High-Performance Computing Using GACRC Resources

Georgia Advanced Computing Resource Center (GACRC)

Enterprise Information Technology Services (EITS)

The University of Georgia

Outline

- What is GACRC?
- What is High Performance Computing (HPC)?
- What is a Cluster and why use a Cluster?
- How do I get started?
- Overview of the Sapelo2 cluster
- How do I learn more?

What is GACRC?

- A high-performance computing (HPC) center at UGA
- A partnership between the Offices of the VP for Information Technology and for Research
- Provide to the UGA research and education community an advanced computing environment:
  - Computing hardware, storage, and networking infrastructure
  - Large collection of math libraries, scientific, and engineering applications
  - Consulting and training services
- Sapelo2 (main cluster, for research use)
- Teaching cluster (for instructional use, student accounts)
- Buy-in program (faculty contribute nodes to Sapelo2)

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

What is High Performance Computing?

1 worker

10 years later……

What is High Performance Computing? (cont.)

2 months later……

What is High Performance Computing? (cont.)

Serial Job with 1 worker

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

Parallel Job with 6 workers

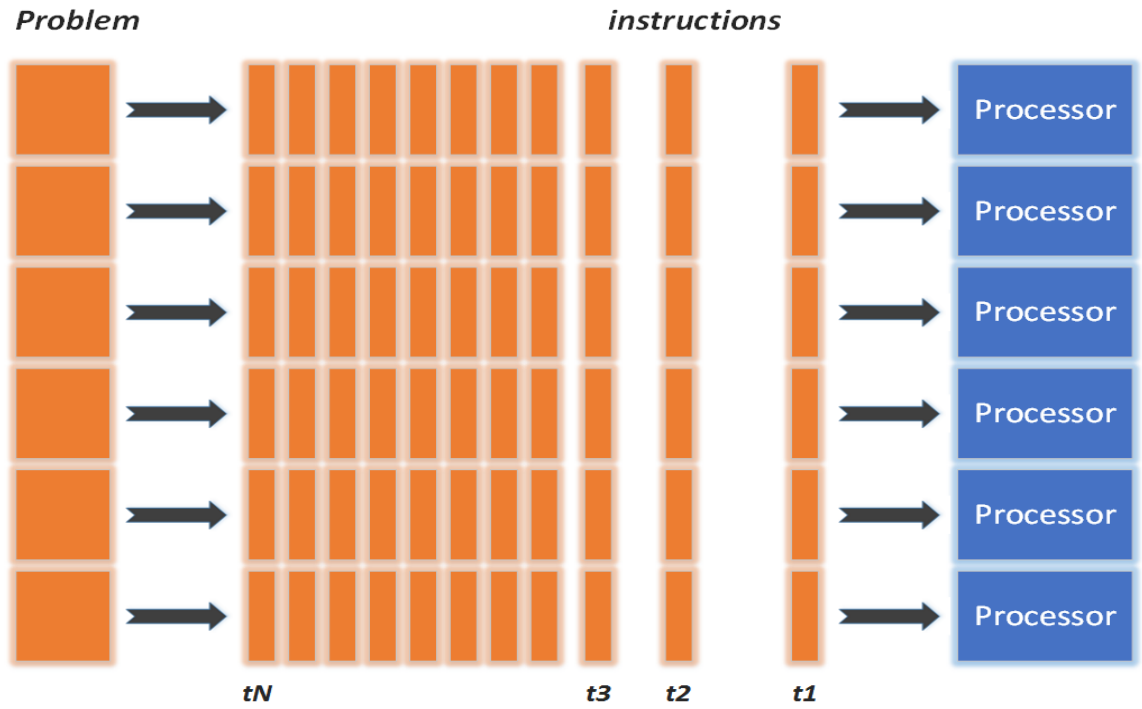

- ✓ 6 instructions executed *simultaneously* on *6* processors (or cores)
- ✓ Shared-memory threaded parallel job or MPI distributed parallel job (Message Passing Interface)

What is a cluster?

