UT Southwestern Medical Center
Lyda Hill Department of Bioinformatics

BioHPC

# Introduce of CryoEM Resources on BioHPC

[web] portal.biohpc.swmed.edu
[email] biohpc-help@utsouthwestern.edu



*06-16-2021*

# Overview CryoEM resources on BioHPC

- **Raw image transfer and on-fly processing, Data Storage**
- **Web Visualization, Interactive job, SLURM job**
- **BioHPC workstation**
- **GPU partitions:**
  - >100 nodes Nvidia - Tesla K20/K40, P4, P40, P100, V100s, 4V100, A100, more will come.
- **Tools:**
  - RELION, cryoSPARC, EMAN2, MotionCorr, Coot, Cryolo, Chimera, IMOD, PEET, SBGrid, et.al.

# Data transfer and Storage

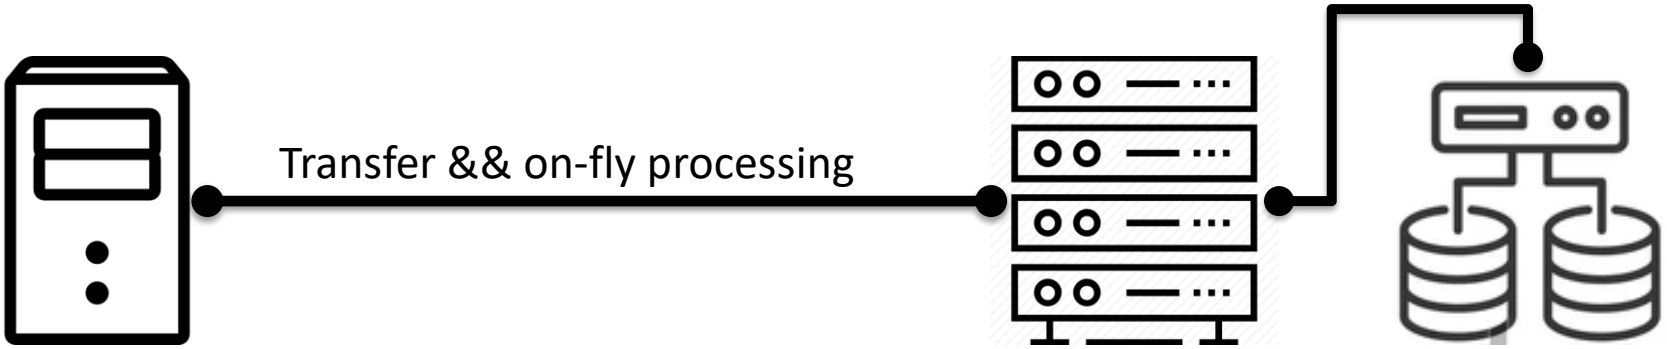


Camera PCs

:\DoseFraction\<**user**>\Images-Disc
:\DoseFraction\<**user**>\KEEP

BioHPC cluster and storage

/project/cryoem/cryoem_transfer/<**krios**>/<**user**>

**…**

**…**

**…**

**Important:** Move to your own space on BioHPC or local server after your cryoEM microscope session.

# Access data storage on BioHPC

**As a BioHPC user, you will have several storage options:**

**BioHPC Standard Storage Limits**

| | |
|---|---|
| /home2/<username> | 50 GB per user |
| /work/<department>/<username> | 5 TB per user |
| /project/<organization>/<group>/<username> | 5 TB per group |
| /archive/<organization>/<group>/<username> | 5 TB per group |


**BioHPC File Exchange (accessible to outside of UTSW)**

https://cloud.biohpc.swmed.edu 50 GB / user

**BioHPC Lamella Cloud Storage (UTSW-only)**


https://lamella.biohpc.swmed.edu 100 GB / user
(FTP, SAMBA or WebDAV*)

**Project and Archive can be increased on PI request with Dept. Chair approval.**
**Archive usage is multiplied by 2/3 (as to encourage use of archive).**
**Data on home2 counts thrice and on work counts twice because of backup.**

**Guides about BioHPC storage**:

1, BioHPC portal Guides/FAQ pages: BioHPC Guides (swmed.edu)

2, BioHPC training slides: BioHPC Storage Overview (swmed.edu)

# Setting up Lamella to access project and work space

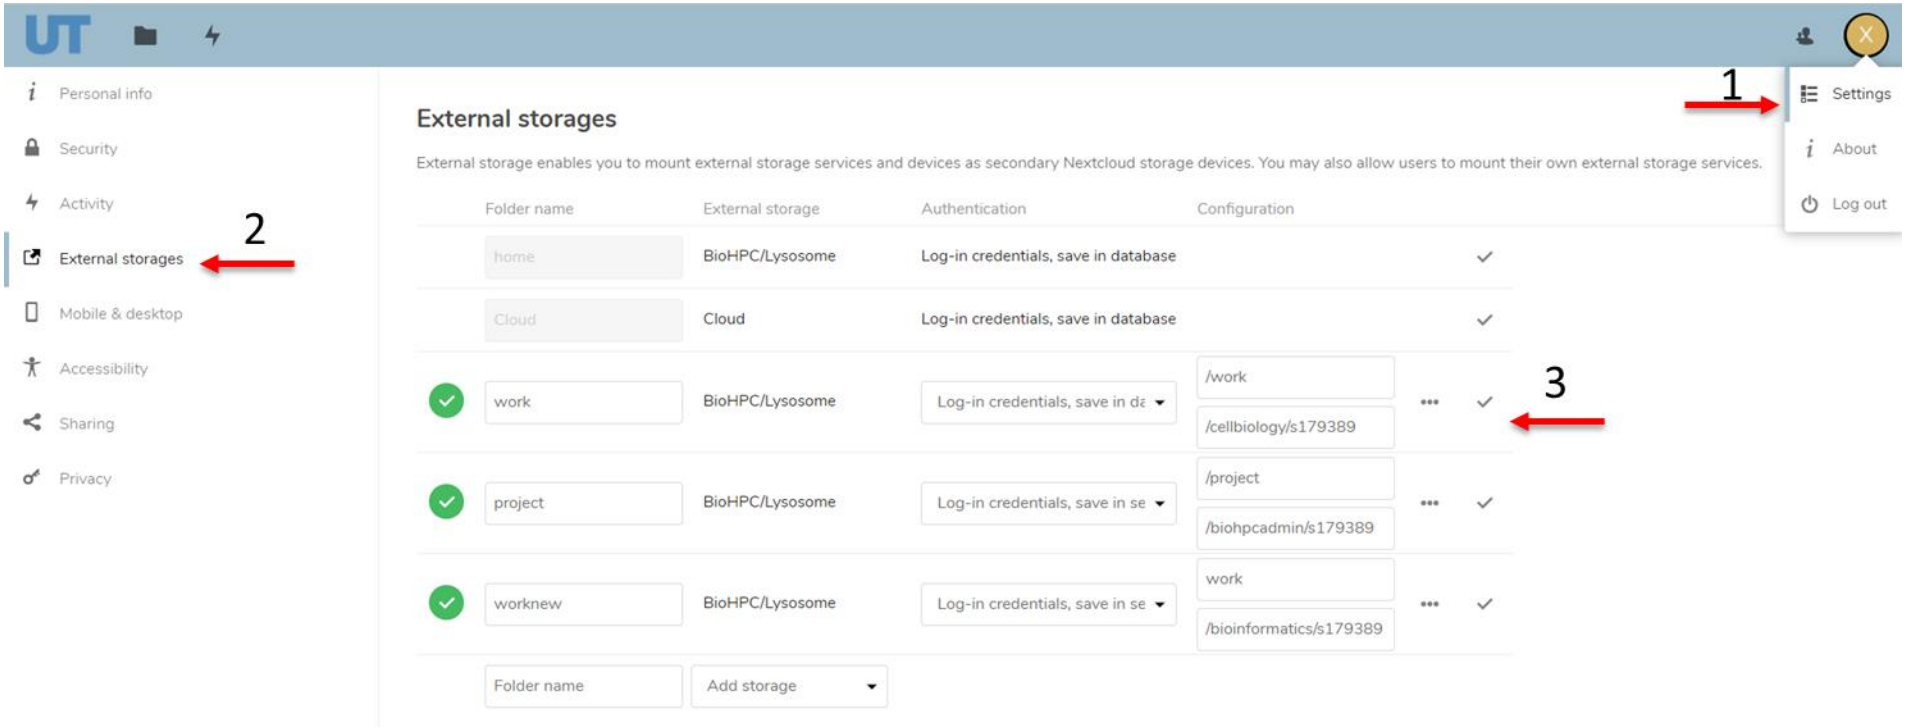

# Setting up Lamella to access project and work

Computer -> Map Network Drive

Folder is: \\lamella.biohpc.swmed.edu*username* *(home dir)*
\\lamella.biohpc.swmed.edu\project
\\lamella.biohpc.swmed.edu\work

Check 'Connect using different credentials'

Enter your BioHPC username and password when prompted.

