


Research Computing Orientation for UF Courses

Matt Gitzendanner
magitz@ufl.edu



UF | Information Technology

OneIT for the #GatorGood

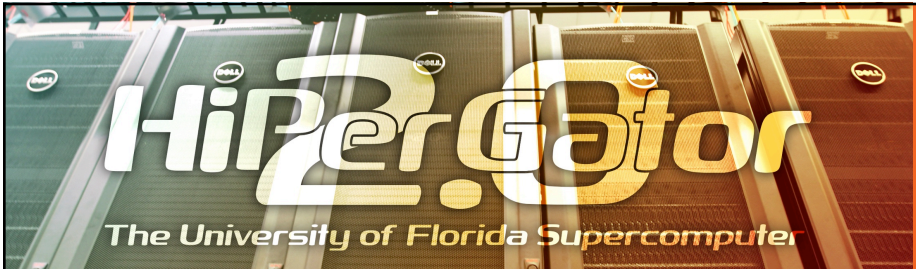


UF Preeminence
INVESTING IN PEOPLE AND PROGRAMS
THAT HELP UF HELP THE WORLD





UF | Information Technology

OneIT for the #GatorGood



HiPerGator
The University of Florida Supercomputer

- 30,000 cores
- #136 in June 2016 Top 500 
 - Top 5 of US public universities
 - HPL RMAX 738 TFlops
- Dell PowerEdge
- Intel Xeon E5-2698v3 2.3GHz CPUs 

UF | Information Technology OneIT for the #GatorGood

Course use of HiPerGator 2.0

- 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

For users with an account

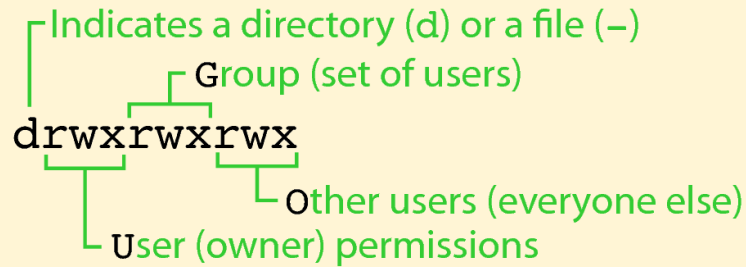
- If you already have a Research Computing account for research:
- **module load class/abe6933**
 - Ensures that jobs run under course allocation
 - Storage is under course allocation

ITAR / EAR / RSICC

- Known controlled software installed:
 - MCNP/MCNPX
 - Serpent
 - Scale
 - Vulcan
- Proof of certification is required to gain access
 - Send proof of certification and email from Dr Goluoglu granting access to support@hpc.ufl.edu
- All code and output data should be in:
/spec/ENU4192/itar/\$USER
- Do not provide access to other individuals

Linux permissions

- Make sure your /scratch/lfs and /home directories remain `drwx-----`



Galaxy: Data intensive biology *for everyone*

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

Galaxy / UFRC

UF Research Computing Information Technology

Galaxy Documentation:
 UFRC Galaxy documentation on the Wiki.
 Large dataset import procedure.
 Tool PBS resource request reference table.
 Log of UF Galaxy changes, fixes and upgrades.
 Known Galaxy problems.

The Galaxy project is supported in part by NSF, NHGRI, and the Huck Institutes of the Life Sciences.

History

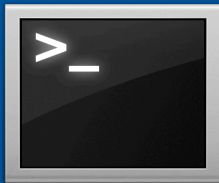
Job ID	Job Name	Status
26	FASTQ Quality Trimmer on data 5	Success
24	FASTQ Quality Trimmer on data 5	Success
23	SPAdes log	Success
22	SPAdes scaffold stats	Success
21	SPAdes scaffolds (fasta)	Success
20	SPAdes contig stats	Success
19	SPAdes contigs (fasta)	Success
18	vevtp on data 14_C ontigs	Success
17	vevtp on data 14_S fasts	Success
16	vevtp on data 13_C ontigs	Success
15	vevtp on data 13_S fasts	Success
14	vevtp on data 8 km er.51	Success
13	vevtp on data 8 km er.21	Success

Galaxy demo

<http://galaxy.rc.ufl.edu>

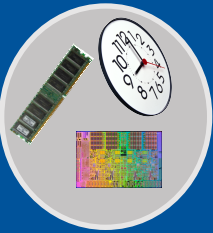
Cluster basics

User
interaction




Login node
(Head node)

Scheduler



Tell the
scheduler what
you want to do

Compute
resources



Your job
runs on
the cluster

UF Information Technology

OneIT for the #GatorGood

Tools

ssh client to connect to
hpg2.rc.ufl.edu



e.g.: Terminal, MobaXterm

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++

SSH Clients

```
ufit-rc-matt:~ matt$ ssh magitz@hpg2.rc.ufl.edu
magitz@hpg2.rc.ufl.edu's password:
Last login: Thu May 26 11:17:26 2016 from 10.243.21.61
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
```

