CryoSPARC Cryo-EM Data Processing on ACES

July 14, 2023
ACES Workshop
Michael Dickens

High Performance
Research Computing
DIVISION OF RESEARCH

ACES
ACCELERATING COMPUTING
FOR EMERGING SCIENCES

CryoSPARC on ACES

- CryoSPARC Academic License ID
- Launch CryoSPARC on ACES Portal app
- Resources for Running CryoSPARC
  - CPU job
  - GPU job
  - SU charge rate
- Resource Monitoring
- Group data directories
- CryoSPARC Tutorial

CryoSPARC Academic License ID

- CryoSPARC provides a personal Academic License ID free of charge for Academic staff and students
- A CryoSPARC Academic License ID is required to launch the CryoSPARC HPRC portal app
- Use your academic email address for the Academic License
  - https://cryosparc.com/download

Initialize Your CryoSPARC Database

ACES CryoSPARC Portal App

portal-aces.hprc.tamu.edu

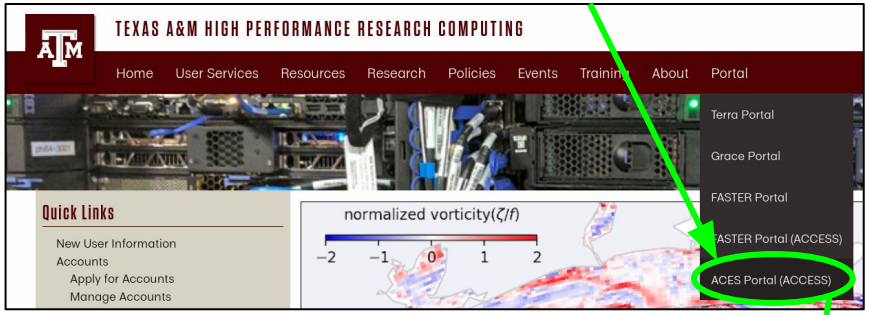

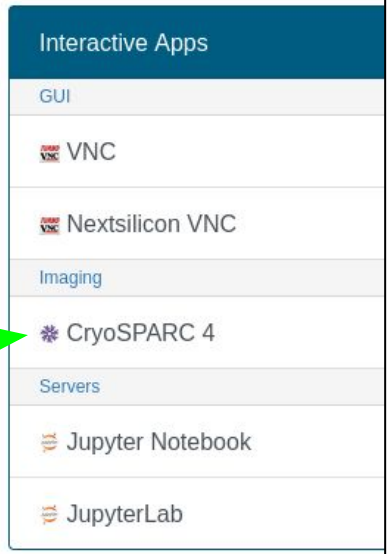

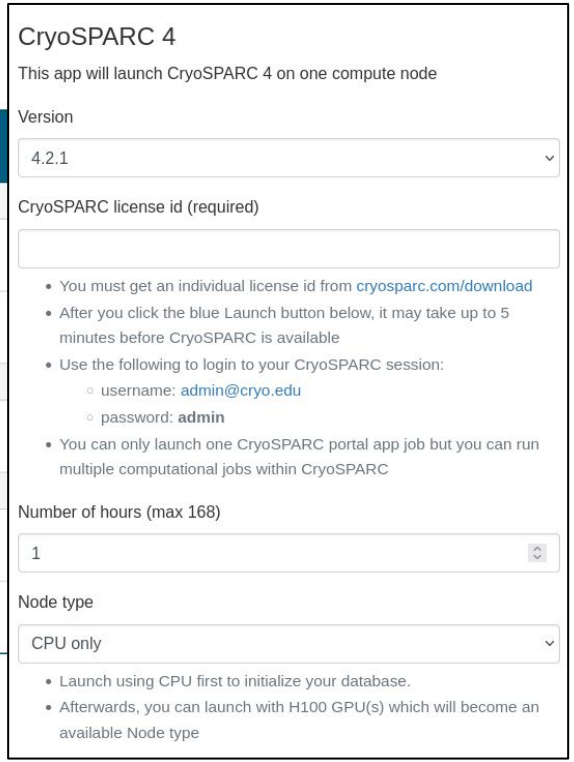

Launch CryoSPARC 4 the first time using 'CPU only' in order to initialize the database.