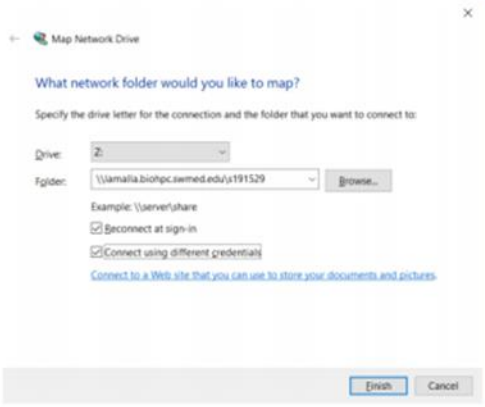


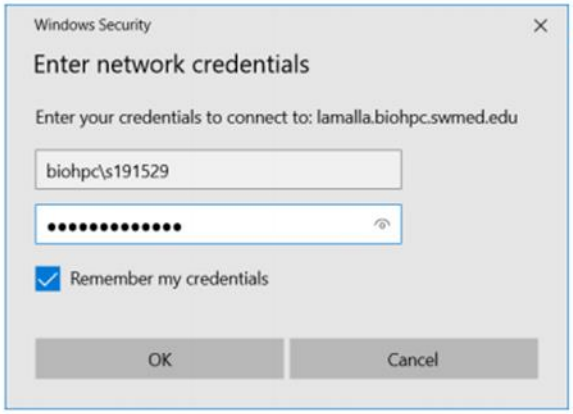


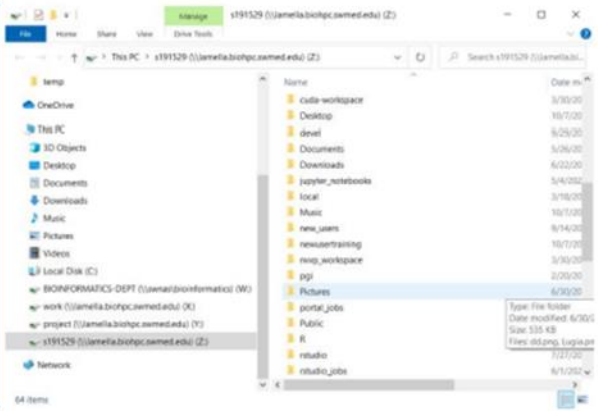

# Web Visualization and Slurm Job

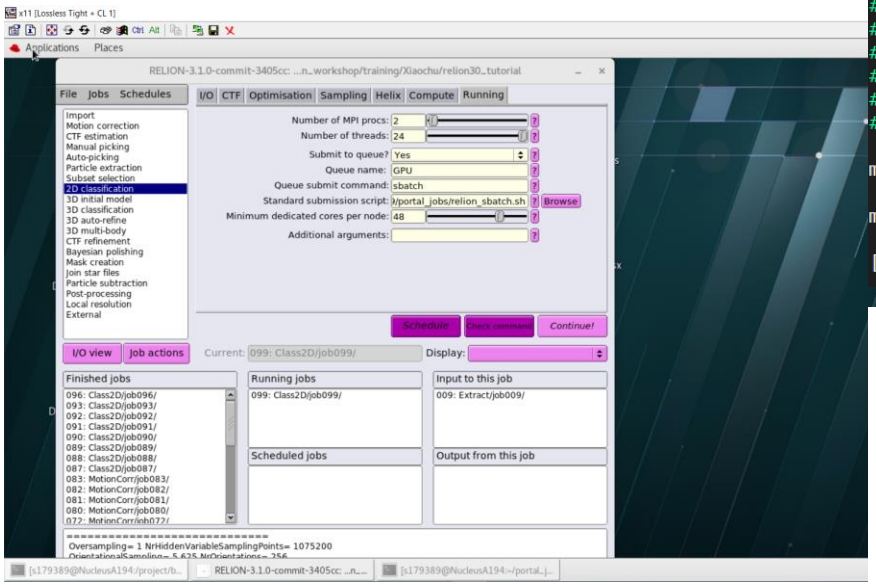


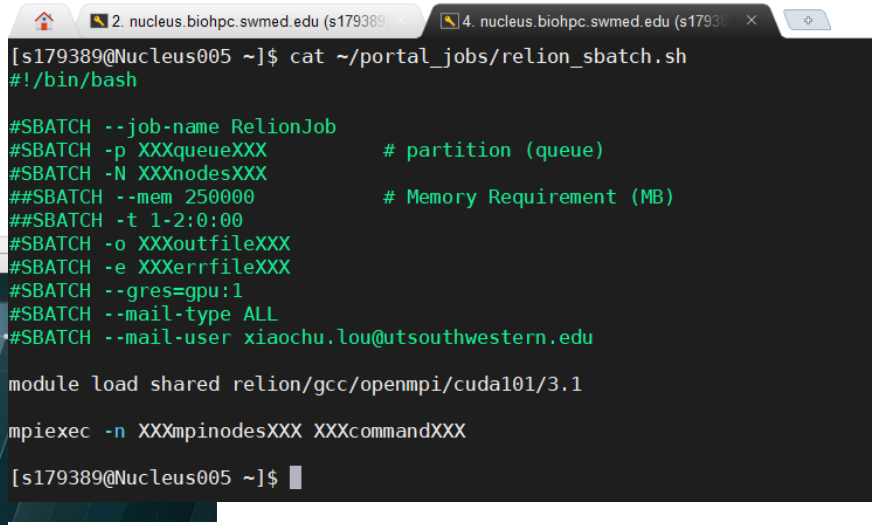


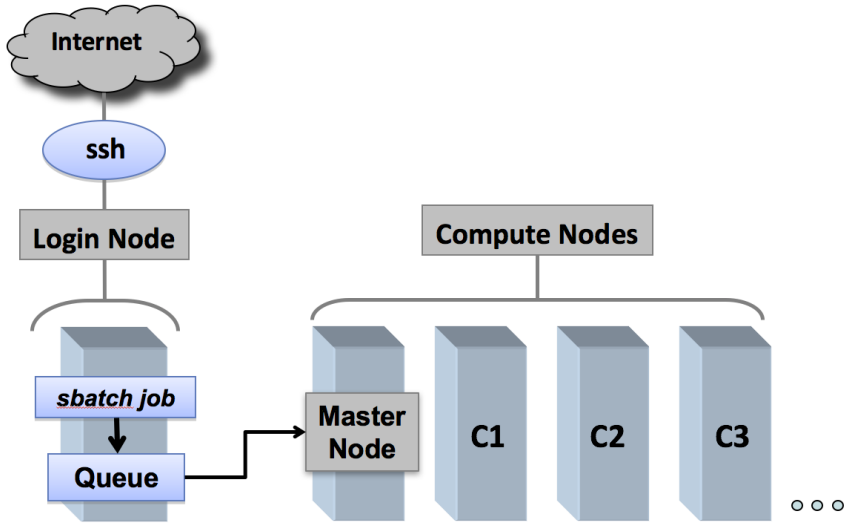

# BioHPC Workstation and Thin Clients

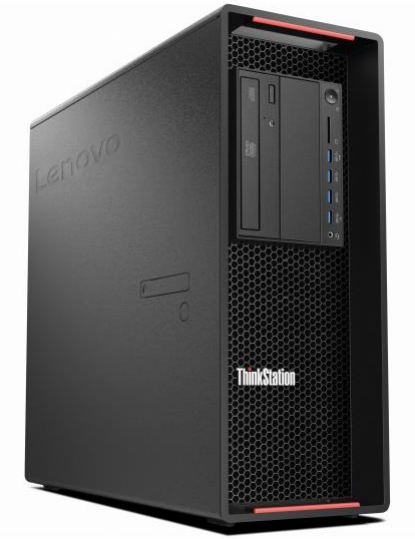



## Why use a BioHPC Workstation or Thin Client?

- They run the same environments as the BioHPC Cluster
  - ReaHat Enterprise 7 (RHEL 7)
  - Features such as bash, slurm, modules are preconfigured
  - No need to apply web visualization session
- It's easier to access large datasets on BioHPC filesystems
  - /home2, /project, /work, /archive are mounted
  - Drag-and-Drop files in the Graphic user interface (GUI)
  - No need to use an FTP client, Samba shares, or web

# GPU nodes

| GPU Partition | Number of CPU/Node | Memory Per Node | Number of GPU/Node | GPU Memory | Number of nodes |
|---|---|---|---|---|---|
| GPU | 32 | 256GB | 1 K20/K40 | 6GB/12GB | 8 |
| GPUp4 | 72 | 384GB | 1 P4 | 8GB | 16 |
| GPUp40 | 72 | 384GB | 1 P40 | 24GB | 16 |
| GPUp100 | 56 | 256GB | 2 P100 | 16GB | 12 |
| GPUv100s | 72 | 384GB | 1 V100S | 32GB | 32 |
| GPU4v100 | 72 | 384GB | 4 V100S | 32GB | 12 |
| GPUA100 | 72 | 1.5TB | 1 A100 | 40GB | 16 |

**Check node availability**

```
[s179389@Nucleus005 ~]$ sinfo -p GPUp4
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
GPUp4        up   infinite      5  alloc NucleusC[002,012-013,016-017]
GPUp4        up   infinite     11   idle NucleusC[003-011,014-015]
```

# RELION GPU Benchmarks

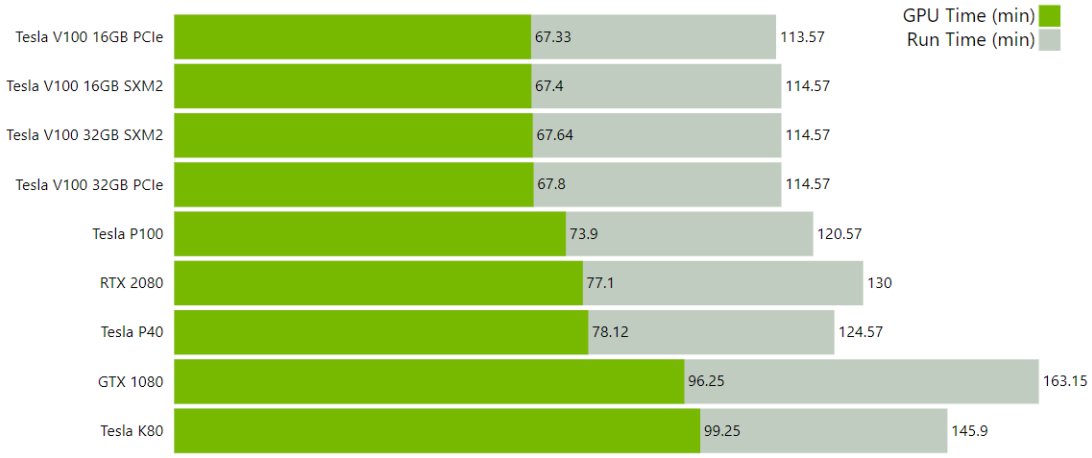


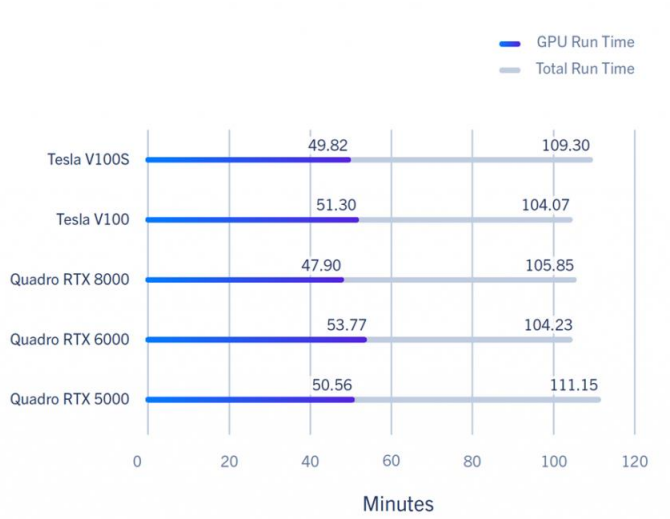


[Ward Lab - RELION GPU Benchmarks (scripps.edu)]

[RELION Cryo-EM Benchmarks and Analysis | Exxact Blog (exxactcorp.com)]

# Modules

```
[s179389@Nucleus005 ~]$ module avail

------------------------------------------------------------ /cm/local/modulefiles ------------------------------
cluster-tools/7.3 dot             gcc/6.1.0        module-git       null             shared
cmd               freeipmi/1.5.2  ipmitool/1.8.17  module-info      openldap         use.own

------------------------------------------------------------ /cm/shared/modulefiles -----------------------------
abra2/2.18                                  intel/mkl/32/2017/6.256
acml/gcc/64/5.3.1                           intel/mkl/64/2017/6.256
acml/gcc/fma4/5.3.1                         intel/mkl/64/2017/current
acml/gcc/mp/64/5.3.1                        intel/mkl/mic/2017/6.256
acml/gcc/mp/fma4/5.3.1                      intel/mpi/32/2017/6.256
acml/gcc-int64/64/5.3.1                     intel/mpi/64/2017/6.256
acml/gcc-int64/fma4/5.3.1                   intel/mpi/mic/2017/6.256
acml/gcc-int64/mp/64/5.3.1                  intel-cluster-checker/2.1.2
acml/gcc-int64/mp/fma4/5.3.1                intel-cluster-runtime/ia32/3.8
afni/20.0.04                                intel-cluster-runtime/intel64/3.8
afni/v17.2.17                               intel-cluster-runtime/mic/3.8
afni/v18.3.03                               intel-tbb-oss/ia32/2017_20170807oss
almost/2.1                                  intel-tbb-oss/intel64/2017_20170807oss
amber/12                                    io_lib/1.13.3
```

```
[s179389@Nucleus005 ~]$ module avail relion

------------------------------------------------------------ /cm/shared/modulefiles -----------------------------
relion/gcc/1.2                      relion/gcc/openmpi/cuda80/2.0-beta    relion/intel/openmpi/1.4
relion/gcc/1.3                      relion/gcc/openmpi/cuda80/2.1.0       relion/intel/openmpi/cuda101/3.1-beta
relion/gcc/mvapich2/1.4             relion/gcc/openmpi/cuda91/2.1.0       relion/intel/openmpi/cuda80/2.0-beta
relion/gcc/openmpi/1.4              relion/intel/1.3                      relion/intel/openmpi/cuda91/3.0.1
relion/gcc/openmpi/cuda101/3.1      relion/intel/mvapich2/1.4
[s179389@Nucleus005 ~]$ module avail cryosparc

------------------------------------------------------------ /cm/shared/modulefiles -----------------------------
cryosparc/2.15.0-6-singularity cryosparc/3.1.0-singularity
[s179389@Nucleus005 ~]$ module avail imod

------------------------------------------------------------ /cm/shared/modulefiles -----------------------------
imod/4.7.5            imod/4.8.50-beta       imod/cuda65/4.8.50-beta imod/cuda80/4.11.0      imod/cuda80/4.9.3
imod/4.8.38           imod/4.9.3             imod/cuda80/4.10.32     imod/cuda80/4.9.12
[s179389@Nucleus005 ~]$
```