Rows with racks of compute nodes

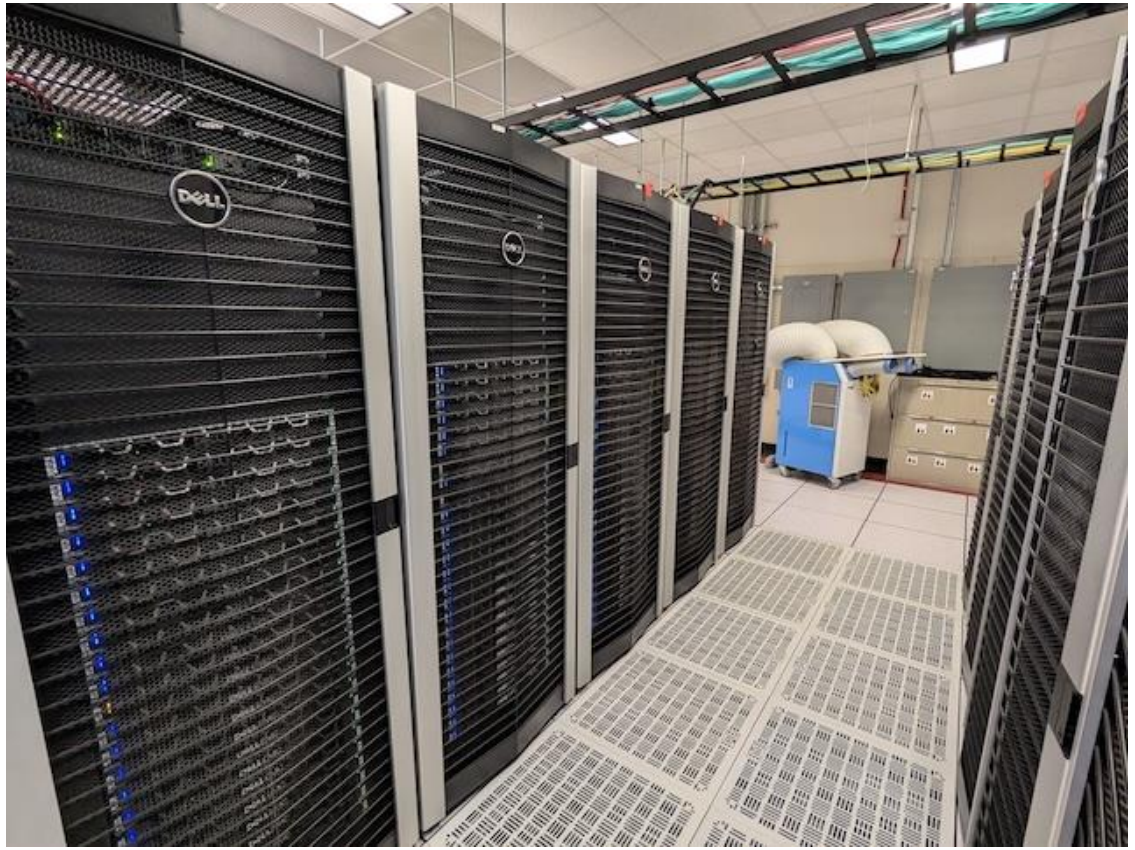

Network cables connecting compute nodes

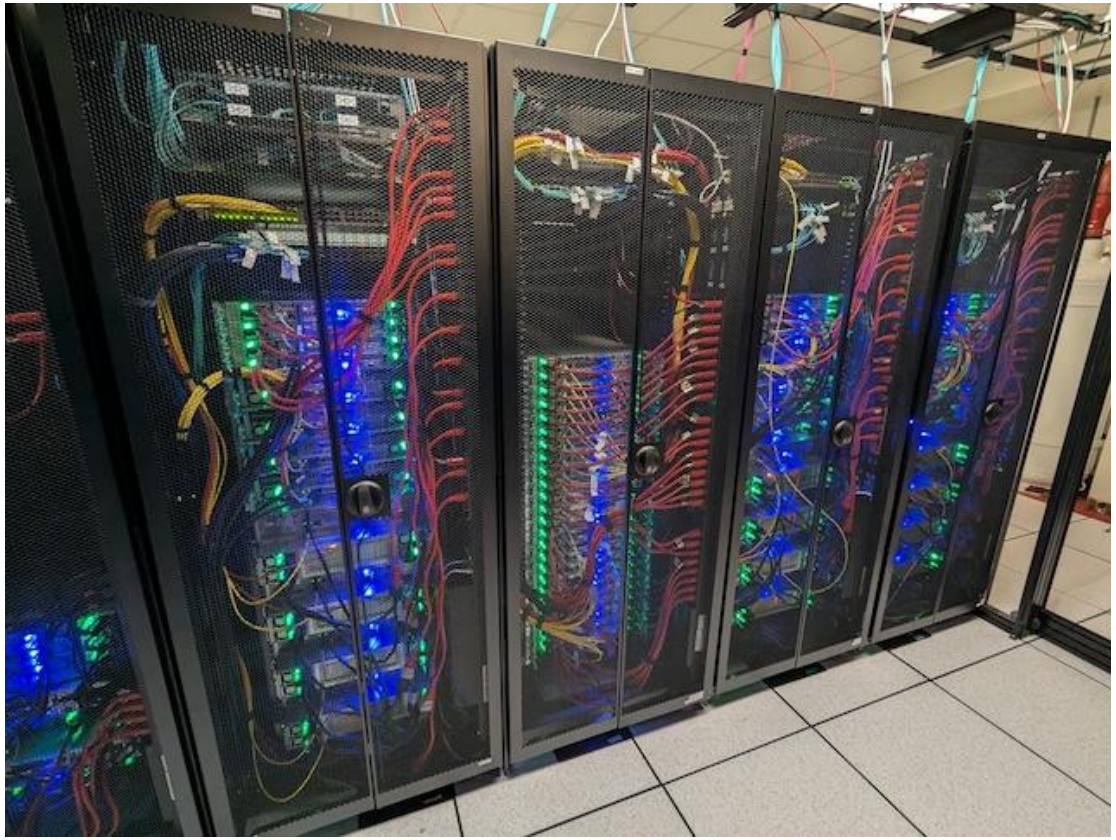

Back of the racks

How do I access a cluster?

vs

Access GACRC's cluster using a Secure Shell (SSH) program from anywhere that has network connectivity to the UGA VPN.

