Matt Gitzendanner

magitz@ufl.edu

UFIT Research Computing Orientation for Courses


HiPerGator
4TH Gen
UNIVERSITY OF FLORIDA
HiPerGator AI
UNIVERSITY OF FLORIDA

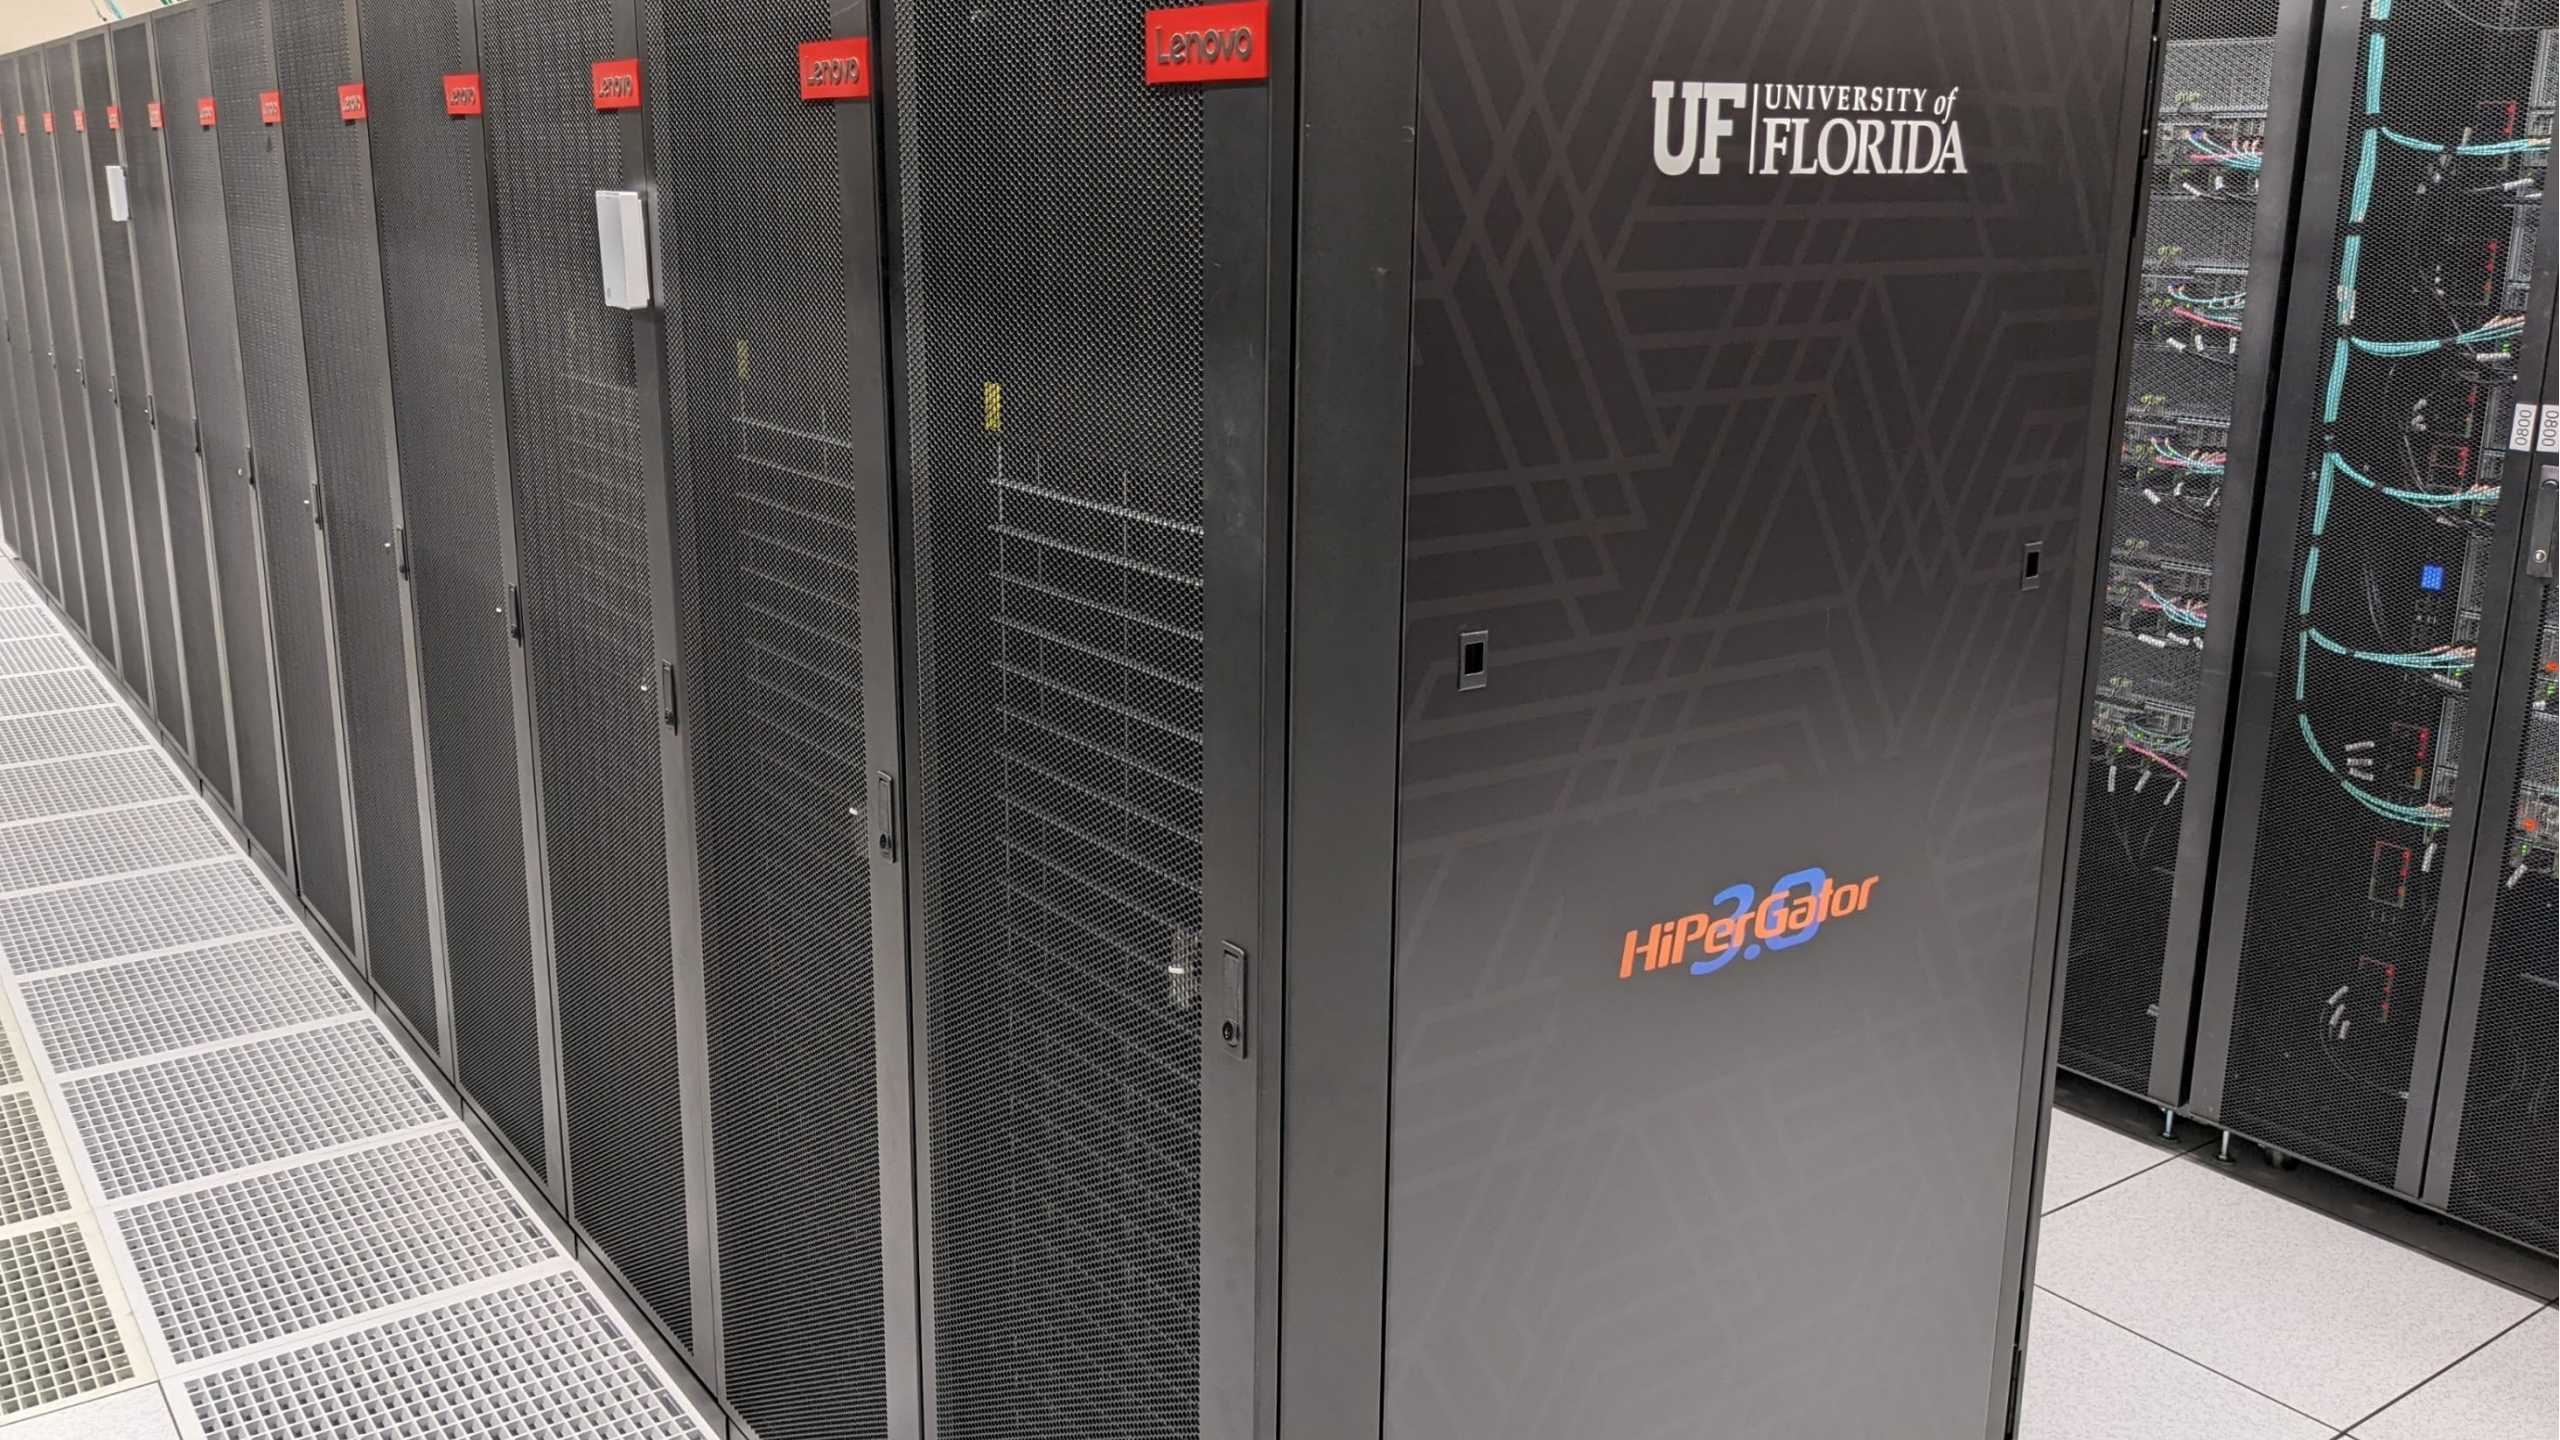
Lenovo
UF UNIVERSITY of FLORIDA
HiPerGator 3.0