# Run RELION and submit RELION batch job for GPU intensive jobs

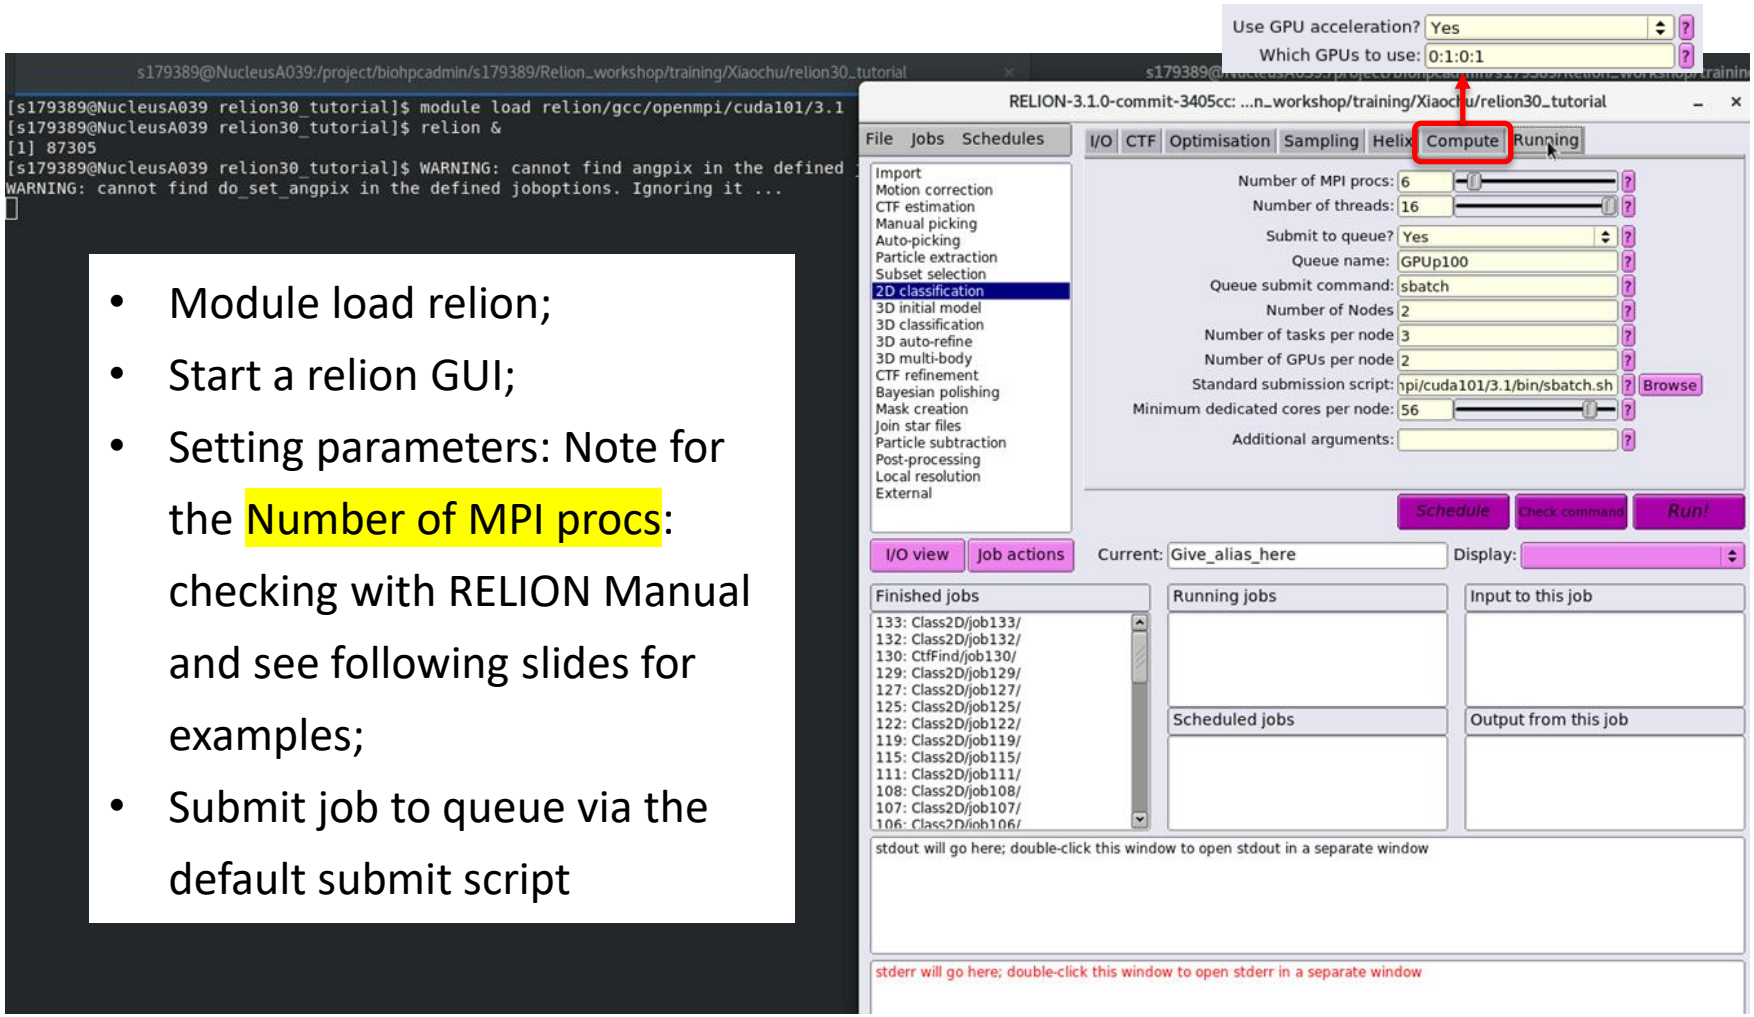


- Module load relion;
- Start a relion GUI;
- Setting parameters: Note for the Number of MPI procs: checking with RELION Manual and see following slides for examples;
- Submit job to queue via the default submit script