Sapelo2 Cluster Overview

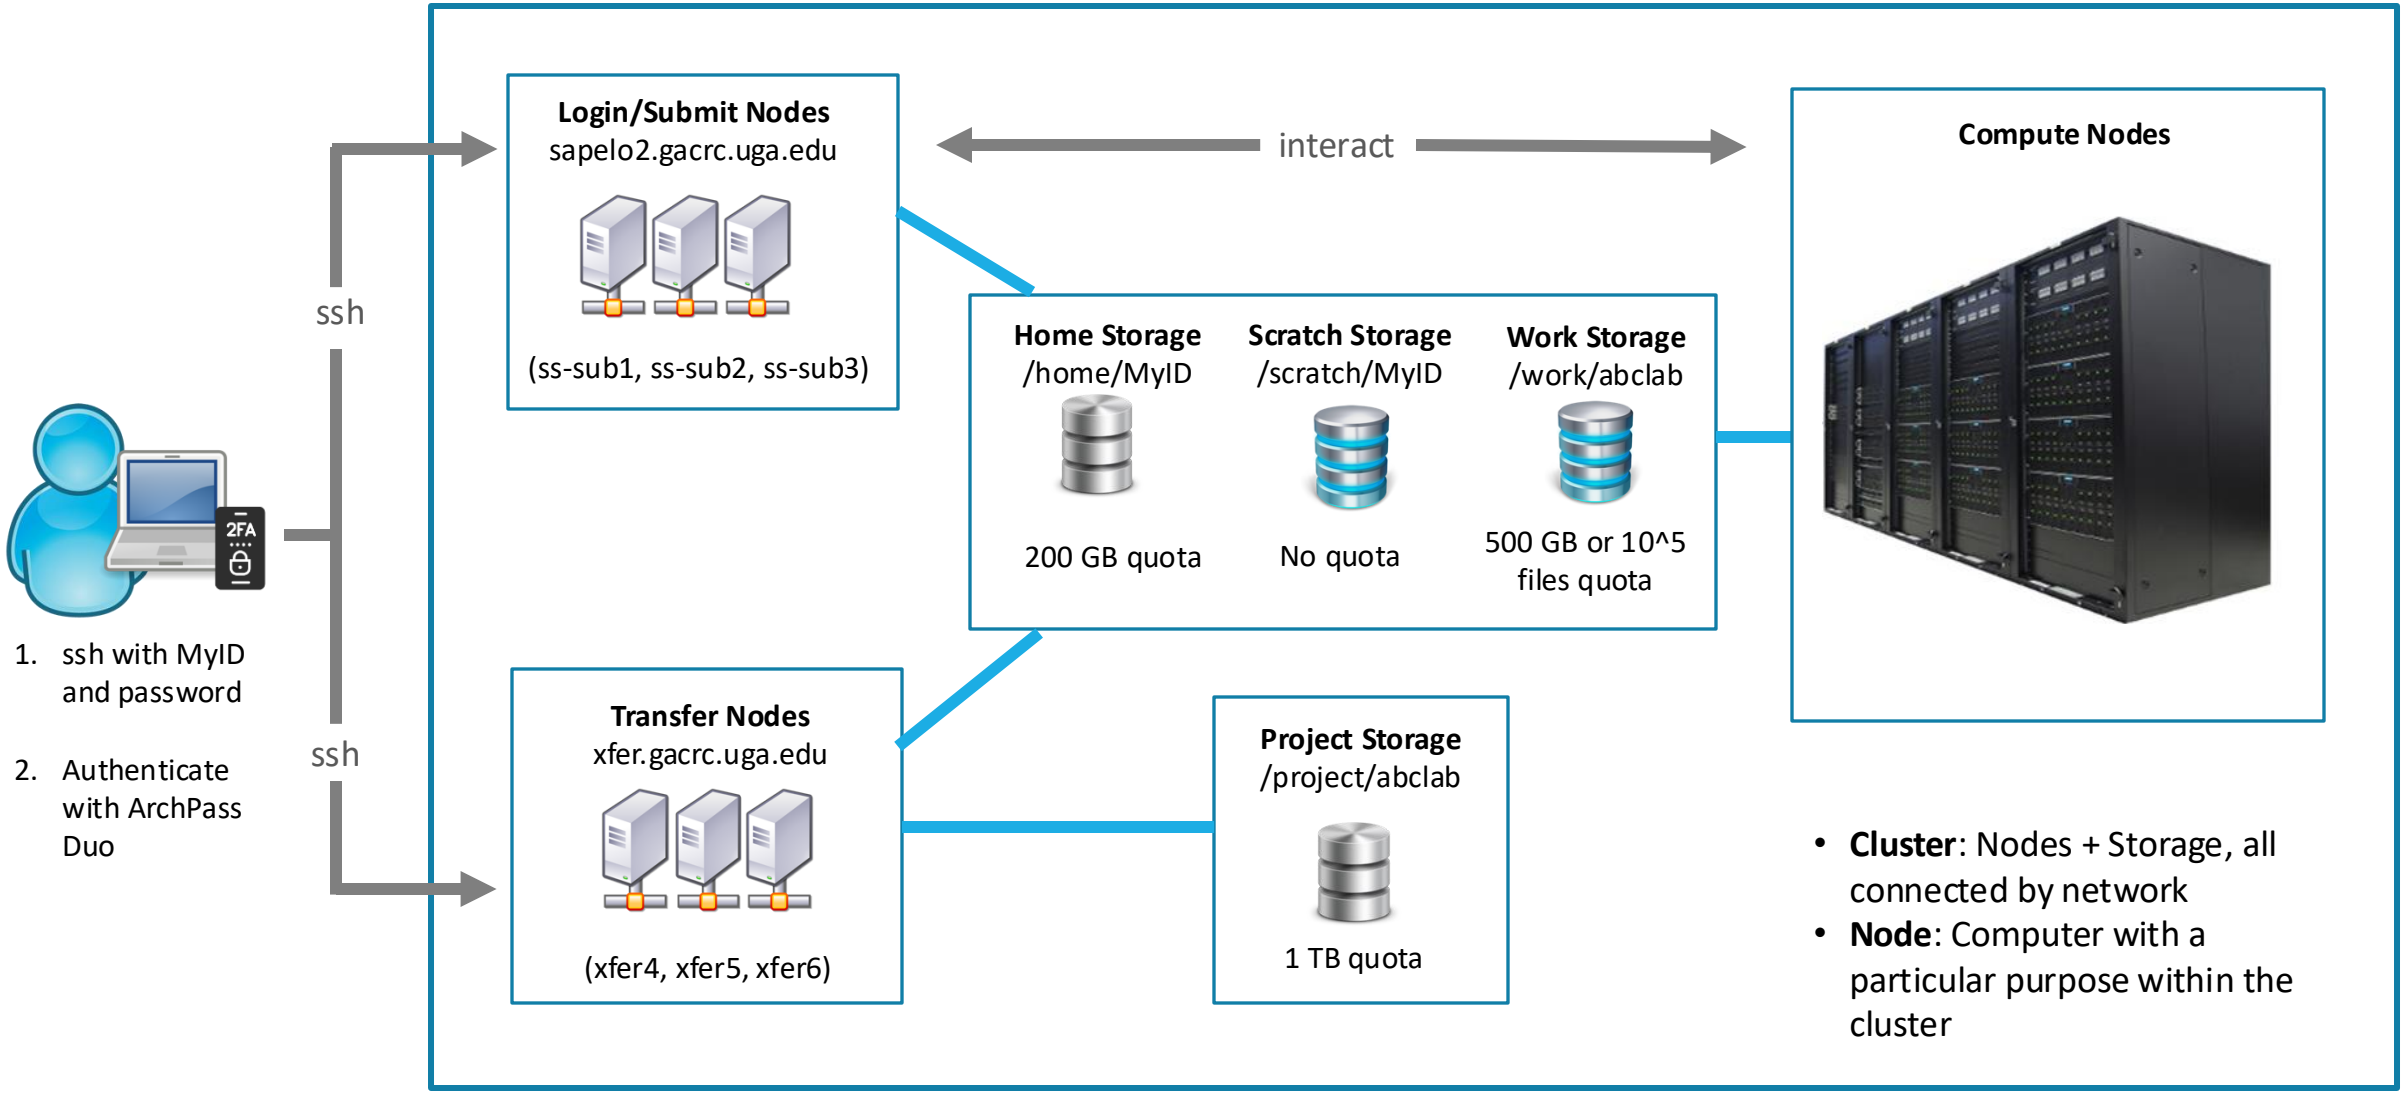

Note: You need to connect to the UGA network using VPN when accessing from outside of the **UGA main campus**.
UGA VPN: https://eits.uga.edu/access_and_security/infosec/tools/vpn/

What is a Cluster? (cont.)

When you work on a cluster, 3 roles are working:

- You: Who submits a job
- Queueing System: Who dispatches your job to run on cluster
- Cluster: Who runs your job