Over 60,000 cores

Over 1,000 GPUs

Over 30 Petabytes of storage

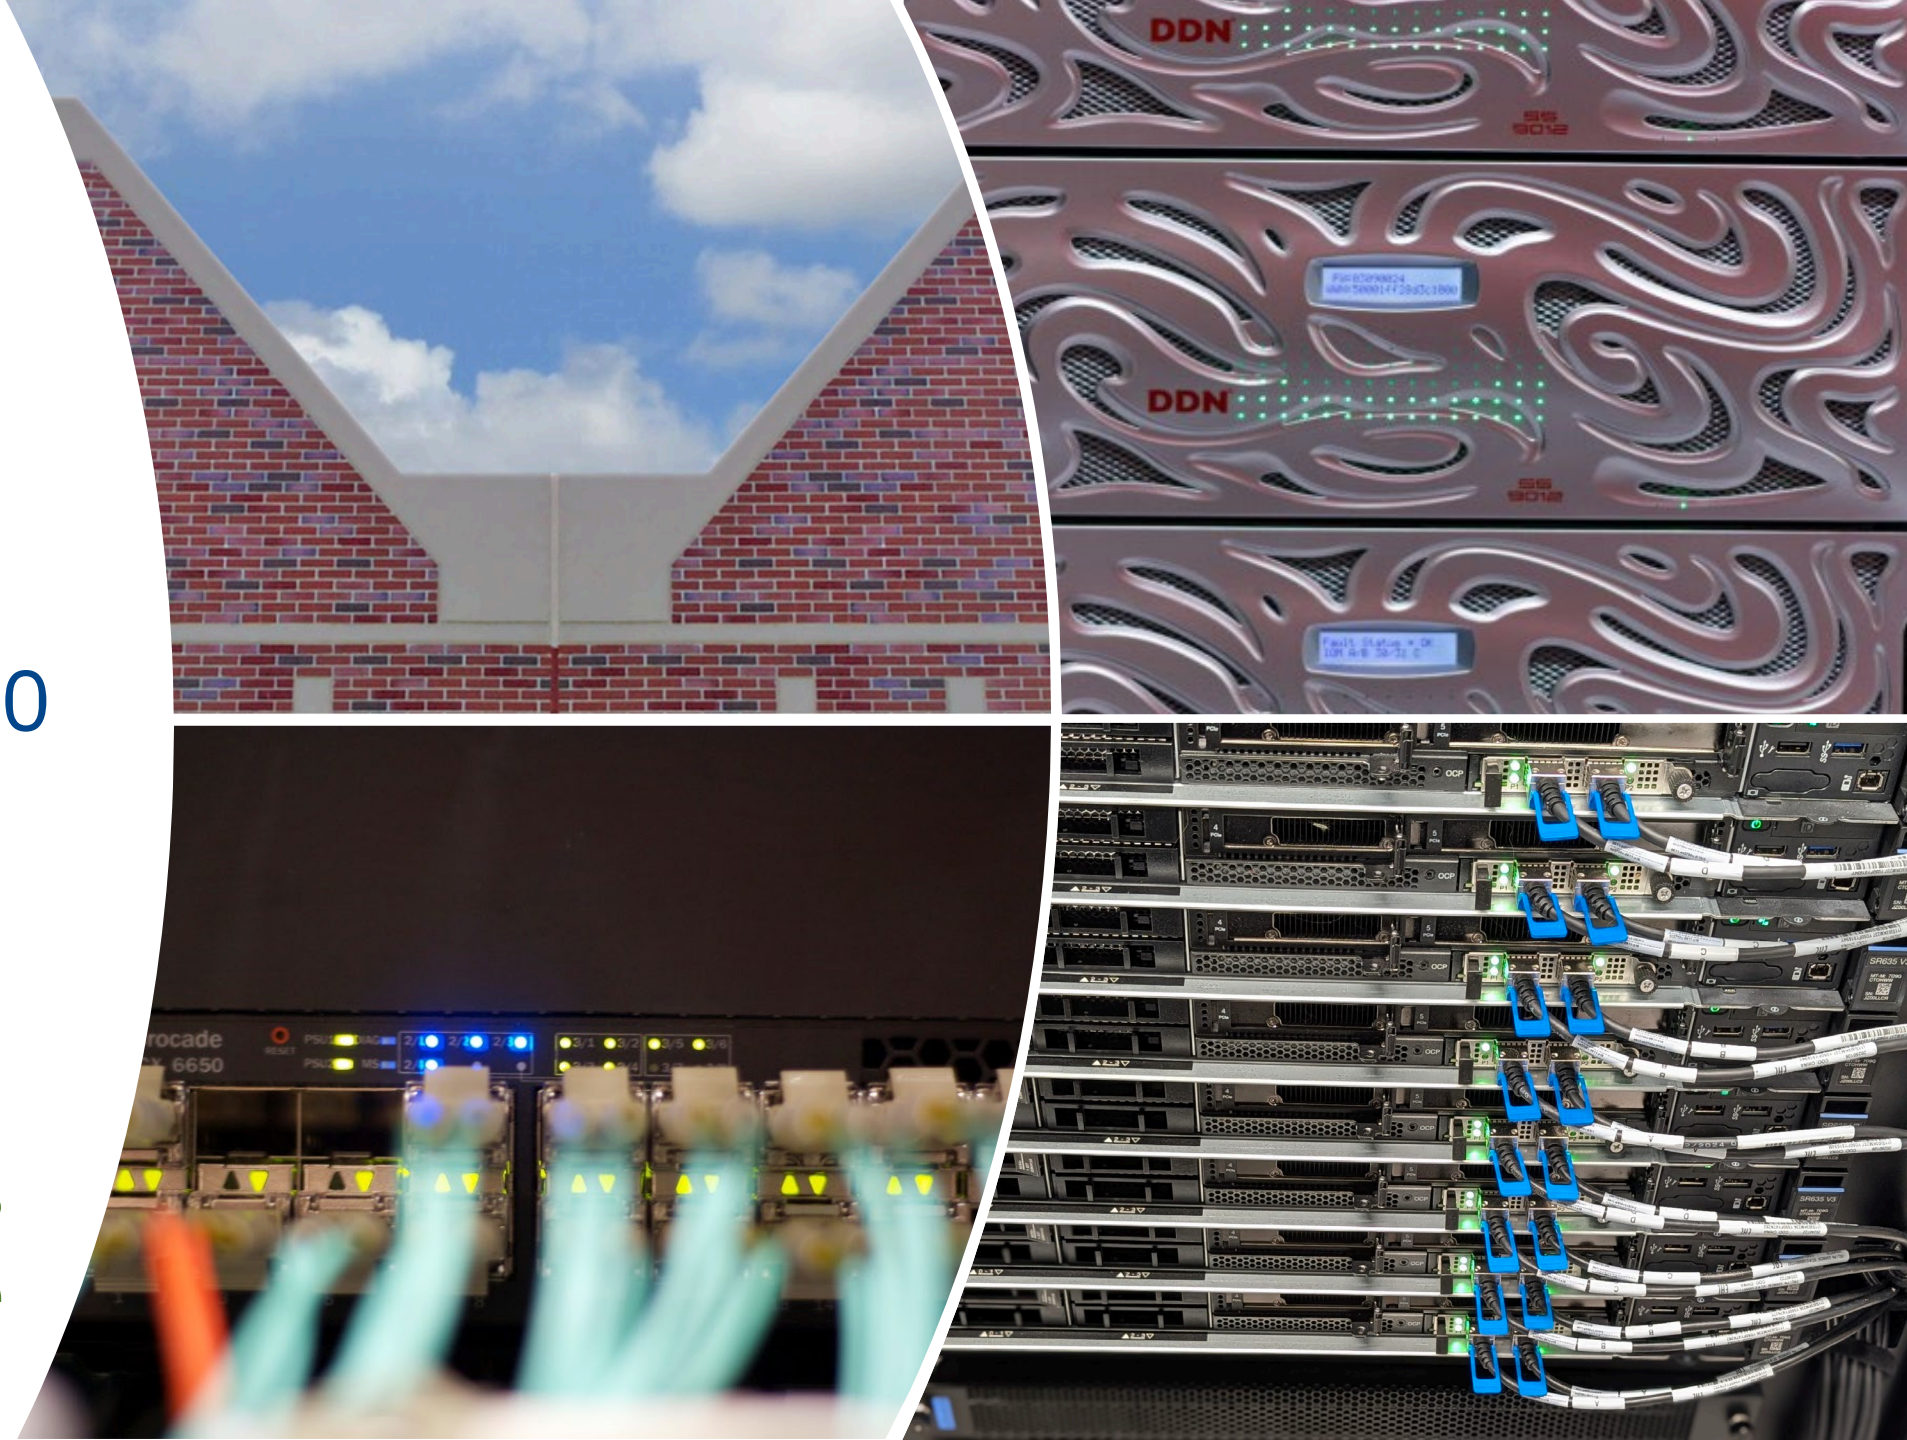

Course use of HiPerGator

The University of Florida Supercomputer for Research

- **Course is allocated 32 cores, 256GB RAM, 2TB Blue storage, GPUs as needed**
  - Design projects with this in mind
  - Time your work with this in mind
  - Use resources efficiently
- **Support requests should go through the course TA**
  - If the 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

HiPerGator Account Training

- Content and links at: https://docs.rc.ufl.edu/training/new_user_training/
  - Page also has additional information for classes at the end

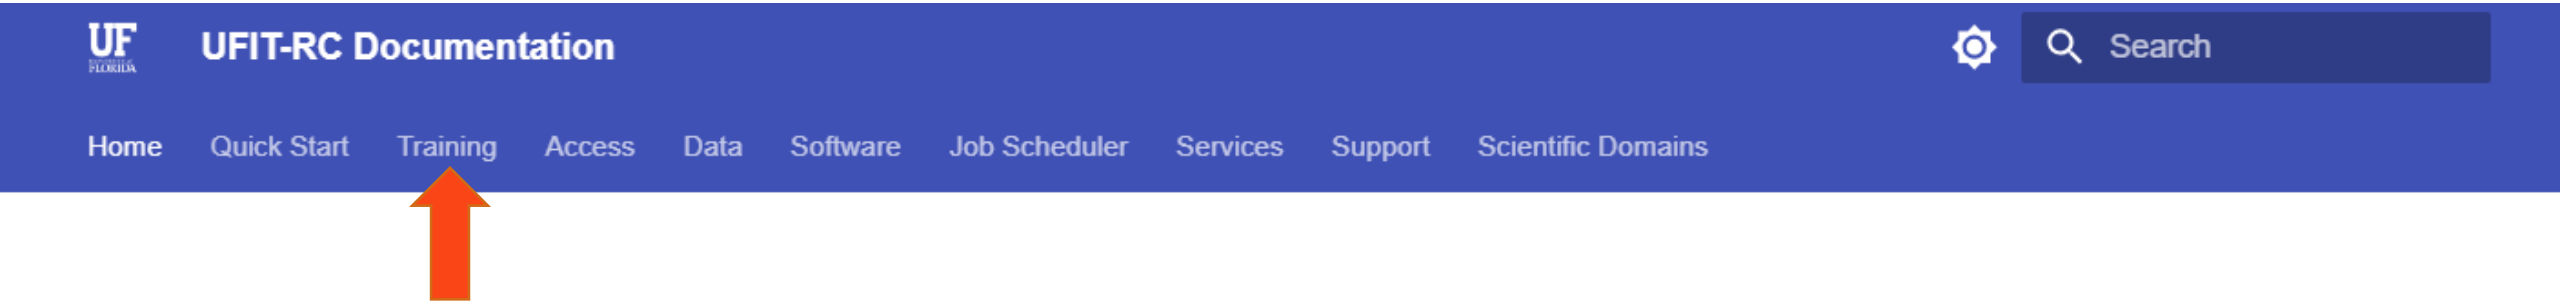

UFIT Research Computing User Documentation

https://docs.rc.ufl.edu/

Welcome to the University of Florida Information Technology Research Computing User Documentation and HiPerGator Compendium. This is a refresh of our original Help Wiki Site.

General information about RC, news, announcements, policies, purchase request and account request forms are located on our main website. All other information you may need to become a proficient and effective HiPerGator (HPG) user, including any extra information about using particular applications on HPG, available reference data, services, and code examples is located in this compendium.