# Examples of submitting RELION job to multiple nodes

**Motion correction**

**1 GPUp100 nodes,**

**2 GPU cards on each nodes**

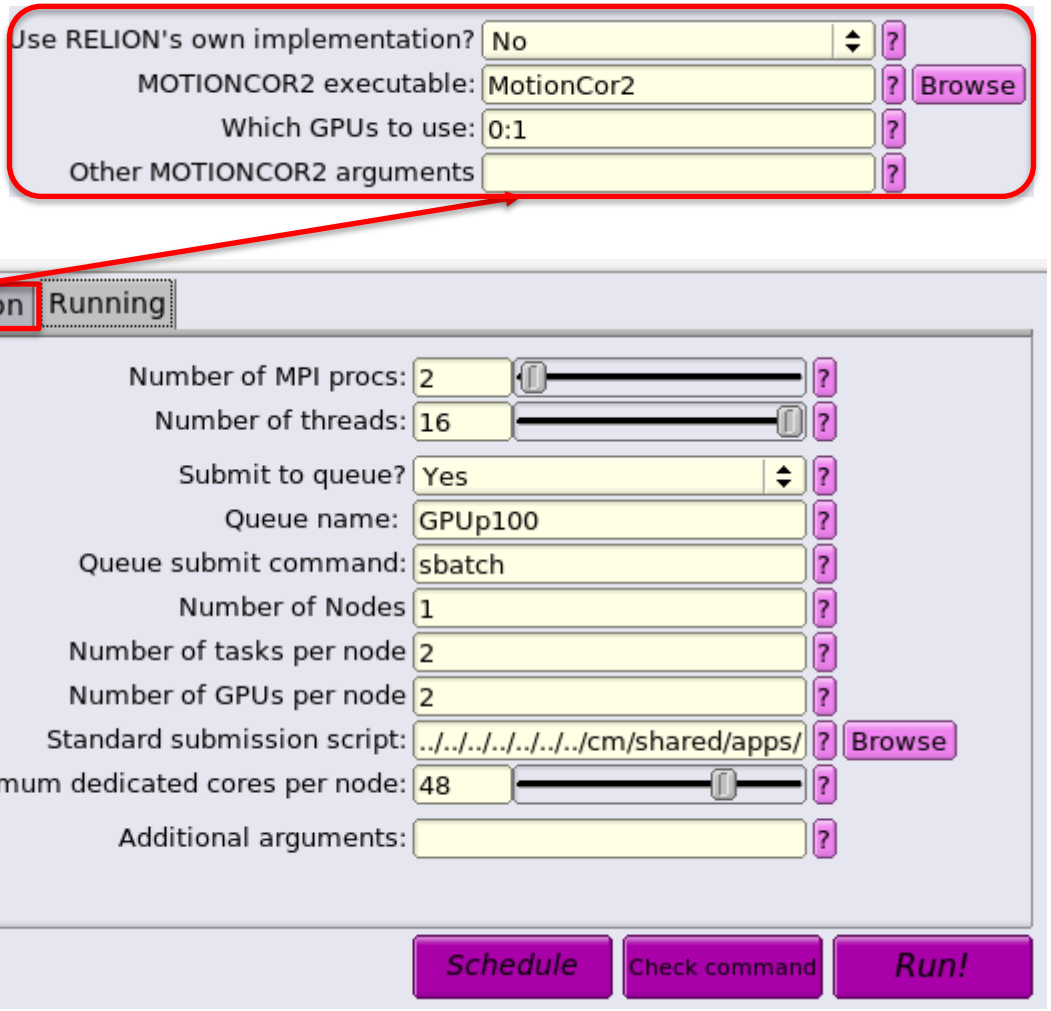

# Examples of submitting RELION job to multiple nodes

**2D classification**

**1 GPUp40 nodes,**

**1 GPU cards on each nodes**

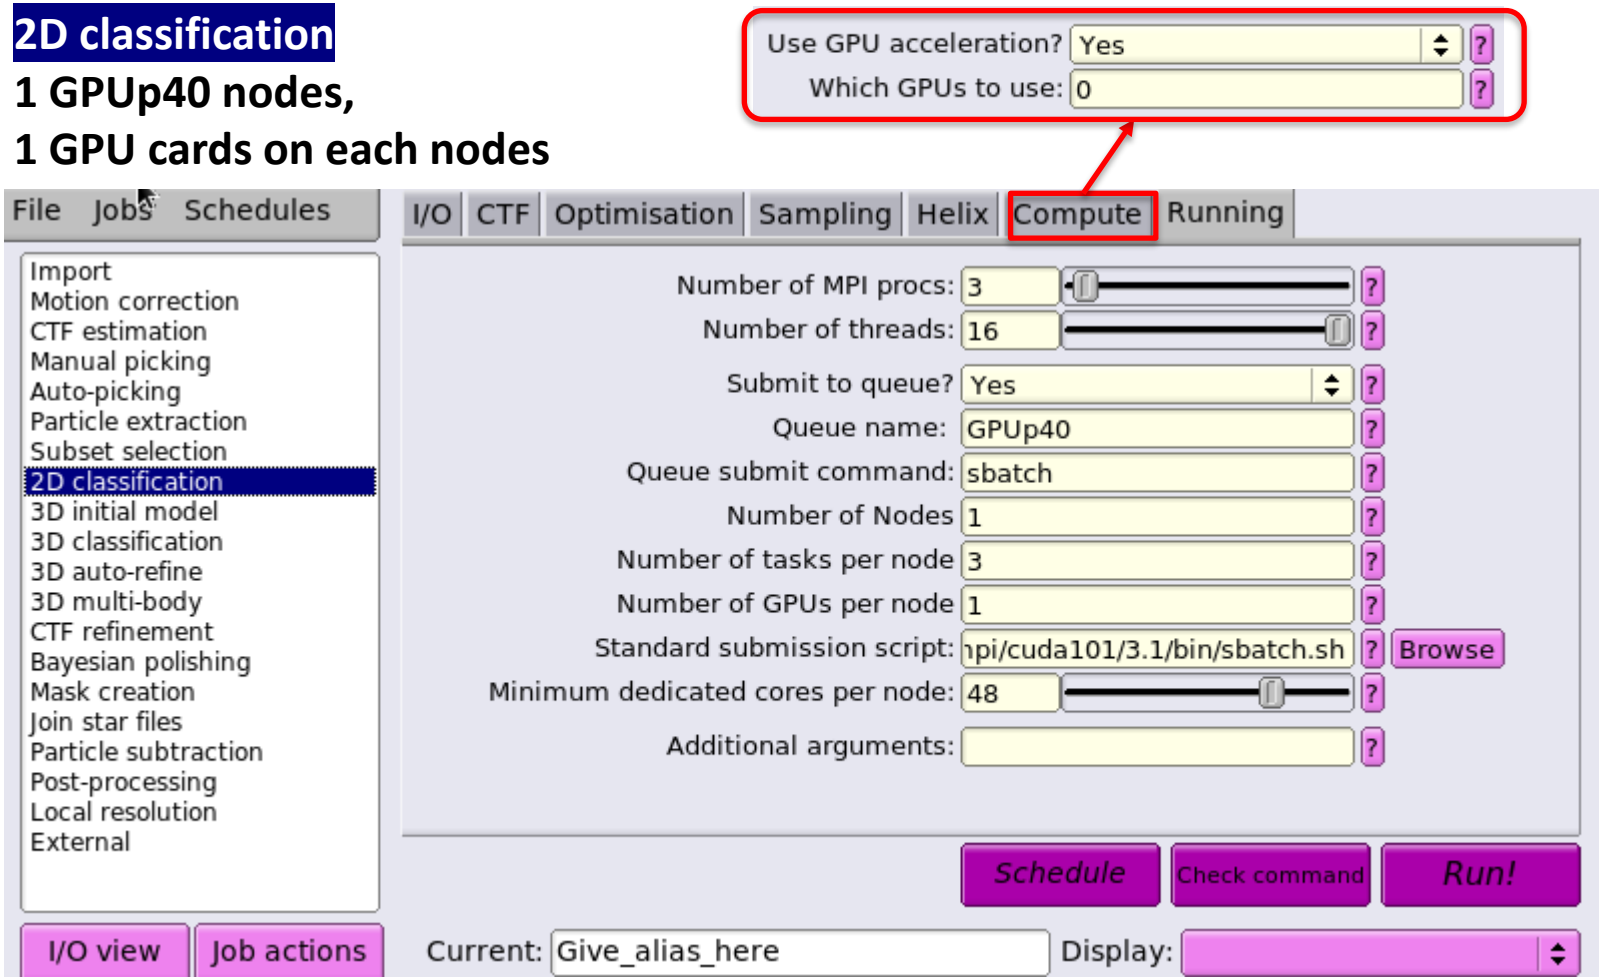

# Examples of submitting RELION job to multiple nodes

**2D classification**

**4 GPUp40 nodes,**

**1 GPU cards on each nodes**

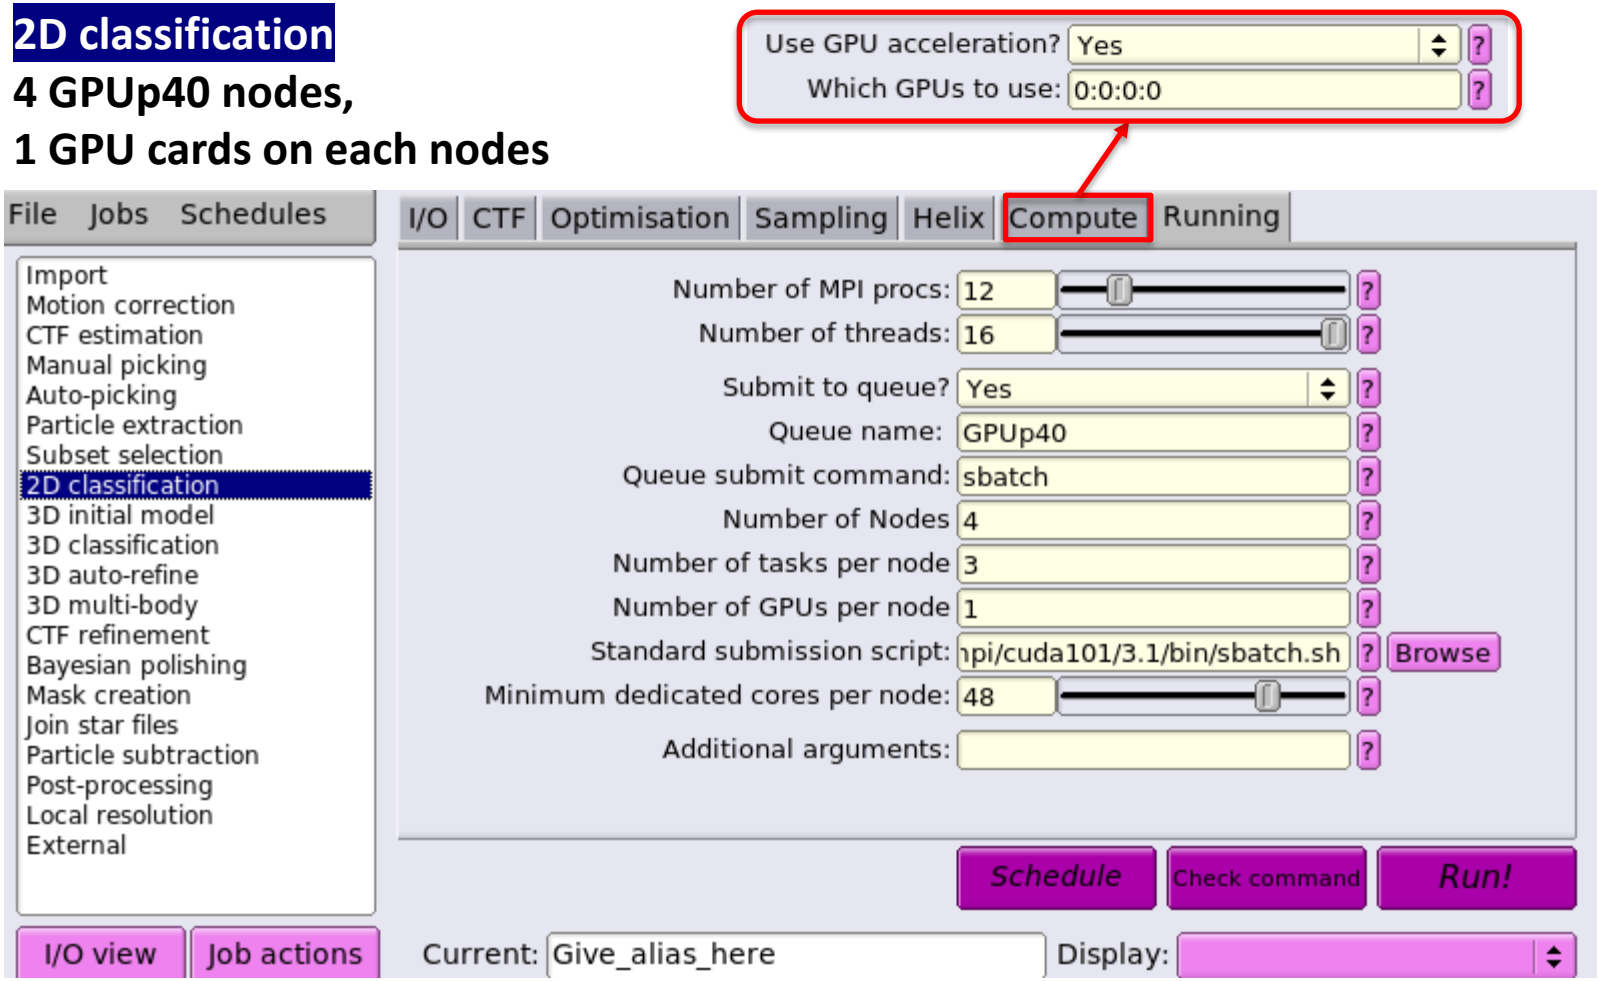

# Examples of submitting RELION job to multiple nodes

**2D classification**

**1 GPUp100 nodes,**

**2 GPU cards on each nodes**

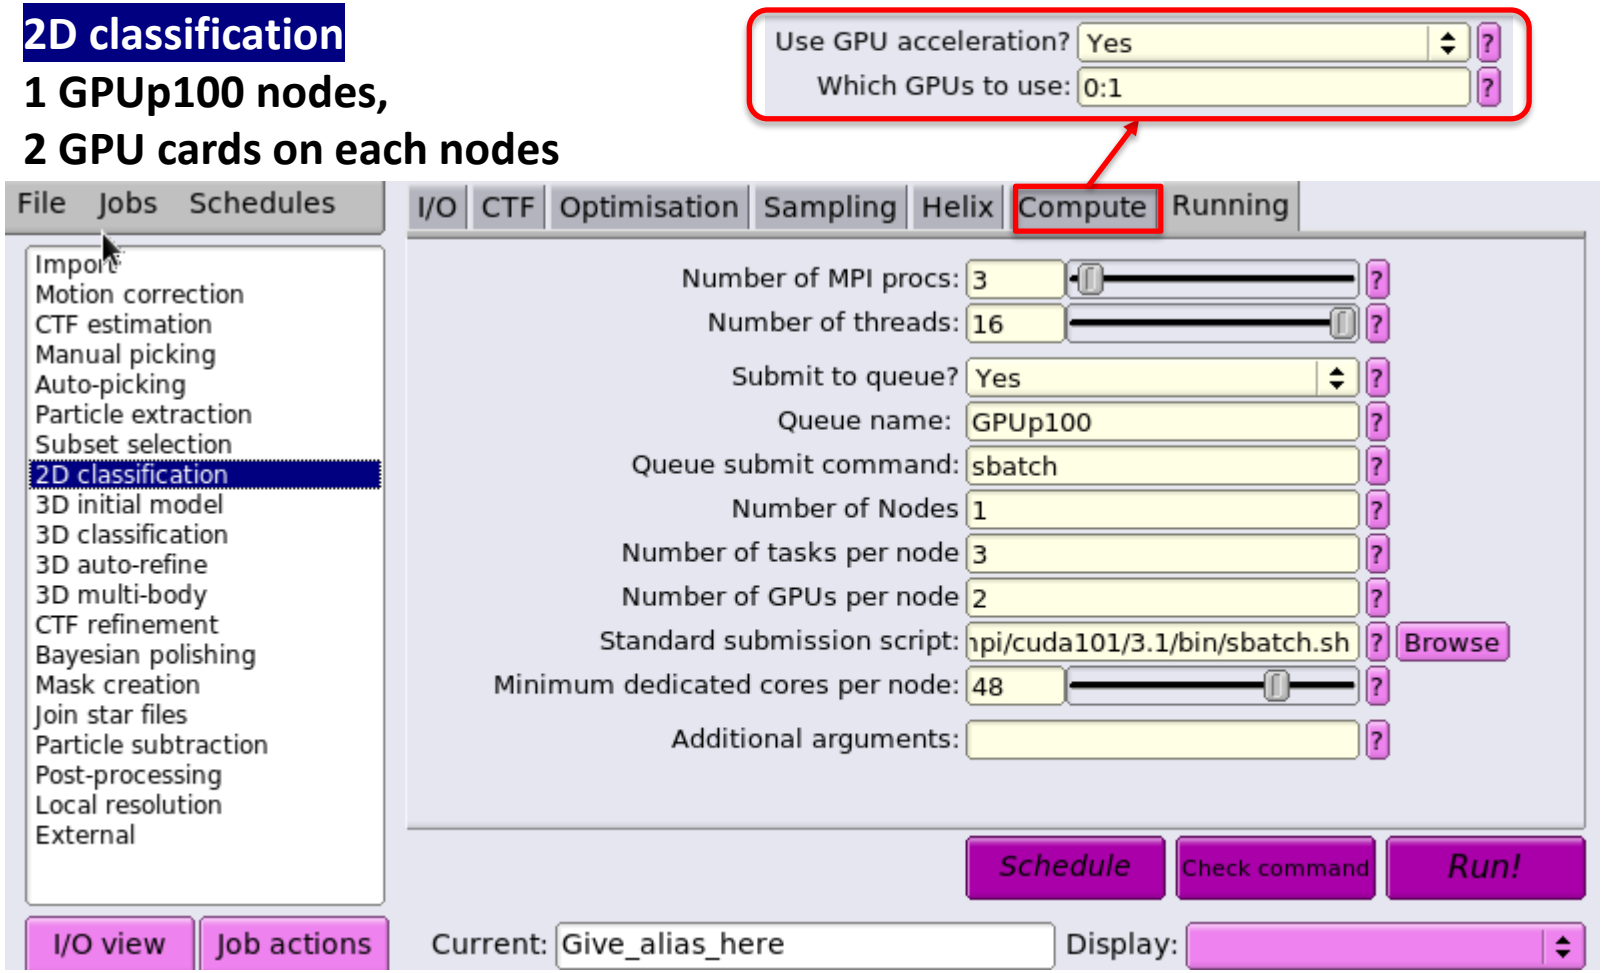

# Examples of submitting RELION job to multiple nodes

**2D classification**

**2 GPUp100 nodes,**

**2 GPU cards on each nodes**

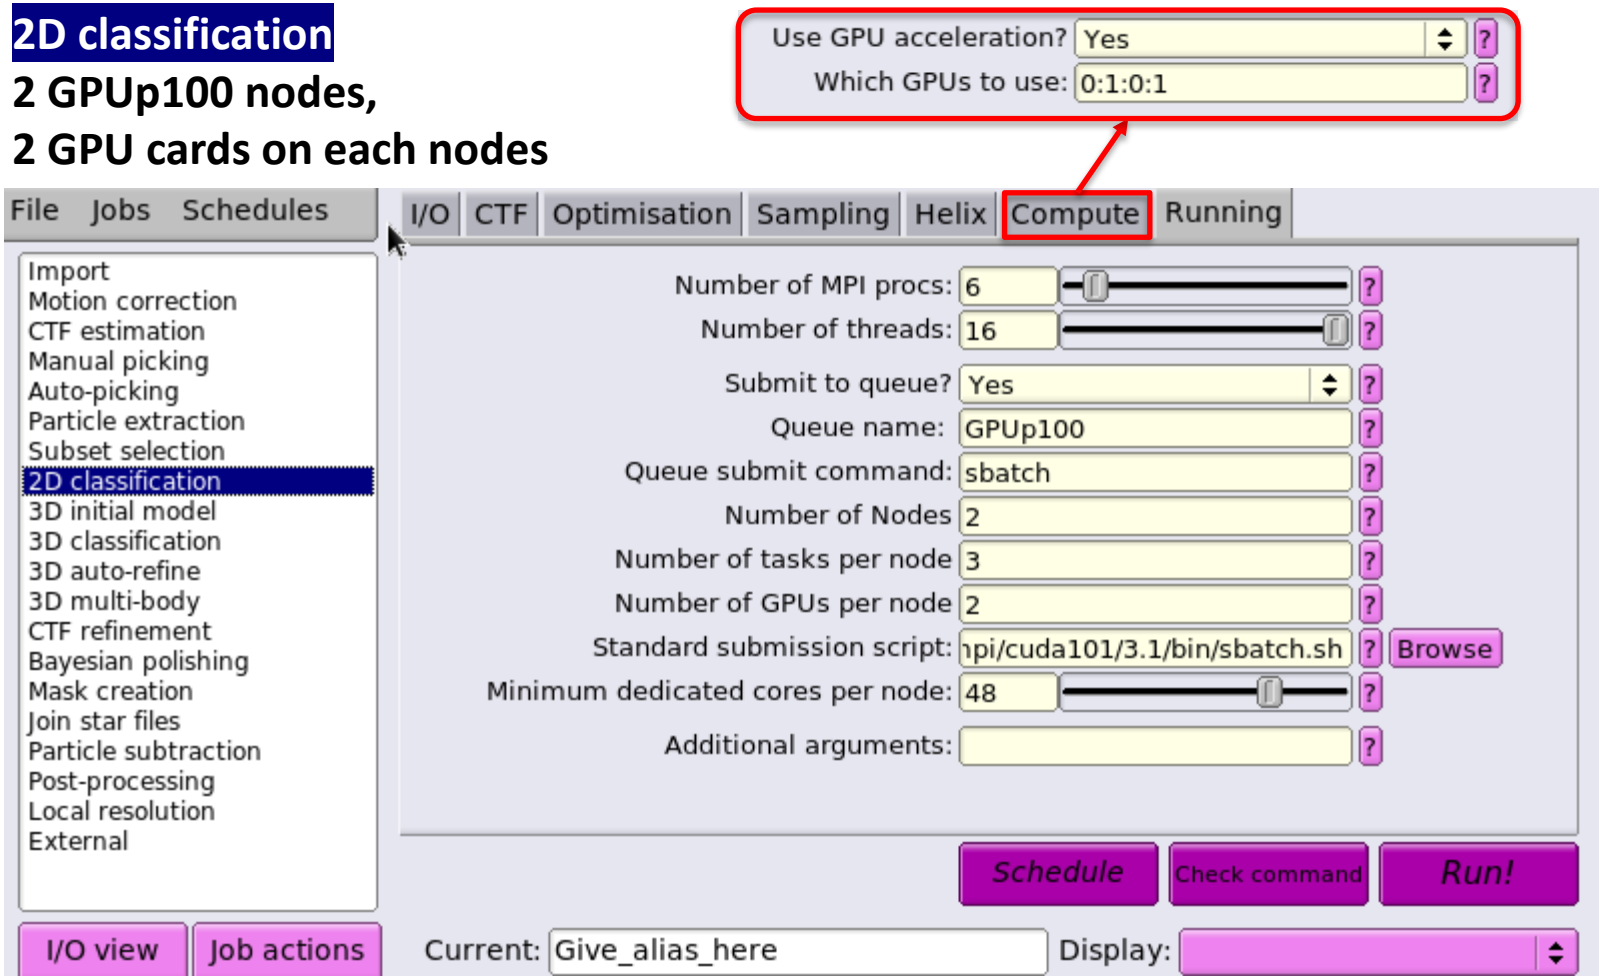

# Examples of submitting RELION job to multiple nodes

**2D classification**

**3 GPUp100 nodes,**

**2 GPU cards on each nodes**

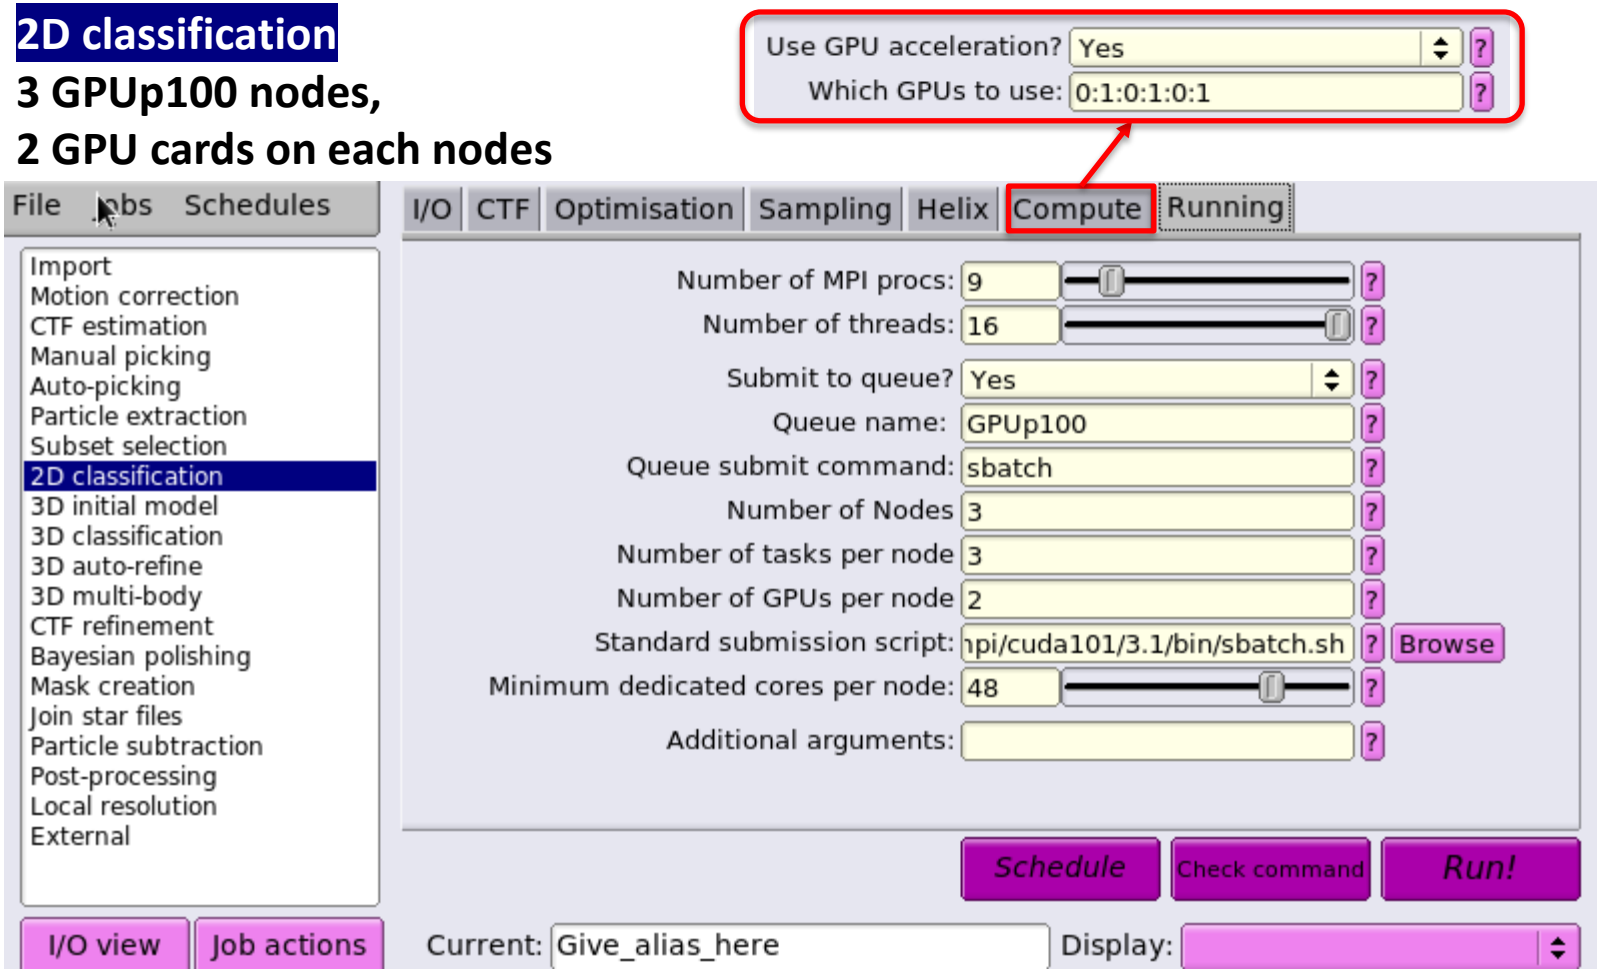

# Examples of submitting RELION job to multiple nodes

**2D classification**

**4 GPUp100 nodes,**

**2 GPU cards on each nodes**

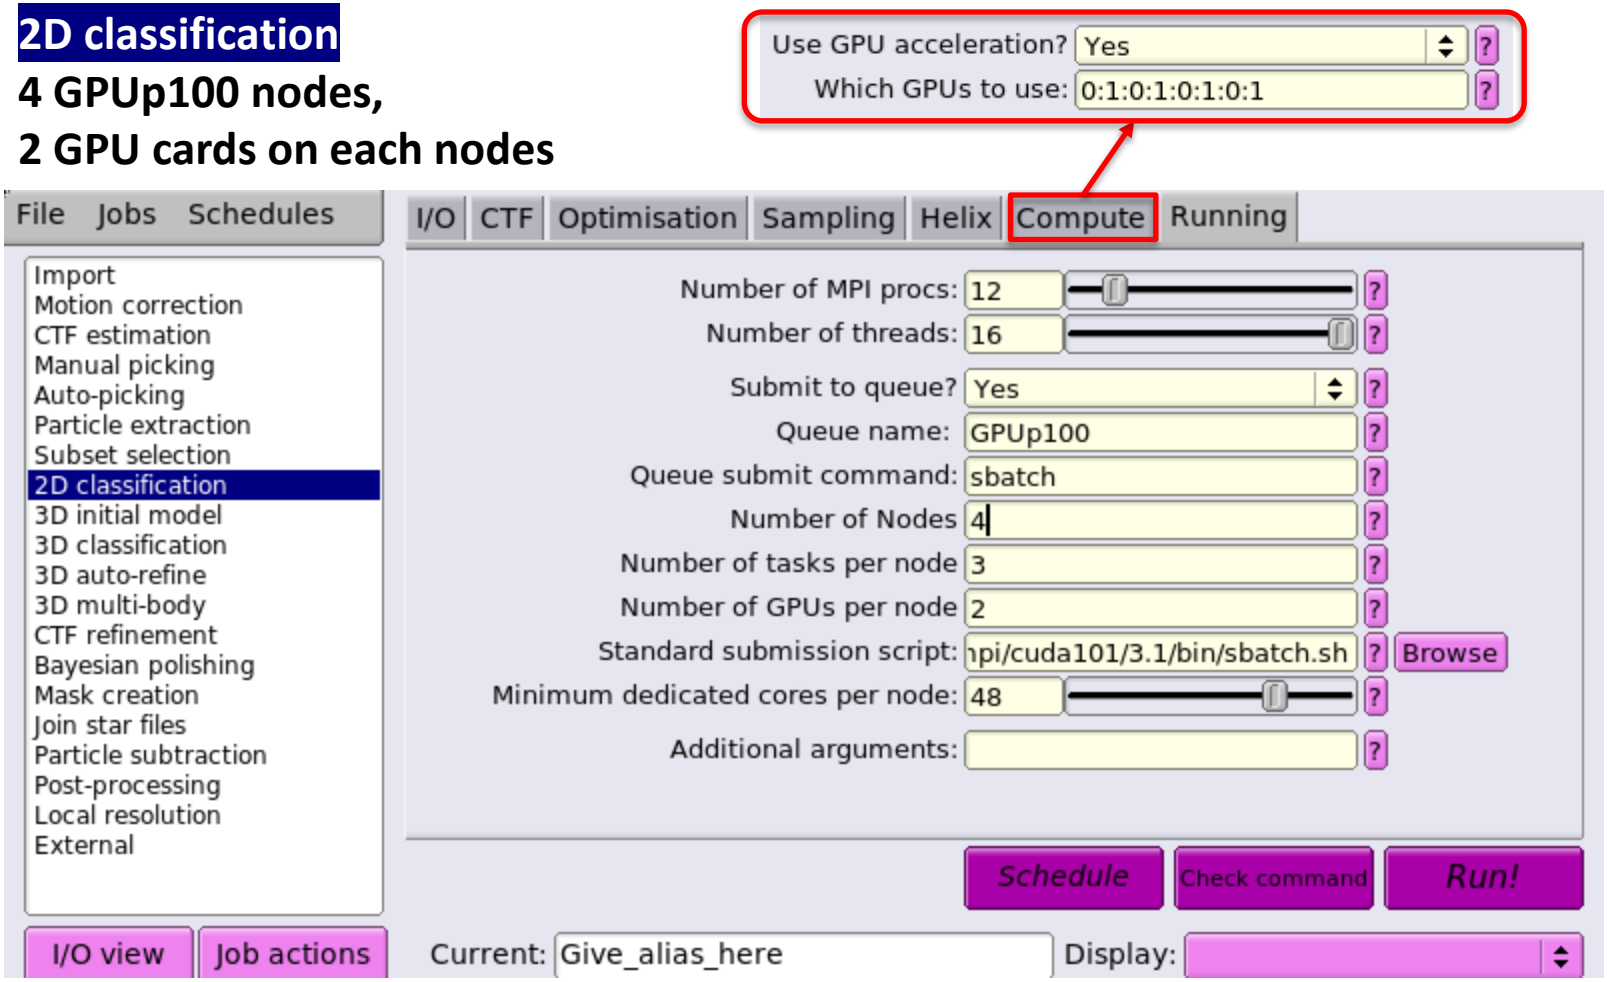

