Georgia Advanced Computing Resource Center
UNIVERSITY OF GEORGIA

High Performance Computing Using Sapelo Cluster

Georgia Advanced Computing Resource Center
EITS/UGA

Zhuofei Hou, Training Advisor
zhuofei@uga.edu

Outline

- GACRC
- What is High Performance Computing (HPC)
- Sapelo Cluster Diagram
- How to Log on
- How to Transfer Files
- How to Use Software
- How to Submit a Job
- How to Find Support and Important Rules

GACRC

- We are a high-performance computing (HPC) center at UGA
- We 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

http://gacrc.uga.edu/about/

What is High Performance Computing

1 worker

10 years later……

What is High Performance Computing

1 month later……

What is High Performance Computing

Serial Job
with 1 worker

- ✓ Only 1 instruction executed at any time point on a single processor

Parallel Job
with 6 workers

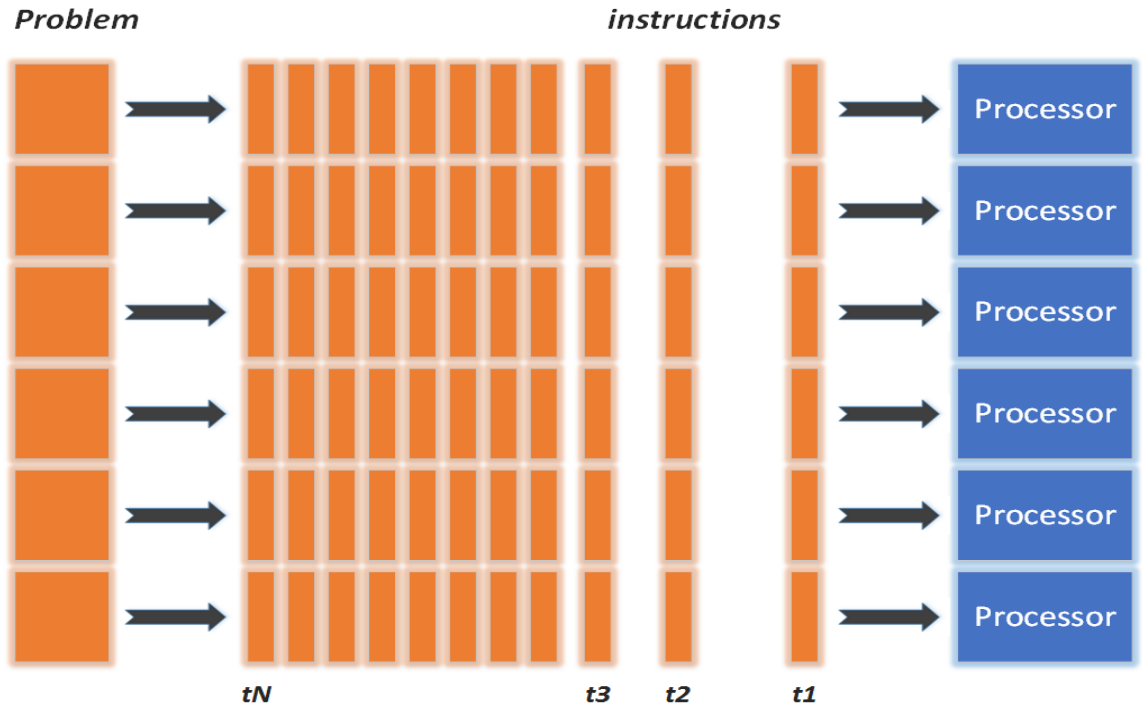

- ✓ 6 instructions executed *simultaneously* on *6* processors
- ✓ Shared-memory threaded job or MPI job (Message Passing Interface)

What is a Cluster?

When you work on cluster, 3 roles are working:

- User: to submit a job
- Queueing System: to dispatch a job to run on cluster
- Cluster: to run a job

