

UF Research Computing: An Introduction

Matt Gitzendanner: magitz@ufl.edu

Alex Moskalenko: om@hpc.ufl.edu

8/28/12

UF Information Technology

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

UF Research Computing

- ▶ Funding
 - Faculty participation (i.e. grant money) provides funds for hardware purchases
 - Matching grant program!
- ▶ Any UF Faculty can use
 - Up to 8 cores at a time
 - 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

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

UF Research Computing




UF Information Technology

www.it.ufl.edu

UF Research Computing

- ▶ Shared Hardware Resources
 - Over **6K cores** AMD and Intel
 - High-speed, low-latency **InfiniBand** interconnects
 - **>1 PB**, high performance Lustre and Nexenta storage
 - **GPUs**—CUDA code, etc.
 - Several large memory (**512GB**) nodes



UF Information Technology

www.it.ufl.edu

UF Research Computing



Where do you start?

UF Information Technology
www.it.ufl.edu

UF Research Computing

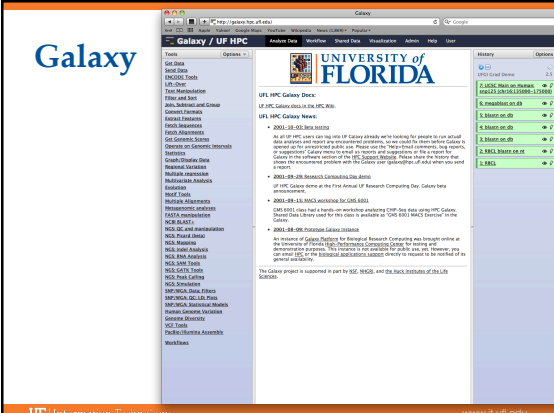
- ▶ Account Policies
 - Personal activities are strictly prohibited on HPC Center systems
 - Class accounts deleted at end of semester
 - **Data are not backed up!**
 - Home directories must not be used for I/O
 - Use /scratch/hpc/\$USER
 - Storage systems may not be used to archive data from other systems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

UF Information Technology
www.it.ufl.edu

What can you run?

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

UF Information Technology
www.it.ufl.edu

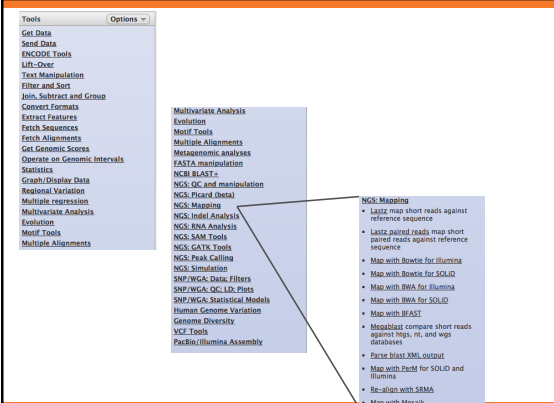


UF Information Technology
www.it.ufl.edu

Galaxy: Data intensive biology for everyone

- ▶ Accessible, reproducible, transparent computational biology
- ▶ galaxy.hpc.ufl.edu
 - Local instance of Galaxy
 - Faster access to storage, easier upload
 - Local compute resources
 - Local control

UF Information Technology
www.it.ufl.edu



NGS Mapping

- Lazzy map short reads against reference sequence
- Lazzy paired reads map short paired reads against reference sequence
- Map with Bowtie for Illumina
- Map with Bowtie for SOLiD
- Map with BWA for Illumina
- Map with BWA for SOLiD
- Map with BEAST
- Mapblast compare short reads against hgps, mt, and wgs databases
- Pave blast XML output
- Map with Pave for SOLiD and Illumina
- Be-align with SBSs
- Map with MosaiK

UF Information Technology
www.it.ufl.edu

Galaxy

Map with Bowtie for Illumina

Will you select a reference genome from your history or use a built-in index?

