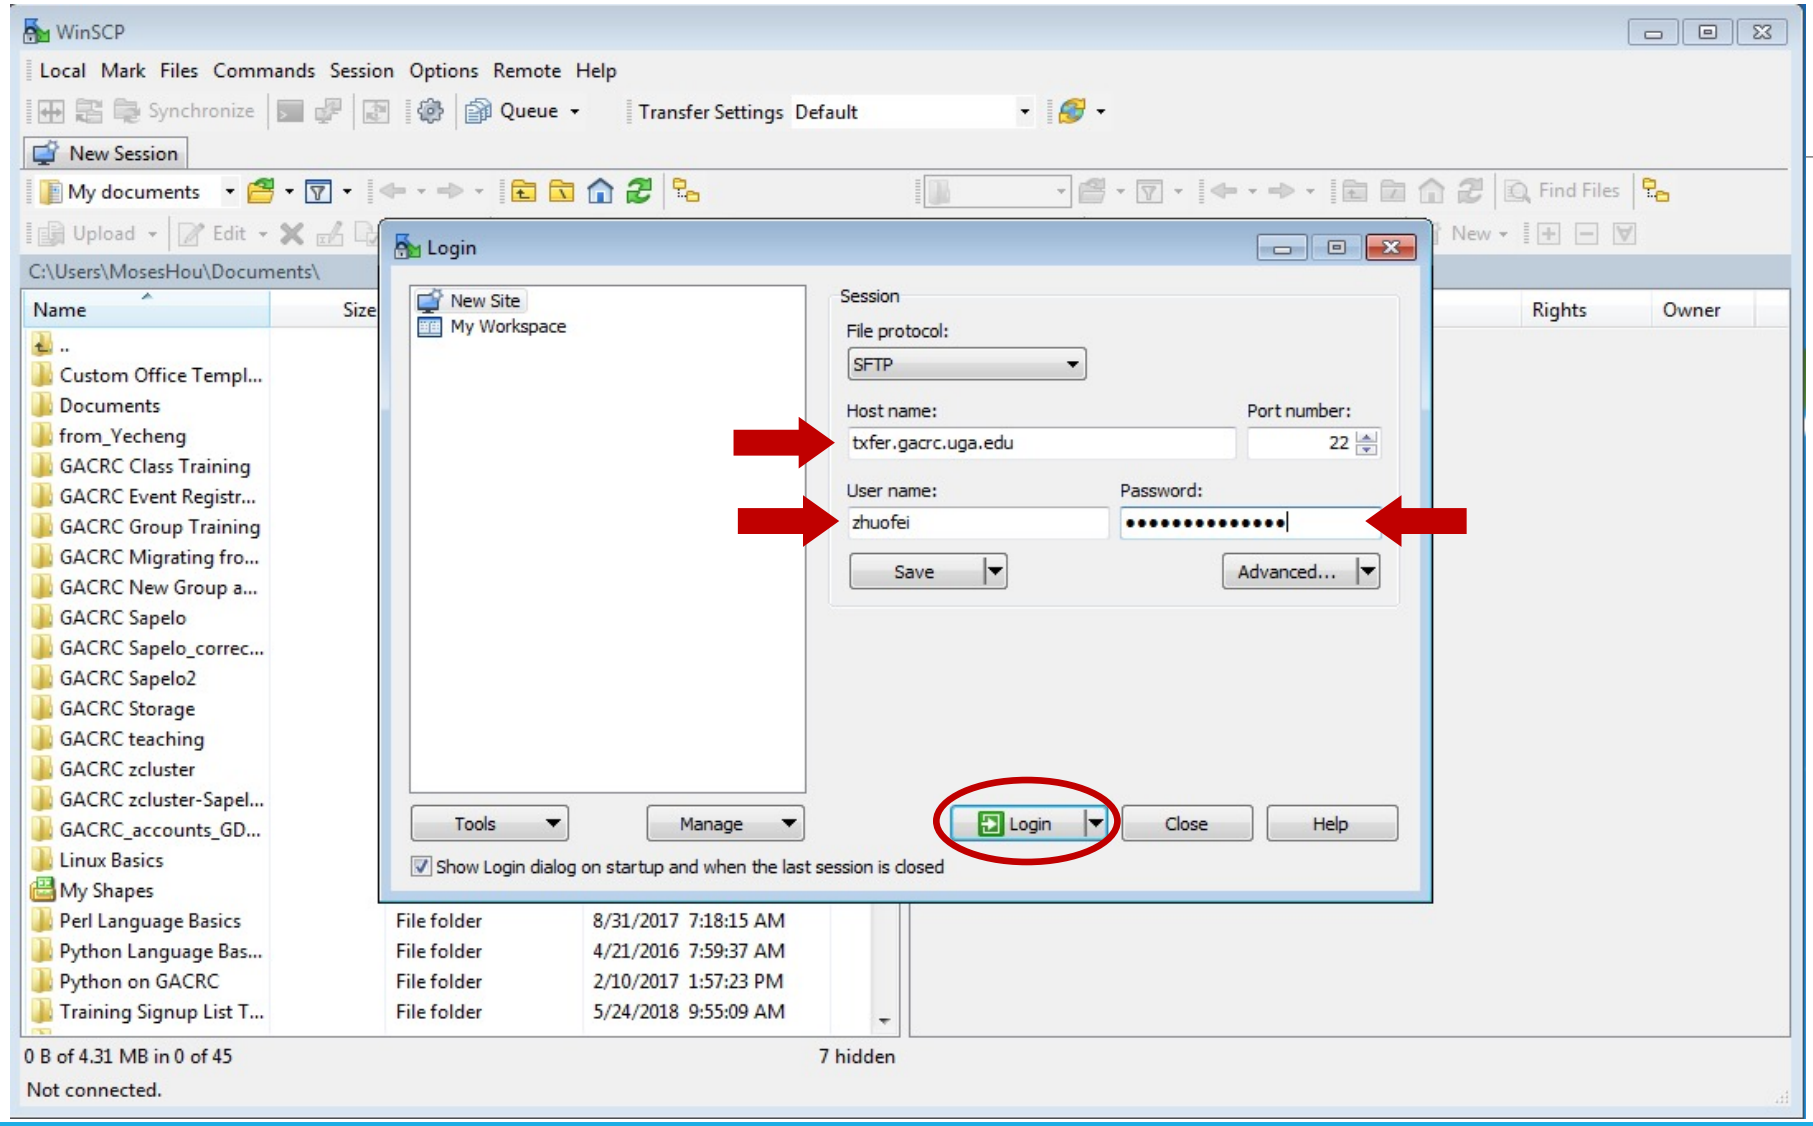

Step 5 (Cont.) - Windows using WinSCP

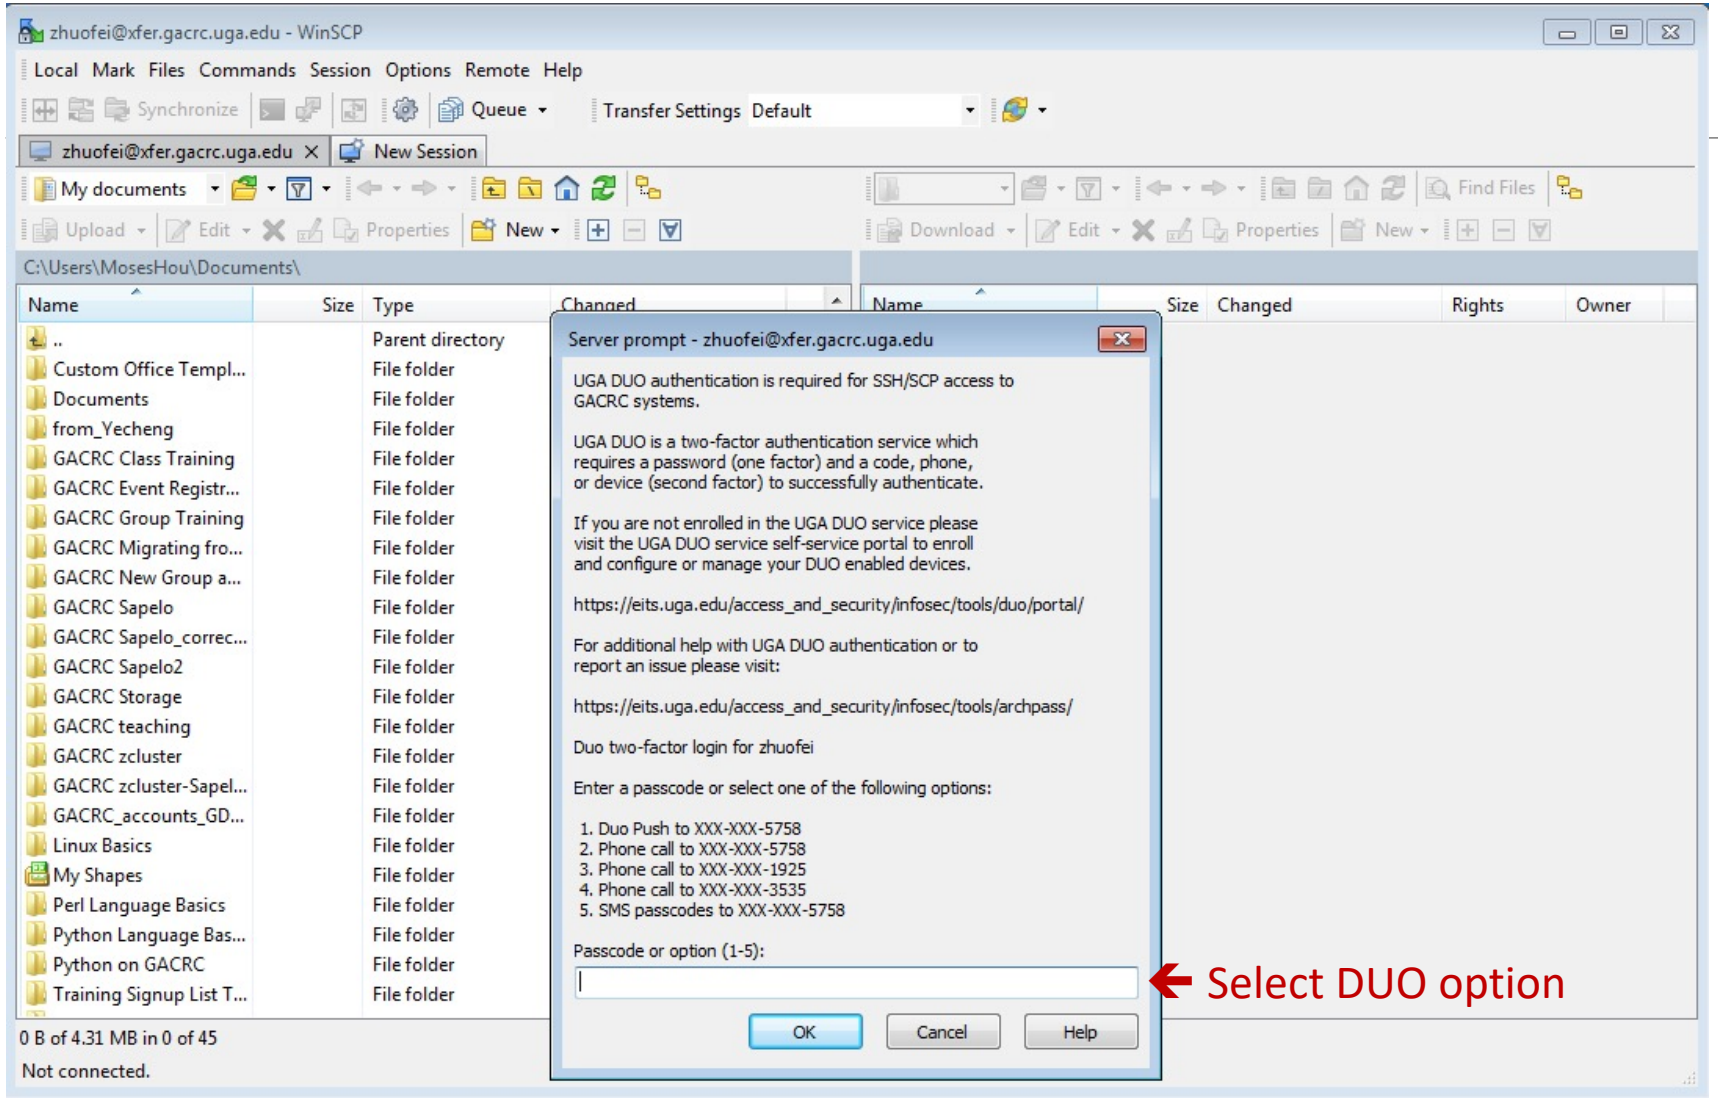

Step 5 (Cont.) - Windows using WinSCP

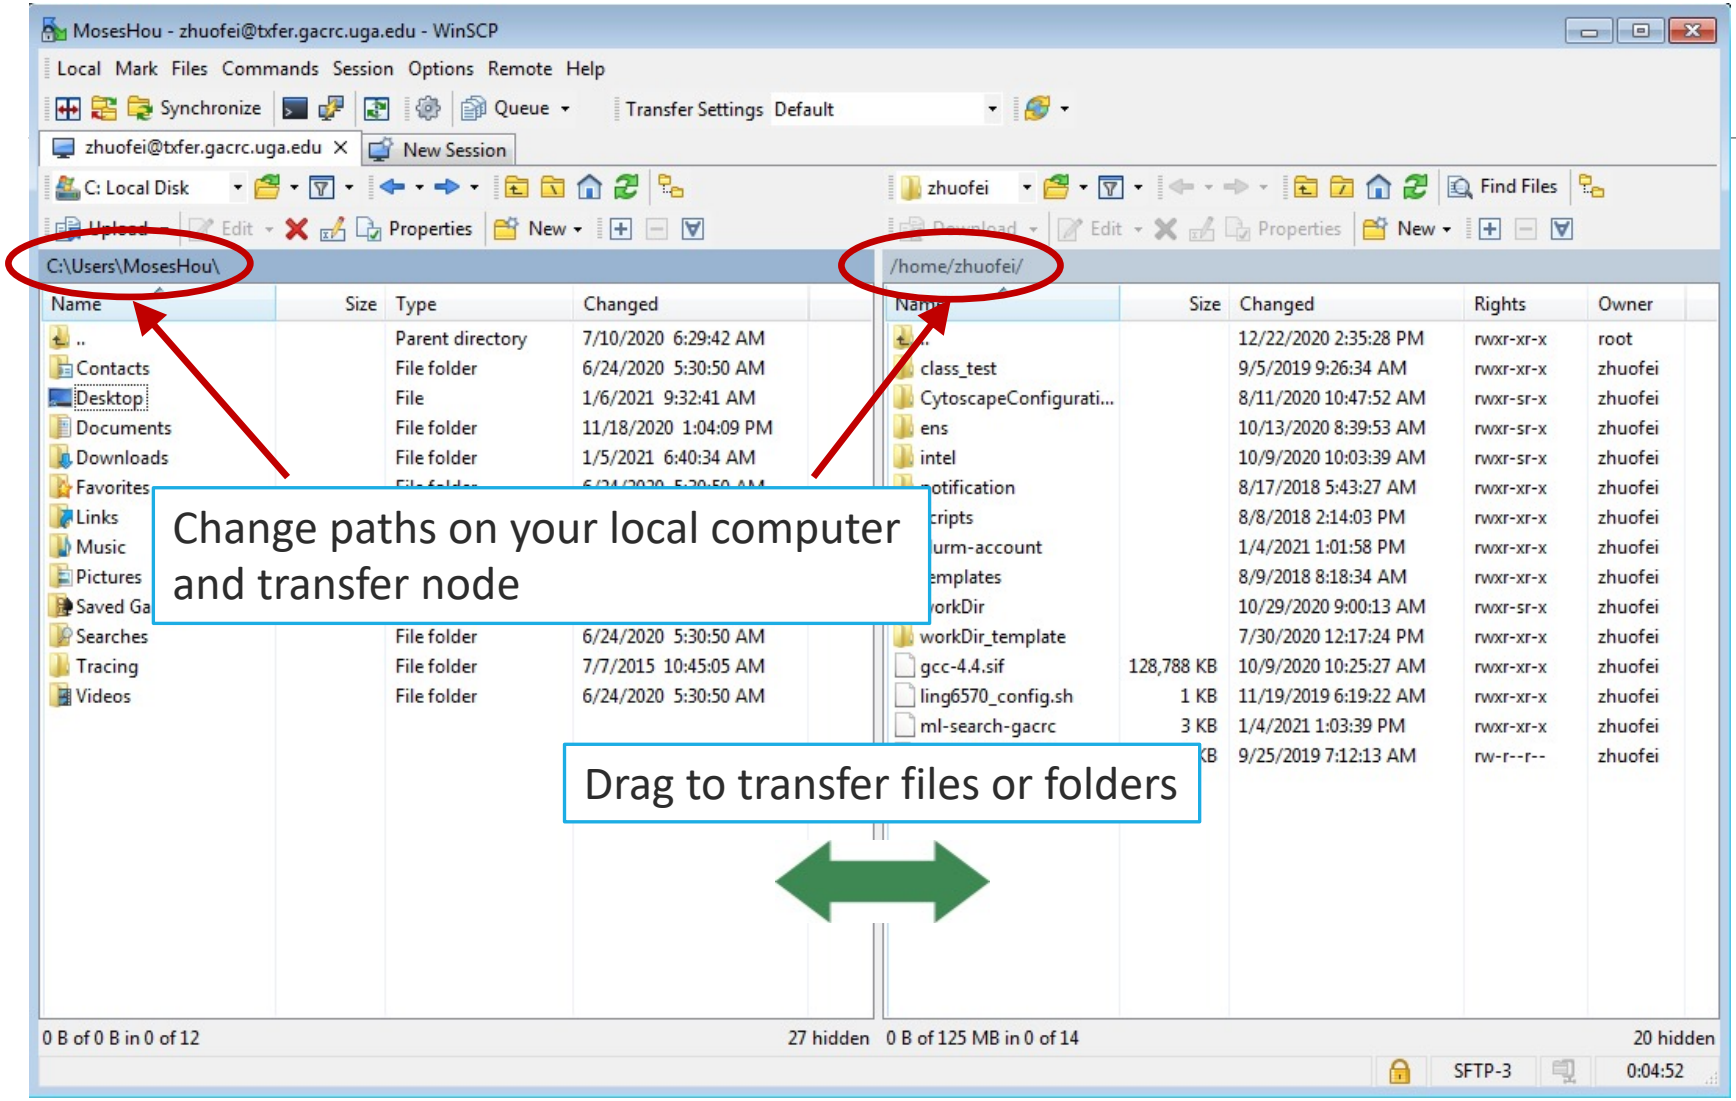

Step 5 (Cont.): Transfer data on cluster to workDir

- Log on to Transfer node (txfer.gacrc.uga.edu)
  - ✓ Mac/Linux: ssh MyID@txfer.gacrc.uga.edu (page 9-10)
  - ✓ Windows: use PuTTY to log in MyID@txfer.gacrc.uga.edu (page 11-13)
- Directories you can access on transfer node:
  1. /home/MyID
  2. /scratch/MyID
  3. /work/CourseID/
- Transfer data between two folders on cluster using cp or mv, e.g.:

```
mv /work/bcmb8330/MyID/datafile /scratch/MyID/workDir
```

Step 6: Make a job submission script in workDir using nano

https://wiki.gacrc.uga.edu/wiki/Sample_batch_job_submission_scripts_on_the_teaching_cluster

```
$ nano sub.sh
```

nano is a simple text editor on Linux. You are welcome to use other editors like vim or emacs.