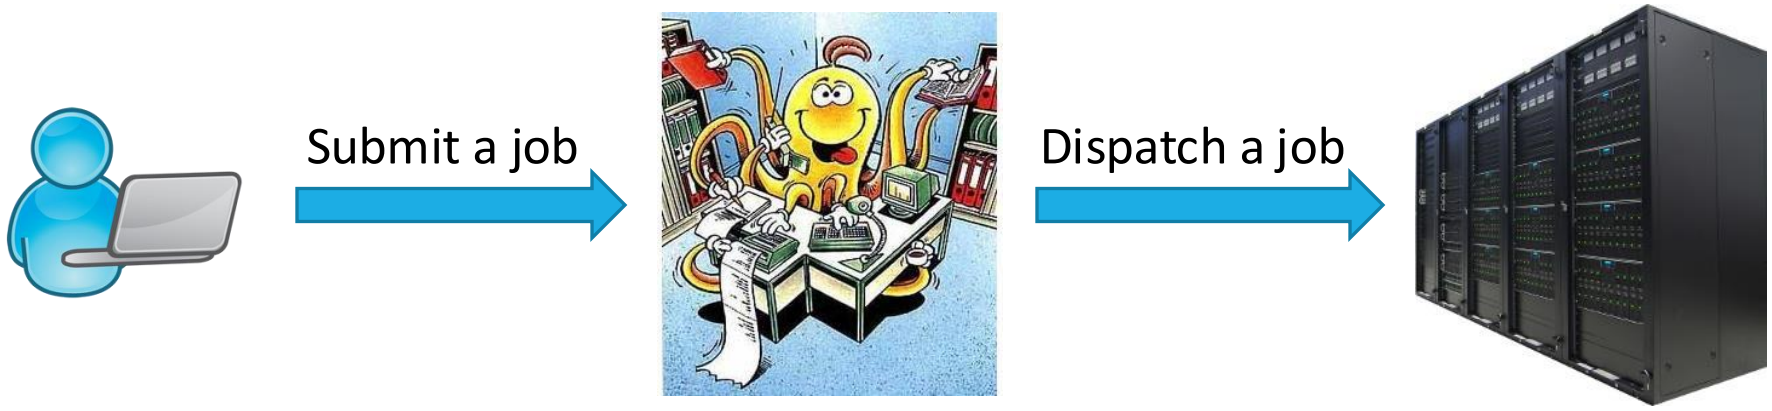

Why use a Cluster?

Research problems can outgrow the capabilities of the desktop or laptop computer.

Examples:

➢ A material science student needs to repeat a stochastic simulation 100 times to determine statistical fluctuations. Each simulation takes 10h using a laptop (total time: over 40 days).

Large number of independent tasks can be done concurrently on the cluster.

➢ A genomic researcher needs to analyze a large sequencing dataset. A laptop does not have enough memory to handle the analysis.

Use large memory nodes on the cluster.

➢ A marine scientist runs a 2D coastal modeling program on a desktop, but going to 3D would require too much memory and computing time.

Run the program in parallel on multiple compute nodes on the cluster.

How do I get started?

➢Ask your PI to request an account for you, using the form https://uga.teamdynamix.com/TDClient/2060/Portal/Requests/ServiceDet?ID=25839

➢Register for training and watch the training videos.

➢Attend Linux training (if needed) and Sapelo2 new user training.

➢Account created during the Sapelo2 training

- User accounts provided using UGA MyID.
- Connection to the cluster via SSH on Linux/Mac OS X, PuTTY on Windows.
- Authenticate with UGA MyID password and Archpass Duo.
- From off campus, connect to the UGA VPN first

Sapelo2 – Computing Hardware

➢**Resources available to all users**

- Regular memory nodes: 32 to 128 cores/node, 2GB to 6GB/core
- High memory nodes: 512GB, 1TB, 2TB, or 3TB/node
- GPU nodes: NVIDIA H100, A100, and L4 devices
- Number of CPU cores: ~ 33,400

➢**Buy-in nodes**

- Various configurations, including GPU nodes.
- Number of CPU cores: ~ 10,300

➢Total number of CPU cores: ~ 43,700

Sapelo2 – Storage

➢Storage to support projects that use the computing cluster only.

- ZFS storage: home directory /home (200GB/user quota, no purge, backup)
- Lustre storage device:

  /scratch (no quota, 30-day purge, no backup)

  /work (500GB and 500k file/group quota, no purge, no backup)

- Local hard drive (SSD, NVMe) on compute nodes: /lscratch (3.5TB of GPU nodes, up to 890GB on other nodes)
- ZFS JBOD: /project (1TB/group, no purge, backup)

Sapelo2 – Software

- 64-bit Linux OS (Rocky 8.8).
- Slurm queueing system to submit jobs to the compute nodes (e.g. sbatch, squeue, sacct).
- Over 1400 environment modules, conda environments, Singularity containers.
- Cannot run applications that need Windows OS.
- Cannot run Docker containers, convert to Singularity (Apptainer) containers.
- Very limited number of commercial software (e.g. MATLAB, Amber, Gaussian), some limited to a few groups that purchased licenses (e.g. SAS, Stata, ANSYS, Schrodinger).

