

Batch Nodes And Queues

In 2021 LCLS switch to the SLURM batch system.

Information on submitting jobs to the SLURM system at LCLS can be found on this page: [Submitting SLURM Batch Jobs](#)

Information on the Automatic Run Processing system (ARP) can be found on this page: [Automatic Run Processing \(ARP\)](#). This is also usable at sites like NERSC and SDF.

A "cheat sheet" showing similar commands on LSF and SLURM can be found here: <https://slurm.schedmd.com/rosetta.pdf>

Refer to the table below for the batch resources available in psana. Submit your job from an interactive node (where you land after doing `ssh psana`). All nodes in the queues listed below run RHEL7. By submitting from an interactive node, also running RHEL7, you will ensure that your job inherits a RHEL7 environment.

Note 1: Jobs for the **current experiment** can be submitted to fast feedback (FFB) queues, which allocate resources for the most recent experiments. The FFB queues in the tables below are for LCLS-II experiments (TMO, RIX and UED). The FEH experiments (LCLS-I, including XPP) can submit FFB jobs to the new [Fast Feedback System](#).



As of February 2023, the offline compute resources have been consolidated into the psanaq. The priority queues have been removed.

sQueue name	Node names on SLURM queues	Number of Nodes	Comments	Throughput [Gbit/s]	Cores/ Node	RAM [GB /node]	Time limit
psanaq	psana15xx psana16xx	34	Primary psana queue	40	16	128	48hrs
psanagpuq	psanagpu113 psanagpu118	6	GPU nodes	10	16	128	48hrs