Ctrl-x to save file and quit from nano

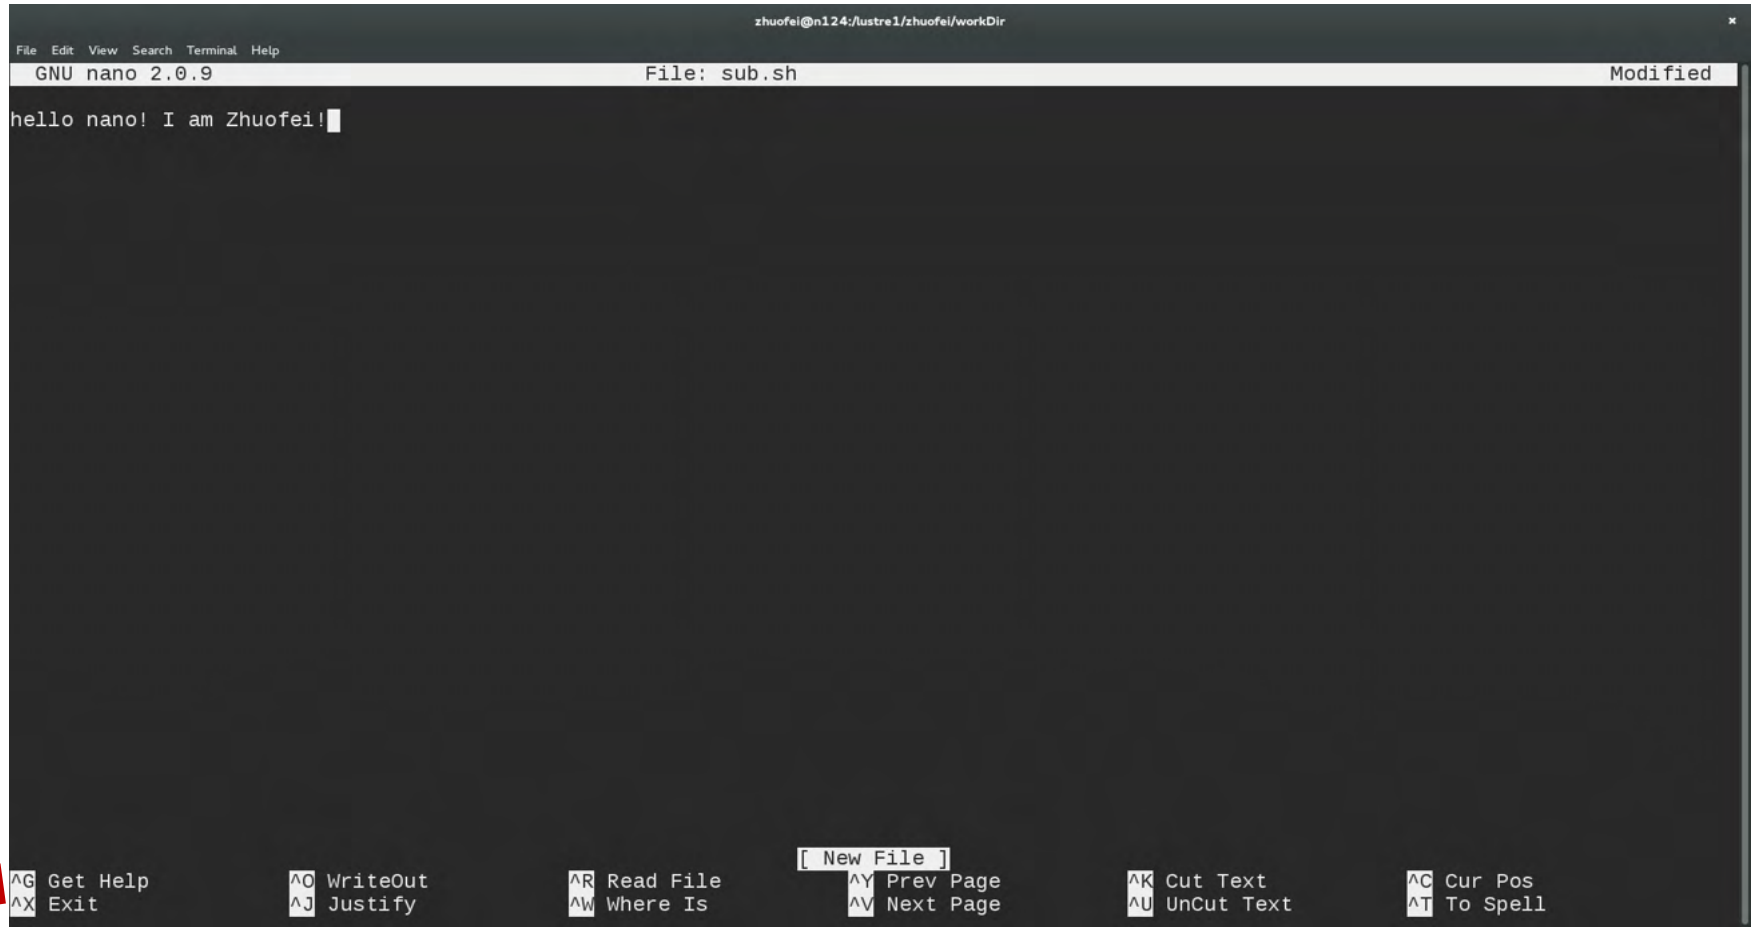

Step 6 (Cont.)

1. Copy sample job to workDir:

cp -r /usr/local/training/bcmb8330/* .

2. Job submission script:

sub_mpi.sh

3. Amber Wiki:

https://wiki.gacrc.uga.edu/wiki/AMBER-Sapelo2

```
#!/bin/bash
#SBATCH --job-name=Amber_MPI              # Job name
#SBATCH --partition=batch                 # Computational partition
#SBATCH --ntasks=10                       # MPI processes/rank number
#SBATCH --cpus-per-task=1                 # Number of CPU cores per task/MPI process
#SBATCH --mem-per-cpu=2gb                 # Memory per CPU core
#SBATCH --time=24:00:00                   # Time limit hrs:mins:secs
#SBATCH --output=log.%j.out               # Standard output log
#SBATCH --error=log.%j.err                # Standard error log
#SBATCH --mail-user=MyID@uga.edu          # Where to send mail
#SBATCH --mail-type=ALL                   # Mail events (BEGIN, END, FAIL, ALL)

cd $SLURM_SUBMIT_DIR                      # cd /scratch/MyID/workDir

ml Amber/24.3-foss-2022a-AmberTools-24.10-CUDA-12.1.1          # Load Amber module
source ${AMBERHOME}/amber.sh              # Config software environment for Amber

# PMEMD: Job1: minimization, solvent
srun ${AMBERHOME}/bin/pmemd.MPI -O \
 -i min_solvent.in \
 -o min_solvent.out \
 -p gfp.parm7 \
 -c gfp.rst7 \
 -ref gfp.rst7 \
 -r gfp_min_solvent.rst7
```

Step 7: Submit a job from workDir using sbatch

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster#How_to_submit_a_job_to_the_batch_queue

```
$ sbatch sub.sh

Submitted batch job 542
```

Tips: sub.sh is a job submission script for

1. specifying computing resources
2. loading software using module load
3. running any Linux commands that you want to run
4. running the software that you loaded