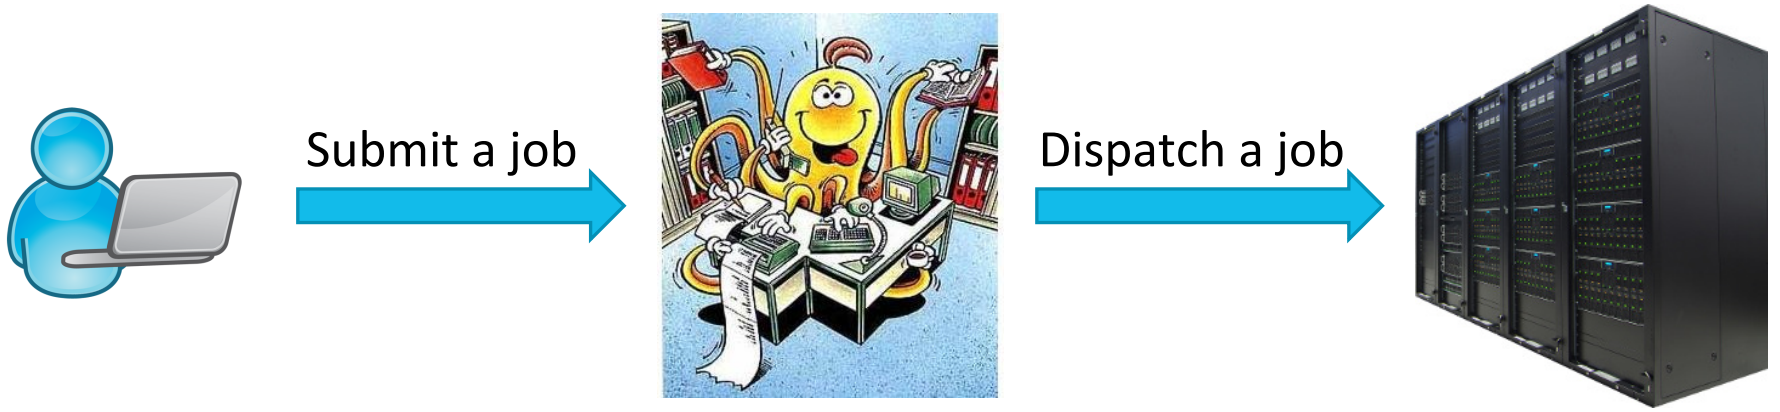

Sapelo Cluster Diagram

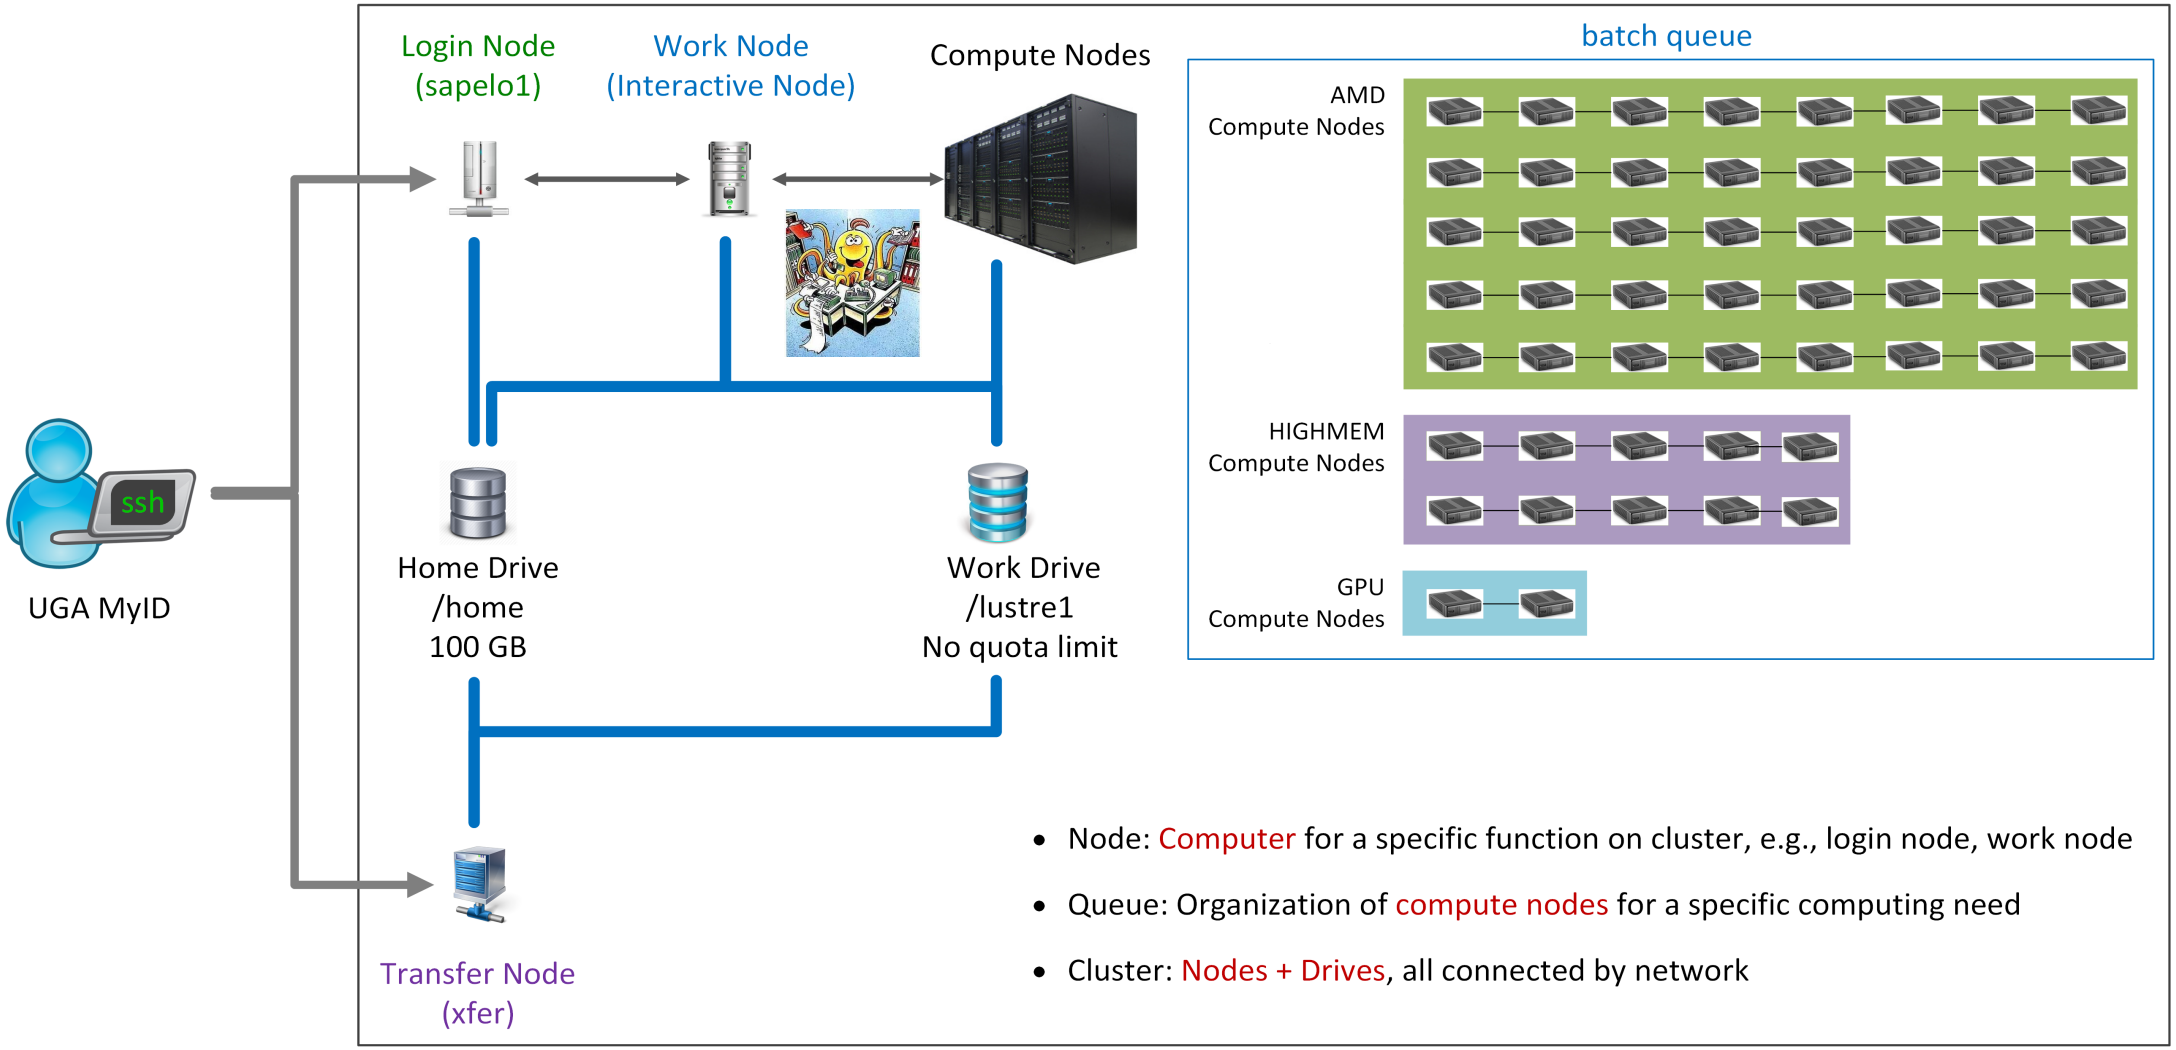

- Node: Computer for a specific function on cluster, e.g., login node, work node
- Queue: Organization of compute nodes for a specific computing need
- Cluster: Nodes + Drives, all connected by network

