

## UF Research Computing: An Introduction

Matt Gitzendanner  
[magitiz@ufl.edu](mailto:magitiz@ufl.edu)

5/14/13

UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

## UF Research Computing



- ▶ Mission
  - Improve opportunities for research and scholarship
  - Improve competitiveness in securing external funding
  - Provide high-performance computing resources **and support** to UF researchers

UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

## UF Research Computing

- ▶ Funding
  - Faculty
  - Matching grant program!
- ▶ Any UF Faculty can use
  - Up to 8 cores
  - Investors gain priority and access to additional resources
- ▶ Comprehensive management
  - Hardware maintenance and 24x7 monitoring
  - Relieve researchers of the majority of systems administration tasks



UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

## Matching Program

Consolidating Resources to Improve Efficiency and Capacity




The Research Computing Matching Program pooled \$642k, thereby creating synergies and improving research infrastructure.

UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

UNIVERSITY OF FLORIDA | High-Performance Computing



**HiPerGator**  
The University of Florida Supercomputer for Research

UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

## UF Research Computing

- ▶ Shared Hardware Resources
  - Over **22K cores**
  - High-speed, low-latency **InfiniBand** interconnects
  - **>3 PB**, high performance Lustre and Nexenta storage
  - **GPGPUs**— 90+, new Kepler-class
  - Several large memory (**512GB to 1TB of RAM**) nodes




UF Information Technology

[www.it.ufl.edu](http://www.it.ufl.edu)

## Investing in the future


- Internet2 Innovation Platform
  - 100 Gbps connectivity
  - Campus Research Network now 200 Gbps
- UF Data Center on Eastside Campus
  - 10,000 sq.ft and 1.75 MW total
  - 5,000 sq. ft. space for Research Computing
- HiPerGator**
  - 16,384 cores
  - Infiniband interconnect
  - 2.1PB fast, high-availability, storage
  - Current ~7,000 core cluster will be integrated—Aug 2013



UF Information Technology

www.it.ufl.edu

## UF Research Computing



Where do you start?

UF Information Technology

www.it.ufl.edu

## UF Research Computing

- User Accounts
  - Qualifications:
    - Current UF faculty sponsor
- Account Policies
  - Personal activities are strictly prohibited on HPC Center systems
  - Class accounts deleted at end of semester
  - Data are not backed up!**
  - Home directories must not be used for I/O
    - Use /scratch/hpc/\$USER
  - Storage systems may not be used to archive data from other systems
  - Passwords expire every 6 months

UF Information Technology

www.it.ufl.edu

## UF Research Computing

- User Accounts
  - Qualifications:
    - Current UF faculty, UF graduate student, and researchers
  - Request at: <http://www.hpc.ufl.edu/support/>
  - Requirements:
    - GatorLink Authentication
    - Faculty sponsorship for graduate students and researchers



UF Information Technology

www.it.ufl.edu

## What can you run?

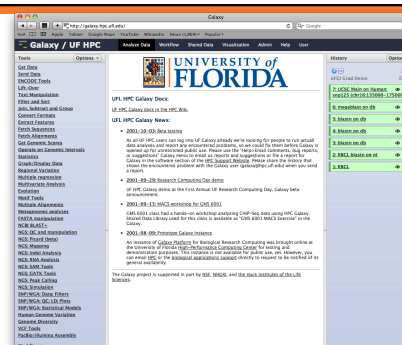


- Linux-based
- Generally command line driven applications
- Galaxy
- Graphical apps can be setup
  - SAS
  - BEAUi

UF Information Technology

www.it.ufl.edu

## Galaxy



UF Information Technology

www.it.ufl.edu



## Cluster login

submit.hpc.ufl.edu

ssh

submit1  
submit2

/home/\$USER

ssh <user>@submit.hpc.ufl.edu

Windows: PuTTY  
Mac/Linux: Terminal

User interaction

Login node (Head node)

UF Information Technology www.it.ufl.edu

## Cluster login

submit.hpc.ufl.edu

ssh

sub: Last login: Mon Jun 11 21:49:41 on ttys000  
Voyager-11: matt\$ ssh magitz@submit.hpc.ufl.edu  
magitz@submit.hpc.ufl.edu's password:  
Last login: Tue Jun 12 16:01:13 2012 from submit.hpc.ufl.edu  
Welcome to the UF HPC Center.

Do not run interactive jobs on the login nodes. If you need to run an interactive job, there are interactive/test nodes for that.

UF HPC Center Account Policies can be found here:  
<http://www.hpc.ufl.edu/users/accounts.php>

Window Mac/Linux

magitz@submit1 ~]\$ pwd  
/home/magitz  
magitz@submit1 ~]\$

User interaction

UF Information Technology www.it.ufl.edu

## Logging in

PuTTY Configuration

Category: Session, Logging, Terminal, Keyboard, Bell, Features, Window, Appearance, Behaviour, Translation, Selection, Colours, Connection

Basic options for your PuTTY session

Specify the destination you want to connect to

Host Name (or IP address): submit.hpc.ufl.edu Port: 22

Connection type: ☐ Raw ☐ Telnet ☒ Rlogin ☒ SSH ☐ Serial

Load, save or delete a stored session

Saved Sessions: Default Settings, Fisher

Close window on exit: ☐ Always ☐ Never ☒ Only on clean exit

UF Information Technology www.it.ufl.edu

## Linux Command Line

>\_

All the Best Linux Cheat Sheets

1. Linux Command Line

- Linux Reference Card - Great reference published on FOSSwire website
- Colin's Linux Manual - Great one page reference to the most popular Linux commands
- Linux Cheat Sheet - An incredibly exhaustive reference for all things Linux
- Forward's Linux Cheat Sheet - A great reference with Don's companions
- Terminal Shortcuts - Cheat sheet for the most common terminal shortcuts
- More Terminal Shortcuts - More shortcuts for history and X