# Examples of submitting RELION job to multiple nodes

RELION 3.1 Manual

**Number of MPI procs** 3

(Note that *when using the EM-algorithm*, 2D classification, 3D classification, 3D initial model and 3D auto-refine use one MPI process as a master, which does not do any calculations itself, but sends jobs to the other MPI processors. Therefore, we often run the EM-algorithm using a single worker MPI process on each of the available GPUs, so we specify 3 here to include the master and one workers on each of the two GPUs.)

The general rule for 2D classification, 3D classification, 3D initial model, and 3D auto-refine would be:

```
IF Num_of_GPU_per_Node == 1:
        Num_of_MPI = 3 * Num_of_Nodes;
IF Num_of_GPU_per_Node >1:
        Num_of_MPI = (Num_of_GPU_per_node + 1) * Num_of_Nodes;

Num_of_threads =  minimum_num_of_cores_per_Node * Num_of_Nodes / Num_of_MPI ;

Num_of_tasks_per_Node = Num_of_MPI / Num_of_Nodes;
```

For other steps, eg, motion correction:

```
Num_of_MPI = Num_of_GPU_per_Node * Num_of_Nodes;
```

# Standard RELION submit script

```
$ cat /cm/shared/apps/relion/gcc/openmpi/cuda101/3.1/bin/sbatch.sh
#!/bin/bash

#SBATCH --job-name RELIONJob