Software Environment

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

➢Software loaded via environment modules (Lmod) – Over 1400 software modules installed!

- Compilers (C/C++, Fortran): GNU/GCC, Intel
- MPI Libraries: OpenMPI, Intel MPI, MVAPICH
- Math Libraries: Intel MKL, GSL, OpenBLAS, LAPACK, FFTW, etc.
- GPU Tools: NVIDIA CUDA Toolkit, cuDNN, NVHPC
- High-Level Languages: Python, Perl, R, Julia, Java
- Plotting: gnuplot, xmgrace, matplotlib, Seaborn, ParaView, etc.
- Other Programming/Visualization Tools: MATLAB, Mathematica
- Machine Learning and Deep Learning Tools: scikit-learn, Tensorflow, Keras, PyTorch, etc.
- **Applications for Engineering, Bioinformatics, Chemistry, Physics, etc.**

Transferring Files

➢UGA has a subscription for Globus, a high-performance data-transfer platform that allows you to perform and/or automate:

- Transfer data from your local machine to GACRC resources.
- Transfer data between file systems on Sapelo2.
- Data transfers between servers in your group or a server and your laptop.
- Sharing data with researchers at UGA and at other institutions.
- Sharing data with the world.

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

https://kaltura.uga.edu/media/t/1_vlprwoc7/176125031

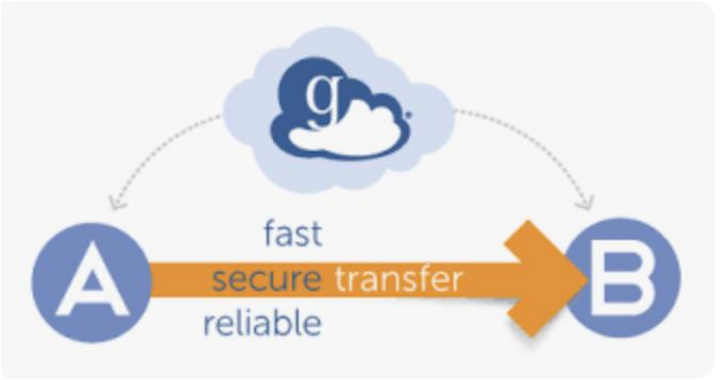

➢Other file transfer programs are available: scp, FileZilla, WinSCP

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

Open OnDemand

- Alternative way to access Sapelo2, connecting from a browser.
- Provides a graphical desktop on the cluster, easy way to run Graphical applications
- Interactive applications (Jupyter, MATLAB, Mathematica, Rstudio, ParaView, Vscode, etc)
- Sessions are not lost when network connection is interrupted or local computer is shutdown. Can start a session from a browser on one computer, and connect to it from a browser on another computer.

https://wiki.gacrc.uga.edu/wiki/OnDemand
https://kaltura.uga.edu/media/t/1_u9d1xrpp/176125031

Consulting and Support

- Install and update software per user request.
- Troubleshoot programs, jobs, and workflows.
- Help users optimize their code.
- Help users implement or automate pipelines.
- Provide user training (Linux, Sapelo2 new user, Python, R, conda environment, etc.).
- Assistance with grant proposal preparation, where GACRC resources are used.

https://gacrc.uga.edu/about/how_we_help_researchers.php

Documentation and Contact

Wiki: http://wiki.gacrc.uga.edu
Web Site: http://gacrc.uga.edu
Help: https://help.gacrc.uga.edu
Videos: https://kaltura.uga.edu/channel/GACRC/176125031

Offices: GACRC staff are located in the Computing Services Building (formerly called Statistics Building), rooms 101 to 108.

Office Hours: https://wiki.gacrc.uga.edu/wiki/Office_Hours

Thank you!