For users with an account

- If you already have a HiPerGator account:
  - Make a folder for yourself at **`/blue/pre1234/<gatorlink>`**
  - When submitting jobs, add:
    `--account=pre1234 --qos=pre1234`

Cluster Overview

User interaction

Login node or web interface

Slurm Scheduler

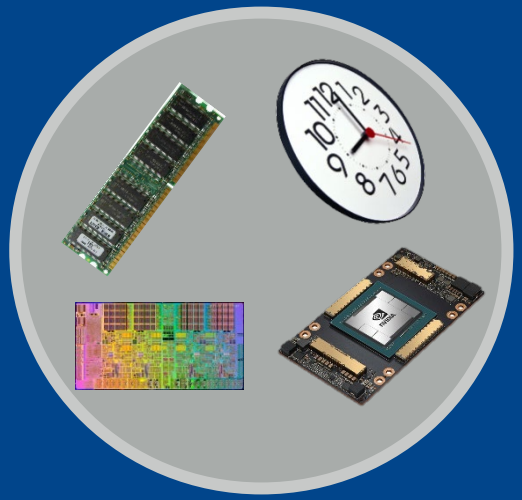

Tell Slurm what you want to do

Compute resources

Your job runs on the cluster

SSH Clients

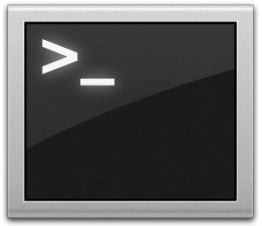

Mac/Linux/Windows: **Terminal**

Windows: Git Bash, PuTTY, Bitvise

```
magitz@login1:~
$ ssh magitz@hpg.rc.ufl.edu
(magitz@hpg.rc.ufl.edu) Password:
(magitz@hpg.rc.ufl.edu) Duo two-factor login for magitz@ufl.edu

Enter a passcode or select one of the following options:

 1. Duo Push to XXX-XXX-4066
 2. Phone call to XXX-XXX-4066
 3. Phone call to XXX-XXX-1960

Passcode or option (1-3): 1
Success. Logging you in...
Last login: Mon Nov  8 08:31:41 2021 from 10.138.154.11

                    Welcome to UF Research Computing
                    ________________________________

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

                https://www.rc.ufl.edu/services/procedures/

                              Backup Policy
                              _____________
```

