Running on psana

The psana system is the default offline system and your data will end up in psana unless you arrange a different destination with your experiment POC. The psana system is also relatively old and it will be retired when more storage becomes available in the SDF system.

Getting an Account

You will need a valid SLAC UNIX account in order to use psana. The instructions for getting a SLAC UNIX account are here.

Getting Access to the System

To get into the LCLS photon computing system you can connect to the pslogin pool from anywhere:

```
ssh pslogin.slac.stanford.edu
```

From `pslogin` you can then reach the analysis nodes (see below).

You can also consider connecting to the LCLS system using NX technology: this approach allows to preserve the status of a connection across multiple sessions and it provides better performance for displaying remote graphics.

Each control room has a number of nodes for local login. These nodes have access to the Internet and are named `psusr<id>`.

The controls and DAQ nodes used for operating an instrument work in kiosk mode so you don't need a personal account to run an experiment from the control room. Remote access to these nodes is not allowed for normal users.

Getting Access to the Data

The path to the experimental data in the psana system is:

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

Data Export

To copy your data out of psana you can connect to the `psexport` pool:

```
ssh psexport.slac.stanford.edu
```

All control rooms and the overflow room in FEH have one or more taps on the Visitor Data Network. These taps can be used to transfer data to a laptop or a storage device. These devices will be automatically assigned an IP address through DHCP as soon as they are connected to the network tap.

Interactive Pool

In order to get access to the interactive nodes, connect to `psana`:

```
ssh psana
```

The `psana` pool is currently made of 18 servers with the following specifications:

- 12-cores, Xeon E5-2620 v3 2.4GHz, 128GB memory, 1 x NVIDIA GPU GTX 1080Ti, 10Gb/s Ethernet, no IB

LCLS provides a few Matlab licenses and one IDL license. Instructions on how to use these tools are in the Matlab and IDL pages.

Batch Queues

There are a number of batch queues (*i.e.* collections of compute nodes) in the psana systems. Instructions describing how to submit jobs can be found on the Submitting Batch Job page.

The main batch farm currently consist of 80 nodes with the following general specifications:

- 12 cores, Xeon X5675, 24GB memory, QDR IB

and 40 nodes with the following specifications:

- 16 cores, Xeon E5-2640 v3, 128GB memory, QDR IB

Remote Visualization

When performing analysis on the psana interactive nodes, it is useful to display plots on your host machine. For host machines near SLAC, using ssh with X-windows forwarding (the -X or -Y options) suffices. X windows forwarding can get slow for host machines in Europe. Some users have found better performance with technology called nomachine, this is documented on the Remote Visualization page.