


Research Computing Orientation for UF Courses


Matt Gitzendanner
magitz@ufl.edu





**Providing
computational
resources for a
top 10 university**





- 51,000 Cores
- 6PB Storage
- High-speed networks

Course use of **HiPerGator**

The University of Florida Supercomputer for Research


- **Course is allocated 32 cores & 112GB RAM**
 - Design projects with this in mind
 - Time your work with this in mind
 - Use resources efficiently
- **Support requests should go through course TA**
 - If TA cannot solve the issue, the TA should open support requests
- **By using your account, you agree to the AUP**
 - <http://www.rc.ufl.edu/about/policies/>
 - No restricted data

For users with an account

- If you already have a Research Computing account for research:
- **module load class/pre1234**
 - Ensures that jobs run under course allocation
 - Storage is under course allocation


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

Tools

ssh client to connect to
hpg.rc.ufl.edu



e.g.: Terminal, MobaXterm

SFTP client to move files
to/from your computer



e.g.: Cyberduck, FileZilla

Text editor

Especially on Windows, be sure to
convert DOS line breaks to Unix, and
don't use Word



e.g.: VS Code



SSH Clients

```
magitz — magitz@login2:~ — ssh magitz@hpg.rc.ufl.edu — 80x30
flmnh-C02N85F3G:~ magitz$ ssh magitz@hpg.rc.ufl.edu
magitz@hpg.rc.ufl.edu's password:
Last login: Fri May 25 12:11:54 2018 from login4.ufhpc

Welcome to UF Research Computing
-----

The user agrees to comply with Research Computing's policies.

https://www.rc.ufl.edu/about/policies

Backup Policy
-----

Data on Research Computing storage systems is NOT backed up without investment
in backup services. Please visit the link below for more information.

https://www.rc.ufl.edu/about/policies/storage/protection/

UFIT Policy Notice
-----

The user understands and acknowledges that the computer and the network are the
property of the University of Florida (UF). The user agrees to comply with the
UF Acceptable Use Policy and Guidelines. UF monitors computer and network
activities without user authorization. UF may provide information about computer
or network usage to UF officials, including law enforcement when warranted.
Therefore, the user should have limited expectations of privacy.

[magitz@login2 ~]$
```