ssh user@hpg.rc.ufl.edu

UF Information Technology UNIVERSITY of FLORIDA

ood.rc.ufl.edu

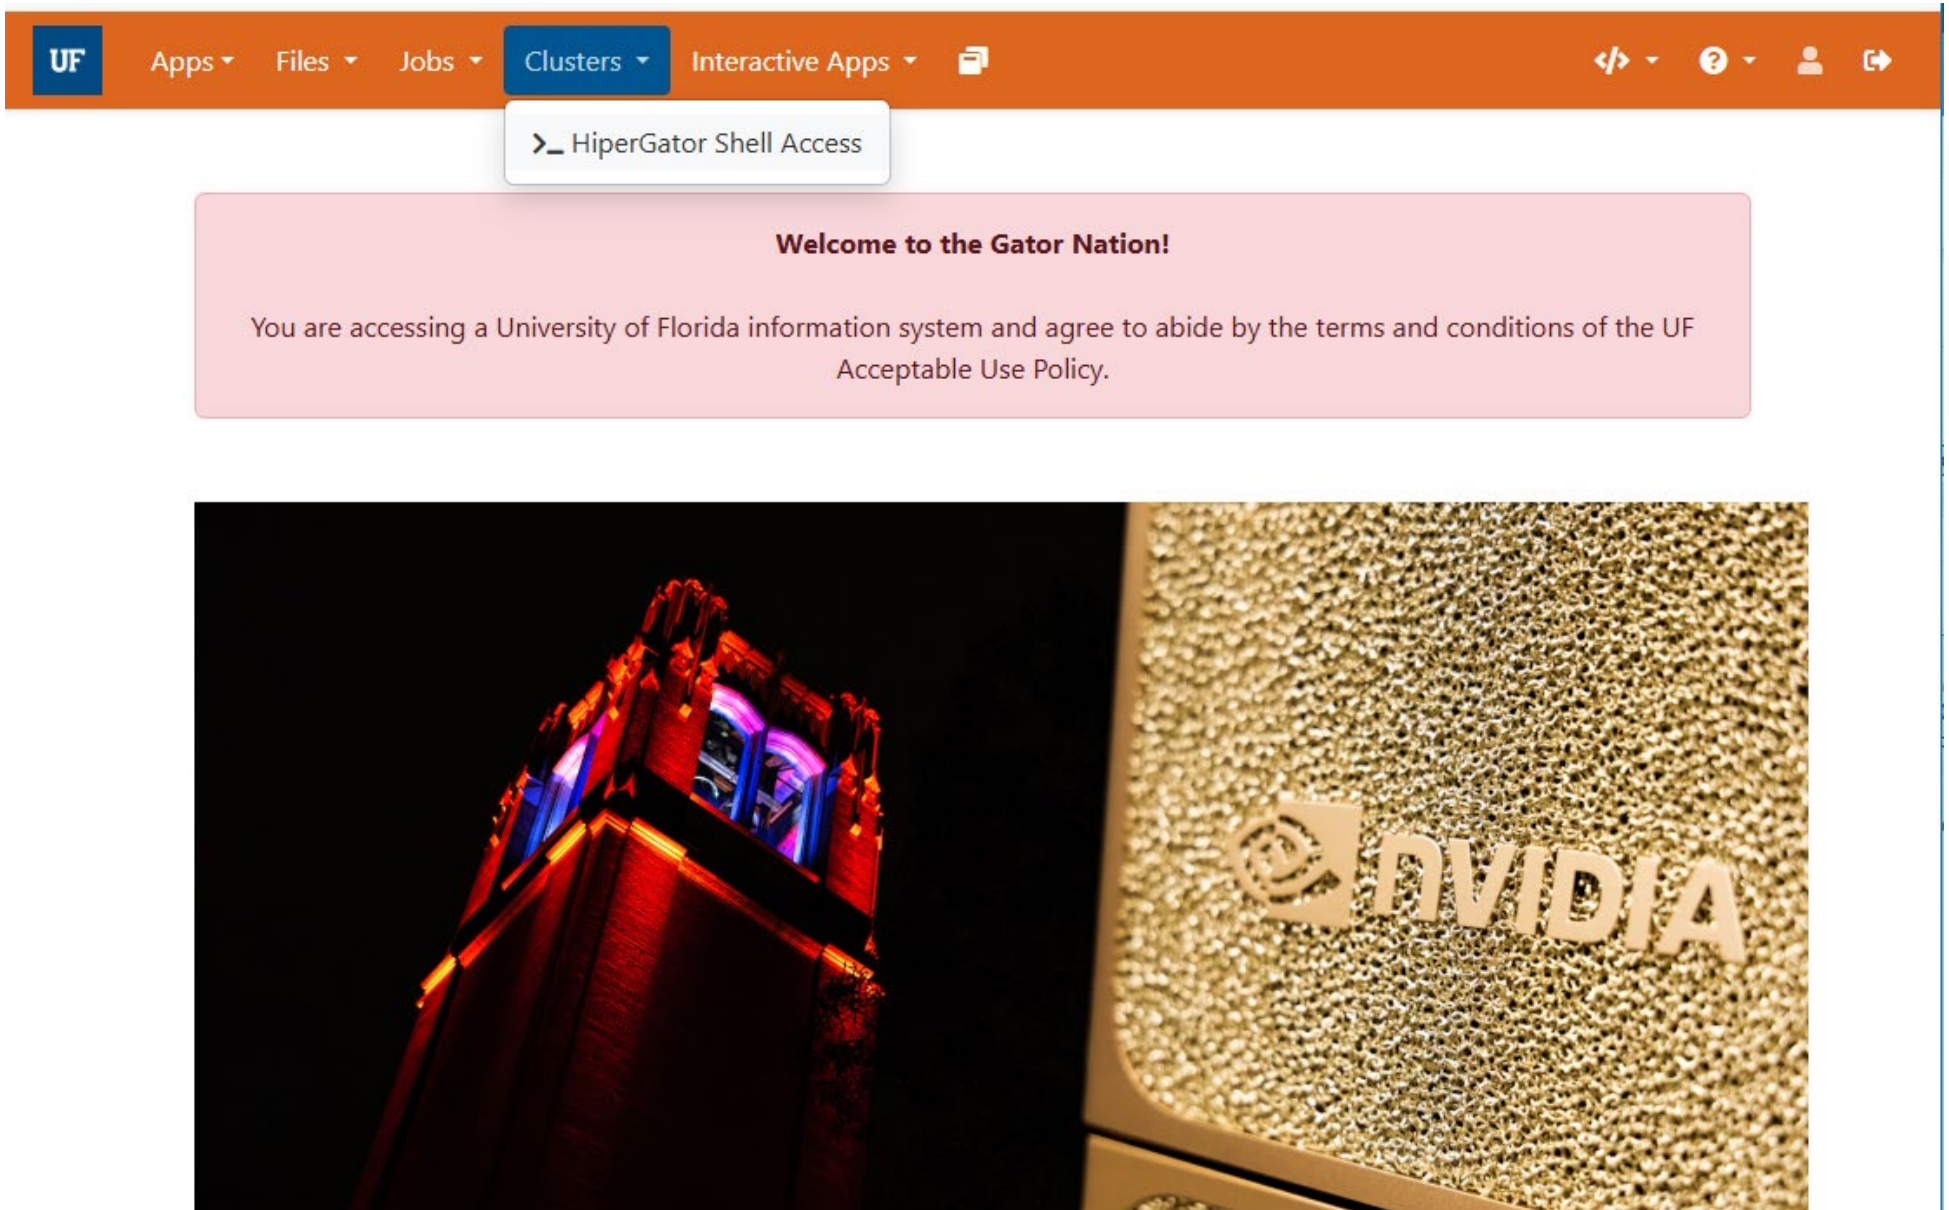

Storage on HiPerGator

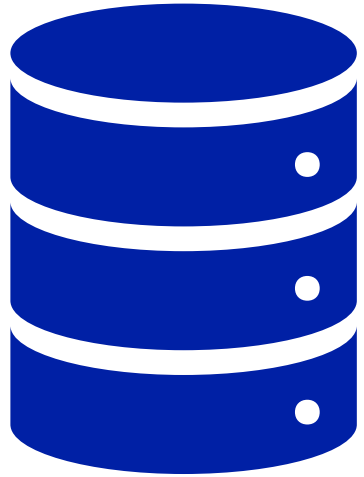

- **Home storage: `/home/<user>`**
  - 40GB limit
  - Scripts, code, compiled applications
  - **Do NOT use for job input/output**
  - **Week of snapshots at `~/.snapshot/`**
- **Blue storage: `/blue/pre1234/<user>`**
  - 2TB limit per class
  - **ALL input/output** from jobs should go here