Built-ins were indexed using default options

Select a reference genome:

If your genome of interest is not listed - contact Galaxy team

Is this library mate-paired?:

FASTQ file:

Must have ASCII encoded quality scores

Bowtie settings to use:

For most mapping needs use Commonly used settings. If you want full control use Full parameter list

Suppress the header in the output SAM file:

Bowtie produces SAM with several lines of header information by default

UF Information Technology www.it.ufl.edu

Galaxy

UF Information Technology www.it.ufl.edu

What can you run?

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

UF Information Technology www.it.ufl.edu

Cluster basics

User interaction

Login node (Head node)

Scheduler

Tell the scheduler what you want to do

Compute resources

Your job runs on the cluster

UF Information Technology www.it.ufl.edu

Cluster login

submit.hpc.ufl.edu

ssh

/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

/home

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

Windows: PuTTY

Mac/Linux: Terminal

User interaction

```

magitz@submit1 ~$ ssh - bash - 67417
Last login: Mon Jun 11 21:49:41 on ttys000
Voyager-II:~ matt$ ssh magitz@submit.hpc.ufl.edu
magitz@submit.hpc.ufl.edu: 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

[magitz@submit1 ~]$ pwd
/home/magitz
[magitz@submit1 ~]$
```

UF Information Technology www.it.ufl.edu

Logging in

UF Information Technology www.it.ufl.edu

Linux Command Line

- ▶ Lots of online resources
 - Google: Linux cheat sheet
- ▶ Training sessions
- ▶ User manuals for applications

UF Information Technology www.it.ufl.edu

Storage at HPC

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 write job output here*
 - **Scratch space:** Lustre File System
 - /scratch/hpc/\$USER
 - 500GB per lab
 - Other options available for purchase and bio investors

/scratch/hpc/ must be used for all file input/output during job execution

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

Tell the scheduler what you want to do

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

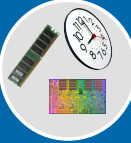
UF Research Computing

► Submission Script

```
#!/bin/bash
#
#PBS -N My_Job_Name
#PBS -M Joe_Shmoie@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
```

Scheduler



Tell the scheduler what you want to do

Compute resources

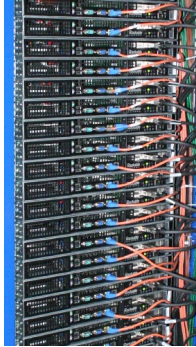



Your job runs on the cluster

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
```

- Lots to consider, but do your best at estimating RAM needed for job
- Over about 2GB of RAM, "costs" toward CPU allocation

Wasted RAM lead to idle CPUs and low job throughput



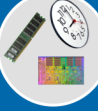
UF Information Technology www.it.ufl.edu

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

Scheduler



Tell the scheduler what you want to do

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

UF Information Technology www.it.ufl.edu

UF Research Computing

► Job Scheduling and Usage

- Test nodes (test01-06) available for interactive use, testing and short jobs (**EL5**: 01, 04, 05; **EL6**: 02,03,06)
 - e.g.: `ssh test01`
- 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 couple GB/core starts counting toward the total PE value of a job

UF Information Technology www.it.ufl.edu

Training Schedule

- ✓ Aug 28: Intro to UFHPC, getting started
- ▶ Sept 10: Modules, RHEL6 Transition, User Q&A
- ▶ Sept 17: The Linux/Unix Shell - An Introduction
- ▶ Sept 24: Running Jobs, Submission Scripts, Modules
- ▶ Oct 1: Galaxy Overview, The Basics
- ▶ Oct 8: NGS Data Techniques: General Methods and Tools
- ▶ Oct 15: NGS Data Techniques: Reference Based Mapping
- ▶ Oct 22: NGS Data Techniques: de Novo Assembly
- ▶ Oct 29: Phylogenetic Analyses
- ▶ Nov 5: Multiprocessing at the HPC Center
- ▶ Nov 12: Using Git and CMake to Organize and Drive Data Analysis Pipelines
- ▶ Nov 19
- ▶ Nov 29
- ▶ Dec 3

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



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

