

UF Research Computing: An Introduction and Getting Started



Matt Gitzendanner
magitz@ufl.edu

9/3/15

UF Research Computing
Information Technology
Home of High-Performance Computing and *HiPerGator*


UF Information Technology www.it.ufl.edu

UNIVERSITY OF FLORIDA'S 2015-2020 STRATEGIC GOALS FOR IT

One IT for the GatorGood

- ▶ Research Computing Strategic Goals
 - Expand HPC, data storage and research network capacity, performance and usability
 - Enhance and expand services that use HPC, data storage and network resources
 - Improve faculty awareness and access to use of Research Computing services

UNIVERSITY OF FLORIDA | High-Performance Computing



HiPerGator

The University of Florida Supercomputer for Research

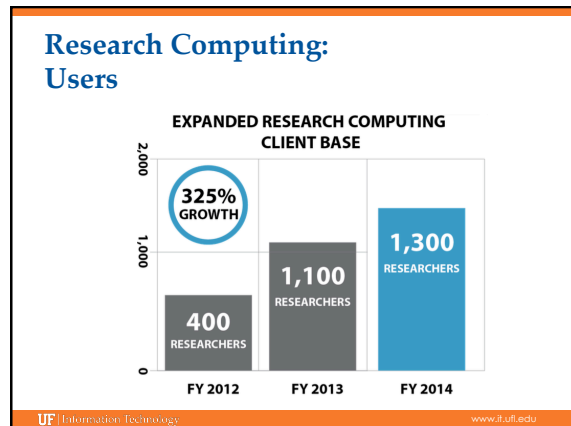
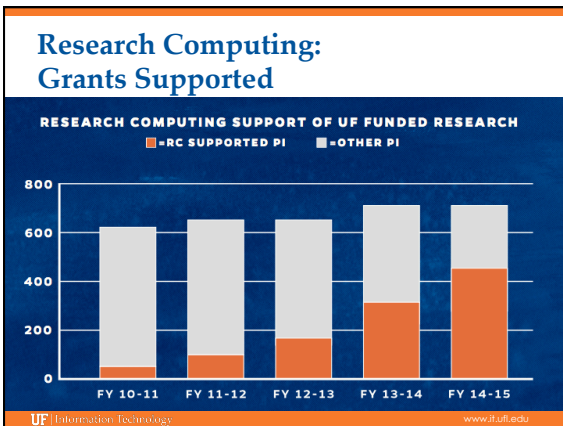
UF Information Technology www.it.ufl.edu



HiPerGator 2

The University of Florida Supercomputer for Research

UF Information Technology www.it.ufl.edu



UF Information Technology www.it.ufl.edu



HiPerGator

The University of Florida Supercomputer for Research

- 16,384 cores—total of about 21,000 cores today
- Infiniband interconnect
- >3PB fast, high-availability, storage
- **GPGPUs**
- Large memory nodes (512GB to 1TB of RAM)

Research Computing

Where do you start?

UF Information Technology www.it.ufl.edu

Research Computing

- ▶ User Accounts
 - Need current UF faculty sponsor

www.rc.ufl.edu


UF Information Technology www.it.ufl.edu

Investor Supported

Investment Option	Total Cost	University Match	Cost to PIs (2015-2016)
Long-Term 1 NCU/5 Years	\$300	\$100	\$200
Billed 1NCU/Hour	\$0.04	-	\$0.04
GPU NVIDIA K80/5 Years	\$2,250	\$1,750	\$500
Storage 1TB/Year	\$125	-	\$125
Replicated Storage 1TB/Year	\$250	-	\$250

UF Information Technology 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

Cluster Basics

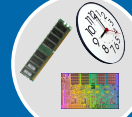
User interaction

Galaxy




Login server (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

gator.rc.ufl.edu


ssh gator1 gator3
gator2 gator4

/home/\$USER

ssh <user>@gator.rc.ufl.edu

Windows: PuTTY 
Mac/Linux: Terminal 

User interaction

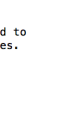


Login server (Head node)