- 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 Overview

User interaction

Login node or web interface

Slurm Scheduler

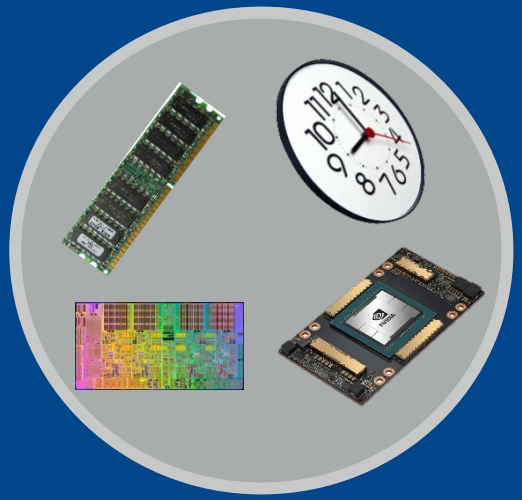

Tell Slurm what you want to do

Compute resources

Your job runs on the cluster

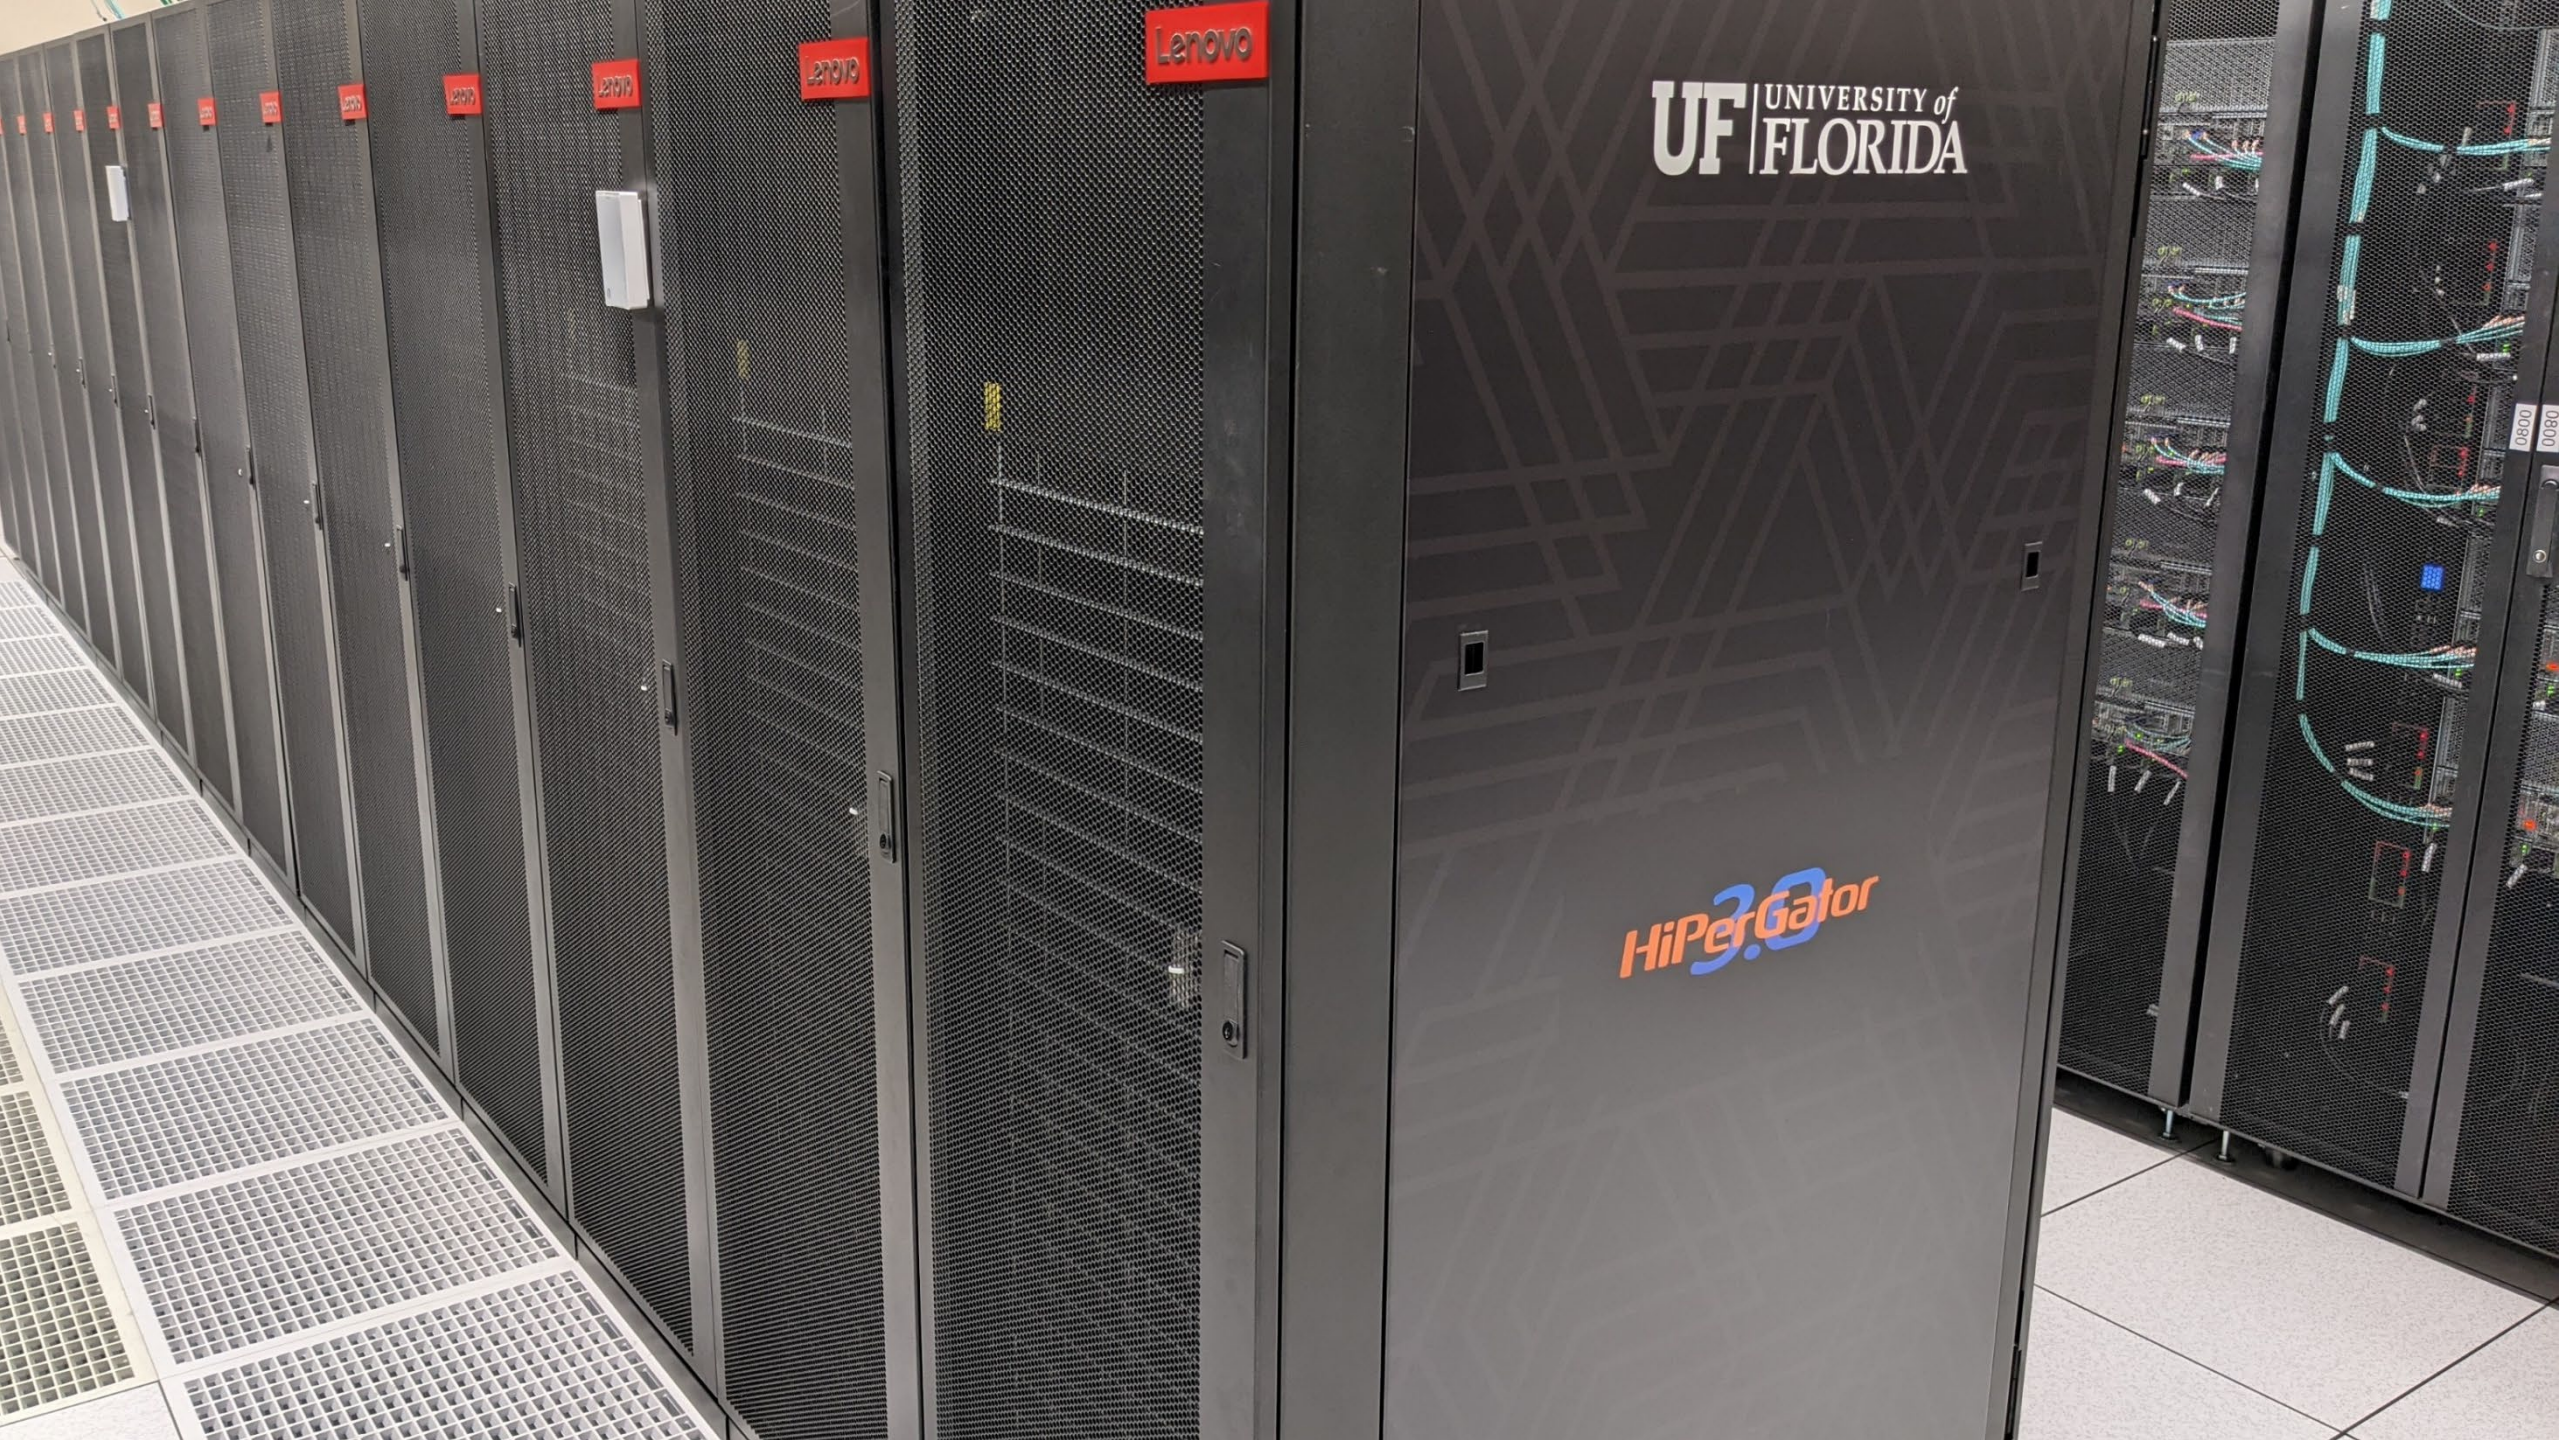
Lenovo
UF UNIVERSITY of FLORIDA
HiPerGator 3.0


