LCLS Data Systems Update: Data Retention Policy and NERSC Collaboration

Amedeo Perazzo, LCLS Controls & Data Systems Division Director

Facility Director 5-way Meeting
October 9th 2017

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

Data Retention Policy: Key Aspects

SLAC

- Only the **members of the experiment group** can access data (raw, intermediate, results) and metadata (elog, run info)
- LCLS provides the **resources to store the data and do the analysis**
  - Most experiments do pre-processing of the data at SLAC
- We advertize **10 years on tape** and **4 months on disk** (with the ability to restore from disk to tape after initial 4 months)
- LCLS provide pool **data transfer nodes** to move the data offsite
  - Tools: bbcp, globus, xrootd

Because of the expected 3 orders of magnitude increase in data throughput, providing resources to the users for LCLS-II will be even more critical after 2020

Data Retention Policy: Folders and Backup Rules

SLAC

Space	Quota	Backup	Lifetime	Comment
xtc	None	Tape archive	4 months	Raw data
usrdaq	None	Tape archive	4 months	Raw data from users' DAQ systems
hdf5	None	Tape archive	4 months	Data translated to HDF5
scratch	None	None	4 months	Temporary data (lifetime not guaranteed)
results	4TB, 10K files	Tape backup	2 years	Analysis results ⭐
calib	None	Tape backup	2 years	Calibration data
User home	20GB	Disk + tape	Indefinite	User code
Tape archive	-	-	10 years	Raw data (xtc, hdf5, usrdaq)
Tape backup	-	-	Indefinite	User home, results and calib folder
Disk backup	-	-	Indefinite	Accessible under ~/.zfs/

https://confluence.slac.stanford.edu/display/PCDS/Data+Retention+Policy

We currently partner with ESnet/NERSC to make use of DOE supercomputing capabilities for LCLS analysis

SLAC

NERSC

ESnet
Energy Sciences Network

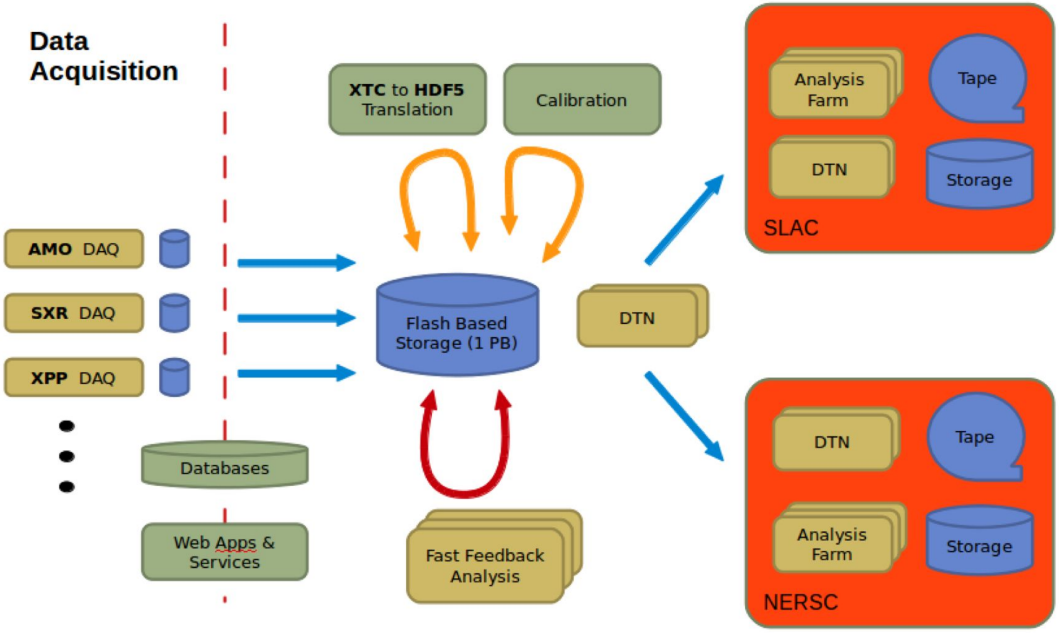

A Photon Science Speedway

Stream science data files on-the-fly from the LCLS beamlines to the NERSC supercomputers via ESnet

Data management handled transparently from the users

Very positive partnership to date, informing our future strategy

Practical Requirements for Offsite Resources for LCLS Computing

- Ability to coordinate planned offsite **outages** to avoid conflict with LCLS experiment schedule
  - Working with NERSC towards cordination API and/or shared calendar of outages
- Ability to access a significant fraction (~20%) of the main machine during a few weeks per year for **bursty jobs** with **high priority** and **short startup time** (within minutes)
  - Reservations can be expensive, preemptive techniques may be preferred
  - **LCLS has 1.5M MPP hours allocation at NERSC for 2017**
- Ability to access small pools of **interactive nodes** for development and debugging