ACES CryoSPARC Running Jobs

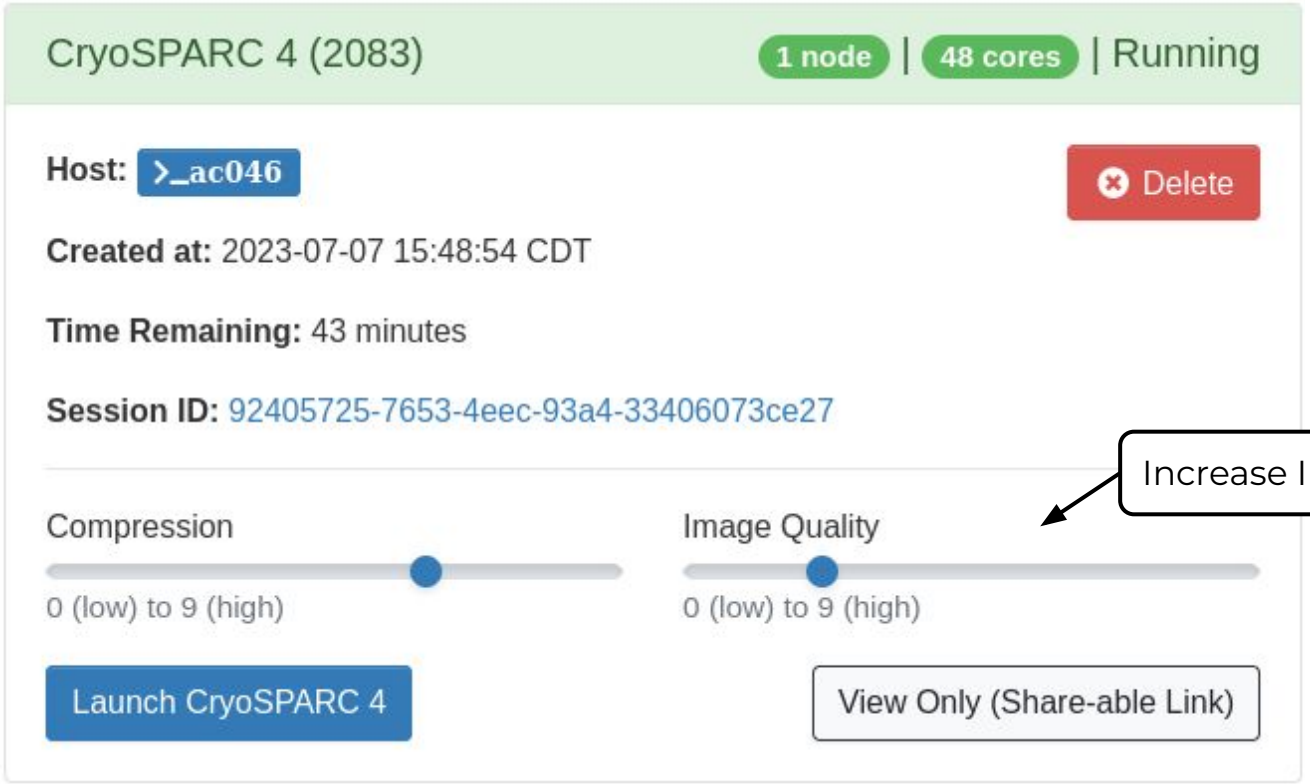

Increase Image Quality to 9

ACES CryoSPARC Portal App

portal-aces.hprc.tamu.edu

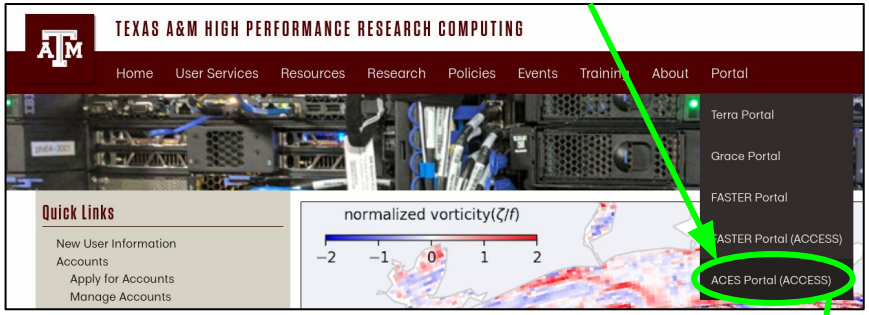

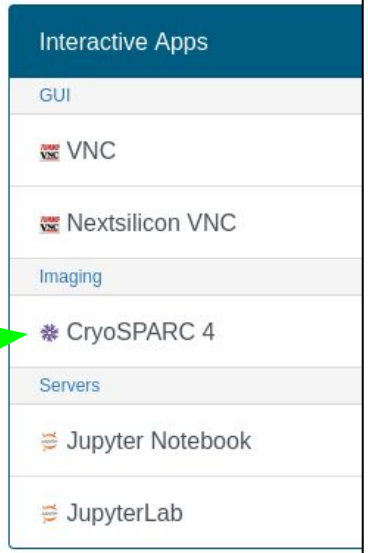

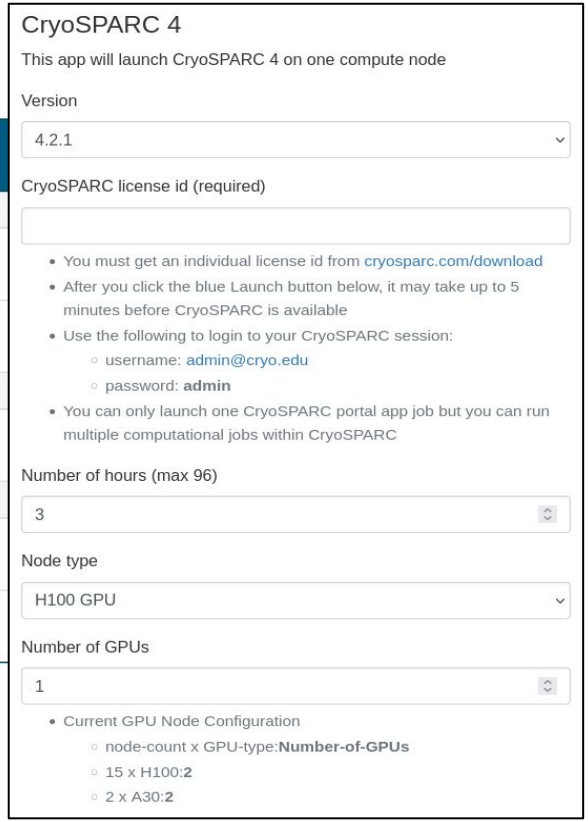

Launch CryoSPARC 4 using 1 x H100 GPU for the tutorial

Monitoring Job Resource Usage

Review CPU and GPU usage for a Job

The **`jobstats`** utility automatically runs for each CryoSPARC portal job and monitors CPU and GPU resource usage and creates a graph for each.

- CPU stats monitors all cores regardless of how many were configured for the job
- GPU stats will create a graph for each GPU that was configured for the job