#SBATCH -p XXXqueueXXX          # partition (queue)
#SBATCH --nodes=XXXextra1XXX
#SBATCH --ntasks-per-node=XXXextra2XXX
##SBATCH --mem 250000           # Memory Requirement (MB)
##SBATCH -t 2-2:0:00
#SBATCH -o XXXoutfileXXX
#SBATCH -e XXXerrfileXXX
#SBATCH --gres=gpu:XXXextra3XXX

module load shared relion/gcc/openmpi/cuda101/3.1

mpiexec -n XXXmpinodesXXX XXXcommandXXX
```

# Create and use user defined submit script

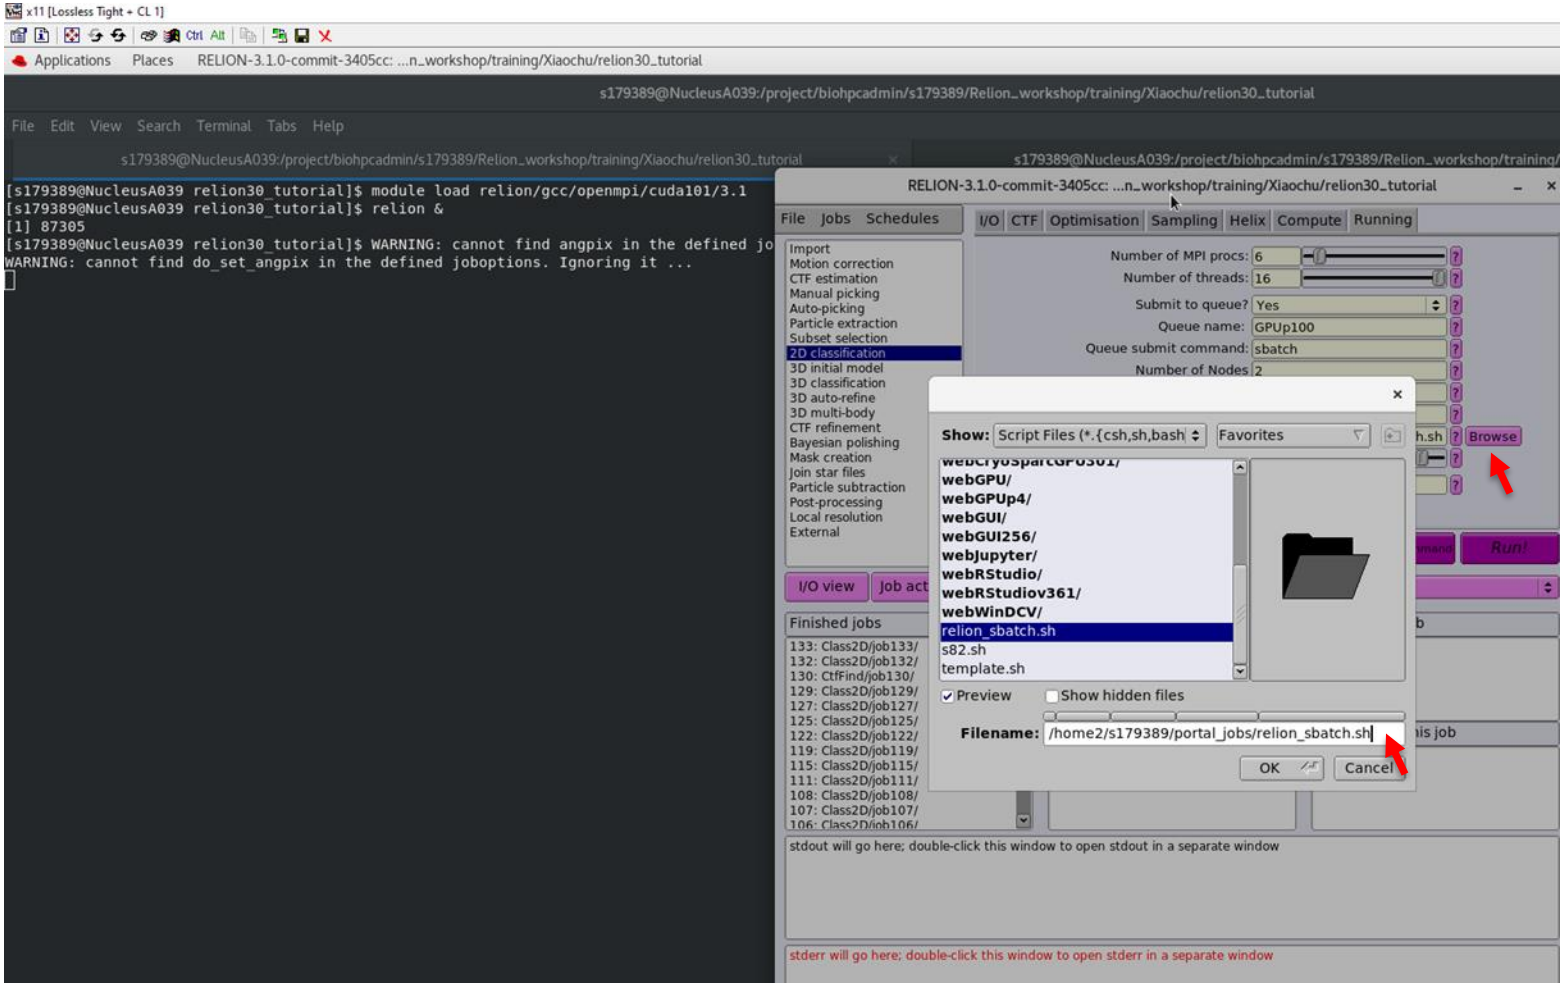

# Customize your RELION submit script

```
$ cat ~/portal_jobs/relion_sbatch.sh
#!/bin/bash

#SBATCH --job-name RELIONJob
#SBATCH -p XXXqueueXXX          # partition (queue)
#SBATCH --nodes=XXXextra1XXX
#SBATCH --ntasks-per-node=XXXextra2XXX
##SBATCH --mem 250000           # Memory Requirement (MB)
##SBATCH -t 1-2:0:00
#SBATCH -o XXXoutfileXXX
#SBATCH -e XXXerrfileXXX
#SBATCH --gres=gpu:XXXextra3XXX
#SBATCH --mail-type ALL
#SBATCH --mail-user xiaochu.lou@utsouthwestern.edu

module load shared relion/gcc/openmpi/cuda101/3.1
mpiexec -n XXXmpinodesXXX XXXcommandXXX
```

# OnDemand cryoSPARC and cryoSPARC Live

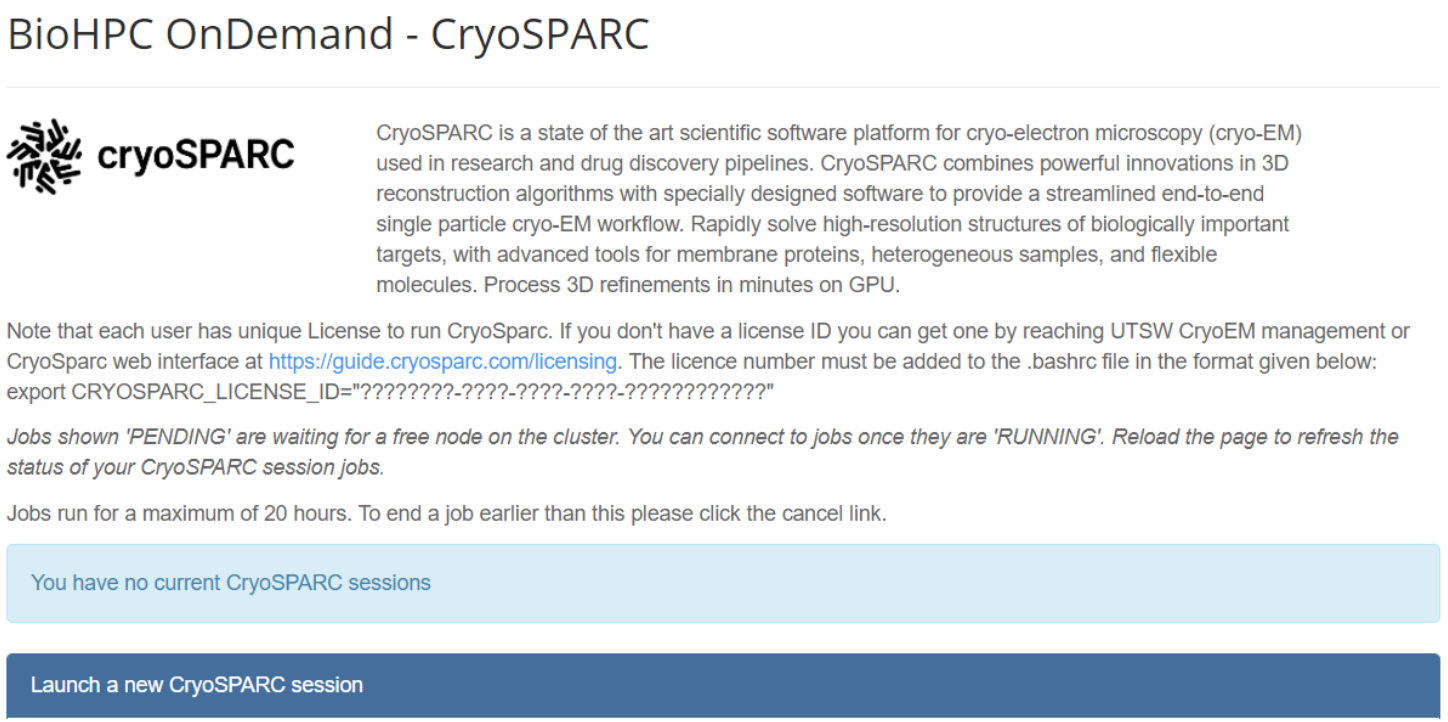


**Guides about cryoSPARC usage**:

1, BioHPC portal Guides page: [CryoEM resource on BioHPC (swmed.edu)]