Step 8: Check job status using sq --me

https://wiki.gacrc.uga.edu/wiki/Monitoring_Jobs_on_the_teaching_cluster

```
sq --me
JOBID NAME PARTITION USER NODES CPUS MIN_MEMORY PRIORITY  TIME  TIME_LIMIT STATE  NODELIST(REASON)
542   Amber_MPI batch zhuofei 1 10   2G         24        1:56  1-00:00:00 RUNNING   rb1-3
```

Step 8 (Cont.): Cancel job using scancel

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster#How_to_delete_a_running_or_pending_job

```
scancel 542


sq --me
JOBID NAME PARTITION USER NODES CPUS MIN_MEMORY PRIORITY TIME TIME_LIMIT STATE    NODELIST(REASON)
<empty output>
```

Step8 (Cont.): Check job details using sacct-gacrc -X and seff

https://wiki.gacrc.uga.edu/wiki/Monitoring_Jobs_on_the_teaching_cluster

```
$ sacct-gacrc -X

JobID     JobName      User    Partition  NNode  NCPUS  ReqMem   CPUTime    Elapsed    Timelimit       State  ExitCode  NodeList
------------------ ---------- -------------- --------- ---------- ----- ----- -------- ---------- ---------- ---------- --------- -------- ---------- ---------- ---------
542   Amber_MPI   zhuofei    batch      1     10      20G   00:38:20  00:03:50  1-00:00:00  CANCELLED+   0:0    rb1-3
543   Amber_MPI   zhuofei    batch      1     10      20G   01:46:00  00:10:36  1-00:00:00  COMPLETED    0:0    rb1-3

seff 543            # Check computing resources used by a COMPLETED job
Job ID: 543
Cluster: gacrc-teach
User/Group: zhuofei/gacrc-instruction
State: COMPLETED (exit code 0)
Nodes: 1
Cores per node: 10
CPU Utilized: 01:39:15
CPU Efficiency: 93.63% of 01:46:00 core-walltime
Job Wall-clock time: 00:10:36
Memory Utilized: 1.58 GB (estimated maximum)
Memory Efficiency: 7.88% of 20.00 GB (2.00 GB/core)
```

Step 8 (Cont.): Check node info using sinfo

https://wiki.gacrc.uga.edu/wiki/Monitoring_Jobs_on_the_teaching_cluster

```
$ sinfo

PARTITION    AVAIL  TIMELIMIT  NODES  STATE NODELIST
batch           up 7-00:00:00      8   idle rb1-[3-10]
gpu             up 7-00:00:00      1  down* c4-23
highmem         up 7-00:00:00      2   idle rb1-[1-2]
interactive     up 7-00:00:00      1    mix rb1-11
interactive     up 7-00:00:00      1   idle rb1-12
franklin_gpu    up 7-00:00:00      1    mix b8-6
franklin_gpu    up 7-00:00:00      1   idle b8-7
fsr4601         up       1:00      8   idle rb1-[3-10]
fsr8602         up      10:00      8   idle rb1-[3-10]
```

idle = no cores in use; mix = some cores are still free; alloc = all cores are allocated

Obtain Job Details

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster#How_to_check_resource_utilization_of_a_running_or_finished_job

Option 1: sq --me for details of running or pending jobs

Option 2: sacct-gacrc -X for details of computing resource usage of a running or finished job

Option 3: seff for details of computing resource usage of a finished job

Option 4: Email notification from finished jobs (completed, canceled, or crashed), if using:

```
#SBATCH --mail-user=username@uga.edu

#SBATCH --mail-type=ALL
```

Run Interactive Jobs

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster#How_to_open_an_interactive_session

https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster#How_to_run_an_interactive_job_with_Graphical_User_Interface_capabilities

Description	Command
Start an interactive session	interact
Start an interactive session with X forwarding	interact --x11

interact	srun --pty --cpus-per-task=1 --job-name=interact --ntasks=1 --nodes=1 --partition=interactive --time=12:00:00 --mem=2GB /bin/bash -l
interact --x11	srun --pty --cpus-per-task=1 --job-name=interact --ntasks=1 --nodes=1 --partition=interactive --time=12:00:00 --mem=2GB --x11 /bin/bash -l

Sample job script for running Amber GPU job

1. Copy sample job to workDir:

cp -r /usr/local/training/bcmb8330/* .

2. Job submission script:

sub_gpu.sh

3. Amber Wiki:

https://wiki.gacrc.uga.edu/wiki/AMBER-Sapelo2

```
#!/bin/bash
#SBATCH --job-name=Amber_GPU          # Job name
#SBATCH --partition=franklin_gpu      # Computational partition (franklin_gpu)
#SBATCH --gres=gpu:1                  # Request 1 GPU device
#SBATCH --ntasks=1                    # Run a single task
#SBATCH --cpus-per-task=4             # Number of CPU cores per task
#SBATCH --mem-per-cpu=2gb             # Memory per CPU core
#SBATCH --time=24:00:00               # Time limit hrs:mins:secs
#SBATCH --output=log.%j.out           # Standard output log
#SBATCH --error=log.%j.err            # Standard error log
#SBATCH --mail-user=MyID@uga.edu      # Where to send mail
#SBATCH --mail-type=ALL               # Mail events (BEGIN, END, FAIL, ALL)

cd $SLURM_SUBMIT_DIR                  # cd /scratch/MyID/workDir

ml Amber/24.3-foss-2022a-AmberTools-24.10-CUDA-12.1.1    # Load Amber module
source ${AMBERHOME}/amber.sh          # Config software environment for Amber

# PMEMD: Job1: minimization, solvent
$AMBERHOME/bin/pmemd.cuda -O \
 -i min_solvent.in \
 -o min_solvent.out \
 -p gfp.parm7 \
 -c gfp.rst7 \
 -ref gfp.rst7 \
 -r gfp_min_solvent.rst7
```

GACRC Wiki http://wiki.gacrc.uga.edu
Kaltura Channel https://kaltura.uga.edu/channel/GACRC/176125031

Connecting: https://wiki.gacrc.uga.edu/wiki/Connecting#Connecting_to_the_teaching_cluster

Running Jobs: https://wiki.gacrc.uga.edu/wiki/Running_Jobs_on_the_teaching_cluster

Monitoring Jobs: https://wiki.gacrc.uga.edu/wiki/Monitoring_Jobs_on_the_teaching_cluster

Transfer File:

https://wiki.gacrc.uga.edu/wiki/Transferring_Files#The_File_Transfer_node_for_the_teaching_cluster_.28txfer.gacrc.uga.edu.29

Sample Job Scripts:

https://wiki.gacrc.uga.edu/wiki/Sample_batch_job_submission_scripts_on_the_teaching_cluster

Linux Command: https://wiki.gacrc.uga.edu/wiki/Command_List

GACRC Support

https://wiki.gacrc.uga.edu/wiki/Getting_Help

➢ Job Troubleshooting:

Please tell us details of your question or problem, including but not limited to:

- ✓ Your user name
- ✓ Your job ID
- ✓ Your working directory
- ✓ The partition name and command you used to submit the job

➢ Software Installation:

- ✓ Specific name and version of the software
- ✓ Download website
- ✓ Supporting package information if have

Please note to make sure the correctness of datasets being used by your jobs!

