

Up and Running at UF Research Computing



Matt Gitzendanner
magitz@ufl.edu

6/18/15

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

UF | Information Technology | www.it.ufl.edu

Research Computing




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

- ▶ 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

UF | Information Technology | www.it.ufl.edu

UNIVERSITY OF FLORIDA | High-Performance Computing



HiPerGator

The University of Florida Supercomputer for Research

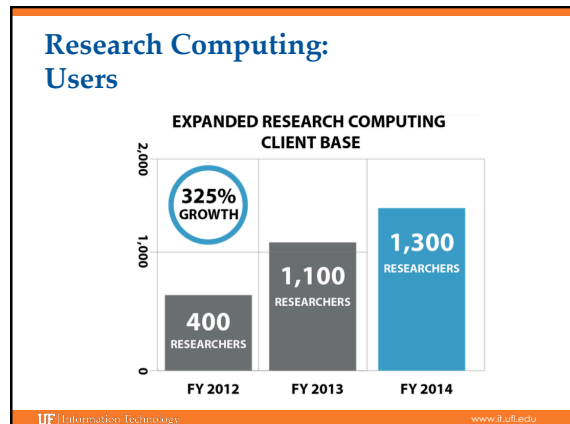
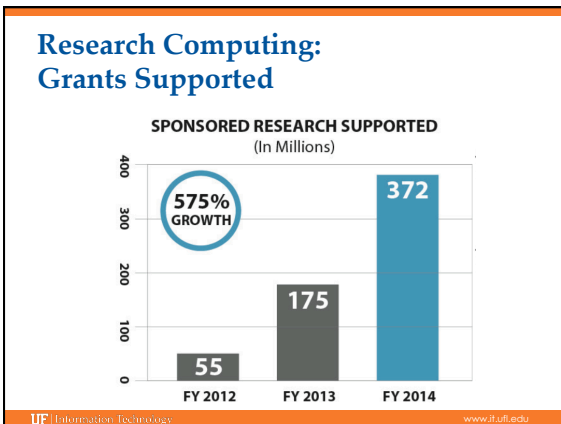
UF | Information Technology | www.it.ufl.edu



HiPerGator 2

The University of Florida Supercomputer for Research

UF | Information Technology | www.it.ufl.edu



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

UF Information Technology www.it.ufl.edu



HiPerGator

The University of Florida Supercomputer for Research

- 16,384 cores—total of about 20,000 cores today
- Infiniband interconnect
- >3PB fast, high-availability, storage
- **GPGPUs**
- Large memory nodes (**750GB to 1TB of RAM**)

UF Information Technology www.it.ufl.edu

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

What can you run?


- ▶ Linux-based
- ▶ Generally command line driven applications
- ▶ Galaxy
- ▶ Graphical apps can be setup
 - SAS
 - MATLAB
 - STATA
 - and others

UF Information Technology www.it.ufl.edu

Cluster Basics

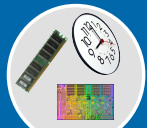
User interaction

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

```

FLMNH-SOL-MAC1:~$ git$ ssh magitz@gator.rc.ufl.edu
magitz@gator.rc.ufl.edu:~$ ssh magitz@gator.rc.ufl.edu
Last login: Mon Jan 19 07:10:57 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

ssh <user>@gator.rc.ufl.edu
UF Research Computing account policies are available at
the following URL.

http://www.rc.ufl.edu/about/policies/account

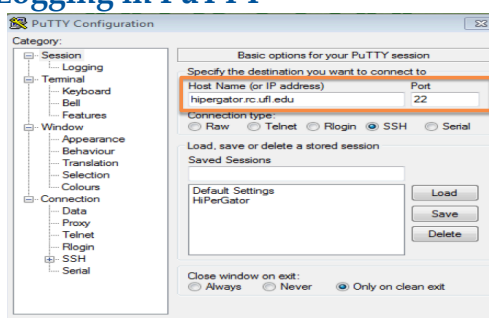
UFIT Policy Notice

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
                    
```

User interaction

UF Information Technology | www.it.ufl.edu

Logging in PuTTY



The screenshot shows the PuTTY Configuration dialog box. The 'Basic options for your PuTTY session' tab is selected. The 'Host Name (or IP address)' field is highlighted with an orange box and contains 'hpergator.rc.ufl.edu'. The 'Port' field contains '22'. The 'Connection type' is set to 'SSH'. The 'Close window on exit' option is set to 'Only on clean exit'.

