

# Introducing



*Matt Gitzendanner*     *magitz@ufl.edu*



**UF** | Information Technology

OneIT for the #GatorGood

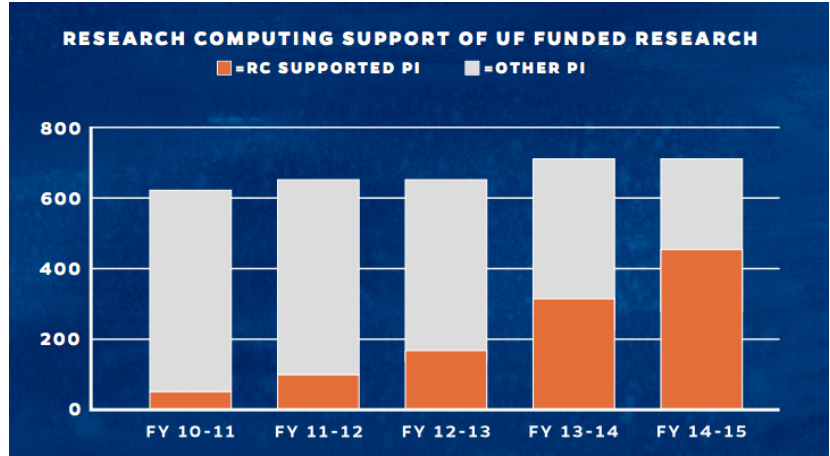
## UNIVERSITY OF FLORIDA'S 2015-2020

# STRATEGIC GOALS FOR IT

One IT for the GatorGood

- Research Computing 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

# Research Computing: Grants Supported



UF | Information Technology

OneIT for the #GatorGood


**HiPerGator**  
The University of Florida Supercomputer



UF | Information Technology

OneIT for the #GatorGood

## Research Computing




Where do you start?

UF | Information Technology

OneIT for the #GatorGood

## Research Computing



- User Accounts
  - Need current UF faculty sponsor
  - GatorLink ID
    - We can assist in creating GatorLink IDs for collaborators

[www.rc.ufl.edu](http://www.rc.ufl.edu)




UF | Information Technology

OneIT for the #GatorGood

## Investor Supported

Investment Option	Total Cost	University Match	Cost to PIs (2016)
Long-Term 1 NCU/5 Years	\$300	\$100	\$200
Billed 1NCU/Hour	\$0.04	-	\$0.04
GPU NVIDIA K80/5 Years	\$2,250	\$1,750	\$500
Storage 1TB/Year	\$125	-	\$125
Replicated Storage 1TB/Year	\$250	-	\$250

## What can you run?



- Linux-based
- Generally command line driven applications



# Galaxy

# wiki.rc.ufl.edu

Log in

Page Discussion
Read View source View history

Go

## UF Research Computing Wiki

Welcome to the University of Florida Research Computing Wiki. The information here augments what is found on our web site with information we think is best provided in a less formal and official context. It is used for information that changes rapidly and might become quickly dated or incorrect on the web site. While there is good and helpful information here, you will find text that has not been proofed to the same standards that we try to maintain on our web site. We hope you find it helpful.

**HiPerGator 2.0 Information**

- HiPerGator 2.0 Information
- SLURM Documentation
- Sample SLURM Scripts
- Moab (PBS) to SLURM command reference

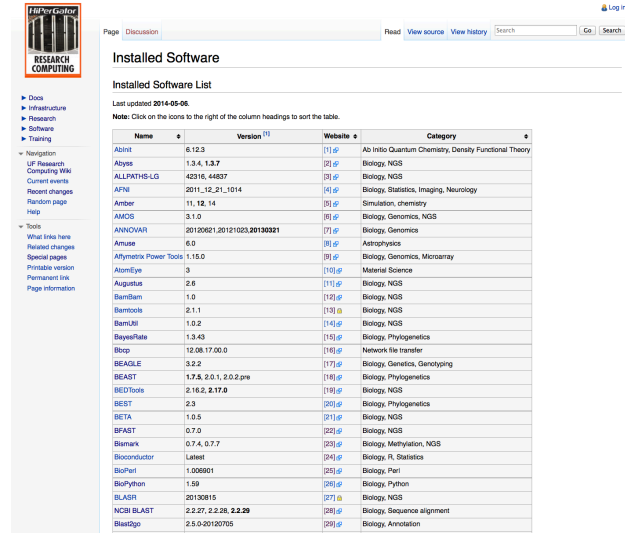
**Getting Started**

- Getting Started
- Mailing Lists
- Changing your Password
- SLURM Commands
- Events\_Calendar
- Open a Support Request
- Non-Batch System Resources
- Training
- View Job Status