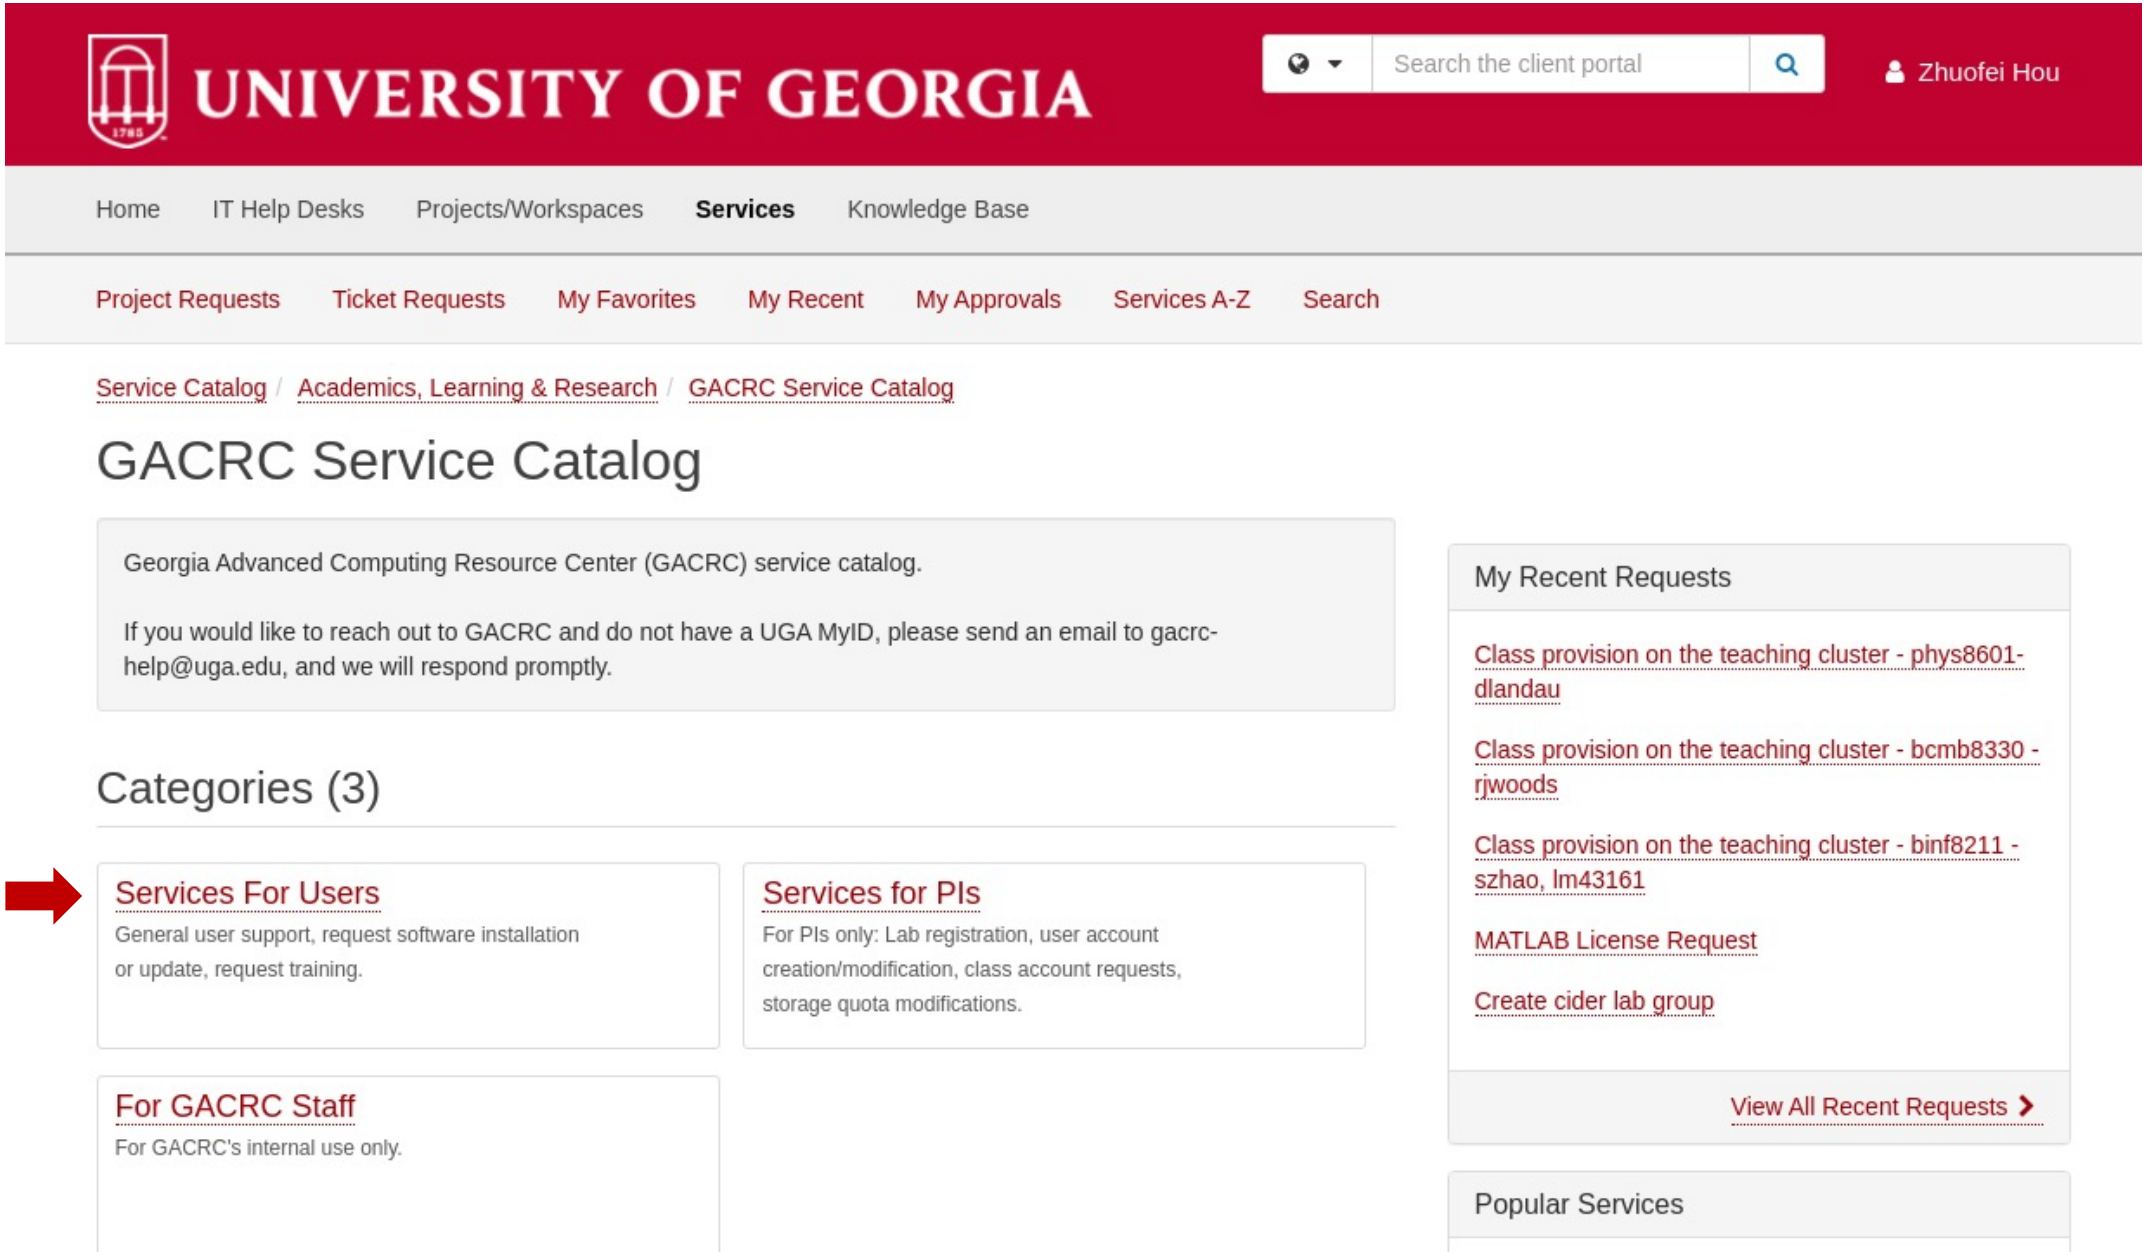
UNIVERSITY OF GEORGIA
Search the client portal
Zhuofei Hou
Home
IT Help Desks
Projects/Workspaces
Services
Knowledge Base
Project Requests
Ticket Requests
My Favorites
My Recent
My Approvals
Services A-Z
Search
Service Catalog / Academics, Learning & Research / GACRC Service Catalog
GACRC Service Catalog
Georgia Advanced Computing Resource Center (GACRC) service catalog.
If you would like to reach out to GACRC and do not have a UGA MyID, please send an email to gacrc-help@uga.edu, and we will respond promptly.
Categories (3)
Services For Users
General user support, request software installation or update, request training.
Services for PIs
For PIs only: Lab registration, user account creation/modification, class account requests, storage quota modifications.
For GACRC Staff
For GACRC's internal use only.
My Recent Requests
Class provision on the teaching cluster - phys8601-dlandau
Class provision on the teaching cluster - bcmb8330 - rjwoods
Class provision on the teaching cluster - binf8211 - szhao, lm43161
MATLAB License Request
Create cider lab group
View All Recent Requests
Popular Services