UF Information Technology | www.it.ufl.edu

Cluster Storage

gator.rc.ufl.edu

ssh gator1 gator3
gator2 gator4

/home/\$USER

/scratch/lfs/\$USER

- /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

Cluster Basics

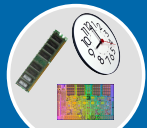
User interaction

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

Don't be intimidated!



```

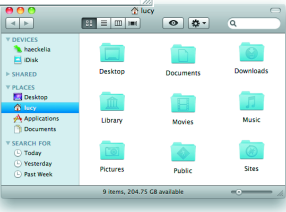
Last login: Thu Jul 25 12:03:00 on tty002
This host: lucy
FUTURE-SQL-MAC11~ gistas
    
```

UF Information Technology | www.it.ufl.edu

File system

```

Terminal — bash — 80x24
Last login: Mon Dec 8 21:44:54 on tty003
Host:~ lucy$ ls
Desktop  Downloads  Movies     Pictures   Sites
Documents Library    Music     Public
Host:~ lucy$
    
```



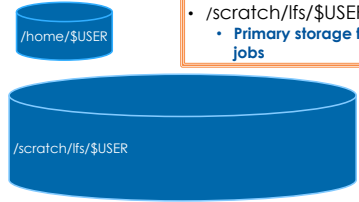
UF Information Technology | www.it.ufl.edu

Cluster Storage

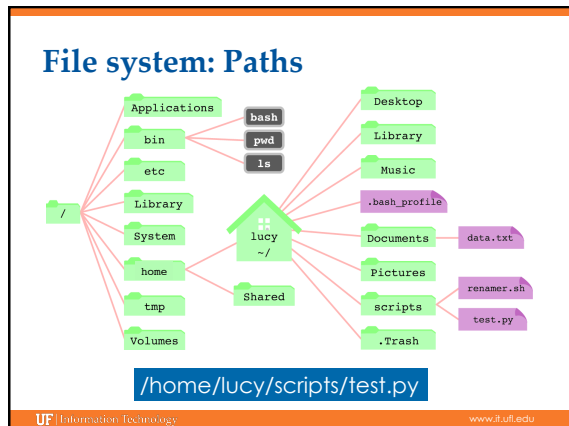
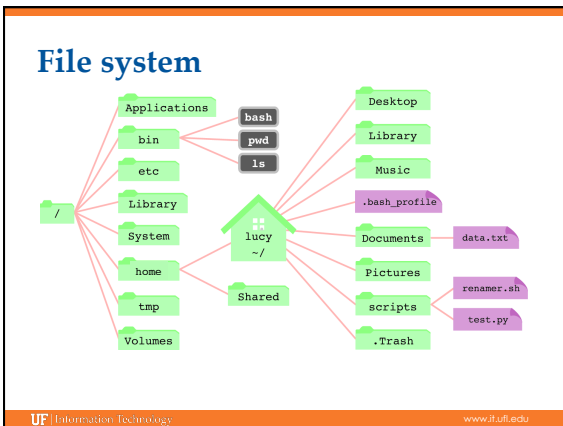
gator.rc.ufl.edu

ssh → gator1 gator3
gator2 gator4

- /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



File system: ls & cd

```
[magitz@gator1 scripts]$ ls
renamer.sh test.py
[magitz@gator1 scripts]$ cd ../Documents/
[magitz@gator1 Documents]$ ls
data.txt
[magitz@gator1 Documents]$
```