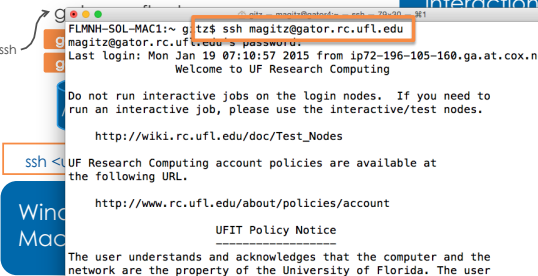
UF Information Technology www.it.ufl.edu

Cluster login

User interaction



Login



FLMNH-SOL-MAC1:~ gitz\$ ssh magitz@gator.rc.ufl.edu
magitz@gator.rc.ufl.edu\$
Last login: Mon Jan 19 07:19:57 2015 from ip72-196-105-160.ga.at.cox.net
Welcome to UF Research Computing

Do not run interactive jobs on the login nodes. If you need to run an interactive job, please use the interactive/test nodes.

http://wiki.rc.ufl.edu/doc/Test_Nodes

UF Research Computing account policies are available at the following URL.
<http://www.rc.ufl.edu/about/policies/account>

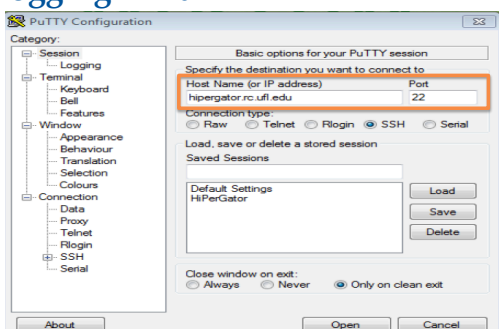
UFIT Policy Notice

The user understands and acknowledges that the computer and the network are the property of the University of Florida. The user agrees to comply with the University of Florida Acceptable Use Policy and Guidelines. The university monitors computer and network activities without user authorization and the university may

Windows: PuTTY
Mac: Terminal

UF Information Technology www.it.ufl.edu

Logging in PuTTY



Category: Session

Basic options for your PuTTY session

Specify the destination you want to connect to

Host Name (or IP address) Port

Connection type:
 Raw Telnet Rlogin SSH Serial

Load, save or delete a stored session

Saved Sessions

Default Settings
HiPerGator

Load Save Delete


Close window on exit:
 Always Never Only on clean exit

About Open Cancel

UF Information Technology www.it.ufl.edu

Development servers

User interaction



Login

- Do not run applications on the login servers
 - Account will be suspended


Do not run interactive jobs on the login nodes.

UF Research Computing account policies are available at the following URL.
<http://www.rc.ufl.edu/about/policies/account>

- Use the development servers for testing and interactive use:
 - ssh dev1 or ssh dev2

UF Information Technology www.it.ufl.edu

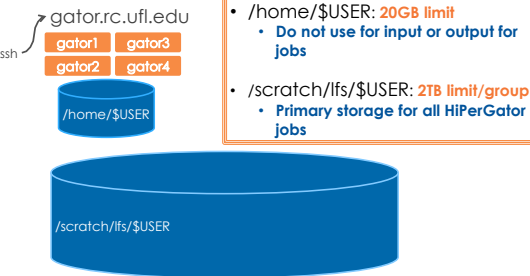
Linux Command Line



- ▶ Lots of online resources
 - Google: Linux cheat sheet
- ▶ Training sessions
 - Sept 10: The Linux/Unix Command Line--An Introduction
- ▶ User manuals for applications

UF Information Technology www.it.ufl.edu

Cluster Storage



- /home/\$USER: **20GB limit**
 - Do not use for input or output for jobs
- /scratch/lfs/\$USER: **2TB limit/group**
 - Primary storage for all HiPerGator jobs