All these aspects will need long term agreement between BES and ASCR

Practical Requirements for Offsite Resources for LCLS Computing (continued)

SLAC

- Ability to store large amount of data O(10PB) on a **fast access file system**
  - Data on this system would be handled accordingly to the LCLS data retention policy
  - **LCLS bought ~1 PB dedicated disk space at NERSC**
- Ability to access **tape** for archival and retrieval of large data sets
  - **Agreement in place with NERSC for tape usage and cost recovery: goal is to have full second copy data archived at NERSC**
- Ability for LCLS users to use their SLAC **accounts** for accessing offsite resources through a federated solution

We just started a conversation with OLCF with the goal to eventually reach an agreement similar to what we have with NERSC which includes dedicated storage and, possibly, data archiving

Data Analytics at the Exascale for Free Electron Lasers

SLAC

Application Project within Exascale Computing Project (DOE/ASCR)

Key idea: stream the data from the LCLS beamlines to exascale HEC via ESnet, do the processing on HPC and provide fast feedback back to the users in quasi real time (< 10 min)

EXASCALE COMPUTING PROJECT

High data throughput experiments	LCLS data analysis framework	Infrastructure
Algorithmic improvements and ray tracing - Example test-cases of Serial Femtosecond Crystallography, and Single Particle Imaging	Porting LCLS code to supercomputer architecture, allow scaling from hundreds of cores (now) to hundred of thousands of cores	Data flow from SLAC to HEC over ESnet