`ssh user@hpg2.rc.ufl.edu`




Mac/Linux: Terminal




Windows:
MobaXterm or PuTTY

MobaXterm



UF


PuTTY



UF | Information Technology

OneIT for the #GatorGood

FileZilla



Host:

sftp.rc.ufl.edu

*Do not use hpg2 login
serve for data
transfer*

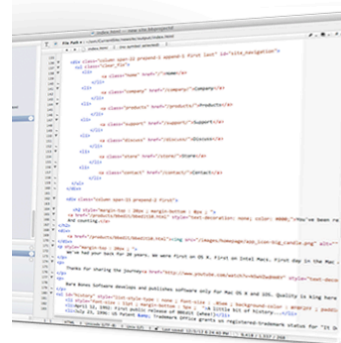
UF | Information Technology

OneIT for the #GatorGood

Text Editors



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

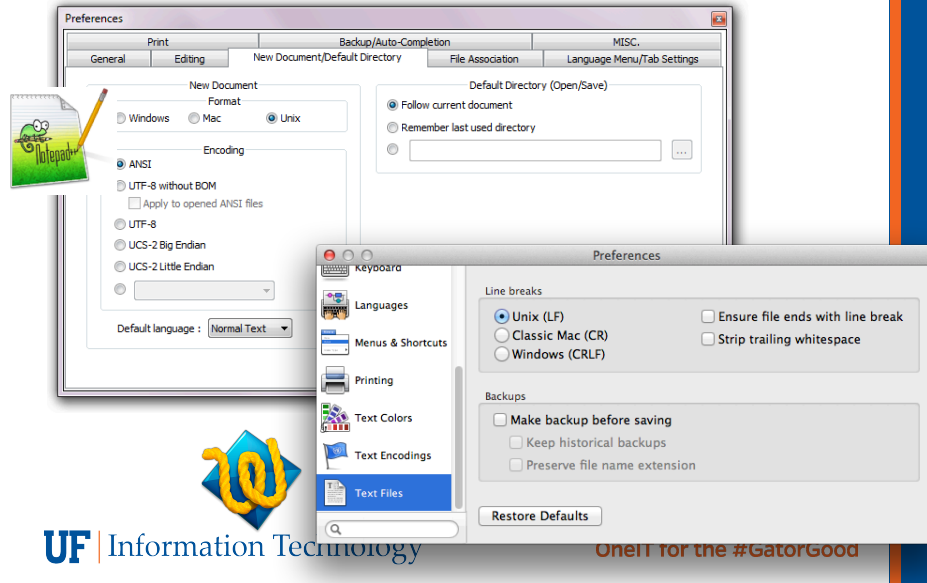


- **Unix line breaks**



UF | Information Technology

Unix line breaks



Storage



- `/home/$USER`
 - 20GB limit
 - scripts, code, small data
 - Do NOT use for job input/output
- `/ufrc/abe6933/$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

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
 - /spec/ENU4192/itar/\$USER
 - For ITAR codes and data
- 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 Information Technology

OneIT for the #GatorGood


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

UF Information Technology

OneIT for the #GatorGood

Development servers

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

```
Welcome to the UF HPC Center.
```

Do not run interactive jobs on the login nodes.

```
http://www.hpc.ufl.edu/about/policies/account
[magitz@gator3 ~]$
```

- Request job in development partition, hpg2-dev
 - **module load ufrc**
 - **srundev -t 120**

UF | Information Technology

OneIT for the #GatorGood



GUI Servers

- gui1.rc.ufl.edu or gui.rc.ufl.edu


```
ssh -Y user@gui1.rc.ufl.edu
```

On Windows, use MobaXterm or PuTTY w/ Xming

- dev1 and dev2


```
ssh -Y user@gator.rc.ufl.edu
ssh -Y dev1
```
- Interactive job



```
ssh -Y user@gator.rc.ufl.edu
qsub -X -I -l
nodes=1:ppn=4,pmem=1gb,walltime=1:00:00
```

UF | Information Technology

OneIT for the #GatorGood

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/bash
#SBATCH --job-name=test      #A name for your job
#SBATCH -o job_%j.out       #Name output file
#SBATCH --mail-type=ALL     #What emails you want
#SBATCH --mail-user=<Email address> #Where
# } #SBATCH --ntasks=1       #Optional—single CPU
# } #SBATCH --mem-per-cpu=100mb #Per processor memory
# } #SBATCH -t=00:01:00      #Walltime in hh:mm:ss
#                               #or d-hh:mm:ss

hostname
module load python
python -v
```


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



HiPerGator 2.0 Compute Servers:
 • 32 cores (2 X 16-core Intel Xeon CPUs)

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



UF | Information Technology

OneIT for the #GatorGood

SLURM CPU Requests


- For threaded applications (single node):

`#SBATCH --nodes=1`

`#SBATCH --ntasks=1`

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

} Optional



UF | Information Technology

OneIT for the #GatorGood

SLURM CPU Requests


- For MPI jobs