Lenovo
UF UNIVERSITY of FLORIDA
HiPerGator
Powered by
Lenovo Neptune

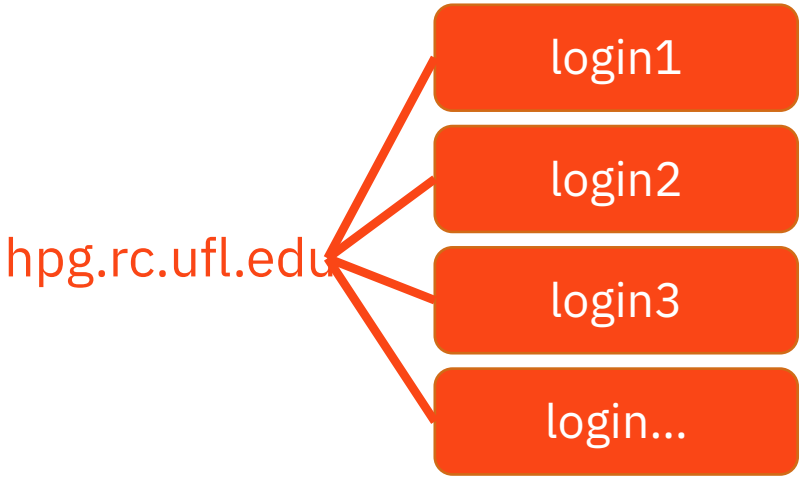

slurm
workload manager

Appropriate use of login nodes

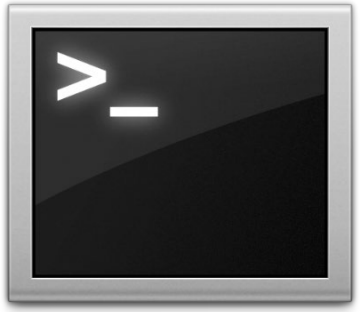

- Login nodes are for:
  - File and job management
  - Short-duration interactive testing and development
- Limit your use to **no more than**:

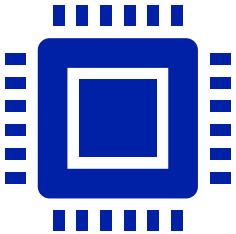

16 cores

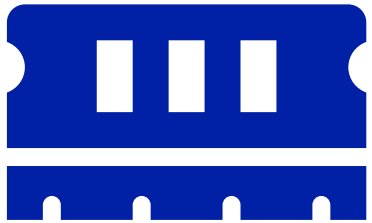

64 GB memory

10 minutes

Slurm Resources

Slurm Scheduler

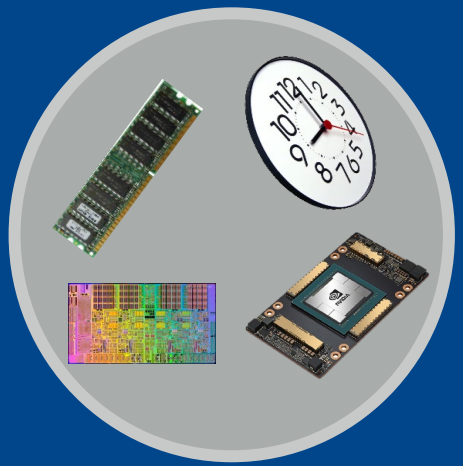

Tell Slurm what you want to do

Development servers

GPU Servers

Galaxy

Jupyter servers

GUI servers

Batch compute resources

UF UNIVERSITY of FLORIDA

Jupyter Hub and on Demand

jhub.rc.ufl.edu

ood.rc.ufl.edu

To setup link to the class blue directory, open a Terminal (File> New > Terminal) and run (e.g. for class **ast4930**):

```
ln -s /blue/ast4930 blue_ast4930
```

Be careful with `pip install`

- Can lead to conflicting versions of packages
- pip installs packages in `~/.local/lib/python3.x/site-packages`

Use conda/mamba

- Create isolated environments
- To use in Jupyter, create custom kernel folder. See help page.
- To use in script:

```
/path/to/bin/python my_script.py
```

Script should start with:
`#!/usr/bin/env python`

scheduling a job

- What resources does your job need?
  - **How many CPUs** do you want, and how do you want them grouped?
  - **How much RAM** your job will use?
  - **How long** your job will run?
  - **How many GPUs**?
  - Also need the commands that will be run to do your work

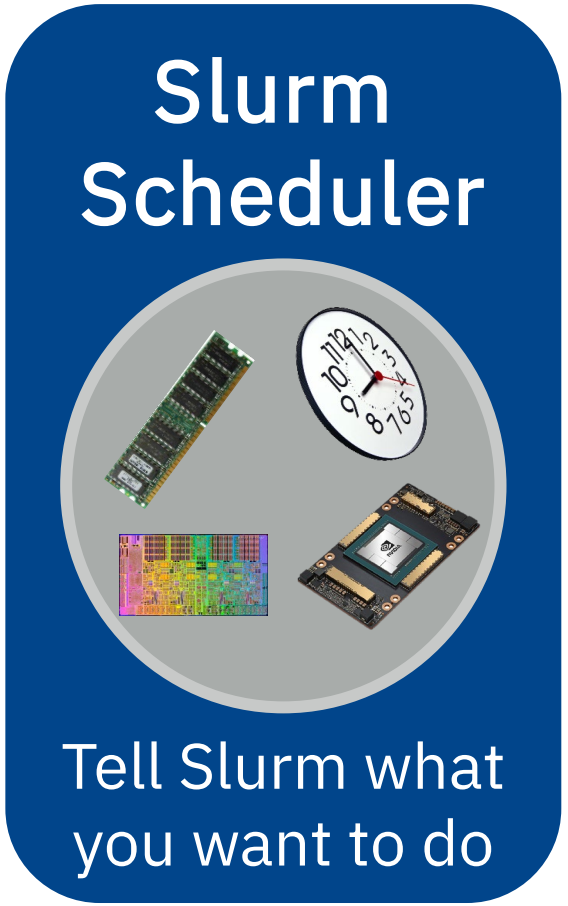

UF Apps Files Jobs Clusters Interactive Apps

Welcome to the Gator Nation!

You are accessing a University of Florida information system and agree to abide by the terms and conditions of the UF Acceptable Use Policy.

Pinned Apps A featured subset of all available apps

- Console — System Installed App
- Matlab — System Installed App
- Jupyter Notebook — System Installed App

powered by OPEN OnDemand

OnDemand version: 4.0.7

Home / My Interactive Sessions / Jupyter Notebook

Interactive Apps

Desktops
- Hipergator Desktop

GUIs
- ANSYS
- Artemis
- Beast
- COMSOL
- Console
- DSI Studio
- Fluent
- Gaussian
- IGV
- ITK-SNAP
- MRIcron
- Matlab
- NetLogo
- QGIS
- Rstudio
- SAS
- SCIRun
- Siril
- Spyder
- Upscayl
- VMD
- VScode
- Visit

HWGUI Apps
- FSLeyes
- Freeview
- MRIcroGL
- MRview

Jupyter Notebook

This app will launch a Jupyter Notebook server on a HiPerGator compute node.
Note: If GPUs are needed for machine learning (pytorch, tensorflow etc.), select the 'gpu' partition and enter a gres request.

Additional jupyter command Arguments

Include additional valid jupyter command arguments. Options must be space-separated.