How to log on – Mac/Linux

1. Open Terminal utility
2. Type command line: `ssh UGAMyID@sapelo1.gacrc.uga.edu`
3. You will be prompted for your UGA MyID password (when you type in the password, the prompt blinks and does not move)

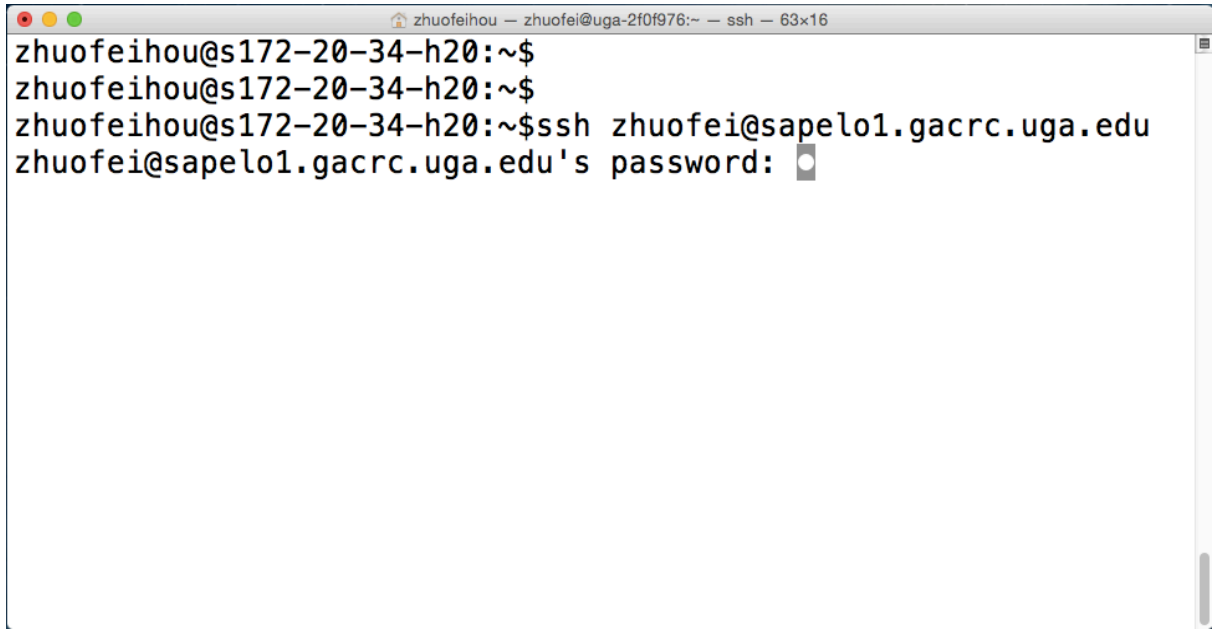

How to log on – Windows

Download to install:

http://eits.uga.edu/hardware_and_software/software/

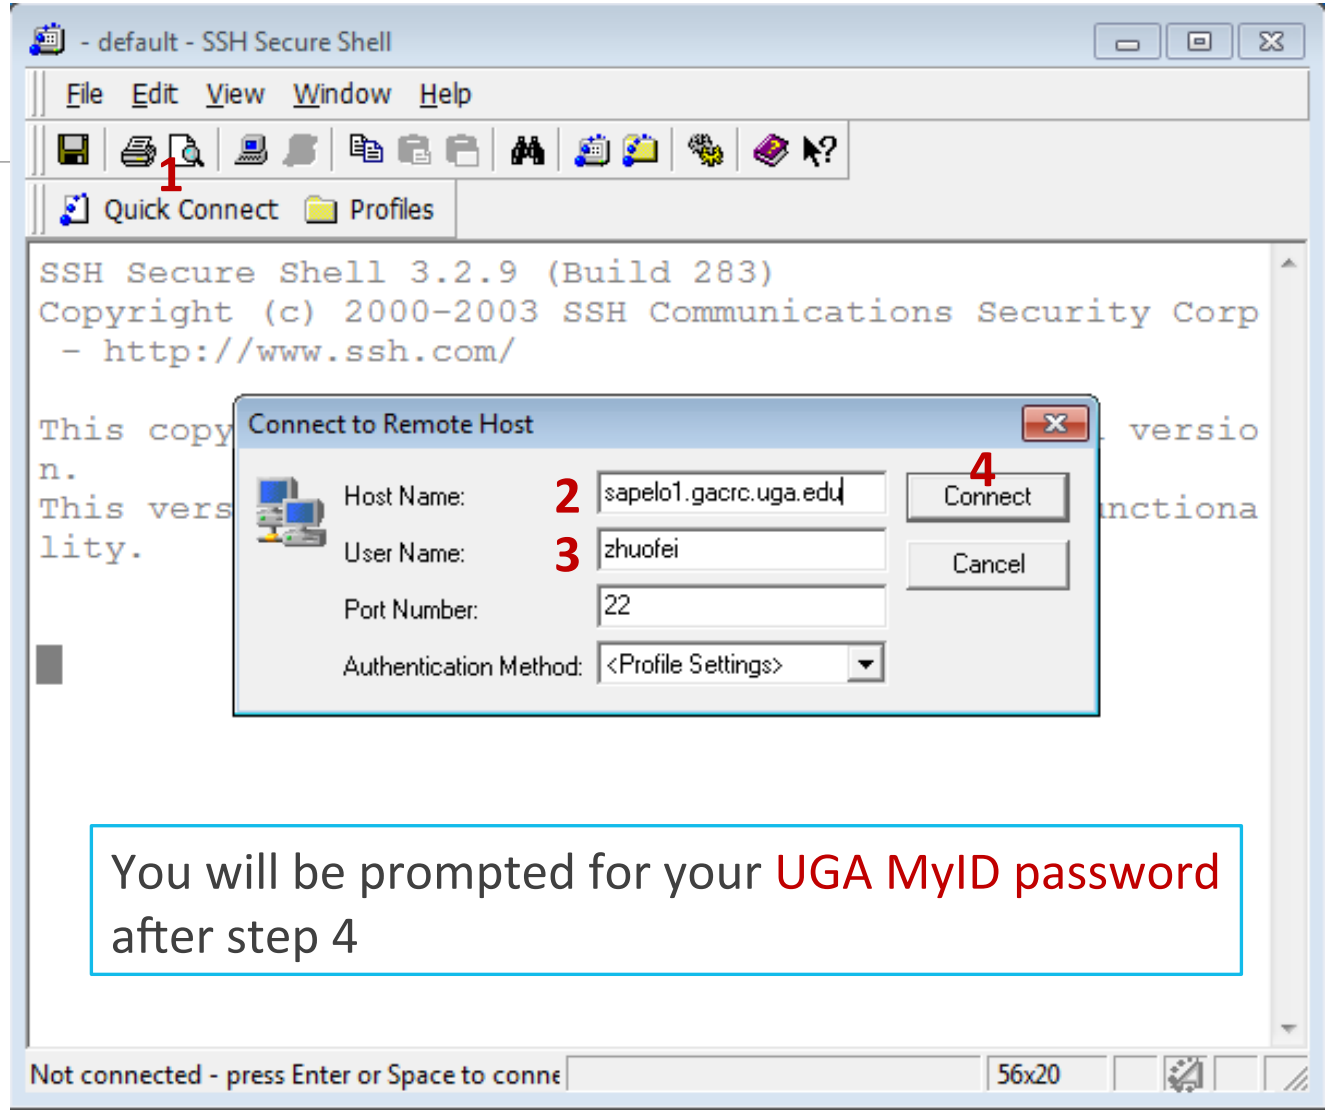

Once You Log on

1. Your working environment is Linux
2. You are on Login Node
3. You are in your home directory: /home/UGAMyID/, for example, /home/zhuofei/

How to Transfer Files – Mac/Linux

Local Computer ➔ Sapelo; Type command in Terminal from your *local computer*

```
scp file zhuofei@xfer.gacrc.uga.edu:/lustre1/zhuofei/
scp -r folder zhuofei@xfer.gacrc.uga.edu:/lustre1/zhuofei/
```

Sapelo ➔ Local Computer; Type command in Terminal from your *local computer*

```
scp zhuofei@xfer.gacrc.uga.edu:/lustre1/zhuofei/file ./
scp -r zhuofei@xfer.gacrc.uga.edu:/lustre1/zhuofei/folder ./
```