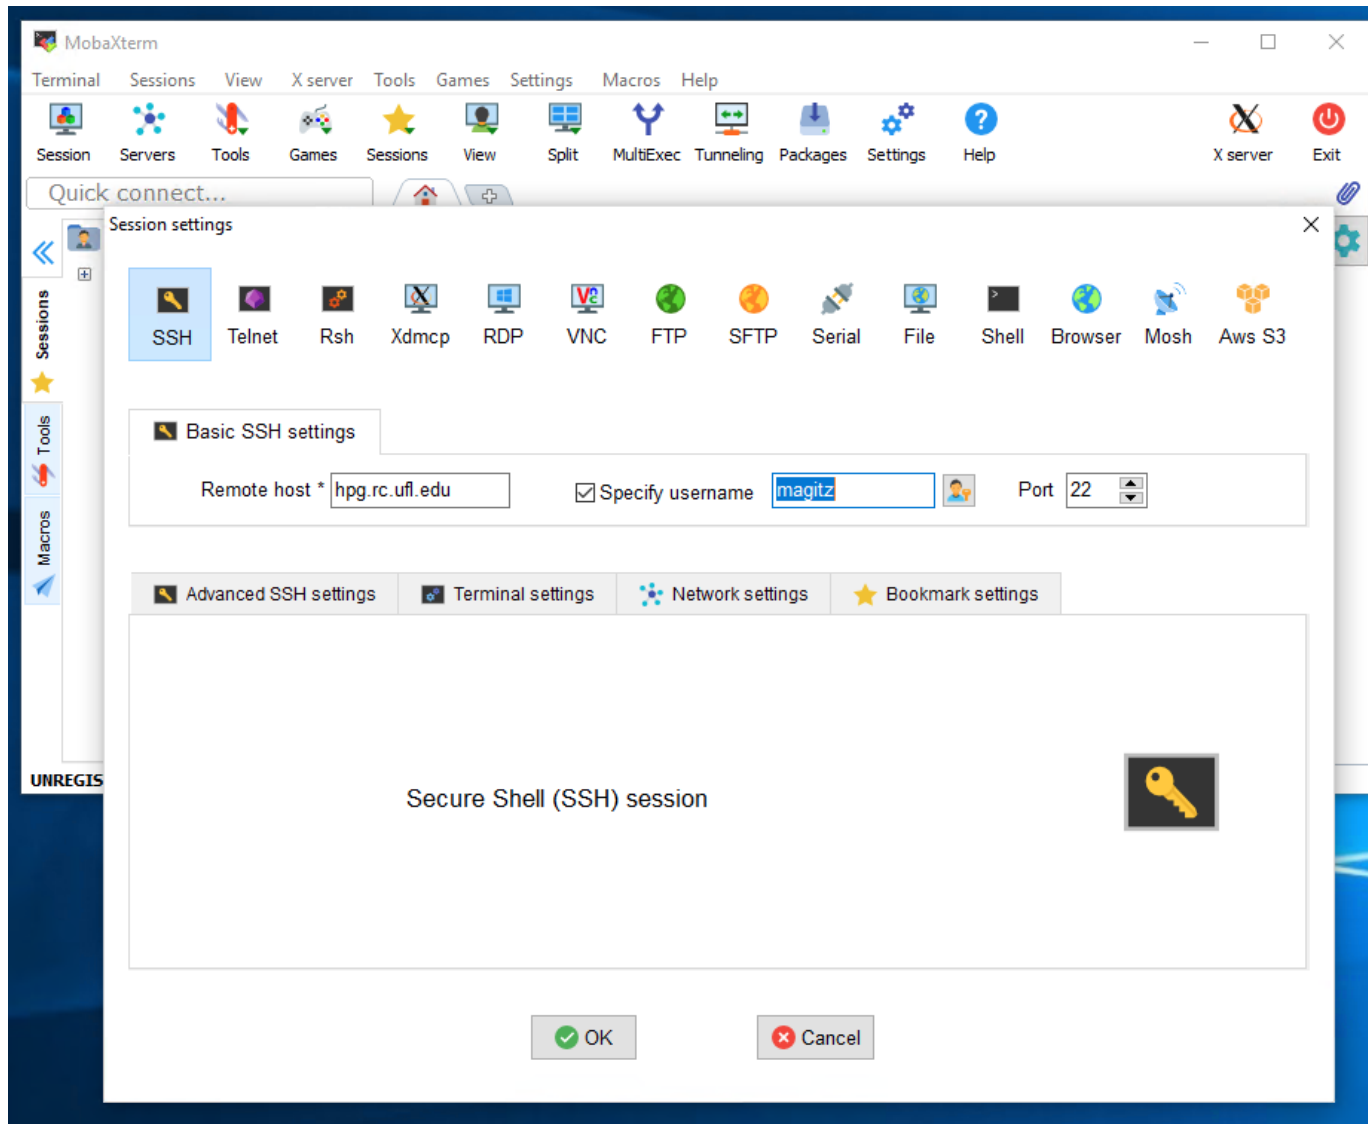
Mac/Linux: Terminal




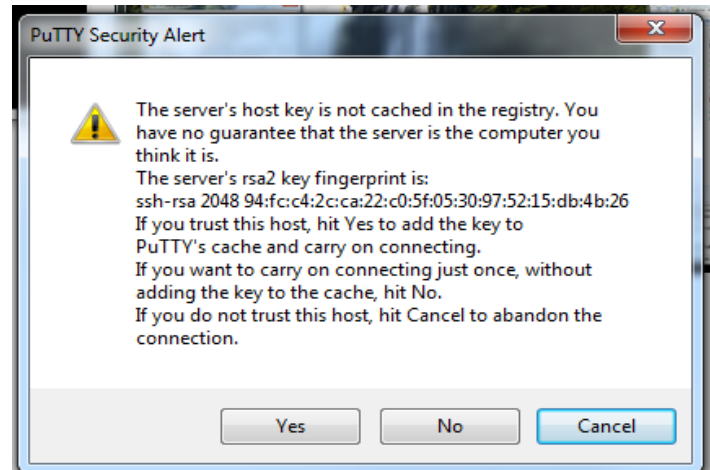
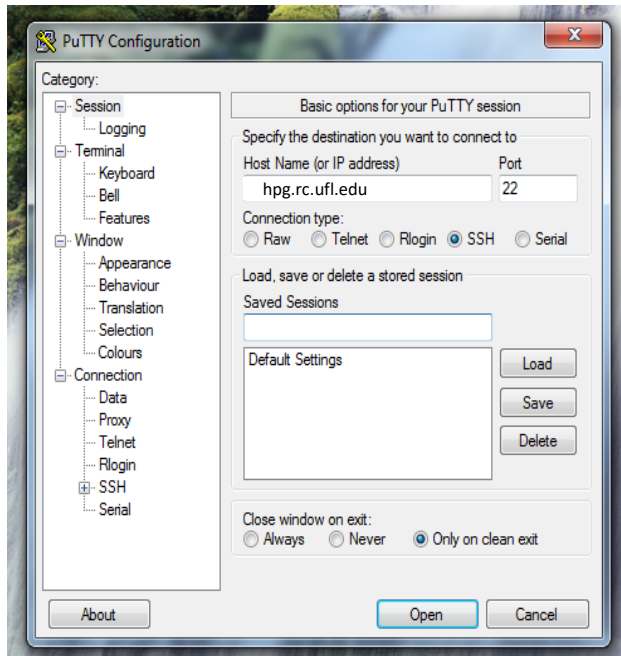
Windows:
MobaXterm or PuTTY

ssh gatorlink@hpg.rc.ufl.edu

MobaXterm



PuTTY



FileZilla



sftp://magitz@hpg.rc.ufl.edu - FileZilla

Host: sftp://hpg.rc.ufl.edu Username: magitz Password: Port: Quickconnect

Status: Connected to hpg.rc.ufl.edu
Status: Retrieving directory listing...
Status: Listing directory /home/magitz
Status: Directory listing of "/home/magitz" successful
Status: Retrieving directory listing of "/ufrc/ufhpc/magitz"...
Status: Listing directory /ufrc/ufhpc/magitz
Status: Directory listing of "/ufrc/ufhpc/magitz" successful

Local site: /Users/magitz/Desktop/ Remote site: /ufrc/ufhpc/magitz

Filename	Filesize	Filetype	Last modified
..			
Desktop Junk		Directory	03/31/2018 21:5...
SAS_test		Directory	04/20/2018 15:2...
data		Directory	12/12/2017 12:4...
.DS_Store	10,244	File	05/30/2018 11:0...
.localized	0	File	01/23/2017 16:2...
Fwunixref.pdf	70,752	pdf-file	01/12/2018 10:1...
Gitzendanner...	1,037,635	pdf-file	03/31/2018 14:5...
INTRO-TO-LL...	251,117	pdf-file	04/04/2018 20:...