**Software and Libraries**

- Installed Software
- Environment Modules System
- GPUs and CUDA

## Wiki: Installed Software



Page Discussion Read View source View history Search Go Search

### Installed Software

Installed Software List

Last updated **2014-05-06**.

Note: Click on the icons to the right of the column headings to sort the table.

Name	Version	Website	Category
Abinit	6.12.3	[1]	Ab Initio Quantum Chemistry, Density Functional Theory
Abys	1.3.4, 1.3.7	[2]	Biology, NGS
ALLPATHS-LG	42316, 44827	[3]	Biology, NGS
AFNI	2011_12_21_1014	[4]	Biology, Statistics, Imaging, Neurology
Amber	11, 12, 14	[5]	Simulation, chemistry
AMOS	3.1.0	[6]	Biology, Genomics, NGS
ANNOVAR	20130621,20121023,20130321	[7]	Biology, Genomics
Armut	6.0	[8]	Astrophysics
Atiyeh's Power Tools	1, 15.0	[9]	Biology, Genomics, Microarray
AtomEye	3	[10]	Material Science
Augustus	2.6	[11]	Biology, NGS
BamBam	1.0	[12]	Biology, NGS
Bamtools	2.1.1	[13]	Biology, NGS
BamLRI	1.0.2	[14]	Biology, NGS
BayesRate	1.3.43	[15]	Biology, Phylogenetics
Bcp	12.08.17.00.0	[16]	Network file transfer
BEAGLE	3.2.2	[17]	Biology, Genetics, Genotyping
BEAST	1.7.5, 2.0.1, 2.0.2.pre	[18]	Biology, Phylogenetics
BEDTools	2.16.2, 2.17.0	[19]	Biology, NGS
BEST	2.3	[20]	Biology, Phylogenetics
BETA	14.5.5	[21]	Biology, NGS
BFAST	0.7.0	[22]	Biology, NGS
Bismark	0.7.4, 0.7.7	[23]	Biology, Methylation, NGS
Bioconductor	Laneet	[24]	Biology, R, Statistics
BioPerl	1.006001	[25]	Biology, Perl
BioPython	1.59	[26]	Biology, Python
BLAST	20130819	[27]	Biology, NGS
NCBI BLAST	2.2.37, 2.2.38, 2.2.29	[28]	Biology, Sequence alignment
BlastSp	2.5.0-20130705	[29]	Biology, Annotation

UF | Information Technology

OneIT for the #GatorGood

## Getting help

▪ [my.it.ufl.edu](http://my.it.ufl.edu)

▪ [support@rc.ufl.edu](mailto:support@rc.ufl.edu)

▪ UF Help Desk

my 



Research Services

UF | Information Technology

OneIT for the #GatorGood

# HiPerGator

The University of Florida Supercomputer



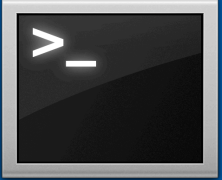
#GATORGOOD

UF | Information Technology

OneIT for the #GatorGood

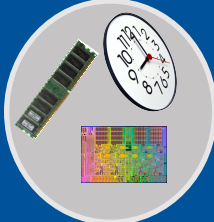
## Cluster basics

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

## 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 Information Technology

OneIT for the #GatorGood


# PuTTY



UF Information Technology

OneIT for the #GatorGood

## Linux Command Line

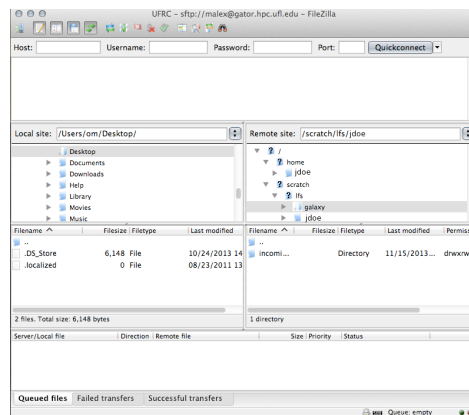


- Lots of online resources
  - Google: Linux cheat sheet
- Training sessions
  - June 8: The Linux/Unix Command Line--An Introduction
- User manuals for applications