# OnDemand cryoSPARC and cryoSPARC Live

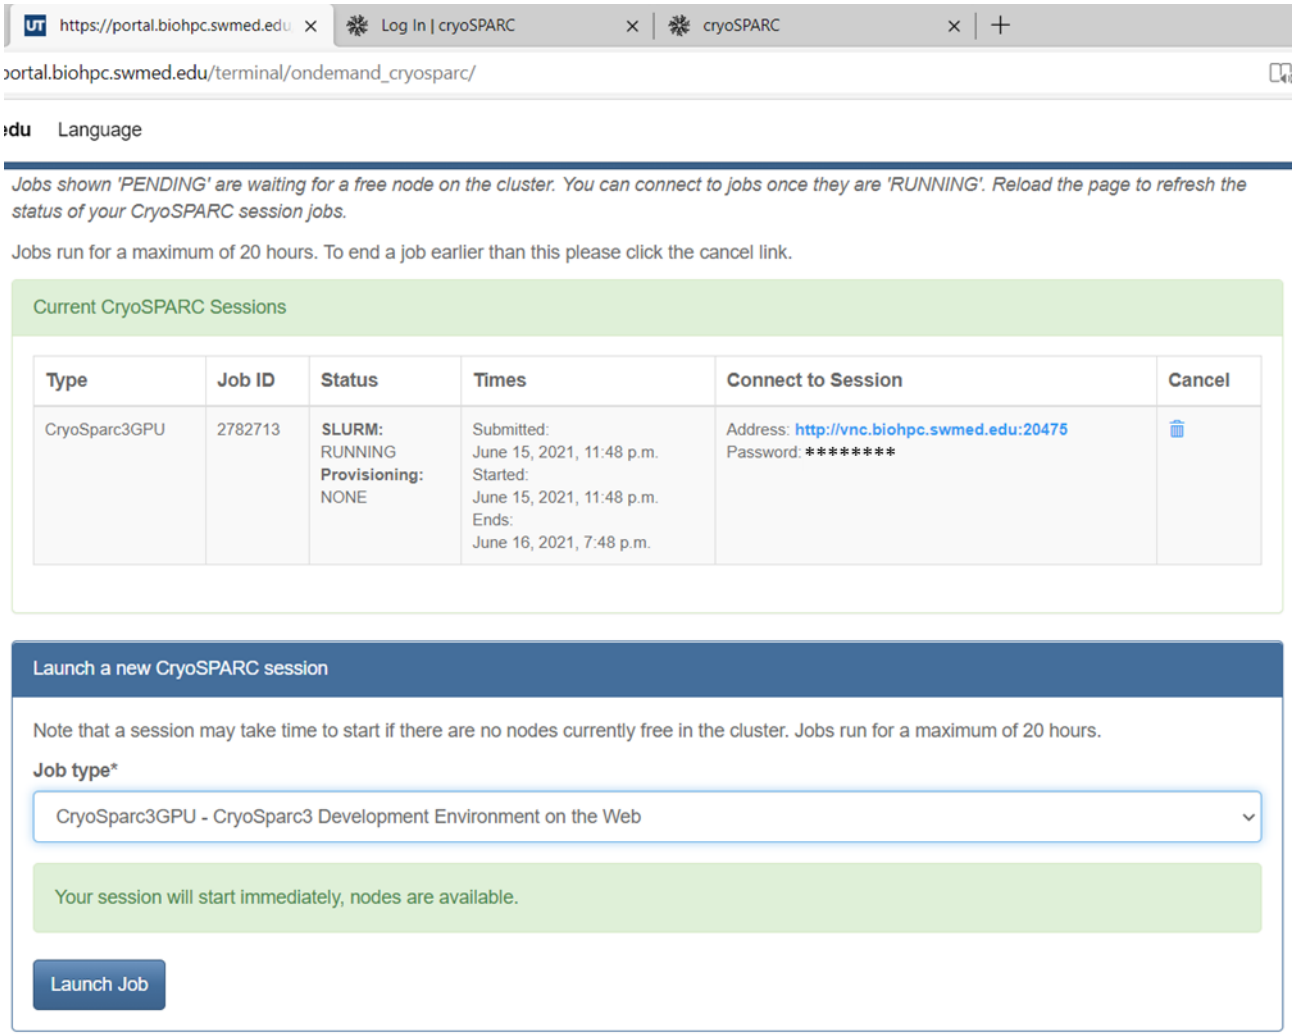

# cryoSPARC and cryoSPARC Live

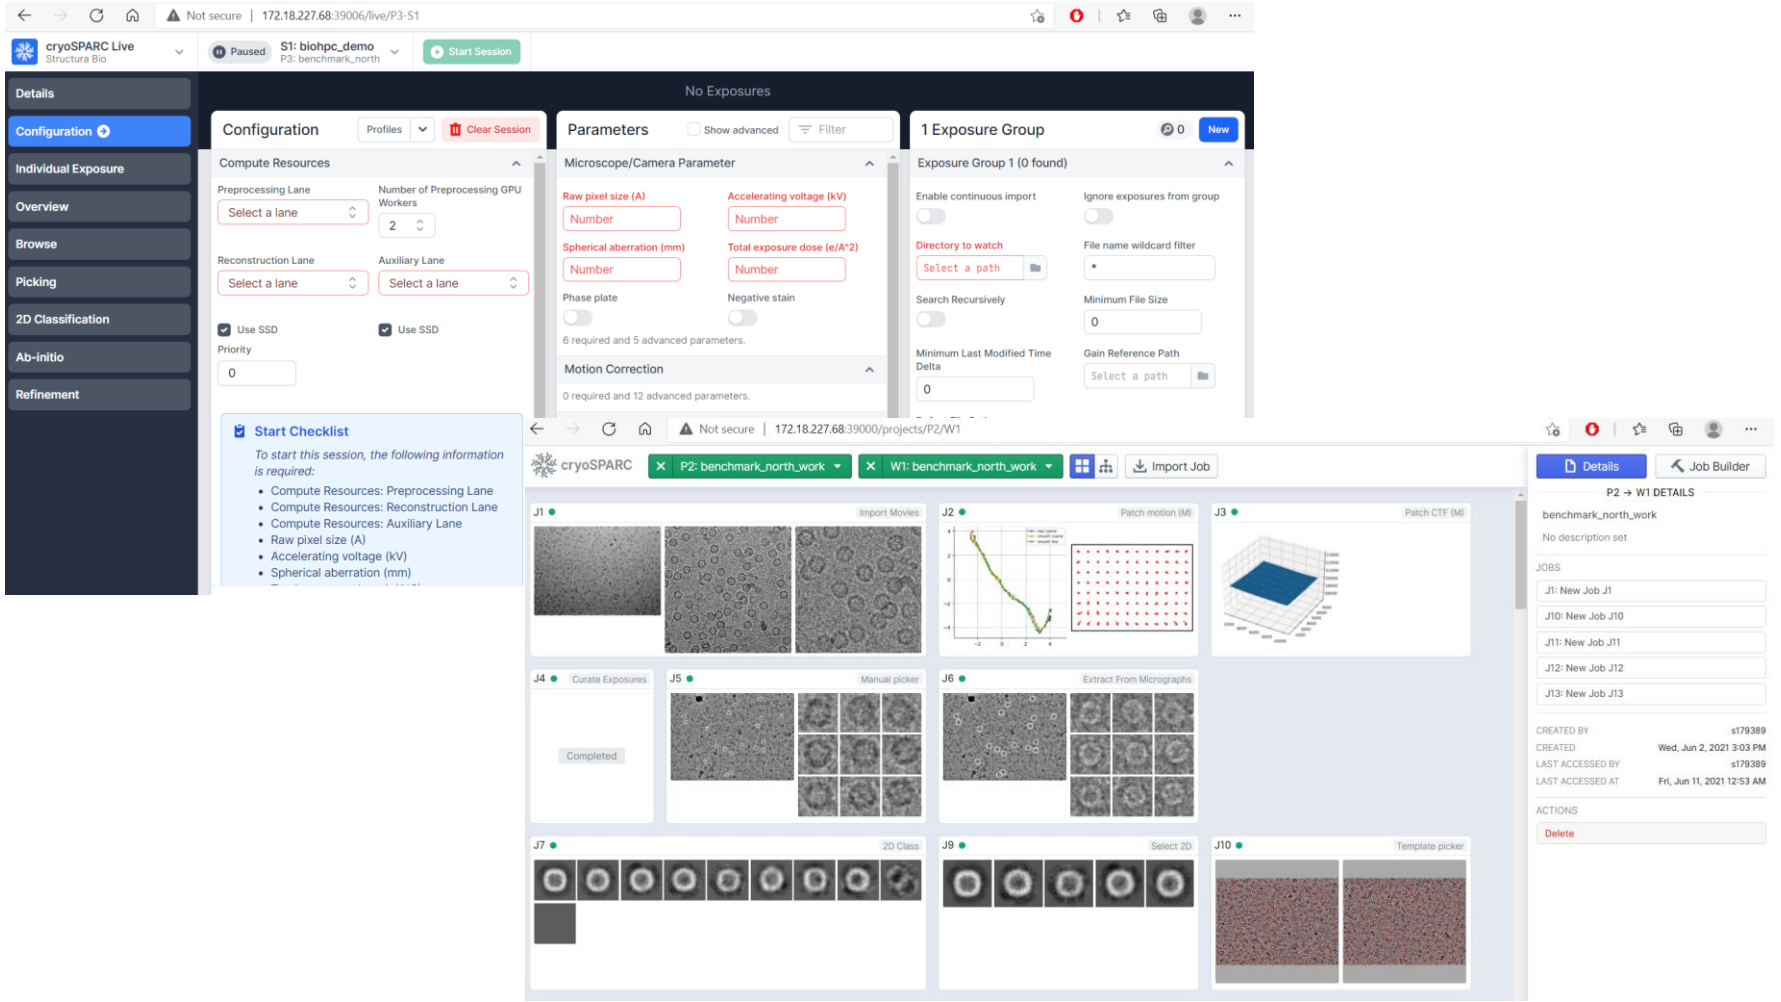

# Submit slurm job to start cryoSPARC and cryoSPARC Live

```
$ cat ~/portal_jobs/cryosparc/cryosparc_sbatch_v100s.sh
#!/bin/bash
#SBATCH --job-name="Cryosparc3"
#SBATCH --partition=GPUv100s
#SBATCH --nodes=1
#SBATCH --gres=gpu:1  # Number of GPUs(per node)
#SBATCH --ntasks=1
#SBATCH --time=2-02:00:00
#SBATCH --output="logs.cryosprac3.%j.%N.txt"
#SBATCH --error=errors.cryosparc3.%j.%N.txt

module load cryosparc/3.1.0-singularity
export no_proxy="localhost"
export CUDA_VISIBLE_DEVICES=0
cryosparc start
tail -f ~/cryosparc-v3/run/command_core.log
```

# Submit slurm job to start cryoSPARC and cryoSPARC Live

```
# Submit to start the job
[s179389@Nucleus005 ~]$ sbatch  ~/portal_jobs/cryosparc/cryosparc_sbatch_v100s.sh
Submitted batch job 2776433

# Cancel the job after finished
[s179389@Nucleus005 ~]$ module load cryosparc/3.1.0-singularity
[s179389@Nucleus005 ~]$ squeue -u s179389
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
           2776433     GPUp4 Cryospar  s179389  R       5:16      1 NucleusC014
[s179389@Nucleus005 ~]$ cryosparc_canceljob  2776433
user confirmed of job: 2776433 running on node: NucleusC014
CryoSPARC is running.
Stopping cryoSPARC
app: stopped
command_core: stopped
command_rtp: stopped
command_vis: stopped
liveapp: stopped
webapp: stopped
database: stopped
Shut down
```

Note: before first time use:
Create a file named ~/.cryopwd and write your Cryosparc password to the first line without any space.