5 files and 3 directories. Total size: 1,369,748 bytes

Filename	Filesize	Filetype	Last modified	Permissions	Owner/G
..					
CodingFun		Directory	10/23/2015 1...	drwxr-xr-x	magitz
CompBio		Directory	11/04/2016 1...	drwxrwxr-x	magitz
DoseResponse		Directory	02/15/2018 1...	drwxrwxr-x	magitz
DoseResponse.old		Directory	09/27/2017 1...	drwxrwxr-x	magitz
ICBR_Download		Directory	08/21/2017 1...	drwxr-xr-x	magitz
PythonWorkshop		Directory	10/29/2017 1...	drwxrwxr-x	magitz
SLURM_examples		Directory	05/29/2018 1...	drwxrwxr-x	magitz
ToolBox		Directory	02/13/2017 2...	drwxr-xr-x	magitz
cli demo		Directory	05/10/2018 1...	drwxrwxr-x	magitz

4 files and 23 directories. Total size: 828,903 bytes

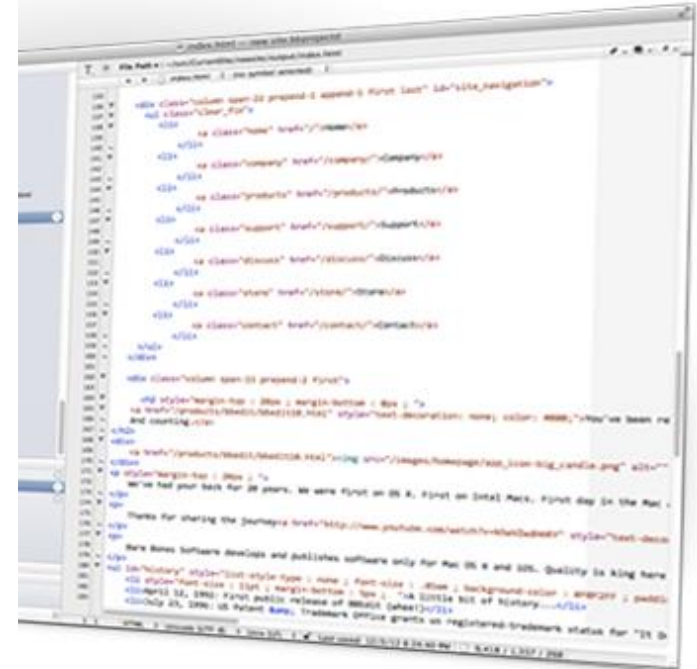

Server/Local file	Direction	Remote file	Size	Priority	Status
-------------------	-----------	-------------	------	----------	--------