UF | Information Technology

OneIT for the #GatorGood

## FileZilla



Host:  
**sftp.rc.ufl.edu**

*Do not use hpg2 login  
serve for data  
transfer*

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

UF | Information Technology

OneIT for the #GatorGood




## Storage



- **/home/<username>**
  - 20GB limit
  - scripts, code, small data
  - **Do NOT use for job input/output**
  
- **/ufrc/<groupname>/<username>**
  - 2TB limit per group
  - **ALL input/output** from jobs should go here

## 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



## 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 core 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
  - **--ntasks=1 (or omit)**
- For parallel jobs on a single node:
  - **--cpus-per-task=8**

## SLURM CPU Requests

- For MPI jobs
  - **--ntasks=32**
  - Gets 32 cores for 32 MPI ranks
  - SLURM will determine node layout
- For Hybrid MPI/OpenMP jobs
  - **--ntasks=4** (4 MPI ranks)
  - **--cpus-per-task=8**
  - **(--nodes=4)** Not needed unless you really want 4 different nodes

## SLURM Memory Requests

- Memory: `--mem-per-cpu=1gb`
  - Can use mb or gb
  - No decimal values: use 1500mb, not 1.5gb

### 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 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 -o 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`

## 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 Task Arrays

- **#SBATCH --array=1-200%10**
- Similar to Moab: range with % to limit number of jobs at a time
- **\$SLURM\_ARRAY\_TASK\_ID**
- Output file naming:
  - %A: job id
  - %a: task id
  - Output.%A\_%a.out



## SLURM environment

- SLURM inherits your environment
  - This includes present working directory
    - Don't need `cd $SLURM_SUBMIT_DIR`
  - Modules that are loaded
- **Be careful of conflicting modules**

## 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
***

```

**UF** | Information Technology

OneIT for the #GatorGood

## Development sessions

- Either:
  - `module load ufrc`
  - Followed by
    - `srundev`
    - `srundev -t 60:00`
- Or
  - `srun -p hpg2-dev --pty -u bash -i`
  - `srun -p hpg2-dev -t 60:00 --pty -u bash -i`

**UF** | Information Technology

OneIT for the #GatorGood

## Checking on jobs

- `squeue`
- `sacct`
  
- See [wiki.rc.ufl.edu/doc/SLURM\\_Commands](http://wiki.rc.ufl.edu/doc/SLURM_Commands)
  
- See <http://slurm.schedmd.com/>

## Example files

```
cd /ufrc/group/user/  
mkdir SLURM_examples  
cd SLURM_examples  
cp /ufrc/data/training/SLURM/*.sbatch .
```



# Satisfaction Survey

- [training.it.ufl.edu](http://training.it.ufl.edu)

The screenshot shows the UF FIT Training website. The top navigation bar includes links for NEWS, CALENDAR, OFFICES & SERVICES, DIRECTORY, GIVING, UF HEALTH, and UF IFAS. The main navigation menu has tabs for TRAINING, CANVAS BASICS, SERVICES (which is underlined), and CALENDAR. Below the navigation, there are links for UF Computing Help Desk and Contact Us. A prominent link for the Satisfaction Survey is displayed. A banner below the navigation states "UFIT Training provides an extensive catalog of" followed by a "NEW AND UPDATED" button. The footer contains the UF Information Technology logo and the slogan "OneIT for the #GatorGood".

## Next Week:

- Open Q&A session
  - 11:00am
  - NPB 2205

The slide content is simple, with a blue vertical bar on the right side. The footer is identical to the screenshot above, featuring the UF Information Technology logo and the slogan "OneIT for the #GatorGood".

# Support

- Support requests



- [Web page](#) and [wiki](#)