```
#SBATCH --ntasks=16
```

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

```
#SBATCH --ntasks-per-socket=8
```

```
#SBATCH --distribution=cyclic:cyclic
```



UF | Information Technology

OneIT for the #GatorGood

SLURM CPU Requests


- For MPI jobs

```
#SBATCH --ntasks=48
```

```
#SBATCH --nodes=4
```

```
#SBATCH --ntasks-per-socket=6
```

```
#SBATCH --distribution=cyclic:cyclic
```




UF | Information Technology

OneIT for the #GatorGood

SLURM CPU Requests

- For Hybrid MPI/OpenMP jobs
 - #SBATCH --ntasks=4 (4 MPI ranks)
 - #SBATCH --cpus-per-task=4
 - #SBATCH --nodes=4
 - #SBATCH --ntasks-per-socket=1
 - #SBATCH --distribution=cyclic:cyclic



SLURM Memory Requests

- Non-MPI jobs:
 - --mem=2gb
 - *Total* memory for the job
- MPI Jobs:
 - --mem-per-cpu=1gb
 - *Per-core* memory
 - Most RAM that any MPI rank will use

SLURM Memory Requests

- Can use **gb** or **mb**
- No decimal values

HiPerGator 2.0 Compute Servers:


- 128 GB total RAM (vs 256 GB on HPG1)
- Diskless servers: OS takes ~8GB 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)

SLURM Time Request

- Limits:
 - Investment QOS: 31 days
 - Burst QOS: 4 days
 - Dev partition: 12 hours
 - GUI partition: 96 hours



As with all resource requests, providing a reasonably accurate request ensures best results

Quality of Service (--qos)

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

SLURM output/error files

```
#SBATCH -o output.file
```

```
#SBATCH -e error.file
```

```
#SBATCH -o output.file #W/o -e  
combined
```

- Can also use `--output` and `--error`

```
#SBATCH --output JobFile.%j.out
```

- Use `%j` instead of `$SLURM_JOBID`

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](#)

Lmod: Environmental Modules System

 Information Technology
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
```

Multiple variants of a version

```
[magitz@submit1 ~]$ module spider mrbayes/3.2.1
Rebuilding cache file, please wait ... Done
```

```
-----
mrbayes: mrbayes/3.2.1
-----
```

```
-
Description:
```

```
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:
```

UF | Information Technology
intel/2012, openmpi/1.6

OneIT for the #GatorGood

Module loading

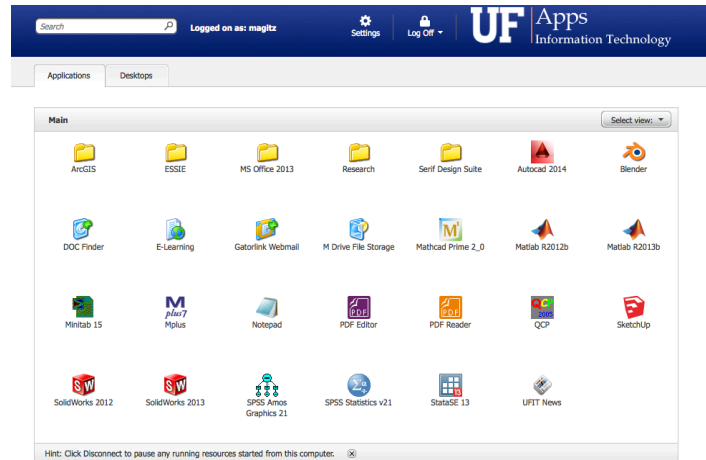
- module load raxml
- module load intel raxml
- module load intel openmpi raxml
- module load intel/12 openmpi/1.6
raxml/3.2
- module unload raxml

UF | Information Technology

OneIT for the #GatorGood

UF Apps

- <https://apps.ufl.edu>

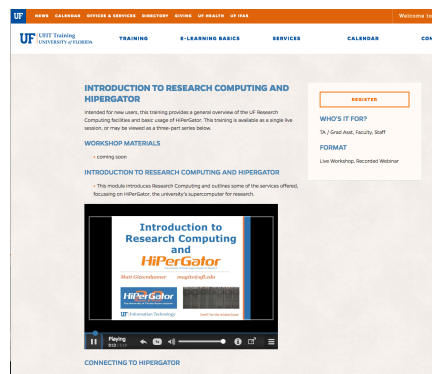


UF | Information Technology

OneIT for the #GatorGood

Training sessions

- Thursdays@ 10:40 and online
<https://training.it.ufl.edu>



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