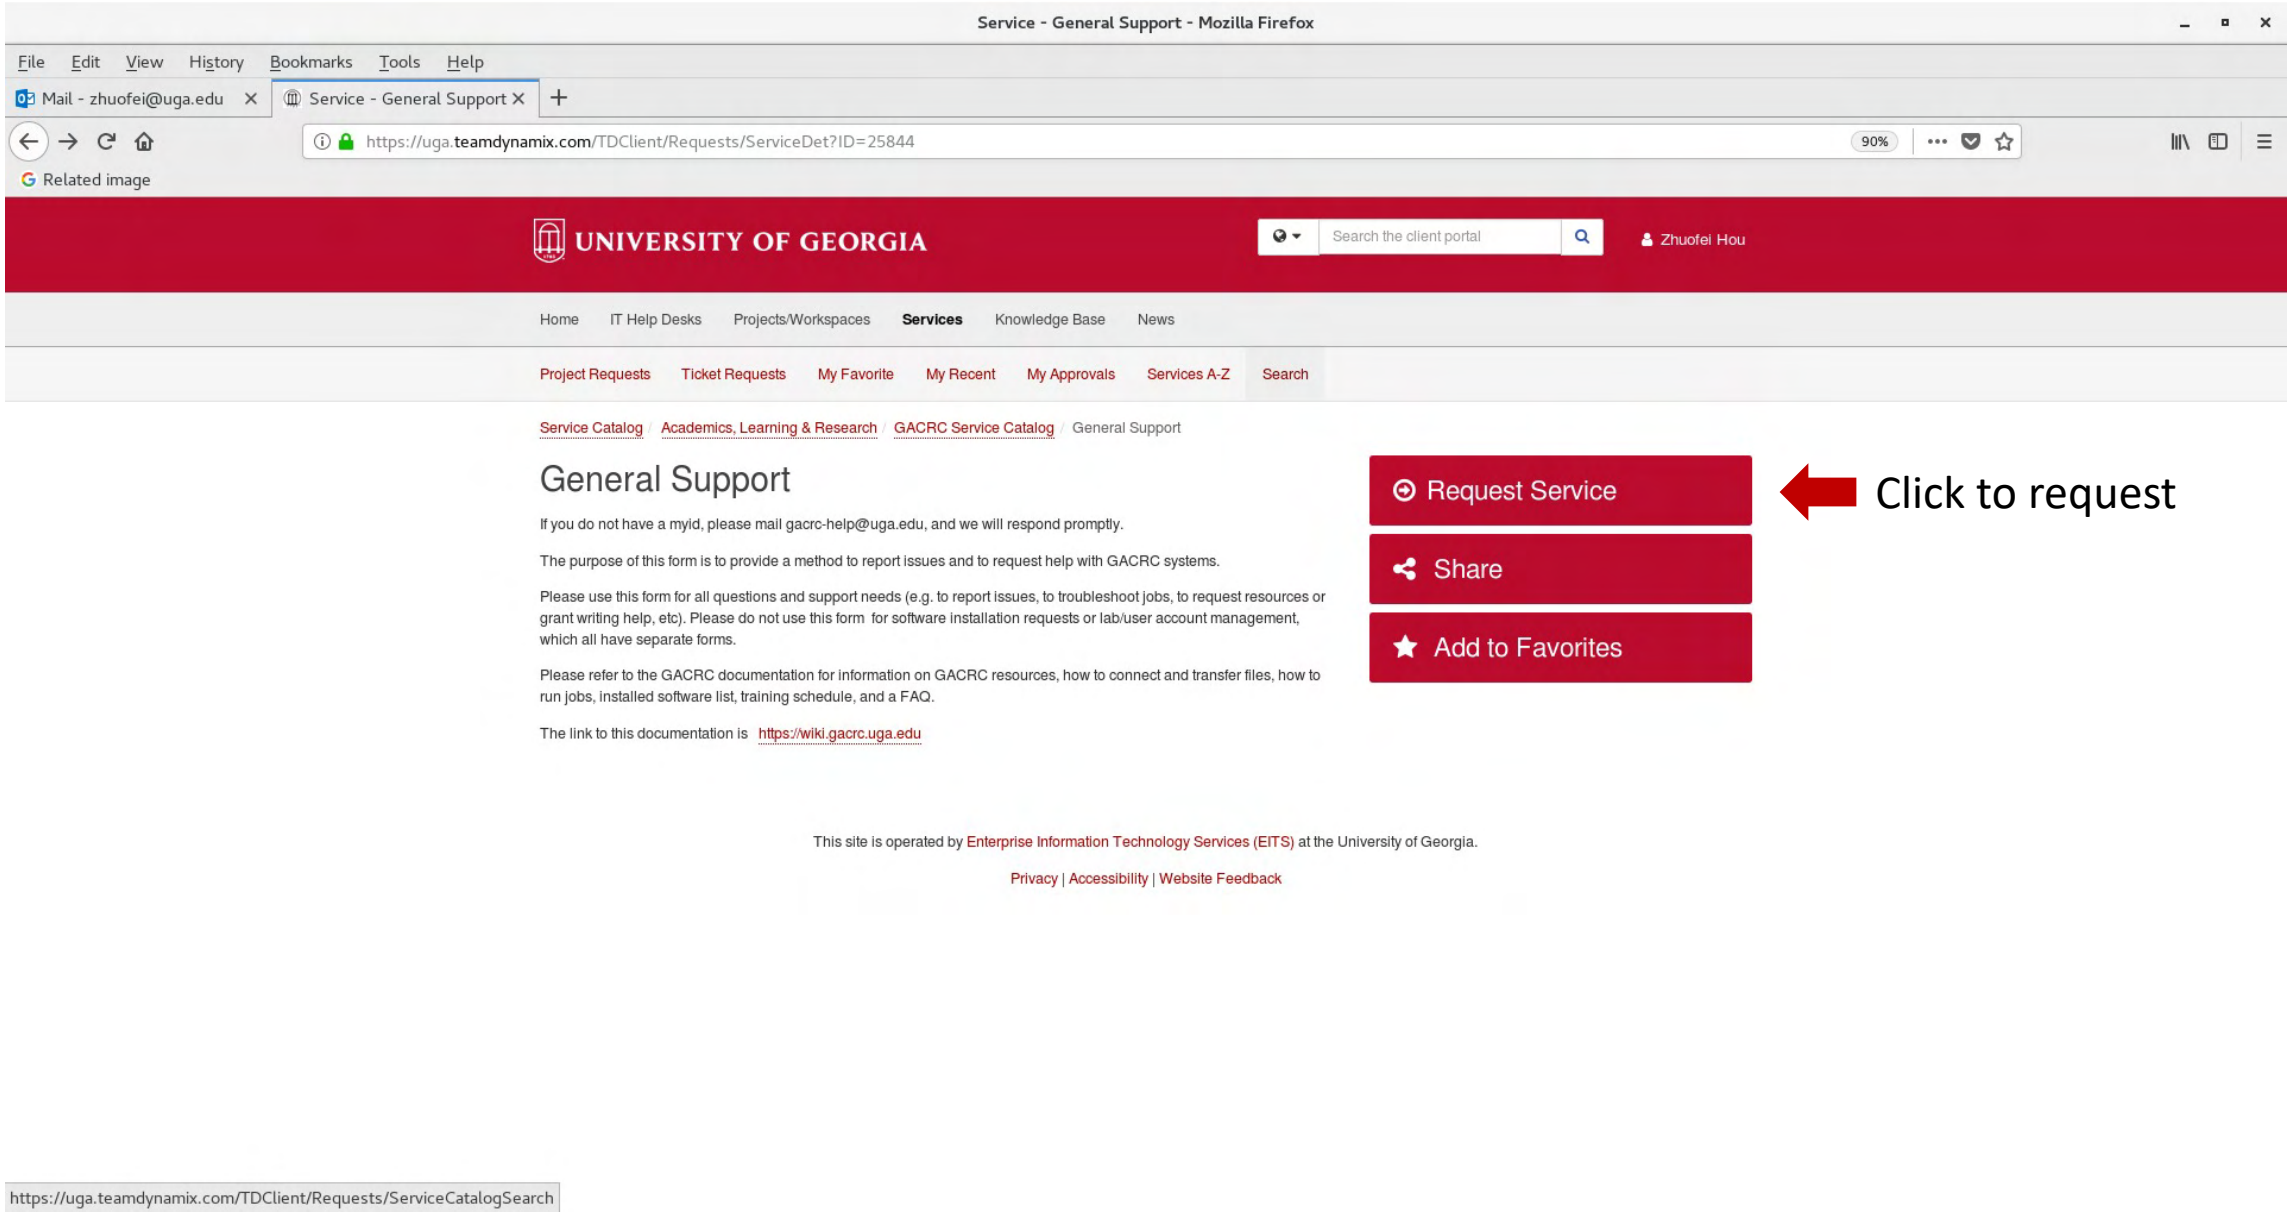
Service - General Support - Mozilla Firefox
File Edit View History Bookmarks Tools Help
Mail - zhuofei@uga.edu
Service - General Support
https://uga.teamdynamix.com/TDClient/Requests/ServiceDet?ID=25844
90%
Related image
UNIVERSITY OF GEORGIA
Search the client portal
Zhuofei Hou
Home IT Help Desks Projects/Workspaces Services Knowledge Base News
Project Requests Ticket Requests My Favorite My Recent My Approvals Services A-Z Search
Service Catalog / Academics, Learning & Research / GACRC Service Catalog / General Support
General Support
If you do not have a myid, please mail gacrc-help@uga.edu, and we will respond promptly.
The purpose of this form is to provide a method to report issues and to request help with GACRC systems.
Please use this form for all questions and support needs (e.g. to report issues, to troubleshoot jobs, to request resources or grant writing help, etc). Please do not use this form for software installation requests or lab/user account management, which all have separate forms.
Please refer to the GACRC documentation for information on GACRC resources, how to connect and transfer files, how to run jobs, installed software list, training schedule, and a FAQ.
The link to this documentation is https://wiki.gacrc.uga.edu
Request Service
Share
Add to Favorites
Click to request
This site is operated by Enterprise Information Technology Services (EITS) at the University of Georgia.
Privacy | Accessibility | Website Feedback
https://uga.teamdynamix.com/TDClient/Requests/ServiceCatalogSearch

Thank You!

Telephone Support

EITS Help Desk: 706-542-3106

Monday – Friday: 7:30 a.m. – 6:30 p.m.

Saturday: closed

Sunday: 1:30 p.m. – 6:30 p.m.

Georgia Advanced Computing Resource Center

101-108 Computing Services building

University of Georgia

Athens, GA 30602

https://gacrc.uga.edu/