Queued files Failed transfers Successful transfers

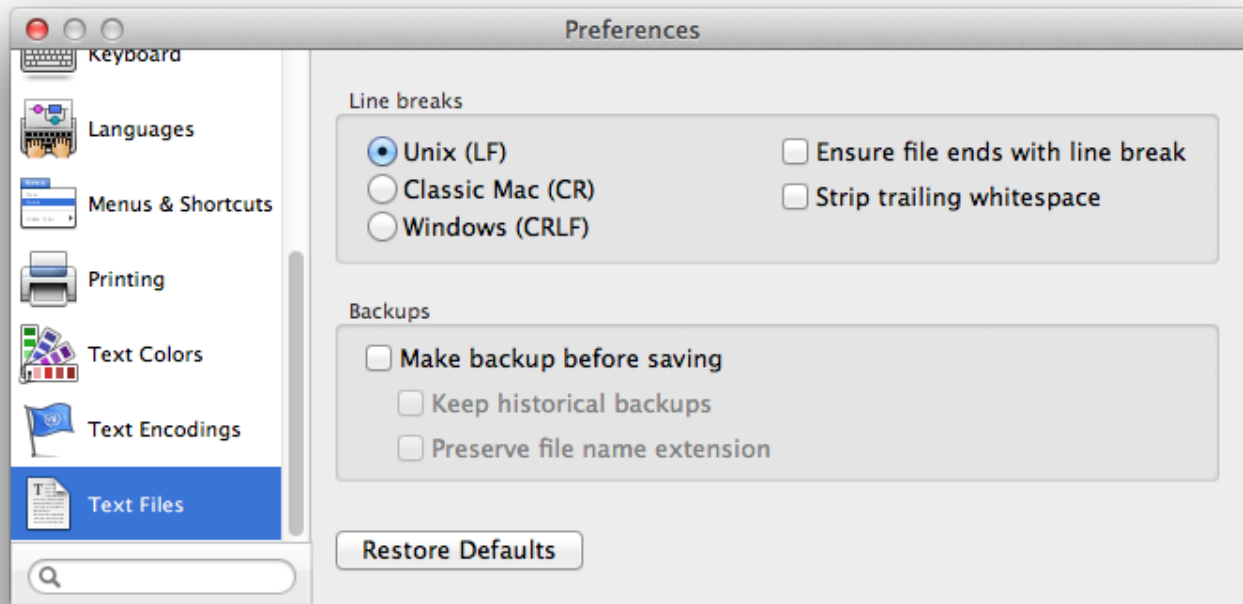

Queue: empty

Text Editors


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



Unix line breaks



Storage




- /home/\$USER
 - 20GB limit
 - scripts, code, small data
 - Do NOT use for job input/output
- /ufrc/pre1234/\$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


Cluster basics

User
interaction



Login node
(Head node)