How to Transfer Files - Windows

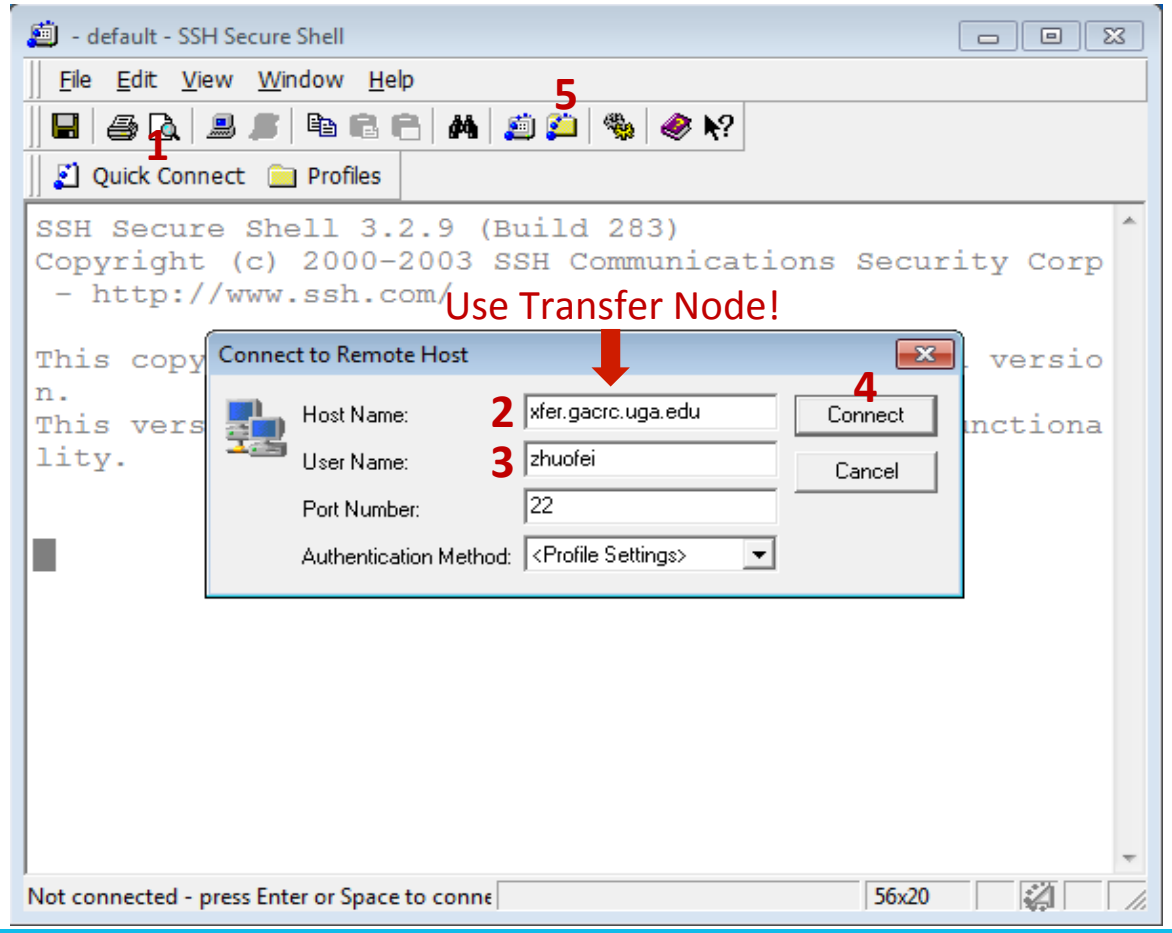

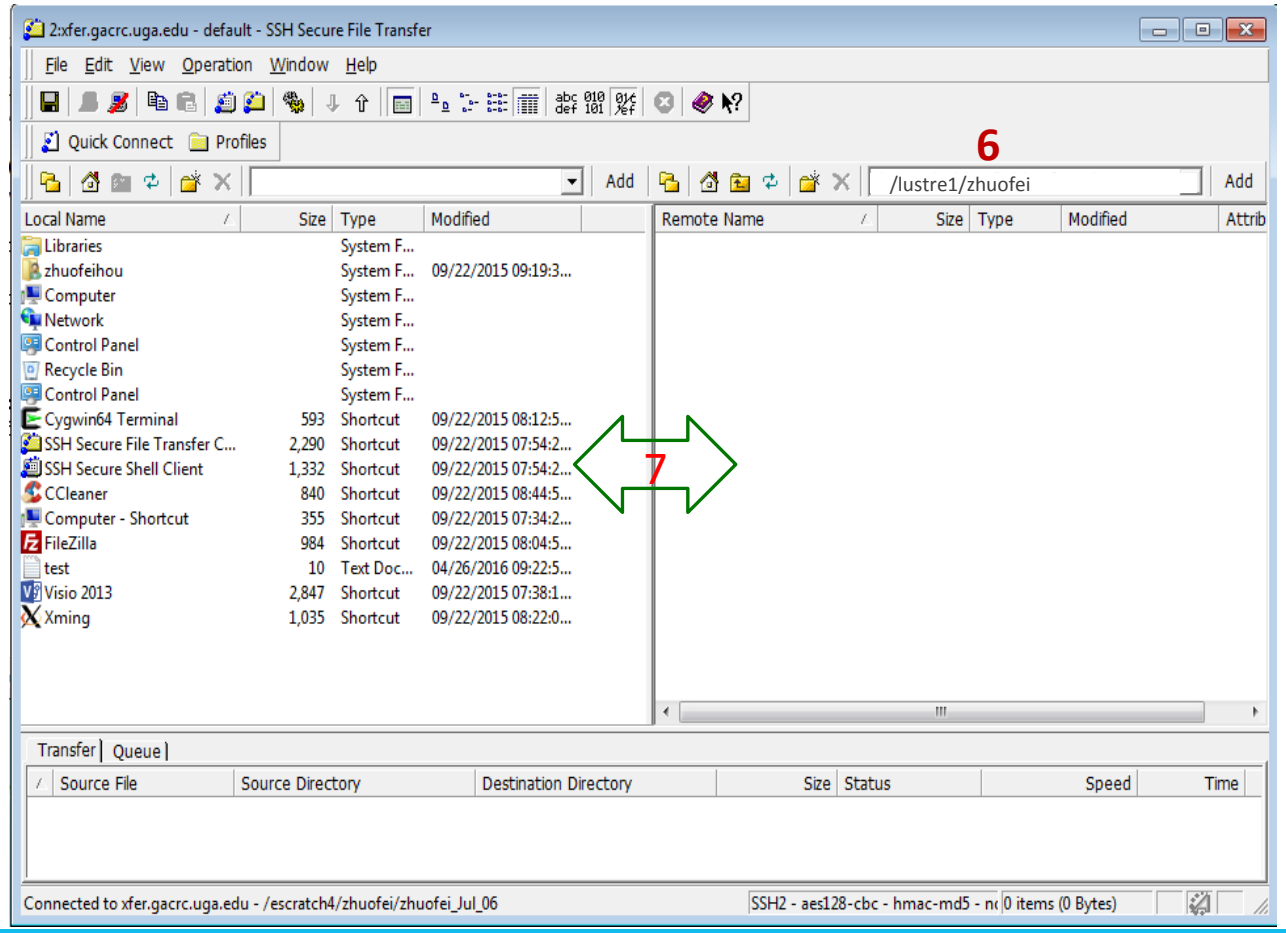

Next

1. Go to Work Node by typing command `qlogin`
2. Change to Work Drive by typing command `cd /lustre1/UGAMyID/`
3. Now you are in your working space and ready to
   - Check software information
   - Create job submission script
   - Submit job

How to Use Software

- On Sapelo ~300 installed software for you to run
- Three module commands to:
  - Have a look on all software : `module avail`
  - Choose a software to run : `module load softwareName`
  - Have a look on loaded software : `module list`

How to Submit a Job

- You need to create a *script* to make a *compute node* to know:
  - Your working directory
  - The software you want to run
  - Any Linux commands you want the node to run
- Wiki Links:

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

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

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

A Common Example – Running Samtools

- Samtools is a common bioinformatics utility https://wiki.gacrc.uga.edu/wiki/Samtools-Sapelo
- Script name: sub.sh
- Working directory: /lustre1/zhuofei/

```
#PBS -S /bin/bash                          # Linux default shell (bash)
#PBS -N testSamtools                       # define job name
#PBS -q batch                              # use Sapelo batch queue
#PBS -l nodes=1:ppn=1:AMD                  # request 1 CPU from 1 AMD node
#PBS -l mem=100gb                          # request 100GB RAM memory
#PBS -l walltime=48:00:00                  # request 48 hours running time

cd $PBS_O_WORKDIR                          # specify and enter (cd) a working directory
module load samtools/1.3.1                 # load samtools
time  samtools <command> [options]         # run samtools
```

A Common Example – Submit Job

- Submit sub.sh to batch queue: Type command on command line from your working directory:

```
$ qsub sub.sh
```

submit a job defined in sub.sh

Note: Once a job is submitted, your can logout (`exit` command) from the cluster and log back in later to check status of job and to retrieve the results.

