

Research Computing Orientation for UF Courses

Matt Gitzendanner: magit@ufl.edu

8/27/15

UF www.ufl.edu

UF www.ufl.edu

SSH Clients

UF www.ufl.edu

PuTTY

UF www.ufl.edu

HiPerGator

The University of Florida Supercomputer for Research

- 16,384 cores (total of ~21,000 today)
- Infiniband interconnect
- 2.1PB fast, high-availability, storage

UF www.ufl.edu

Course use of HiPerGator

- Coursework should be done through course account
 - All course accounts are deleted at the end of the semester
- <http://www.rc.ufl.edu/help/accountrequest/>
- Course is allocated 32-cores
 - Design projects with this in mind
 - Time your work with this in mind
- Support requests should go through course TA
 - If TA cannot solve the issue, the TA should open support requests

UF www.ufl.edu

FileZilla

Tutorial at: <http://wiki.rc.ufl.edu/doc/f9e/f9e>

UF www.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 www.ufl.edu

Galaxy: Data intensive biology for everyone

- Accessible, reproducible, transparent computational biology
- galaxy.rc.ufl.edu

UF www.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 www.ufl.edu

Unix line breaks

UF www.ufl.edu


Storage

- `/home/$USER`
 - 20GB limit
 - scripts, code, small data
 - Do NOT use for job input/output
- `/scratch/lfs/$USER`
 - 2TB limit per group
 - ALL input/output from jobs should go here

- All storage systems are for research and coursework data only
- Nothing is backed up
- All course accounts are deleted at the end of the semester

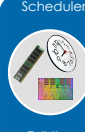
UF www.ufl.edu

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

UF

Development servers

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

Do not run interactive jobs on the login nodes.

UF BPC Center Account Policies can be found here:
<http://www.bpc.ufl.edu/about/policies/account>
 (needs password) - 15

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

UF

HiPerGator

The University of Florida Supercomputer for Research

- Sharing Resources:**
 - Remember courses are limited to 32 cores total
 - showq -w group=ams6001




UF

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

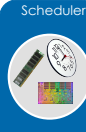
Wasted RAM leads to idle CPUs and low job throughput



UF

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

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

End-of-job emails:

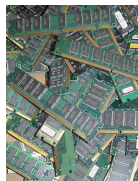
#PBS -M Joe_Shmoe@ufl.edu
#PBS -m abe

```
PBS Job Id: 358634.moab.ufhpc
Job Name: NR.25.nex
Exec host: c7a-sl/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

RAM- bigmem queue

- For jobs asking for over 16GB per core (pmem)
- #PBS -q bigmem
- ITB node



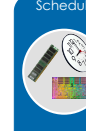
UF

Research Computing

- Submission Script

```
#!/bin/bash
#PBS -M My_Job_Name
#PBS -M Joe_Shmoe@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=01:05:00

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



Scheduler

Tell the scheduler what you want to do


UF

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



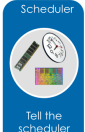
UF

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
Courses	7 days	<12 hrs	3 days



Scheduler

Tell the scheduler what you want to do

UF


UF Research Computing

- Job Management
 - qsub <file_name>: job submission
 - qstat -u <user>: check queue status
 - showq -f -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

So what is this “module” thing?

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



TACC

TEXAS ADVANCED COMPUTING CENTER
Enabling Discovery Through Data Change the World

UNIVERSITY OF TEXAS
AT AUSTIN

[Home](#) [About](#) [Resources](#) [User Services](#) [Research & Development](#) [Partnerships](#) [Education & Outreach](#) [News](#)

Lmod: Environmental Modules System


UP [for more information](#)

[www.tacc.utexas.edu](#)

The standard way

```
PATH=$PATH:$some/long/path/to/application
export $PATH
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$some/long/path/to/
place/1/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



UP University of Pittsburgh



www.cs.cmu.edu


Training sessions

► Thursdays@ 12:50 and online
<http://wiki.hpi.uni-edu/doc/Training>

Training Resources and


The most recent of online, technical and knowledge from training sessions are listed below

- **► of Training Resources** (in this slide and Training Portal)
- Download the slides from the slides
- 
- **► of Training Resources** (in this slide and Training Portal)
- Download the slides from the slides
- Download an example (in this slide and Training Portal)
- 

UP  www.up.uni-paderborn.de www.hpi.uni-edu

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



UF University of Florida Research Computing www.rc.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

```
(mag@elsturn) ~$ module spider gaussian  
Reloading config file, please wait ... done
```

gaussian:

Description:
A software for electronic structure modeling

Version:
gaussian/e01
gaussian/g0
gaussian/g09

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

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

Multiple variants of a version

```
[magit]#submit -j$ module spider mrbayes/3.2.1
Rebuilding cache file, please wait ... Done
-----
mrbayes: mrbayes/3.2.1
-----
```

Descriptions:

- Bayesian inference of phylogeny

This module can be loaded directly: module load mrbayes/3.2.1

Additional variants of this module can also be loaded after the loading the following modules:

```
intel/2012, openmpi/1.6
```

11P

more info

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
```