Scheduler



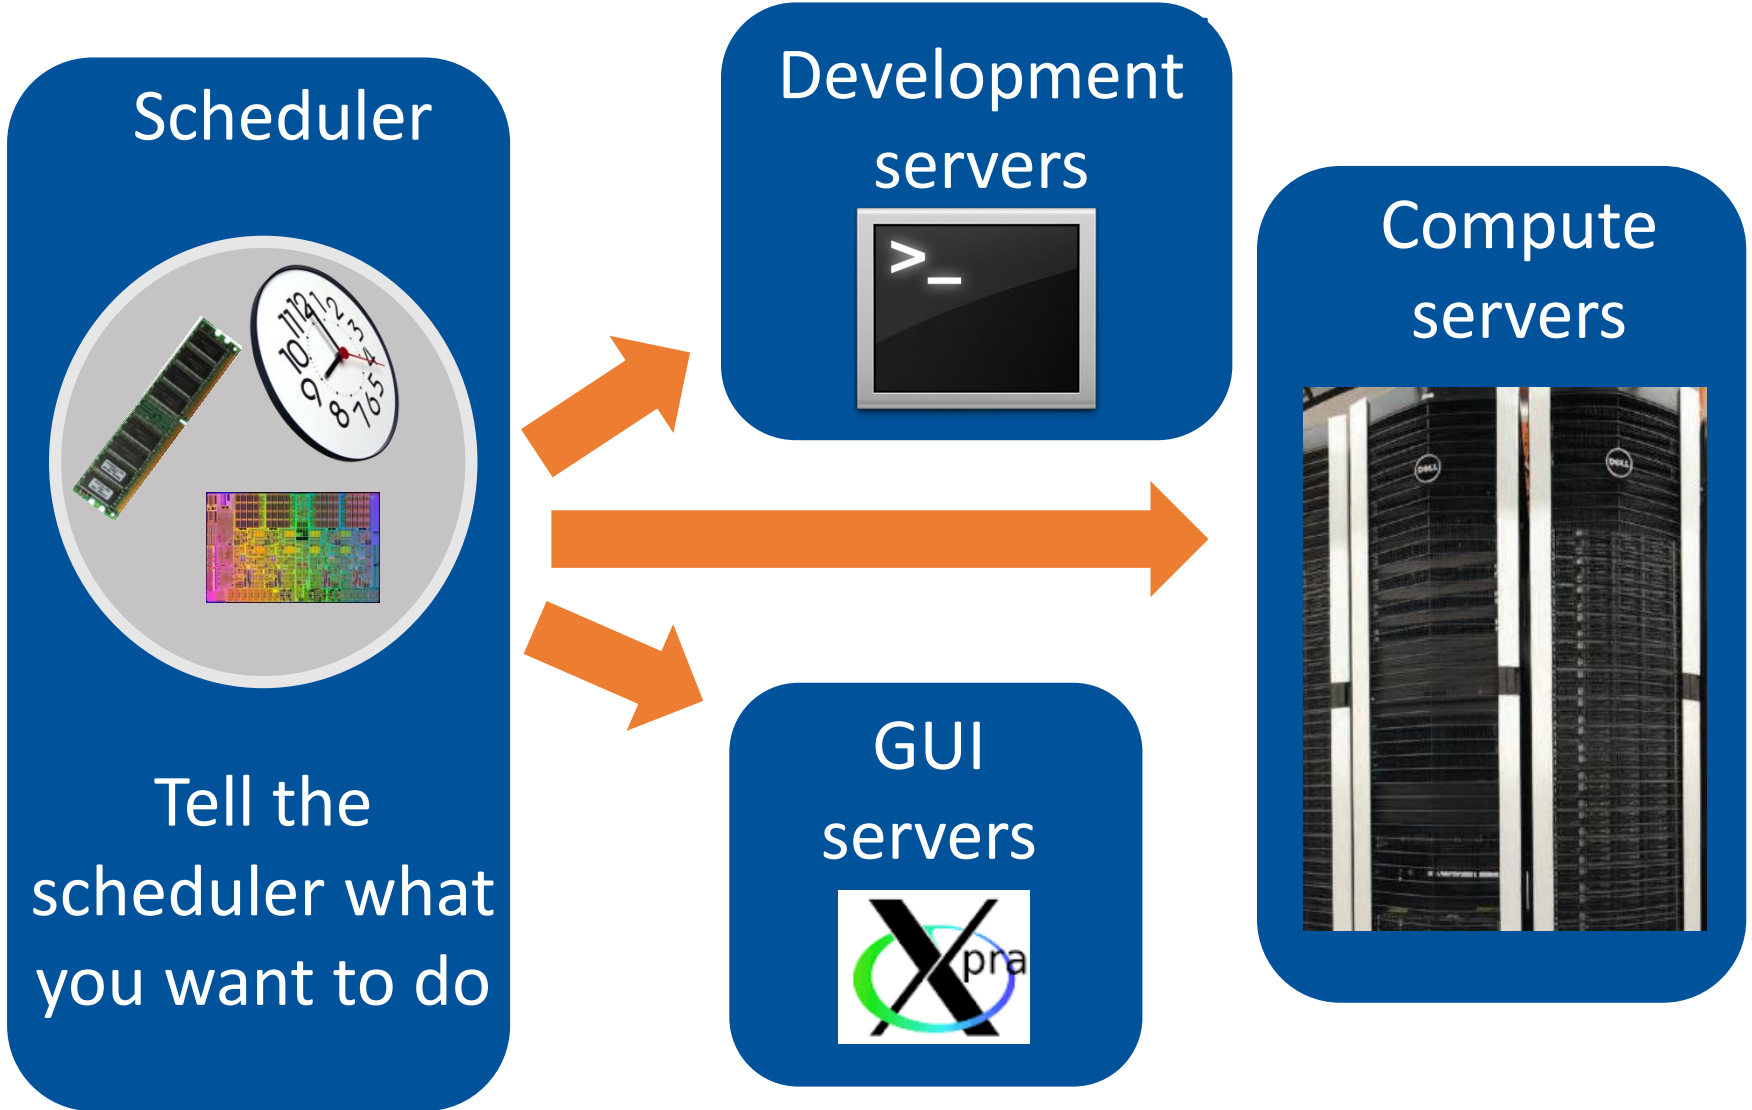
Tell the
scheduler what
you want to do

Compute
resources



Your job
runs on
the cluster


Cluster basics



Scheduling a job

- Need to tell scheduler what you want to do
 - **How many CPUs** you want and how you want them distributed
 - **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