UF Information Technology www.it.ufl.edu

The command prompt

```
[magitz@gator1 scripts]$ ls
renamer.sh test.py
[magitz@gator1 scripts]$ cd ../Documents/
[magitz@gator1 Documents]$ ls
data.txt
[magitz@gator1 Documents]$
```

user host current directory

UF Information Technology www.it.ufl.edu

File system: Basic commands

- pwd: Present Working Directory
- ls: List--show contents of directory
- cd: Change Directory
- mkdir: Make directory
- rm: Remove a file (rm -r for directories)
- rmdir: Remove directory

There is no trashcan!
(or backups)

- cp: Copy a file: cp data.txt ../analysis/
- mv: Move a file: mv data.txt ../analysis/

UF Information Technology www.it.ufl.edu

Command arguments

- Modify how a command works
 - Used from the most basic command like ls, to most applications you will run

```
[magitz@gator1 scripts]$ ls
renamer.sh test.py
[magitz@gator1 scripts]$ ls -l
total 0
-rw-rw-r-- 1 magitz ufhcp 0 Jul 25 12:22 renamer.sh
-rw-rw-r-- 1 magitz ufhcp 0 Jul 25 12:22 test.py
[magitz@gator1 scripts]$ ls -la
total 8
drwxrwxr-x 2 magitz ufhcp 4096 Jul 25 12:22 .
drwxrwxr-x 4 magitz ufhcp 4096 Jul 25 12:25 ..
-rw-rw-r-- 1 magitz ufhcp 0 Jul 25 12:22 renamer.sh
-rw-rw-r-- 1 magitz ufhcp 0 Jul 25 12:22 test.py
[magitz@gator1 scripts]$
```

UF Information Technology www.it.ufl.edu

Linux permissions

Indicates a directory (d) or a file (-)

Group (set of users)

Other users (everyone else)

User (owner) permissions

drwxrwxrwx

- Change with: chmod o-rwx file
 - ls -l shows permission details

```
-rwxrwxr-- 1 magitz ufhcp 0 Jul 25 12:22 renamer.sh
  u   g   o      user   group
```

UF Information Technology www.it.ufl.edu

Viewing text files

- more and less: show contents of a file, less has more features (search, scroll up, etc.)
- head: show start of file (10 lines by default)
 - head -n 20 file.txt
- tail: show end of file (10 lines by default)
- grep: search for text in file
 - grep Gainesville weather.txt
- Text editors: nano, vi, vim, emacs, etc.

UF Information Technology www.it.ufl.edu

Redirecting output

- ▶ The ">" can be used to redirect output:
 - `grep GNV weather.txt > Gainesville.txt`
 - `ls -l > file_list.txt`
- **Caution:** > will overwrite the specified file!
- ▶ You can append to a file with ">>"

UF Information Technology www.it.ufl.edu

Connecting commands

- ▶ The pipe, "|", can be used to send, or pipe, the output of one command to the input of another:
 - `ls -l | grep Jun > files_modified_in_June`

Tip: Don't use spaces in file or directory names

UF Information Technology www.it.ufl.edu

Tools

- ssh client to connect to `hipergator.rc.ufl.edu`
e.g.: Terminal, PuTTY
- SFTP client to move files to/from your computer
e.g.: Cyberduck, FileZilla
- Text editor to prepare files
Especially on Windows, be sure to convert DOS line breaks to Unix, and *don't use Word*. Both have SFTP built in
e.g.: TextWrangler, Notepad++

UF Information Technology www.it.ufl.edu

SSH Clients

Mac/Linux: Terminal

Windows: PuTTY

`ssh user@gator.rc.ufl.edu`

UF Information Technology www.it.ufl.edu

PuTTY

UF Information Technology www.it.ufl.edu

Text Editors

- ▶ **Not** Microsoft Word or other word processor
- ▶ Contextual coloring
- ▶ Built-in SFTP Client
- ▶ Regular expression find/replace
- ▶ **Unix line breaks**


UF Information Technology www.it.ufl.edu

Unix line breaks

The image shows a 'Preferences' dialog box for a text editor. The 'Line breaks' section is highlighted, showing three radio button options: 'Unix (LF)', 'Classic Mac (CR)', and 'Windows (CRLF)'. The 'Unix (LF)' option is selected. Other options include 'Ensure file ends with line break' and 'Strip trailing whitespace'. There are also checkboxes for 'Follow current document' and 'Remember last used directory'.