ExaFEL Data Flow

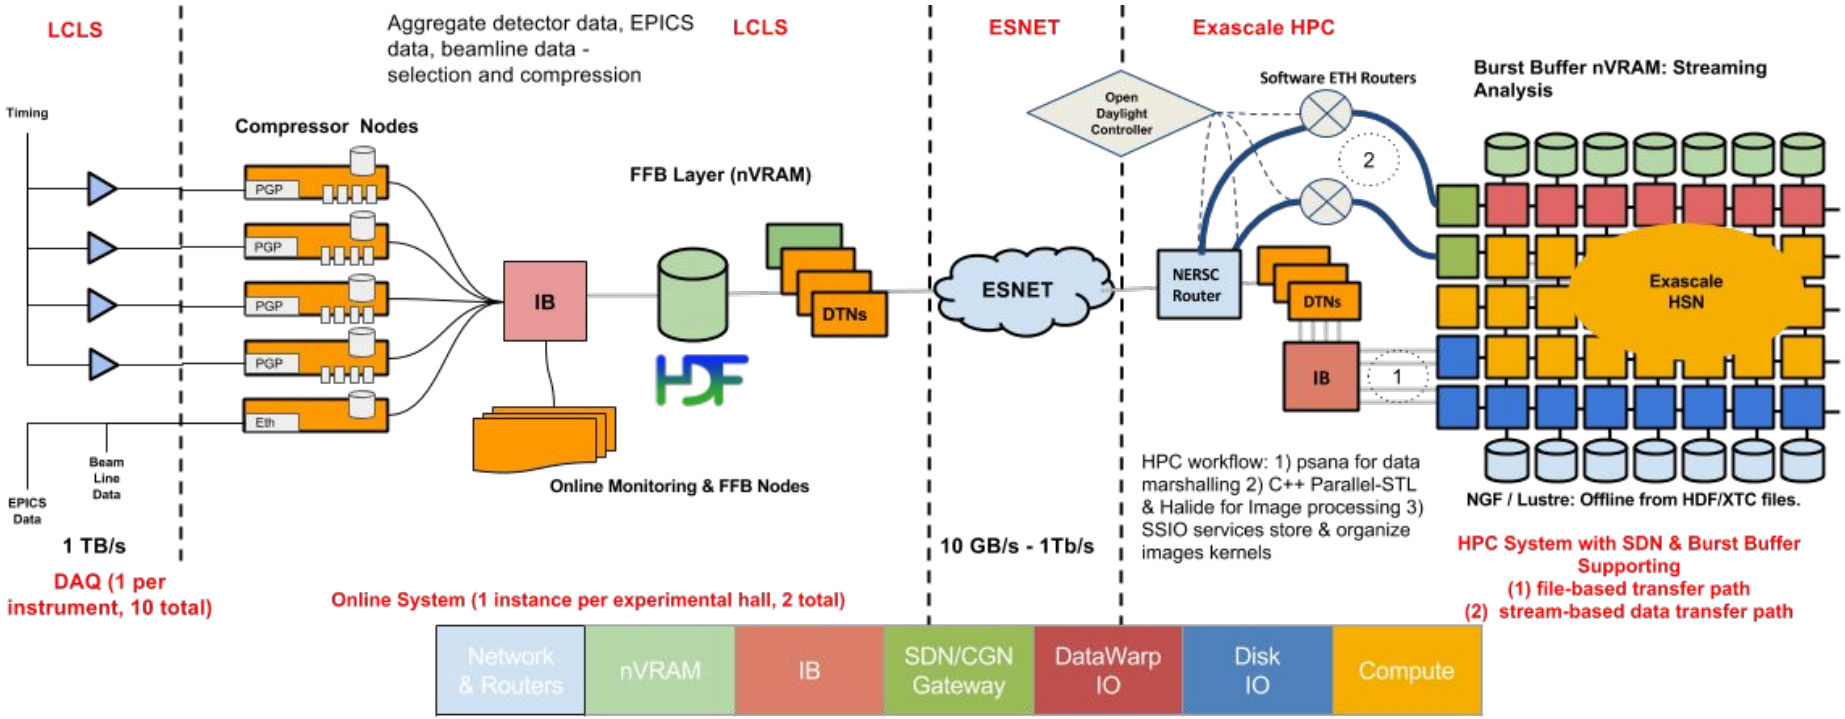

The Role of ESnet

By 2022 we'll need > **200-Gb/s** capabilities between the LCLS beamlines and ESnet

By 2026 we'll require **Tb/s** capabilities

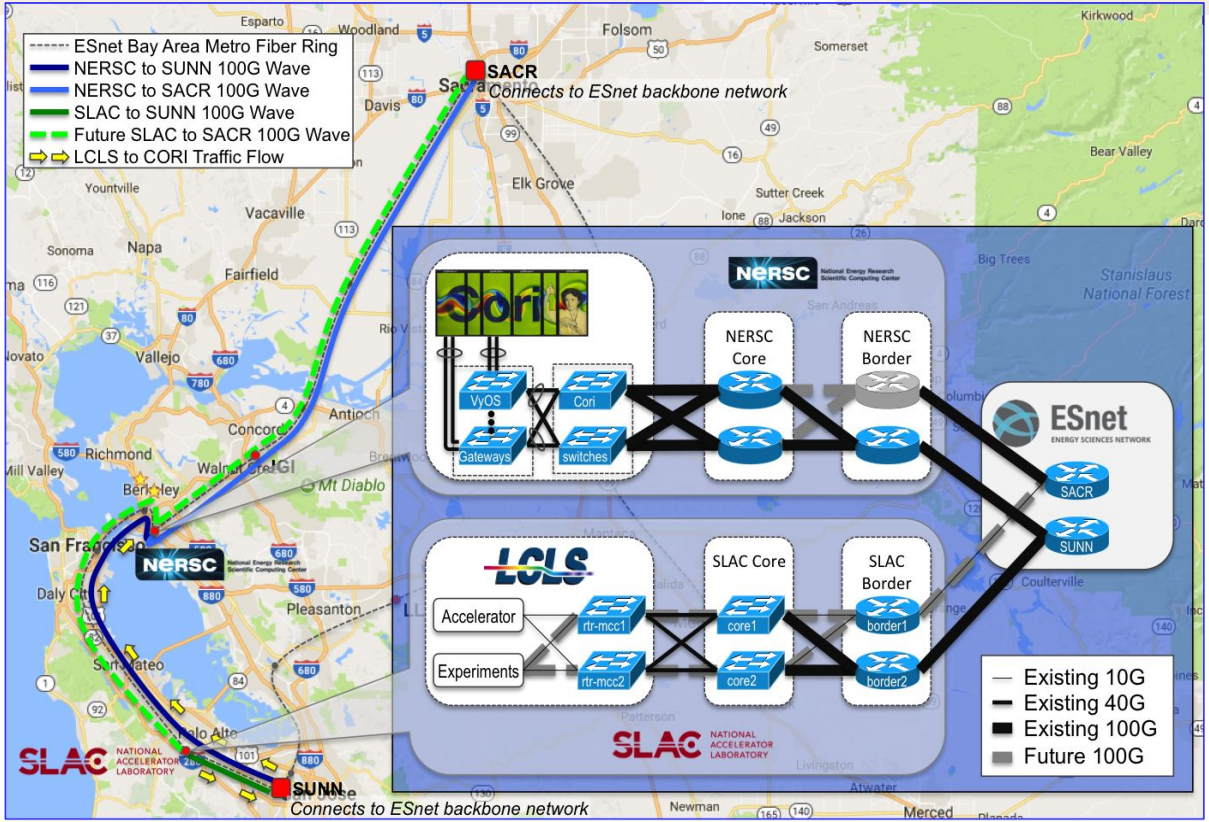

ESnet will be instrumental in providing LCLS-II access to HEC