



# Fiji Compute Cluster

Jonathan DeMasi  
Matt Hynes-Grace  
BioFrontiers IT



# Basic Facts about Fiji

- ~ 4000 Cores
- ~ 55 Standard compute nodes (64 Threads (32 physical cores) - 512 GB of RAM)
- ~ 3 high memory compute nodes (48 physical cores - 1TB of RAM)
- 3 Types of GPU nodes (nVidia Titan (x4), Tesla k20 (x4), Tesla k40 (x1))
- 1 PB high-speed scratch filesystem (BeeGFS)
- 650 TB /Users home filesystem - with ZFS snapshotting and offsite replication
- 1 visualization node for (Jupyterhub, RStudio, and X2Go remote X sessions)
- Fiji uses the Slurm scheduler



# Storage best practices

- Run in scratch, save to /Users for backup
- /scratch directory is NEVER backed up and can't be recovered
- /Users is very susceptible to interruption or problems if a user creates I/O with thousands of cores
- Standard user share in /Users is 1T but that can be increased with approval from the PI
- /scratch does struggle with millions of tiny files so sometimes a local copy is useful



# Job best practices

- ALWAYS submit compute work as a compute job - if it isn't file staging/transfer or editing scripts, it should be a job!
- Test your batch job with a small subset before requesting hundreds of cores - BIT is always available to help get your jobs running properly during normal business hours
- Scratch is NEVER backed up but you should use it to compute
- You can track your job with the **sacct** and **scontrol** commands to verify what resources were used / needed
  - Efficiently submitted jobs make the cluster better for EVERYONE - be a good neighbor



# Hints to be a good neighbor

- Only use what you request
  - The system will kill your job when it goes over memory but we aren't so strict on cores
- Only reserve cores that are needed for your job
- Accurately request walltime
- Generally keep use under 1000 cores in short queue and 500 cores in long
  - This is not a firm limit and can be adjusted in both directions.
  - If you have a large computational need/deadline - let us know in advance and we'll do our best to accommodate
- Clean up /tmp or /localscratch if you use and realize it's limited to ~100G



# Other good ideas

- Keep your sbatch scripts and any custom code in a Git repository - if you don't know about Git, we frequently host Git workshops and encourage you to attend one.
- If you're using someone else's scripts, ensure you change the parameters/paths in the job to reflect your own environment - take the time to understand what it's doing, don't just copy-paste!
- Separate data into raw inputs, intermediates, and final outputs
  - When you're done with a project, delete your intermediates and any outputs that could easily be recomputed, then you aren't needlessly paying for storage



# Visualization Resources

- All visualization tools run on [fiji-viz.colorado.edu](http://fiji-viz.colorado.edu)
  - 144 cores, 512G of RAM
  - Fiji-viz is a SHARED resource - we ask users to self police when using the system
- Jupyterhub
  - The only resource on Fiji-viz that can actually submit jobs to the compute cluster
  - Can run jobs for up to 24 hours with 1, 16, or 64 cores
  - Support for custom Python 3 virtual environments so you can run whatever you need
- RStudio-server
  - Graphical Rstudio session but with access to cluster mounts / resources
- X2Go
  - Used to spawn an entire X session for using tools like IGV



# Jupyterhub quick demo

- Describe what the “hub” is actually doing
- Basic usage example, what to expect at login
- Can use custom virtual environments for extensibility and reproducibility
- Questions





# Rstudio-server

- What does it do?
- How do you log in?
- Current limitations
  - Not easy to maintain virtual environments like Python
  - Does not submit jobs to compute cluster like Jupyterhub, so must be mindful of usage



# X2Go

- What does it do and why would you need to use it?
  - Primarily used by users who need to utilize IGV
  - Can be used for graphical data staging / movement
- Quick demo



# Basic Submission Example

- First! Check it out from git:
  - `git clone https://biof-git.colorado.edu/jode5573/fiji-examples.git`
- Make sure to change the email address and resources requests, etc.
- `sbatch slurm-cpuonly.sbatch`
- `queue -u jode5573` (this is my username)



# How to get help

- If you are unsure of how to do something on Fiji, consult the user guide!
- Email to [bit-help@colorado.edu](mailto:bit-help@colorado.edu) - include your job number(s), locations or copies of sbatch scripts, and any other pertinent information
- This helps us track ongoing issues but we will always be willing schedule time if needed to investigate and discuss your issue
- Talk to us when you are having issues, you don't have to be frustrated first!



# Resources

Jon's example-filled Git repo: <https://biof-git.colorado.edu/jode5573/fiji-examples>

Fiji user guide and other documentation:

<http://bficores.colorado.edu/biofrontiers-it/cluster-computing/fiji>