# Initializing and Use SBGrid on BioHPC

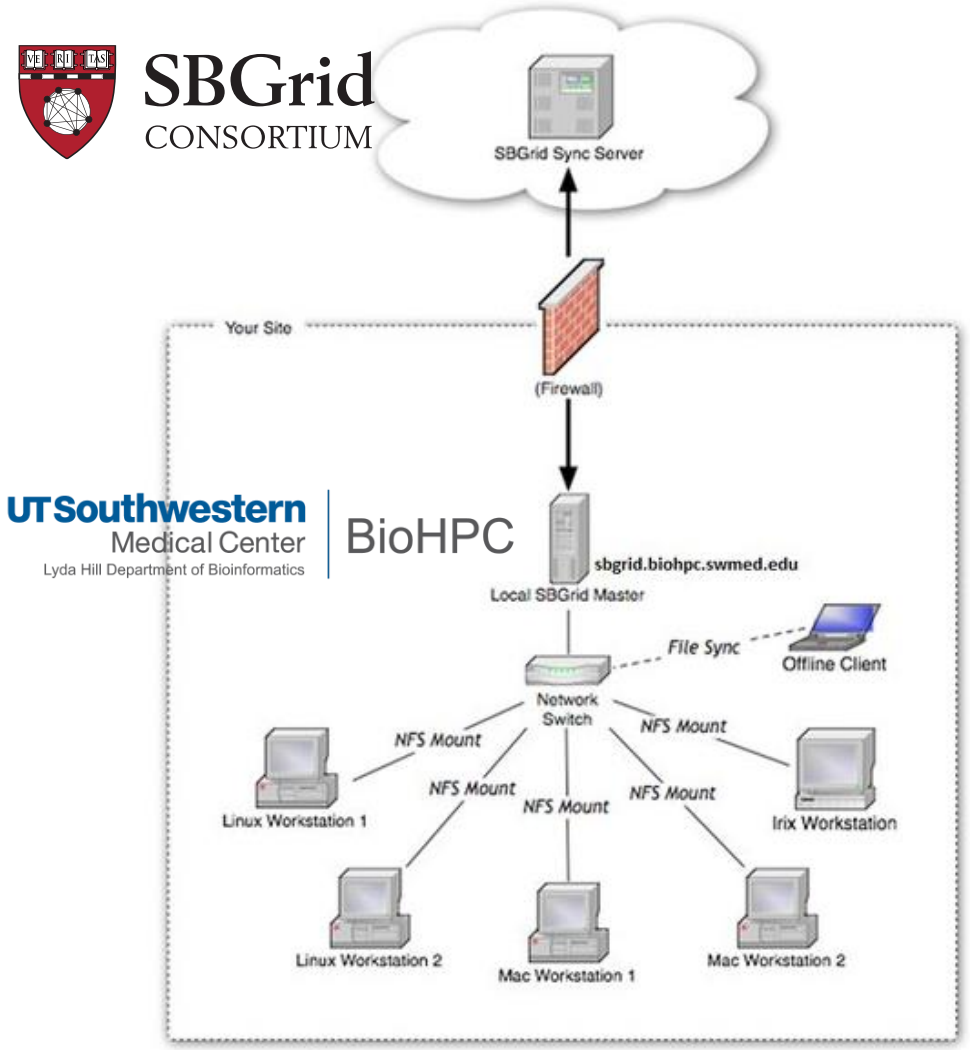


## License / Access

SBGRID is available to SBGRID member labs under the license below:

**********
**********

Your use of the applications contained in the /programs directory constitutes acceptance of the terms of the SBGrid License Agreement included in the file /programs/share/LICENSE. The applications distributed by SBGrid are licensed exclusively to member laboratories of the SBGrid Consortium.

**********
**********

Information about joining SBGRID can be obtained from the SBGRID consortium via: https://sbgrid.org/join/

**Guides about SBGrid usage**:

1, BioHPC portal Guides page: SBGrid on BioHPC (swmed.edu)

# Initializing and Use SBGrid on BioHPC

```
[s179389@Nucleus005 ~]$ cat ~/.bashrc
# .bashrc

# Source global definitions
if [ -f /etc/bashrc ]; then
        . /etc/bashrc
fi

# User specific aliases and functions
module load slurm/16.05.8 shared

alias data='cd /project/biohpcadmin/s179389'

#sbgrid initializing
. /programs/sbgrid.shrc
```

# Initializing and Use SBGrid on BioHPC

```
********************************************************************************
                  Software Support by SBGrid (www.sbgrid.org)
********************************************************************************
 Your use of the applications contained in the /programs  directory constitutes
 acceptance of  the terms of the SBGrid License Agreement included  in the file
 /programs/share/LICENSE.  The applications  distributed by SBGrid are licensed
 exclusively to member laboratories of the SBGrid Consortium.
             Run sbgrid-accept-license to remove the above message.
********************************************************************************
 SBGrid was developed with support from its members, Harvard Medical School,
 HHMI, and NSF. If use of SBGrid compiled software was an important element
 in your publication, please include the following reference in your work:

 Software used in the project was installed and configured by SBGrid.
 cite: eLife 2013;2:e01456, Collaboration gets the most out of software.
********************************************************************************
 SBGrid installation last updated: 2021-03-05 (Update available)
 Please submit bug reports and help requests to:       <bugs@sbgrid.org>  or
                                                       <http://sbgrid.org/bugs>
           For additional information visit https://sbgrid.org/wiki
********************************************************************************
                              SBGrid Announcements
 - There are known issues with MacOS 11.0 “Big Sur” and some SBGrid
   applications. We recommend not upgrading Apple computers to 11.0
   until these issues can be addressed.
     see https://sbgrid.org/wiki/big_sur for more info.
********************************************************************************
[s179389@Nucleus006 ~]$ which relion
/programs/x86_64-linux/system/sbgrid_bin/relion
[s179389@Nucleus006 ~]$ module load relion/gcc/openmpi/cuda101/3.1
[s179389@Nucleus006 ~]$ which relion
/cm/shared/apps/relion/gcc/openmpi/cuda101/3.1/bin/relion
```

# Use IMOD and PEET on BioHPC

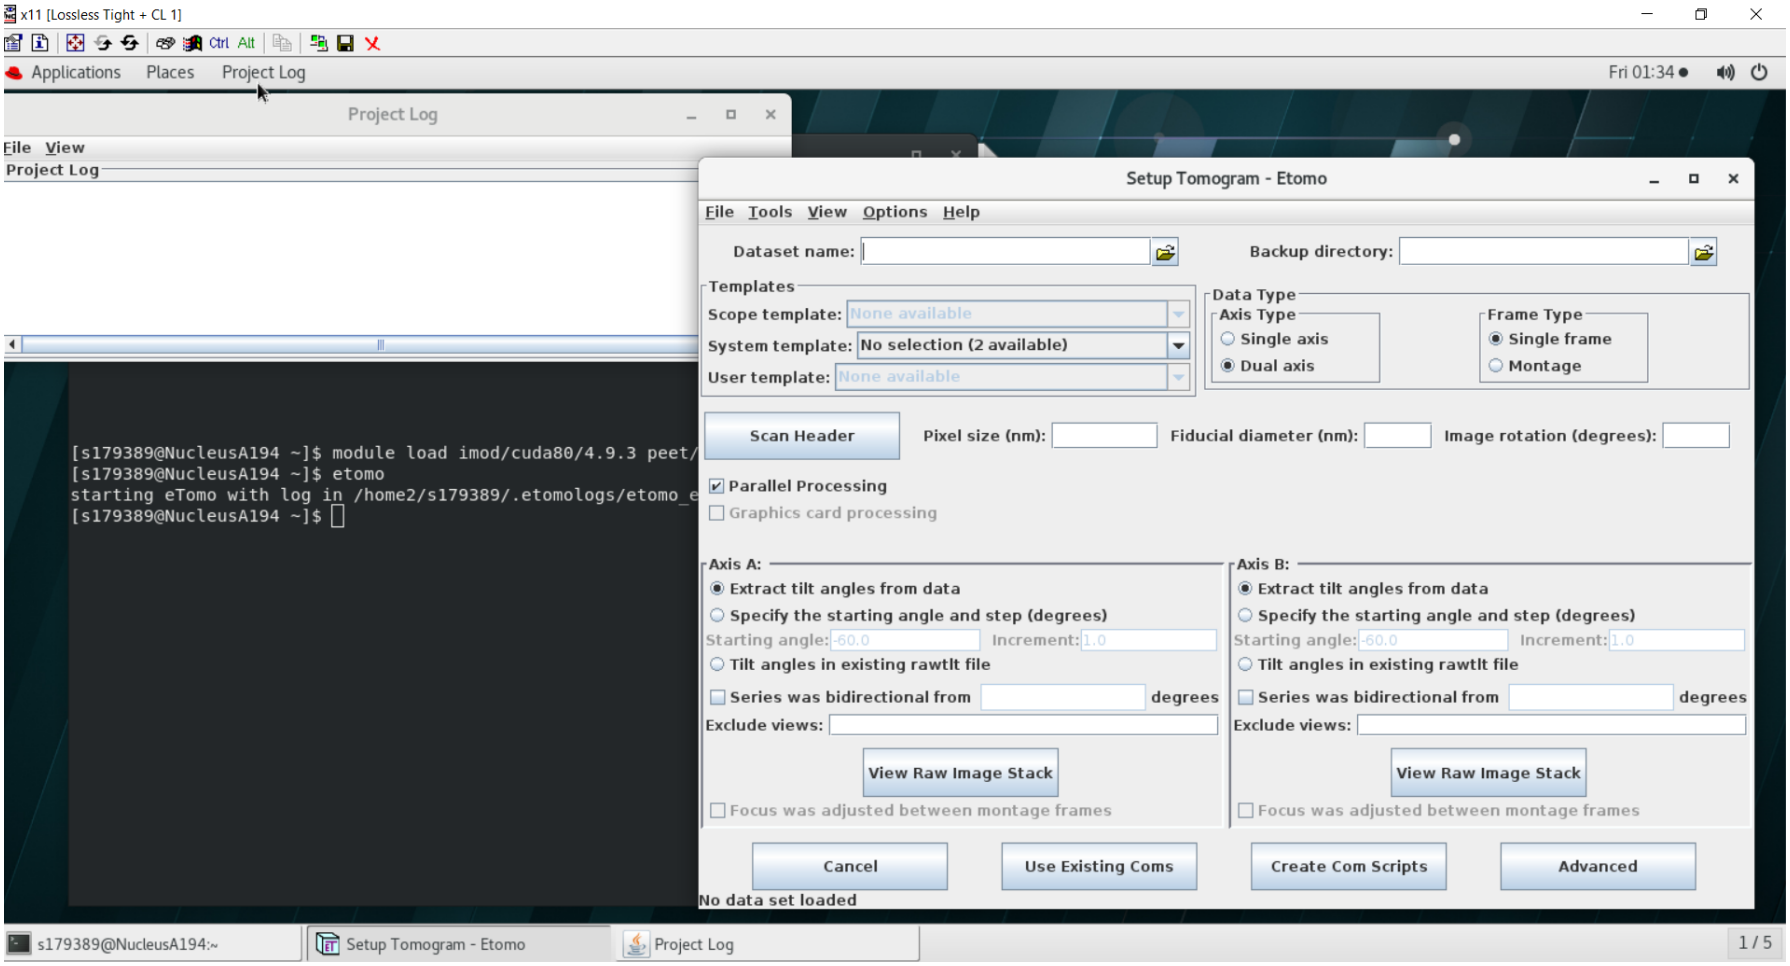

UT Southwestern Medical Center
Lyda Hill Department of Bioinformatics

BioHPC

*Questions? Comments?*

*Email:* biohpc-help@utsouthwestern.edu

*Thanks!*