Environment Modules

Load additional environment modules needed by your kernels

Number of CPU cores

8

Number of CPU cores on the compute node requested for the job, (default = 1).

Maximum memory requested for this job in Gigabytes (--mem, -m)

64

Maximum amount of memory to be used by the job (blank/default = 4GB per node or cpu core). If you are using advanced memory options in **Additional SLURM Options** below, then leave this blank.

Slurm Account (--account, -A)

Enter an alternative Slurm account if required. (default = same as your primary group)

QOS (Required if custom account is set above, --qos, -q)

Enter an alternative QoS, **required if alternative account is entered above**. Note: if you use the burst qos by appending '-b' to the account name, your jobs may take longer to start. There is no burst QOS for GPU jobs. (default = investment QOS)

Job Time Limit in hours (--time, -t).

7

Time in hours requested from SLURM for this job to run. (default = 1) **Please note, interactive gpu jobs will be limited to 12 hours.**

Cluster partition (--partition, -p)

default

Select a specific cluster partition for job. (default = first available compute partition)

Generic Resource Request (--gres).

gpu:1

This is the Generic resource request string to request GPU resources. See also https://docs.rc.ufl.edu/scheduler/gpu_access

Additional Slurm Options

Accepts valid srun options per https://slurm.schedmd.com/srun.html. Each additional option should be separated by a space.

Launch

* The Jupyter Notebook session data for this session can be accessed under the data

Basic Slurm job script

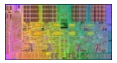

```
#!/bin/sh
#SBATCH --cpus-per-task=1              # Run on a single CPU
#SBATCH --mem=1gb                      # Memory limit
#SBATCH --time=00:05:00                # Time: hr:min:sec

#SBATCH --job-name=job_test            # Job name
#SBATCH --mail-type=ALL                # Mail events
#SBATCH --mail-user=email_address      # Where to send mail
#SBATCH --output=serial_%j.out         # Output and error log

pwd; hostname; date    # Print some information
module load python     # Load needed modules
echo "Running plot script on a single CPU core"
python /data/training/SLURM/plot_template.py
date                   # Print ending time
```

Slurm CPU Requests

- For threaded applications (single node):

```
#SBATCH --nodes=1   # Physical servers
#SBATCH --ntasks=1  # MPI ranks or processes
#SBATCH --cpus-per-task=8
```

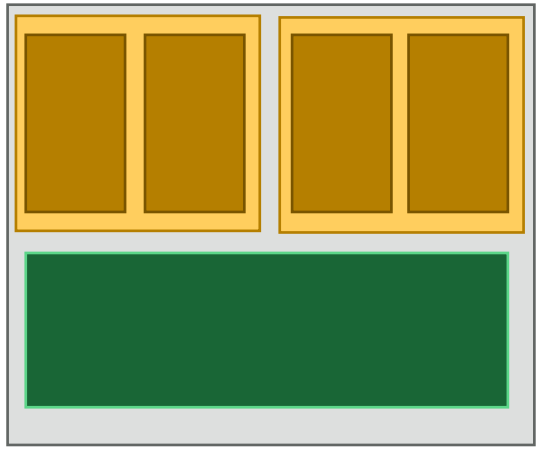

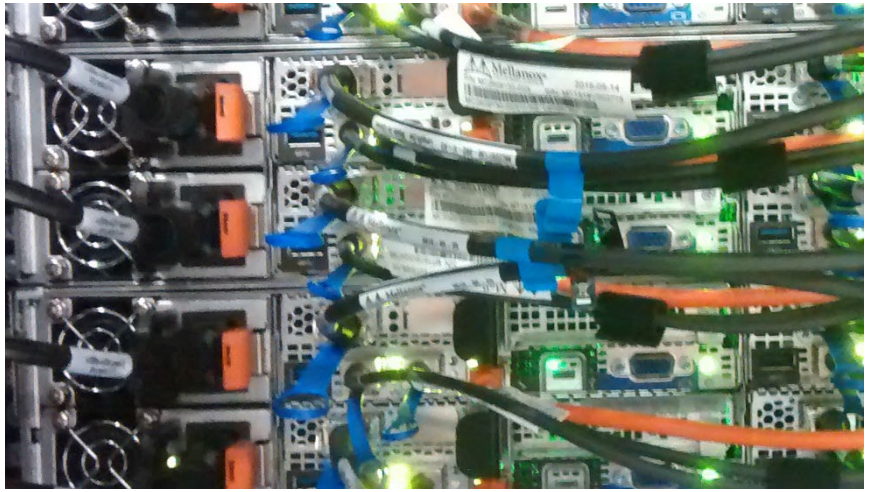

SLURM GPU Requests

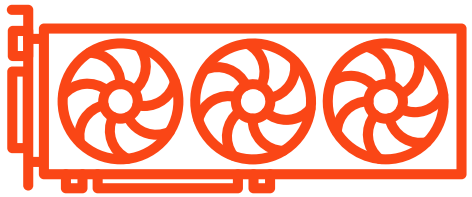

- In a script:

  #SBATCH --gpus=1

- In OOD:

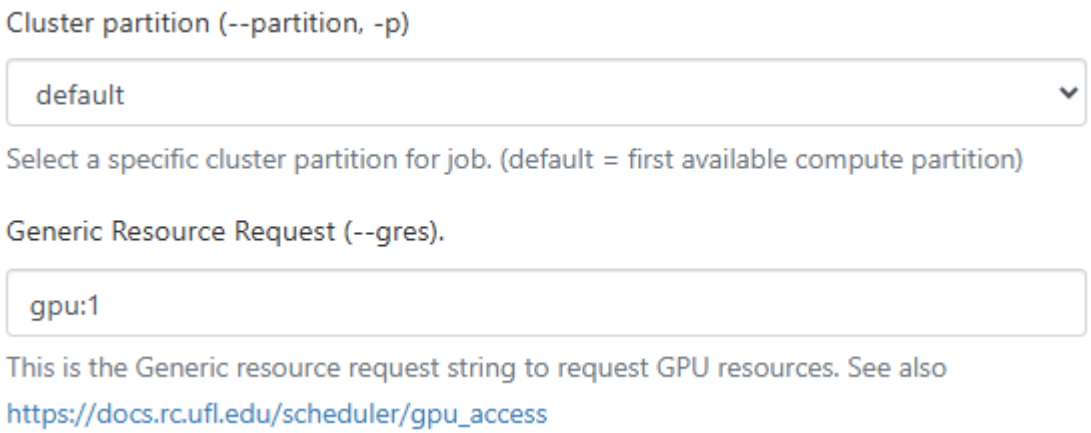

- Courses should use **ONLY** L4 GPUs
- See: https://docs.rc.ufl.edu/doc/GPU_Access

Slurm MPI Job CPU Requests

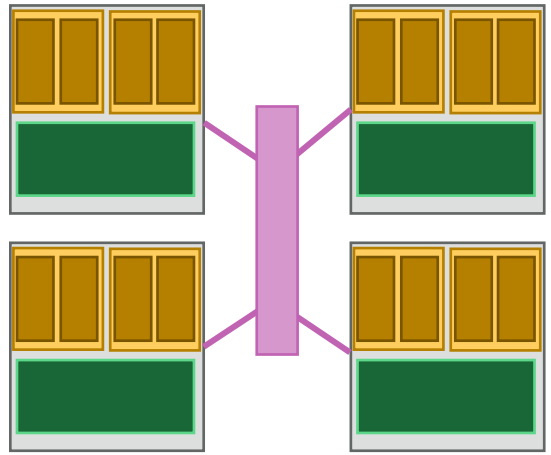

```
#SBATCH --nodes=1
#SBATCH --ntasks=16
#SBATCH --cpus-per-task=1
#SBATCH --ntasks-per-node=16
#SBATCH --ntasks-per-socket=8

#SBATCH --distribution=cyclic:cyclic

ml gcc openmpi
srun --mpi=$HPC_PMIX myApp …
```

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

Servers have
512 GB to 1TB
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)

Quality of Service (--qos)

- Each group has two QOS options
  - Investment QOS: `--qos=group`
  - 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
- Cannot include a GPU
- 4-day time limit

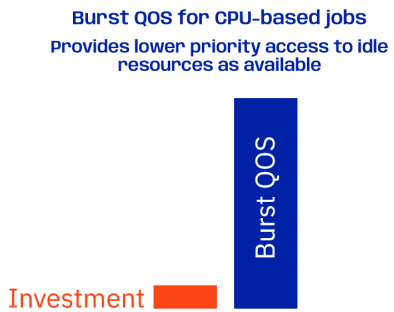

Submit your job

```
[magitz@login3 SLURM_examples]$ sbatch single_job.sh
Submitted batch job 30592170
[magitz@login3 SLURM_examples]$ squeue --me
   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]$
```

Helpful commands

- **`slurmInfo`** – Get information on compute resources in your group
- **`blue_quota, orange_quota, home_quota`** – Get information on quotas and use
- **`ncdu`** – Calculate and show storage usage by folder. Can take time to calculate—most helpful in `/home`
- **`showAllocation`** – Get information on investments
- **`jobhtop, jobnvtop`** – View real-time job performance data
- **`qos_to_burst, qos_to_main`** – Move pending jobs to burst/main QOS

docs.rc.ufl.edu

UFIT Research Computing User Documentation

Welcome to the University of Florida Information Technology Research Computing User Documentation and HiPerGator Compendium. This is a refresh of our original Help Wiki Site.