Check Job Status

- Type command on command line: qstat

```
[zhuofei@n78 zhuofei]$ qstat
Job ID                    Name             User            Time Use S Queue
------------------------------------------------------------------------
481929.pbs                testJob1         zhuofei         900:58:0 C batch
481931.pbs                testJob2         zhuofei         04:00:03 R batch
481934.pbs                testJob3         zhuofei                0 Q batch
```

R : Job is running
C : Job is completed (canceled/crashed) and is no longer running; stays on list for 1 hour
Q : Job is pending, waiting for resources to become available

Cancel a Job

- Type command on command line: qdel *JobID*

```
[zhuofei@n78 zhuofei]$ qdel 481934
[zhuofei@n78 zhuofei]$ qstat
Job ID                    Name             User            Time Use S Queue
------------------------------------------------------------------------
481929.pbs                testJob1         zhuofei         900:58:0 C batch
481931.pbs                testJob2         zhuofei         04:00:03 R batch
481934.pbs                testJob3         zhuofei                0 C batch
```

Job is canceled and status is changed to C

Check Job Memory Usage

- Type command on command line: qstat -f *JobID*

```
[zhuofei@75-104 MPIs]$ qstat -f 699847
Job Id: 699847.pbs.scm
    Job_Name = testJob
    Job_Owner = zhuofei@uga-2f0f976.scm
    .
    resources_used.vmem = 6548528kb
    resources_used.walltime = 07:01:36
    job_state = C
    .
    Resource_List.mem = 5gb
    Resource_List.walltime = 10:00:00
```

Sapelo Workflow

1. Linux/Mac user:

```
ssh username@sapelo1.gacrc.uga.edu
```

Windows user:

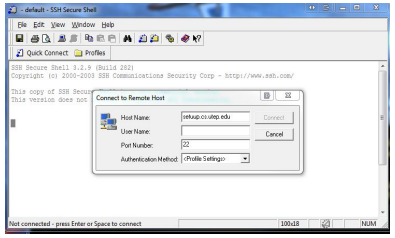

Login

2. `qlogin`

Interactive

3. `cd /lustre1/username`

4. `mkdir ./workDir`

5. `cd ./workDir`

7. `nano ./sub.sh`

```
#PBS -S /bin/bash
#PBS -q batch
#PBS -N testBlast
#PBS -l nodes=1:ppn=1:AMD
#PBS -l mem=20gb
#PBS -l walltime=48:00:00
cd $PBS_O_WORKDIR
module load ncbiblast+/2.2.29
blastn [options] > outputfile
```

8.

```
$ qsub sub.sh
```

lustre1

6. Linux/Mac user:

```
scp file username@xfer.gacrc.uga.edu:/lustre1/username/workDir
```

Windows user:

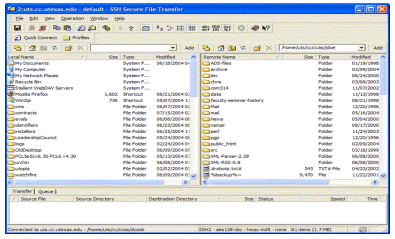

Note: `-r` option of scp command will recursively copy a directory

Sapelo Workflow

1. Log on to Sapelo Login node: `ssh username@sapelo1.gacrc.uga.edu`
2. From Sapelo Login node, transfer to Interactive node by issuing the command: `qlogin`
3. From Interactive node: Change directory to global scratch: `cd /lustre1/username`
4. Create a working subdirectory on global scratch: `mkdir ./workDir`
5. Change directory to `workDir`: `cd ./workDir`
6. Transfer data to `workDir` using `scp` or SSH Client File Transfer (with `tar` or `gzip`)
7. Make a Sapelo job submission script: `nano ./sub.sh`
8. Submit job: `qsub ./sub.sh`

Useful PBS headers: `#PBS -l mem=200gb` : use 256GB high-RAM AMD nodes
`#PBS -l nodes=1:ppn=24:AMD` : request 24 cores for 24 threads, max 48!

How to Find Support

GACRC Support: http://help.gacrc.uga.edu/

- Job Troubleshooting: Tell us details of your problem, please also include:
  - ✓ UserID and JobID
  - ✓ Working directory and script name
  - ✓ Any command you used
  - ✓ Brief of your problem

Note:
It's USER's responsibility to make sure the correctness of data being used by jobs!

Important Rules

- No Sharing User Account! ➔ Use your own user account
- No Job Running Directly on Login node! ➔ Submit job to batch queue
- No Data Transferring into/out of Cluster Using Login node! ➔ Use Transfer node

Useful Links

- GACRC Wiki: https://wiki.gacrc.uga.edu/wiki/Main_Page
- GACRC Software: https://wiki.gacrc.uga.edu/wiki/Software
- GACRC Support: http://gacrc.uga.edu/help/
- GACRC Training: https://wiki.gacrc.uga.edu/wiki/Training

Thank You!

Georgia Advanced Computing Resource Center
4098C Stegeman Coliseum
University of Georgia

Telephone Support
EITS HELPDESK: 706-542-3106
MONDAY – THURSDAY: 8AM – 10PM
FRIDAY: 8AM – 6PM
SATURDAY – SUNDAY: 1PM – 7PM