


Research Computing Orientation for UF Courses

Matt Gitzendanner
magitz@ufl.edu



UF | Information Technology

OneIT for the #GatorGood



Computational Resources for a
 Preeminent University

UF | Information Technology

OneIT for the #GatorGood



- 51,000 Cores
- 3PB Storage
- High-speed networks

UF | Information Technology

OneIT for the #GatorGood

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/account-request/>
- 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 | Information Technology


OneIT for the #GatorGood

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

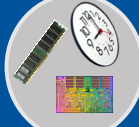
Cluster basics

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



SFTP client to move files
to / from your computer

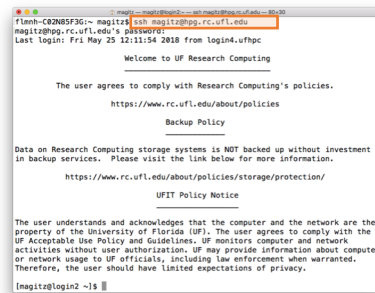


Text editor to prepare files

Especially on Windows, be sure to convert DOS
line breaks to Unix, and *don't use Word*
Both have SFTP built in



SSH Clients



Mac/Linux: Terminal



Windows:
MobaXterm or PuTTY

ssh user@hpg.rc.ufl.edu

MobaXterm

The screenshot shows the MobaXterm 'Session settings' dialog box. The 'Basic SSH settings' tab is active, showing a 'Remote host' of '192.168.1.1' and a 'Port' of '22'. The 'Advanced SSH settings' tab is also visible, showing 'Terminal settings' and 'Bookmark settings'. The main window displays a 'Secure Shell (SSH) session'.

UF Information Technology OneIT for the #GatorGood

PuTTY

The screenshot shows the PuTTY configuration dialog box with the 'Session' tab selected, showing 'Host Name (or IP address)' as '192.168.1.1' and 'Port' as '22'. A 'PuTTY Security Alert' dialog box is overlaid on top, warning that the server's host key is not cached in the registry and asking if the user wants to add it to the cache. The background shows the PuTTY terminal window.

UF Information Technology OneIT for the #GatorGood

FileZilla

The screenshot shows the FileZilla interface with two panes: 'Local site' and 'Remote site'. The local site shows a file system with folders like 'Desktop Junk', 'Downloads', and 'Public'. The remote site shows a file system with folders like 'C:\Program Files', 'C:\Program Files (x86)', and 'C:\Users\mag12'. A file list table is visible at the bottom.

Filename	Filesize	Filetype	Last modified	Permissions	Owner
Desktop Junk	Directory	03/30/2018 21:5...			
Downloads	Directory	04/20/2018 16:2...			
Public	Directory	12/12/2017 02:4...			
123.docx	10,244	File	03/20/2018 16:5...		
Justified	0	File	01/28/2017 18:2...		
FontColor.pdf	10,716	pdf file	01/20/2018 16:5...		
GitAndDocker...	1,837,895	pdf file	03/30/2018 14:5...		
INTRO-TO-LI...	25,117	pdf file	04/04/2018 20:...		

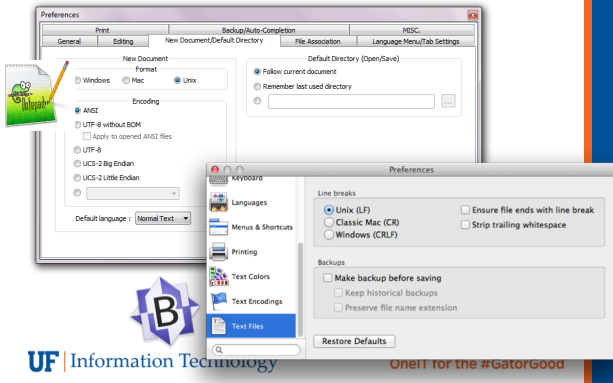
UF Information Technology OneIT for the #GatorGood

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

Unix line breaks



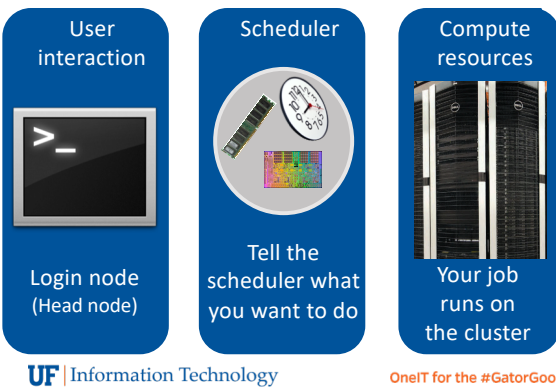
Storage



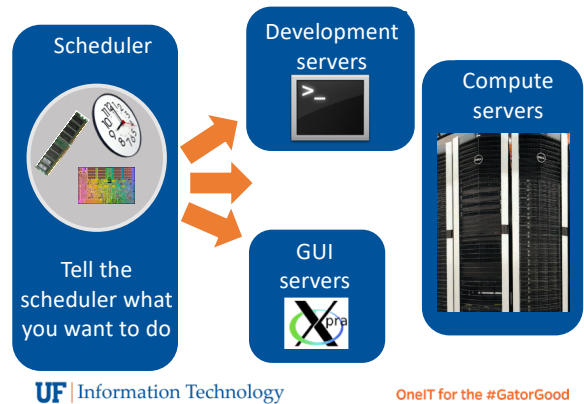
- /home/\$USER
 - 20GB limit
 - scripts, code, small data
 - Do NOT use for job input/output
- /ufrc/pre1234/\$USER
 - 2TB limit per group
 - ALL input/output from jobs should go here

- 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




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



Basic SLURM job script

```
#!/bin/sh
#SBATCH --job-name=serial_job_test # Job name
#SBATCH --mail-type=ALL # Mail events
#SBATCH --mail-user=<email_address> # Where to send mail
#SBATCH --ntasks=1 # Run on a single CPU
#SBATCH --mem=1gb # Memory limit
#SBATCH --time=00:05:00 # Time: hrs:min:sec
#SBATCH --output=serial_test_1j.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
```

SLURM CPU Requests

- 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




UF | Information Technology

OneIT for the #GatorGood

SLURM CPU Requests

- For threaded applications (single node):

```
#SBATCH --nodes=1 } Optional
#SBATCH --ntasks=1 }
#SBATCH --cpus-per-task=8
```



UF | Information Technology

OneIT for the #GatorGood

Node information

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

UF | Information Technology

OneIT for the #GatorGood

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

UF | Information Technology

OneIT for the #GatorGood

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
***

```

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)

Quality of Service (--qos)

- Each group has two QOS options
 - Investment QOS:
 - The NCUs 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`

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 Powering Discoveries That Change The World

THE UNIVERSITY OF TEXAS AT AUSTIN

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

UF | Information Technology **Lmod: Environmental Modules System** OneIT for the #GatorGood

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
UF Information Technology
```

OneIT for the #GatorGood

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`

UF Information Technology

OneIT for the #GatorGood

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/>
 - Frequently Asked Questions
 - Account set up and maintenance



UF Information Technology

OneIT for the #GatorGood