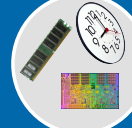
Cluster basics

User interaction




Login server (Head node)

Scheduler



Tell the scheduler what you want to do


Compute resources



Your job runs on the cluster

Learn by doing...

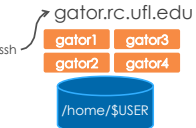
Handouts also available on Wiki



Software you need for this tutorial:

ssh Client	Terminal	✓	putty.org
SFTP Client	cyberduck.io		filezilla-project.org
Text Editor	barebones.com		notepad-plus-plus.org


Cluster login



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

Windows: PuTTY
Mac/Linux: Terminal

User interaction



Login server (Head node)

Cluster login

User interaction

```

FLMNH-SOL-MAC11:~ g$ ssh magitz@gator.rc.ufl.edu
magitz@gator.rc.ufl.edu:~$
Last login: Mon Jan 19 07:10:57 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

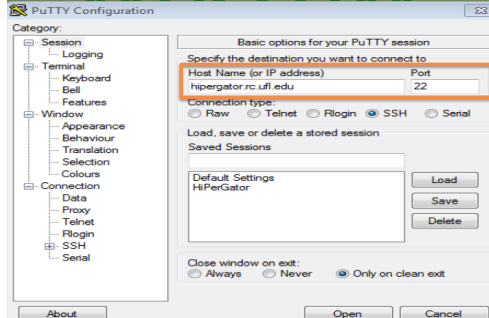
ssh <user>@gator.rc.ufl.edu
UF Research Computing account policies are available at
the following URL.

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
    
```

Logging in PuTTY



The image shows the 'PuTTY Configuration' dialog box. The 'Session' category is selected. Under 'Basic options for your PuTTY session', the 'Host Name (or IP address)' is set to 'gator.rc.ufl.edu' and the 'Port' is set to '22'. The 'Connection type' is set to 'SSH'. There are 'Load', 'Save', and 'Delete' buttons for saved sessions.



Development servers

- 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

Learn by doing...

Handouts also available on Wiki




Software you need for this tutorial:

ssh Client	Terminal	✓	putty.org
SFTP Client	cyberduck.io		filezilla-project.org
Text Editor	barebones.com		notepad-plus-plus.org

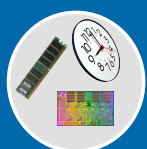
Cluster basics

User interaction




Login server (Head node)

Scheduler



Tell the scheduler what you want to do

Compute resources

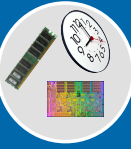


Your job runs on the cluster

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

Scheduler



Tell the scheduler what you want to do

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.

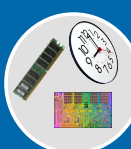
Research Computing

- Submission Script

```
#!/bin/bash
#PBS -N My_Job_Name
#PBS -M Joe_Shmoee@ufl.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
```

Scheduler



Tell the scheduler what you want to do

Compute resources



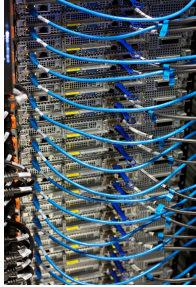
Your job runs on the cluster

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




UF Information Technology www.it.ufl.edu

HiPerGator

The University of Florida Supercomputer for Research

- **64 cores per node**
 - If RAM allows, MPI jobs under ~32 cores, should use `nodes=1:ppn=##`
 - Some older nodes have 4-16 cores

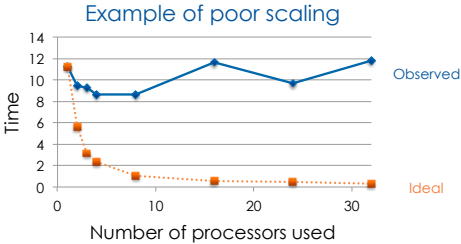


UF Information Technology www.it.ufl.edu

Parallel efficiency

▶ How well does your application scale?