UF Information Technology www.it.ufl.edu

Research Computing



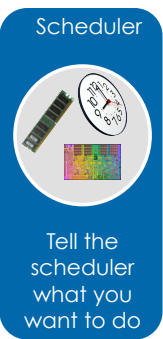
- ▶ Storage
 - **Home:** /home/\$USER
 - For code compilation and user file management only
 - **Do not use for job input/output!**
 - Include `cd $PBS_O_WORKDIR` or similar in scripts
 - **Scratch space:** Lustre File System
 - /scratch/lfs/\$USER

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



Tell the scheduler what you want to do

UF Information Technology www.it.ufl.edu

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

Research Computing

- ▶ Submission Script

```
#!/bin/bash
#PBS -N My_Job_Name
#PBS -M Joe_Shmoefufl.edu
#PBS -m abe
#PBS -o My_Job_Name.log
#PBS -e My_job_Name.err
#PBS -l nodes=1:ppn=1
#PBS -l pmem=900mb
#PBS -l walltime=00:05:00

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



Tell the scheduler what you want to do

UF Information Technology www.it.ufl.edu

Nodes and processors

Single processor apps:


```
#PBS -l nodes=1:ppn=1
```

Threaded (& MPI) apps:

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

MPI apps:

```
#PBS -l nodes=2:ppn=64
```



RAM

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

- ▶ *Per-processor* memory request
- ▶ Lots to consider, but do your best at estimating memory needed for job
- ▶ Over about 4GB of RAM, "costs" toward CPU allocation



Wasted RAM leads to idle CPUs and low job throughput

End-of-job emails: `#PBS -M Joe_Shmo@ufl.edu` `#PBS -m abe`

```
PBS Job Id: 358634.moab.ufhpc
Job Name: NR.25.nex
Exec host: c7a-s1/60
Execution terminated
Exit_status=0
resources_used.cput=07:16:09
resources_used.mem=251348kb
resources_used.vmem=318916kb
resources_used.walltime=07:16:52
```


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

	Maximum	Short	Long
Investor	31 days	<12 hrs	7 days
Other	7 days	<12 hrs	3 days

Scheduler



Tell the scheduler what you want to do


Research Computing

- ▶ Job Management
 - `qsub <file_name>`: job submission
 - `qstat -u <user>`: check queue status
 - `showq -r -u <user>`: shows job efficiency
 - `qdel <JOB_ID>`: job deletion
 - `checkjob -v <job number>` (shows PE value)
 - `pbs_info -f my_job.pbs` (get job PE and group resources before submitting a job)



UF Preeminence
INVESTING IN PEOPLE AND PROGRAMS
THAT HELP HELP THE WORLD







UF Information Technology

www.it.ufl.edu

GatorCloud

- ▶ Microsoft Office and OneDrive@UF
- ▶ GatorBox
- ▶ UF Research Apps
- ▶ GatorVault
- ▶ GatorCloud Virtual Systems
- ▶ GatorCloud Portal

UF Information Technology

www.it.ufl.edu

Training Schedule

- ✓ Sept 3: Intro to UFHPC, getting started
- ▶ Sep 10: The Linux/Unix Shell - An Introduction
- ▶ Sep 17: HiPerGator: Running Jobs, Submission Scripts, Modules
- ▶ Sep 24: MATLAB at Research Computing: A Hands-on Tutorial to Running Your MATLAB Code at Research Computing

More training sessions
will be added

UF Information Technology

www.it.ufl.edu

UF Research Computing

- ▶ Help and Support
 - <https://my.it.ufl.edu>
 - For any kind of question or help requests
 - <http://wiki.rc.ufl.edu>
 - Documents on hardware and software resources
 - Various user guides
 - Many sample submission scripts
 - <http://rc.ufl.edu>
 - Frequently Asked Questions
 - Account set up and maintenance



UF Information Technology

www.it.ufl.edu