- Lots of online resources
  - Google: Linux cheat sheet
- Training sessions
  - May 21: The Linux/Unix Command Line - An Introduction
- User manuals for applications

UF Information Technology www.it.ufl.edu

## Storage at HPC

submit.hpc.ufl.edu

ssh

submit1  
submit2

/home/\$USER

/scratch/hpc/\$USER

\$ cd /scratch/hpc/\$USER/

Copy your data to submit using **scp** or a SFTP program like Cyberduck or FileZilla

UF Information Technology www.it.ufl.edu


## UF Research Computing

- Storage
  - Home:** /home/\$USER
    - For code compilation and user file management only
    - Do not use for job input/output!**
    - Include cd \$PBS\_o\_WORKDIR or similar in scripts
  - Scratch space:** Lustre File System
    - /scratch/hpc/\$USER
    - 500GB per lab
- Other storage options available for purchase

UF Information Technology www.it.ufl.edu

## Scheduling a job

- Need to tell scheduler what you want to do
  - How many CPUs** you want and how you want them grouped
  - How much RAM** your job will use
  - How long** your job will run
  - The commands that will be run



UF Information Technology

www.it.ufl.edu

## UF Research Computing

- Ordinary Shell Script

```
#!/bin/bash
date
module load test_app
test_app -i file.txt
```

Read the manual  
for your application

Commands typed  
on the command  
line can be put in a  
script

UF Information Technology


www.it.ufl.edu

## UF Research Computing

- Submission Script

```
#!/bin/bash
#
#PBS -N My_Job_Name
#PBS -M Joe_Shmoe@ufl.edu
#PBS -m abe
#PBS -o My_Job_Name.log
#PBS -j oe
#PBS -l nodes=1:ppn=1
#PBS -l walltime=00:05:00
#PBS -l pmem=900mb

cd $PBS_O_WORKDIR
date
module load test_app
test_app -i file.txt
```




UF Information Technology

www.it.ufl.edu

## Nodes and processors

```
#PBS -l nodes=1:ppn=4
#PBS -l nodes=2:ppn=8
```



UF Information Technology

www.it.ufl.edu

## RAM

```
#PBS -l pmem=900mb
```

- Per-processor** RAM request
- Lots to consider, but do your best at estimating RAM needed for job
- Over about 3GB of RAM, "costs" toward CPU allocation

Wasted RAM leads  
to idle CPUs and  
low job throughput



UF Information Technology

www.it.ufl.edu

## Walltime

```
#PBS -l walltime=00:50:00
```

- Fairly straight forward
- As with all resource requests, accuracy helps ensure **your** jobs and all other jobs will run sooner



UF Information Technology

www.it.ufl.edu

## UF Research Computing

- ▶ Job Management
  - `qsub <file_name>`: job submission
  - `qstat -u <user>`: check queue status
  - `qdel <JOB_ID>`: job deletion
  - `qdelmine`: delete ALL of your current jobs

UF Information Technology

www.it.ufl.edu

## UF Research Computing

- ▶ Job Scheduling and Usage
  - Job scheduler selects jobs based on priority
    - Priority is determined by several components
    - Investors have higher priority
    - Non-investor jobs limited to 8 processor equivalents (PEs)
    - RAM: requests beyond a few GB/core start counting toward the total PE value of a job
  - Test nodes (test01-06) available for interactive use, testing and short jobs
    - Connect from submit node:
 

```
[magitz@submit1 ~]$ssh test01
```

UF Information Technology

www.it.ufl.edu

## Training Schedule

- ✓ Jan 14: Intro to UFHPC, getting started
- ▶ Jan 28: The Linux/Unix Shell - An Introduction
- ▶ Feb 4: Running Jobs, Submission Scripts, Modules
- ▶ Feb 11: Dr. Dhruva Chakravorty: Amber
- ▶ Feb 18: Galaxy Overview, The Basics
- ▶ Feb 25: Dr. David Ostrov: Molecular Docking
- ▶ Mar 11: NGS Data Techniques: General Methods and Tools
- ▶ Mar 18: NGS: Reference Based Mapping & de Novo Assembly
- ▶ Mar 25: Phylogenetic Analyses
- ▶ Apr 1: Multiprocessing at the HPC Center
- ▶ Apr 8: Introduction to GPU nodes
- ▶ Apr 15:
- ▶ Apr 22:

UF Information Technology

www.it.ufl.edu

## UF Research Computing

- ▶ Help and Support
  - Help Request Tickets
    - <https://support.hpc.ufl.edu>
    - For any kind of question or help requests
    - Searchable database of solutions
  - We are here to help!
    - [support@hpc.ufl.edu](mailto:support@hpc.ufl.edu)



UF Information Technology

www.it.ufl.edu

## UF Research Computing

- ▶ Help and Support (Continued)
  - <http://wiki.hpc.ufl.edu>
    - Documents on hardware and software resources
    - Various user guides
    - Many sample submission scripts
  - <http://hpc.ufl.edu/support>
    - Frequently Asked Questions
    - Account set up and maintenance



UF Information Technology

www.it.ufl.edu