General information about RC, news, announcements, policies, purchase request and account request forms are located on our main website. All other information you may need to become a proficient and effective HiPerGator (HPG) user, including any extra information about using particular applications on HPG, available reference data, services, and code examples is located in this compendium.

Our documentation is split into sections you can explore via the links in the navigation panel at the top or by following the Table of Contents navigation at the bottom.

New users are highly encouraged to take the New User Training course to familiarize themselves with using a supercomputer for research, save time, and avoid many beginner mistakes.

Table of Contents

Policies and Procedures

For users and sponsors using HiPerGator

→ Policies and Procedures

Quick Start

Introduction to HiPerGator

→ Interfaces

→ From Zero to HiPerGator

Training

HiPerGator Training Resources

→ Training Overview

→ HiPerGator User Training

Access

Connect to and interact with HiPerGator

→ SSH Connections

→ Open OnDemand

Data

Transferring & Accessing Data

→ Transfer data to and from HiPerGator

→ Reference data

Software

Using Modules and Apps on HiPerGator

→ View installed programs and tools

→ Accessing Modules

UF Information Technology UNIVERSITY of FLORIDA

Training

UFIT-RC Documentation

Search

Home | Quick Start | Training | Access | Data | Software | Job Scheduler | Services | Support | Scientific Domains

Training
- Introduction
- New User Training
- HiPerGator Training
- How-To Videos
- AI Training
- By Users for Users

UFIT Research Computing User Training

Table of contents
- Training Resources

UFIT Research Computing provides a wide variety of training sessions, materials, and events to enable our customers to make the most of the HiPerGator resources and services available to them and become proficient and effective HiPerGator users as soon as possible.

Note

- All new users are encouraged to take our free HiPerGator training Canvas course. It is available at no cost and gives you an introduction to HPC, SLURM, Modules, etc. You can view the training objectives and register for the course at the HiPerGator User Training page.

- Live training sessions are available throughout the year. **See the** Training Schedule
- Many of the sessions are available as pre-recorded videos linked below.
- AI Education and Training lists selected AI education and learning materials.

Training Resources

Training category	Additional details
General HiPerGator Training	General Use of HiPerGator: topics include Introduction, basic Linux commands, Jupyter, Open on Demand, and the SLURM scheduler.
AI Training, including Practicum AI	Resources for AI training, including Practicum AI, NVIDIA DLI, and more.
Quick How-To Videos	A set of short, how-to videos on topics like logging into HiPerGator, transferring data, and more.
Additional Training Options	How to schedule customized training and resource from other sources.
By Users for Users	Training and documentation produced **by** users **for** users.

UF Information Technology
UNIVERSITY of FLORIDA

A note about IDEs

Visual Studio Code

Spyder

PC

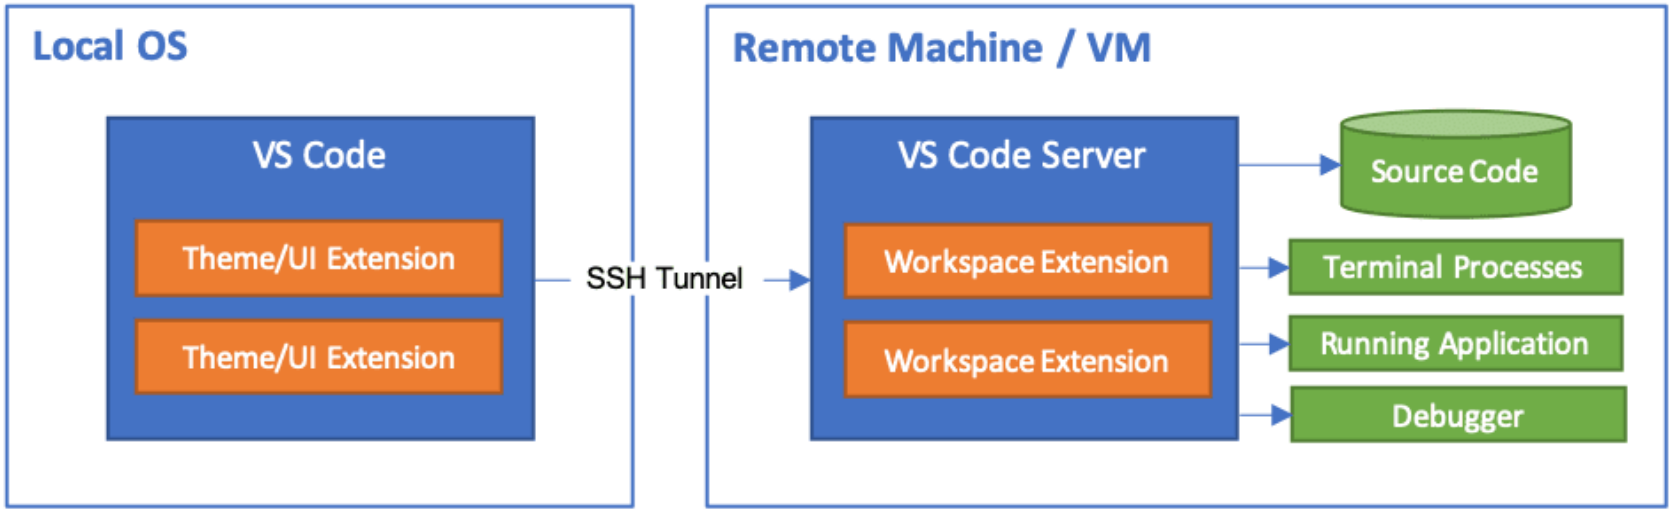

This is a HiPerGator login server!!

- Follow limits
- No GPUs

See VS Code Server: https://code.visualstudio.com/docs/remote/vscode-server
UFIT-RC Help Page: https://docs.rc.ufl.edu/domain/vscode_development/

VSCode Remote Tunnel

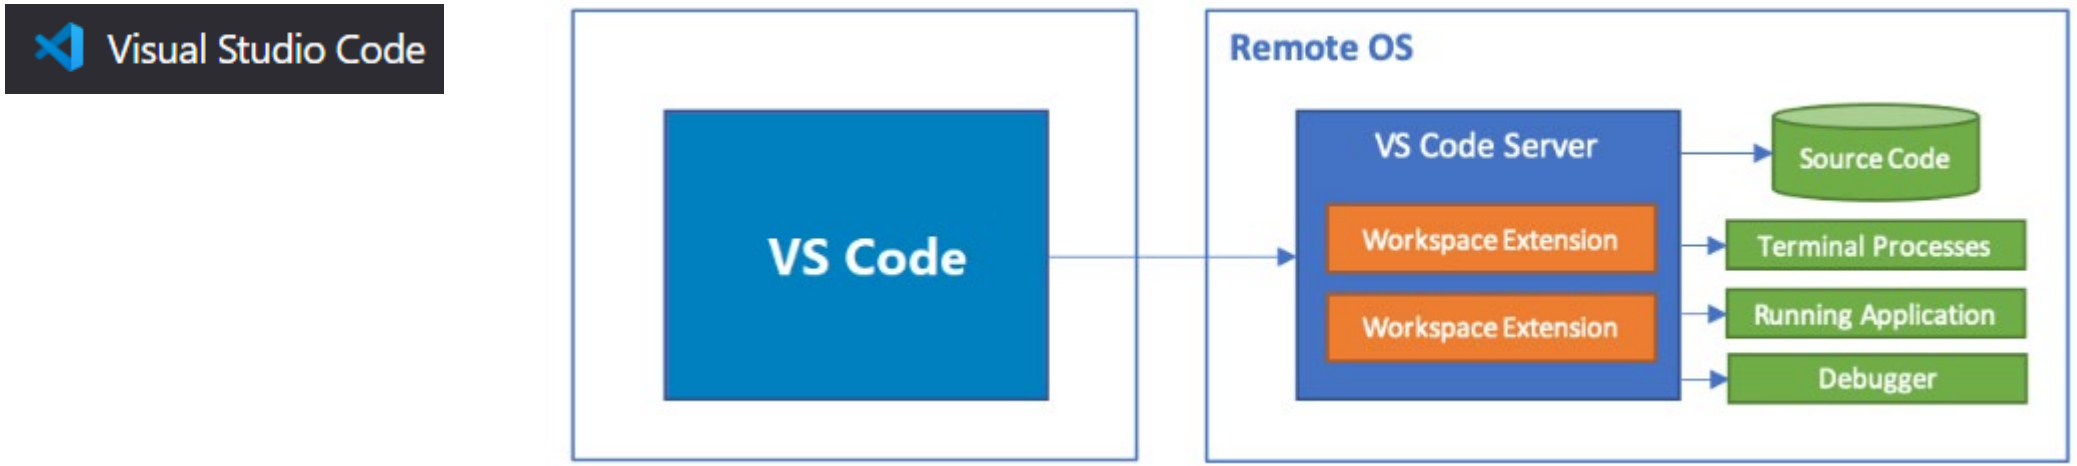

See VS Code Remote Tunnels: https://code.visualstudio.com/docs/remote/tunnels
UFRC Help Page: https://docs.rc.ufl.edu/domain/vscode_development/

Questions?

Matt Gitzendanner
magitz@ufl.edu