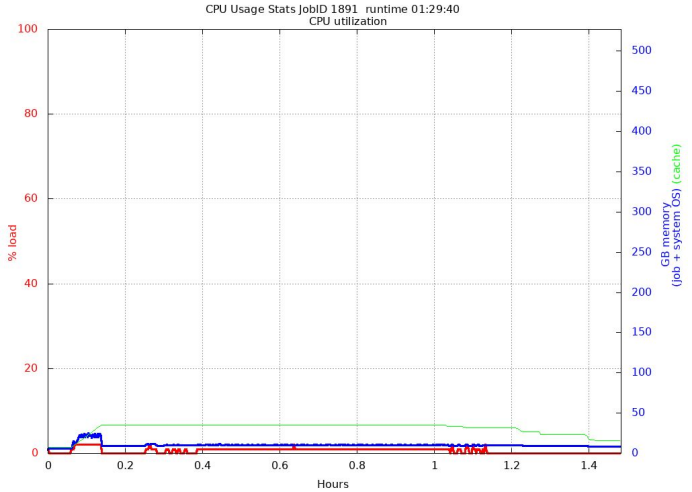

`stats_cpu.1891.png`

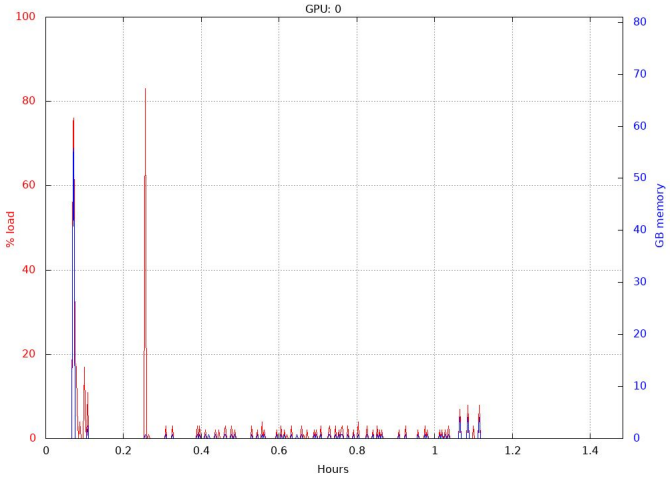

`stats_gpu.1891.png`

Review CPU and GPU usage for a CryoSPARC Job

1. At the bottom of the CryoSPARC portal app, access the CryoSPARC data root directory
2. Click 'output'
3. Click 'Modified at' to sort newest session on top
4. Click the Name of the session session you want to view based on the date
5. Click the .png image files to view the CPU and GPU usage stats

1.

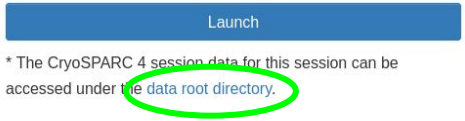

2.

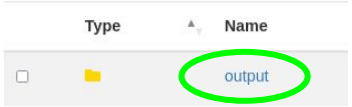

3.

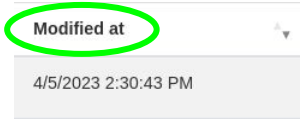

4.

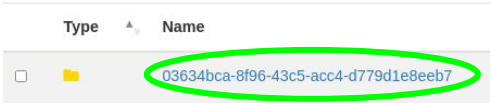

5.

Group Data Directories

- Group data directories can be used to share input files or to have a group work space
- Send a request to the HPRC helpdesk to create a group
  - help@hprc.tamu.edu
  - provide a group name and NetIDs of members
  - group directories have their own disk quotas separate from individual users
- You will need to mount your group directory when launching the CryoSPARC portal app

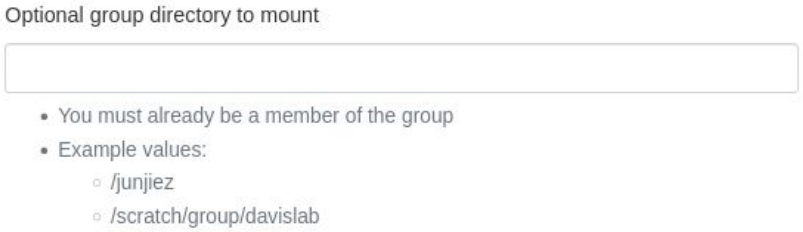

Resources for Running CryoSPARC

- CryoSPARC is available as an Interactive App on the HPRC ACES Portal
  - https://portal-aces.hprc.tamu.edu
- ACES has a total of 110 compute nodes
  - 15 x H100 GPU nodes (2 x H100 GPUs per node)
  - 2 x A30 GPU nodes (2 x A30 GPUs per node)
  - 72 x CPU only nodes
  - 21 x other specialized nodes not for CryoSPARC
- SUs are charged based on type of node selected
  - CPU only = 1 SU per hour per core or per 5GB memory whichever is greater
  - GPU = same as CPU only plus additional charge for number and type of GPUs selected
- Specify enough time to allow your processing to complete
  - can request an time extension for a job at help@hprc.tamu.edu
- If you launch a job for 24 hours and you finish your work in 12 hours and *Delete* the portal job, you will be reimbursed SUs for the unused 12 hours.

Resource Limitations

- Can only launch one CryoSPARC portal session at a time
  - can run multiple CryoSPARC jobs within a portal session
- CPUs
  - Some CryoSPARC jobs do not require a GPU
- GPUs
  - The GPU queue can get busy at times and it may take over an hour for your job to launch
  - Cancel your pending job if you will be away from your computer and it hasn't started yet
- Submitted jobs can have delayed start times due to unavailability of GPUs or a scheduled system maintenance
  - use the **`maintenance`** command to see if there is a scheduled maintenance

CryoSPARC Tutorial

High Performance
Research Computing
DIVISION OF RESEARCH

https://hprc.tamu.edu

HPRC Helpdesk:

help@hprc.tamu.edu
Phone: 979-845-0219

Help us help you. Please include details in your request for support, such as, Cluster (ACES, FASTER, Grace, Terra, ViDaL), NetID (UserID), Job information (JobID(s), Location of your jobfile, input/output files, Application, Module(s) loaded, Error messages, etc), and Steps you have taken, so we can reproduce the problem.