Basic SLURM job script

```
#!/bin/sh
```

```
{ #SBATCH --ntasks=1           # Run on a single CPU
  #SBATCH --mem=1gb           # Memory limit
  #SBATCH --time=00:05:00     # Time: hrs:min:sec
  #SBATCH --job-name=serial_job_test # Job name
  #SBATCH --mail-type=ALL     # Mail events
  #SBATCH --mail-user=email_address # Where to send mail
  #SBATCH --output=serial_test_%j.out # Output and error log
```

```
pwd; hostname; date
```

```
module load python
```

```
echo "Running plot script on a single CPU core"
```

```
python /ufrc/data/training/SLURM/plot_template.py
```

```
date
```


SLURM CPU Requests

```
#!/bin/sh
```

```
#SBATCH --ntasks=1
```

- Nodes: `--nodes` or `-N`
 - Request a certain number of physical servers
- Tasks: `--ntasks` or `-n`
 - Total number of tasks job will use
- CPUs per task: `--cpus-per-task` or `-c`
 - Number of CPUs per task

SLURM CPU Requests

- For single processor jobs

`#SBATCH --nodes=1`

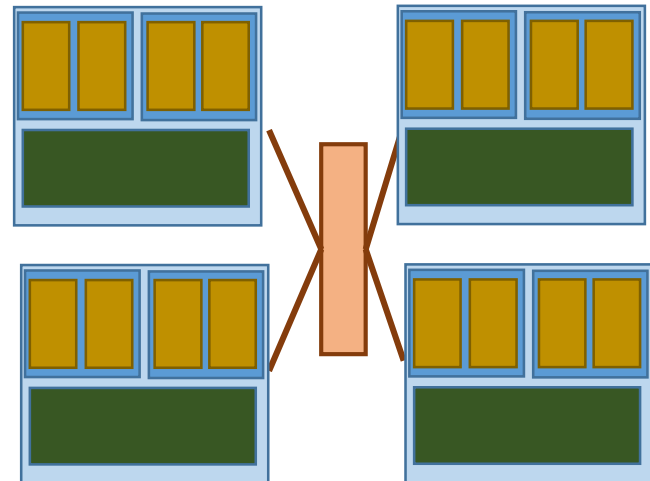

`#SBATCH --ntasks=1`

`#SBATCH --cpus-per-task=1`

} All
Optional

SLURM CPU Requests

- Parallel applications
 - OpenMP, Threaded, Pthreads
 - All cores on one sever, shared memory
 - MPI
 - Can use multiple servers
 - See:
https://wiki.rc.ufl.edu/doc/Sample_SLURM_Scripts



SLURM CPU Requests


- For threaded applications (single node):

```
#SBATCH --nodes=1
```

```
#SBATCH --ntasks=1
```

```
#SBATCH --cpus-per-task=8
```

} Optional



Node information

- HiPerGator 2.0 Compute Servers
 - 32 cores
 - Two 16-core sockets
- HiPerGator 1.0 Compute Servers
 - 64 cores
 - Four 16-core sockets

SLURM Memory Requests

- `--mem=1gb` (total memory)
- `--mem-per-cpu=1gb` (memory per core)
 - Can use mb or gb
 - No decimal values: use 1500mb, not 1.5gb

HiPerGator 2.0 Compute Servers:

- ~ 120 GB available RAM

HiPerGator 1.0 Compute Servers:

- ~ 250 GB available RAM

SLURM Time Request

- Time: `--time` or `-t`
 - 120 (minutes)
 - 2:00:00 (hh:mm:ss)
 - 7-0 (days-hours)
 - 7-00:00 (days-hh:mm)
 - 7-00:00:00 (days-hh:mm:ss)

Emails

Job ID: 94392

Cluster: hipergator

User/Group: magitz/ufhpc

State: COMPLETED (exit code 0)