Example of poor scaling




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 4GB of RAM, "costs" toward CPU allocation

Wasted RAM leads to idle CPUs and low job throughput



UF Information Technology www.it.ufl.edu


Processor Equivalents

- ▶ Accounts for large RAM requests
- ▶ Average ~4GB RAM/core

1 core, 10GB RAM: ~2.5 PEs
 1 core, 60GB RAM: ~15 PEs

- ▶ Non-investor's limit: 8 PEs
- ▶ Investor limits are based on PEs

▶ `pbs_info -f <job_file>`



UF Information Technology www.it.ufl.edu

End-of-job emails:

```
#PBS -M Joe_Shmoie@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
```

UF Information Technology www.it.ufl.edu

RAM– bigmem queue

- ▶ For jobs asking for over 16GB per core (pmem)
- ▶ #PBS -q bigmem
- ▶ 1TB and 750GB nodes



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

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



UF Information Technology

www.it.ufl.edu

Job Management

- ▶ Commands:
 - 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 Information Technology

www.it.ufl.edu

Learn by doing...

Handouts also available on Wiki



Software you need for this tutorial:

ssh Client	Terminal ✓	putty.org
SFTP Client	cyberduck.io	filezilla-project.org
Text Editor	barebones.com	notepad-plus-plus.org

UF Information Technology

www.it.ufl.edu

So what is this “module” thing?

- ▶ **lmod**—Implementation of Environment Modules developed at TACC
- ▶ Allows easy management of user's environment



Home About Resources User Services Research & Development Partnerships Education & Outreach News

Lmod: Environmental Modules System

UF Information Technology

www.it.ufl.edu

The standard way

```
PATH=$PATH:/some/long/path/to/application
export $PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/long/path/to/
place/I/probably/cant/find
export $LD_LIBRARY_PATH
```

- ▶ Need to track down paths to applications, libraries, etc.
- ▶ Multiple compilers, and MPI implementations
- ▶ Manage dependencies
- ▶ Multiple versions of apps



UF Information Technology

www.it.ufl.edu

Module discovery

- ▶ `module spider`
 - List everything
- ▶ `module spider cl`
 - List applications that have cl in name
- ▶ `module spider amber/12`
 - List details about this version of AMBER
- ▶ `module key molecular`
 - Keyword search for applications

Multiple versions

```
[magitz@gator1 ~]$ module spider gaussian
Rebuilding cache file, please wait ... done
```

```
-----
gaussian:
```

```
Description:
```

```
A software for electronic structure modeling
```

```
Versions:
```

```
gaussian/e01
gaussian/g03
gaussian/g09
```

```
-----
To find detailed information about gaussian please enter the full name.
For example:
```

```
$ module spider gaussian/g09
-----
```

Module loading

- ▶ `module load raxml`
- ▶ `module load intel raxml`
- ▶ `module load intel openmpi raxml`
- ▶ `module load intel/12 openmpi/1.6 raxml/3.2`
- ▶ `module unload raxml`

Training sessions

- ▶ Online <http://wiki.hpc.ufl.edu/doc/Training>

Training Resources [edit]

The most recent set of slides, handouts and recordings from training sessions are listed below.

- [UF Research Computing: An Introduction and Getting Started](#)
 - Download the slides from the presentation.

- [View a recording of the session](#)



- [An Introduction to the Linux Command Line](#)
 - Download the slides from the presentation.
 - Download the handout with the exercises.
 - Download an example Linux cheat sheet.

- [View a recording of the session](#)



UF Research Computing

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

