

Fast Feedback System

- [Access to data](#)
- [FFB SLURM partitions](#)
- [Directories and Lifetime of data on the FFB](#)
- [FFB File Permissions](#)
- [Access to Lustre ana-filesystems](#)

Access to data

The FFB system is designed to provide dedicated analysis capabilities during the experiment.

- The FFB currently offers the fastest file system (WekaIO on NVME disks via IB HDR100) of all LCLS storage systems however it size is only about 500 TB.
- The raw data will be kept on the FFB a week after an experiment ends however for data intensive experiments files might be purged even before an experiment ends.
Files deleted from the FFB will be available only on one of the offline systems (psana, SDF, or NERSC).
- The raw data are copied to the offline storage system and to tape immediately, i.e. in quasi real time during the experiment, not after they have been deleted from FFB.
- The users generated data created in the *scratch/* folder are [moved to the offline storage](#) when the experiment is deleted from the FFB.
- **When running on the FFB the *xtc/* and *scratch/* folder should be used for reading and writing** (below */cds/data/drpsrcf/...*). The Lustre ana-filesystems should (must) not be used (only exception is *calib/*, see below).
- The LCLS Jupyterhub allows to start notebooks on the psffb nodes which will have access to the data and the FFB scratch folder of an experiment.

You can access the FFB system from pslogin, psdev or psnx with:

```
% ssh psffb
```

The experiment data will be available under:

```
/cds/data/drpsrcf/<instrument>/<experiment>
```

There are two options for telling psana which directory the data is in. One can add the "dir=" keyword to the psana DataSource, like this:

```
dsource = DataSource('exp=cxilu9218:run=20:smd:dir=/cds/data/drpsrcf/cxi/cxilu9218/xtc')
```

or one can set the following environment variable:

```
export SIT_PSDM_DATA=/cds/data/drpsrcf
```

Besides the *xtc/* folder for the raw data the *scratch/* folder allows user to write their processing output. This folder will be [moved to the offline filesystem](#) after an experiment is done. The *calib/* is a link to the offline calib folder. A ":live" is often added to the DataSource string in order to process xtc files while they are written (currently only works for LCLS1-style DAQ/analysis).

FFB SLURM partitions

You can submit your fast feedback analysis jobs to one of the queues shown in the following table. The goal is to assign dedicated resources to up to four experiments for each shift. Please contact your POC to get one of the high priority queues, 1, 2, 3 or 4, assigned to your experiment.

Queue	Comments	Throughput [Gbit/s]	Nodes	Cores/ Node	RAM [GB /node]	Default Time Limit
anaq	For the week after the experiment	100	62	61	128	12hrs
ffbl1q	Off-shift queue for experiment 1	100	5	61	128	12hrs
ffbl2q	Off-shift queue for experiment 2	100	14	61	128	12hrs
ffbl3q	Off-shift queue for experiment 3	100	20	61	128	12hrs
ffbh1q	On-shift queue for experiment 1	100	5	61	128	12hrs
ffbh2q	On-shift queue for experiment 2	100	14	61	128	12hrs
ffbh3q	On-shift queue for experiment 3	100	20	61	128	12hrs

ffbgpuq	Nvidia RTX A5000	100	2	32	128	12 hrs
---------	------------------	-----	---	----	-----	--------

Note that jobs submitted to ffbh<n>q will preempt jobs submitted to anaq and jobs submitted to ffbh<n>q will preempt jobs submitted to ffbh<n>q and anaq. Jobs that are preempted to make resources available to higher priority queues are paused and then are automatically resumed when resources become available.

The FFB system uses SLURM for submitting jobs - information on using SLURM can be found on the [Submitting SLURM Batch Jobs](#) page.

Directories and Lifetime of data on the FFB

xtc folder

- xtc files are immediately copied to the offline filesystem
- the lifetime on the ffb is dictated by how much data is generated
 - typically files stay on the ffb a week after an experiment ends
 - however if space is need the oldest files will be purged from the FFB even before an experiment has finished
 - after an experiment is done the ffb should not be used anymore except if discussed with the POC

scratch folder

Once an experiment has been complete the ffb **scratch** folder is moved to the experiments scratch folder on the offline filesystems. The following rules are applied:

1. scratch folder is made non accessible to users.
2. files and directories below the ffb *scratch/* are moved to the *scratch/ffb/* on the offline filesystem:
 - /cds/data/psdm/<instr>/<expt>/scratch/ffb/*
 - except for hdf5 files in the *smalldata* folder (see next).
 - Once the data are on the offline scratch the [Data Retention Policy](#) applies. The transfer preserves the files mtime which is used by the cleanup.
3. hdf5 files below *scratch/hdf5/smalldata/* are moved the the *hdf5/smalldata/* folder on the offline filesystem, e.g.
 - /cds/data/drpsrcf/mfx/mfx123456/scratch/smalldata/*.h5 -> /cds/data/psdm/psdm/mfx/mfx123456/hdf5/smalldata/*
 - a. **Only** h5 files (and directories) below *scratch/hdf5/smalldata/* are copied to the *hdf5/smalldata* folder on the offline storage.
 - b. Files that don't match the rule in a. will be moved to the offline *scratch/ffb/hdf5* folder.
 - c. If a h5 file already exists on the offline storage and is newer than the one on the ffb the ffb file will not be copied but just removed.

FFB File Permissions

The same permission, based on ACLs, as used for the Lustre analysis file-systems are used for the FFB. However, there is an issue with the current version of the file system:

- the umask is applied when creating files and directories which violates the ACL specs. As the default umask is 022 the group write permission will be removed. We recommend to set ones umask to:

```
% umask 0002
```

Access to Lustre ana-file systems

The batch nodes also have access to the ana Luster filesystems. For example the calib folder */cds/data/drpsrcf/<instrument>/<experiment>/calib* is a link to the folder on the ana filesystem.



The bandwidth to the ana-file systems is very limited and must be used for only light load e.g.: reading calibration constants.

The following figure shows the connectivity of the nodes:

- Each FFB file server (16 of them) has a 100Gb/s IB connection
- Each batch node has a 100Gb/s IB connection
- Batch nodes have a 10Gb/s Ethernet connection
- All Ethernet Luster access goes eventually through psnslustre02 and is limited 10Gb/s
- The figure also shows which network is used for the different file path