Nodes: 1

Cores per node: 4

CPU Utilization: 00:00:44

CPU Efficiency: 52.38% of 00:01:24 core-walltime

Memory Utilization 1.52 MB

Memory Efficiency: 0.04% of 4.00 GB

Emails

```
Job ID: 5019
Cluster: hpg1
User/Group: magitz/ufhpc
State: CANCELLED (exit code 0)
Cores: 1
CPU Utilization: 00:00:00
CPU Efficiency: 0.00% of 00:00:00 core-walltime
Memory Utilization 1.26 MB
Memory Efficiency: 126.17% of 1.00 MB
```

Job error file:

```
slurmstepd: Job 5019 exceeded memory limit (1292 > 1024), being
killed
slurmstepd: Exceeded job memory limit
slurmstepd: *** JOB 5019 ON dev1 CANCELLED AT 2016-05-16T15:33:27
***
```

Quality of Service (--qos)

- Each group has two QOS options
 - Investment QOS:
 - The NCU's the group has purchased
 - `--qos=pre1234` (or leave off as this is default)
 - Burst QOS: **9X the allocation**
 - The burst capacity, available when idle resources are available on the cluster
 - `--qos=pre1234-b`
- Users can choose higher priority, or larger pool of resources

SLURM

- Note that multi-letter directives are double-dash:

- `--mail-type` `sbatch: error: distribution type 'ail-type=ALL' is not recognized`
- `--ntasks`
- `--mem-per-cpu`


- Do not use spaces with =

- `--mail-user=magitz@ufl.edu` ✓
- `--mail-user magitz@ufl.edu` ✓
- not: `--mail-user= magitz@ufl.edu`

Examples

- See example job scripts at:
help.rc.ufl.edu

Batch System

- SLURM Commands
 - Using Variables in SLURM
 - SLURM Job Arrays
 - Account and QOS Limits Under SLURM
 - GUI Partition
 - Big Memory Partition
 - Annotated Job Script Walk-through
 - Sample SLURM Scripts
 - UFRC environment module
- 

- Or: `/ufrc/data/training/SLURM/`

Submit your job

```
[magitz@login3 SLURM_examples]$ sbatch single_job.sh
```

```
Submitted batch job 30592170
```

```
[magitz@login3 SLURM_examples]$ squeue -u magitz
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST(REASON)
30592170	hpg2-comp	serial_j	magitz	R	0:30	1	c24b-s15

```
[magitz@login3 SLURM_examples]$
```

So what is this “module” thing?

- **Imod**—Implementation of Environment Modules developed at TACC
- Allows easy management of user’s environment



TEXAS ADVANCED COMPUTING CENTER

Powering Discoveries That Change The World



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

Lmod: Environmental Modules System

The standard way

```
PATH=$PATH:/some/long/path/to/application
```

```
export $PATH
```

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/long/path/to/pl  
ace/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



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@submit1 ~]$ 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 resources

<https://training.it.ufl.edu>

The screenshot shows the UF IT Training website. The top navigation bar includes links for NEWS, CALENDAR, OFFICES & SERVICES, DIRECTORY, GIVING, UF HEALTH, and UF IPAS. The main navigation bar includes TRAINING, E-LEARNING BASICS, SERVICES, CALENDAR, and CONTACT. The main content area features a section for 'INTRODUCTION TO RESEARCH COMPUTING AND HIPERGATOR'. It includes a 'REGISTER' button, a 'WHO'S IT FOR?' section listing 'TA / Grad Asst, Faculty, Staff', and a 'FORMAT' section listing 'Live Workshop, Recorded Webinar'. Below this is a video player showing a slide titled 'Introduction to Research Computing and HiPerGator' by Matt Gitzendanner (mugitz@ufl.edu). The slide also features the HiPerGator logo and the text 'UF Information Technology' and 'Gift for the #GatorGood'. The video player shows it is playing at 0:13 / 5:19. Below the video player is a section for 'CONNECTING TO HIPERGATOR' with a bullet point: 'This module covers connecting the HiPerGator with ssh and SFTP clients from Mac and'.

UF Research Computing

- Help and Support

- <http://help.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

