Common Commands

Connect to the HPC

```
ssh [USER]@hpc-login.rcc.fsu.edu
```

Check storage quota

```
$ pan_quota ~
```

Show my running/pending jobs

```
$ squeue -u $(whoami)
```

Show my partitions (login required)

```
$ rcctool my:partitions
```

Show node information in a partition

```
$ sinfo -p [partition]
```

Estimate job start time

```
$ squeue –start -j [job_id]
```

Show available Slurm node features

```
$ sinfo -o %f
```

General Access Partitions

genacc_q	2 weeks	General Access
condor	90 days	No MPI / Long jobs
backfill	4 hours	Short jobs
backfill2	4 hours	Pre-emption enabled
gpu_q	2 days	gpGPU jobs
quicktest	10 min	For testing

Job States

PD	Pending
CF	Configuring
R	Running
CG	Completing
PR	Preempted
S	Suspended
TO	Timeout
F	Failed
NF	Node Failed
CA	Cancelled

Submitting Jobs

Submit a job to the Slurm Scheduler

```
$ sbatch MY_SCRIPT.sh
Submitted batch job 1234567
```

Check job status

```
$ squeue -j 1234567
```

Cancel job

```
$ scancel 1234567
```

View job stats while running

```
$ sstat -j 1234567
```

Example Submit Script

```
#!/bin/bash
#SBATCH --job-name="my_job"        Job Name
#SBATCH -n 32                      Number of cores / tasks
#SBATCH -t 2-10:30:00              Maximum job execution time (Days-H:M:S)
#SBATCH -p genacc_q                Partition to submit job to
#SBATCH --mem-per-cpu=7.8G         Memory per CPU (default is 3.9GB)
#SBATCH -o job_output.txt          Optionally specify output file name
#SBATCH -e job_errors.log          Optionally separate error output file
#SBATCH -N 16                      Number of physical nodes to allocate
#SBATCH -C "YEAR2012,intel"        Run job on nodes with specific features
#SBATCH --mail-type="ALL"          Email ("NONE","BEGIN","END","REQUEUE")

module load intel                  Load kernel modules and run program
srun my_program.sh < input.txt     MPI jobs should use srun, not mpirun
```

FLORIDA STATE UNIVERSITY
INFORMATION TECHNOLOGY SERVICES
Research Computing Center

https://rcc.fsu.edu @fsurcc support@rcc.fsu.edu