

Present and Future Computing Requirements for the Advanced Light Source (ALS)

Banda

ALS

Berkeley Lab

NERSC BES Requirements for 2017
October 8-9, 2013
Gaithersburg, MD

Project Description

Senior Investigators:

Roger Falcone, ALS Director

Michael Banda, ALS Operations Deputy


Problem: “The growing consensus within light source scientific communities is that scientific insight and discovery at BES facilities are now being limited by computational and computing capabilities much more than by detector or accelerator technology.” (Hexemer et al., white paper)

Environment:

- ALS enables the science of ~2000 users/year who work at 37 user beamlines.
- The data generated is heterogeneous.
- Some beamlines produce large data sets that require storage and computational resources beyond those that users can provide or have access to.

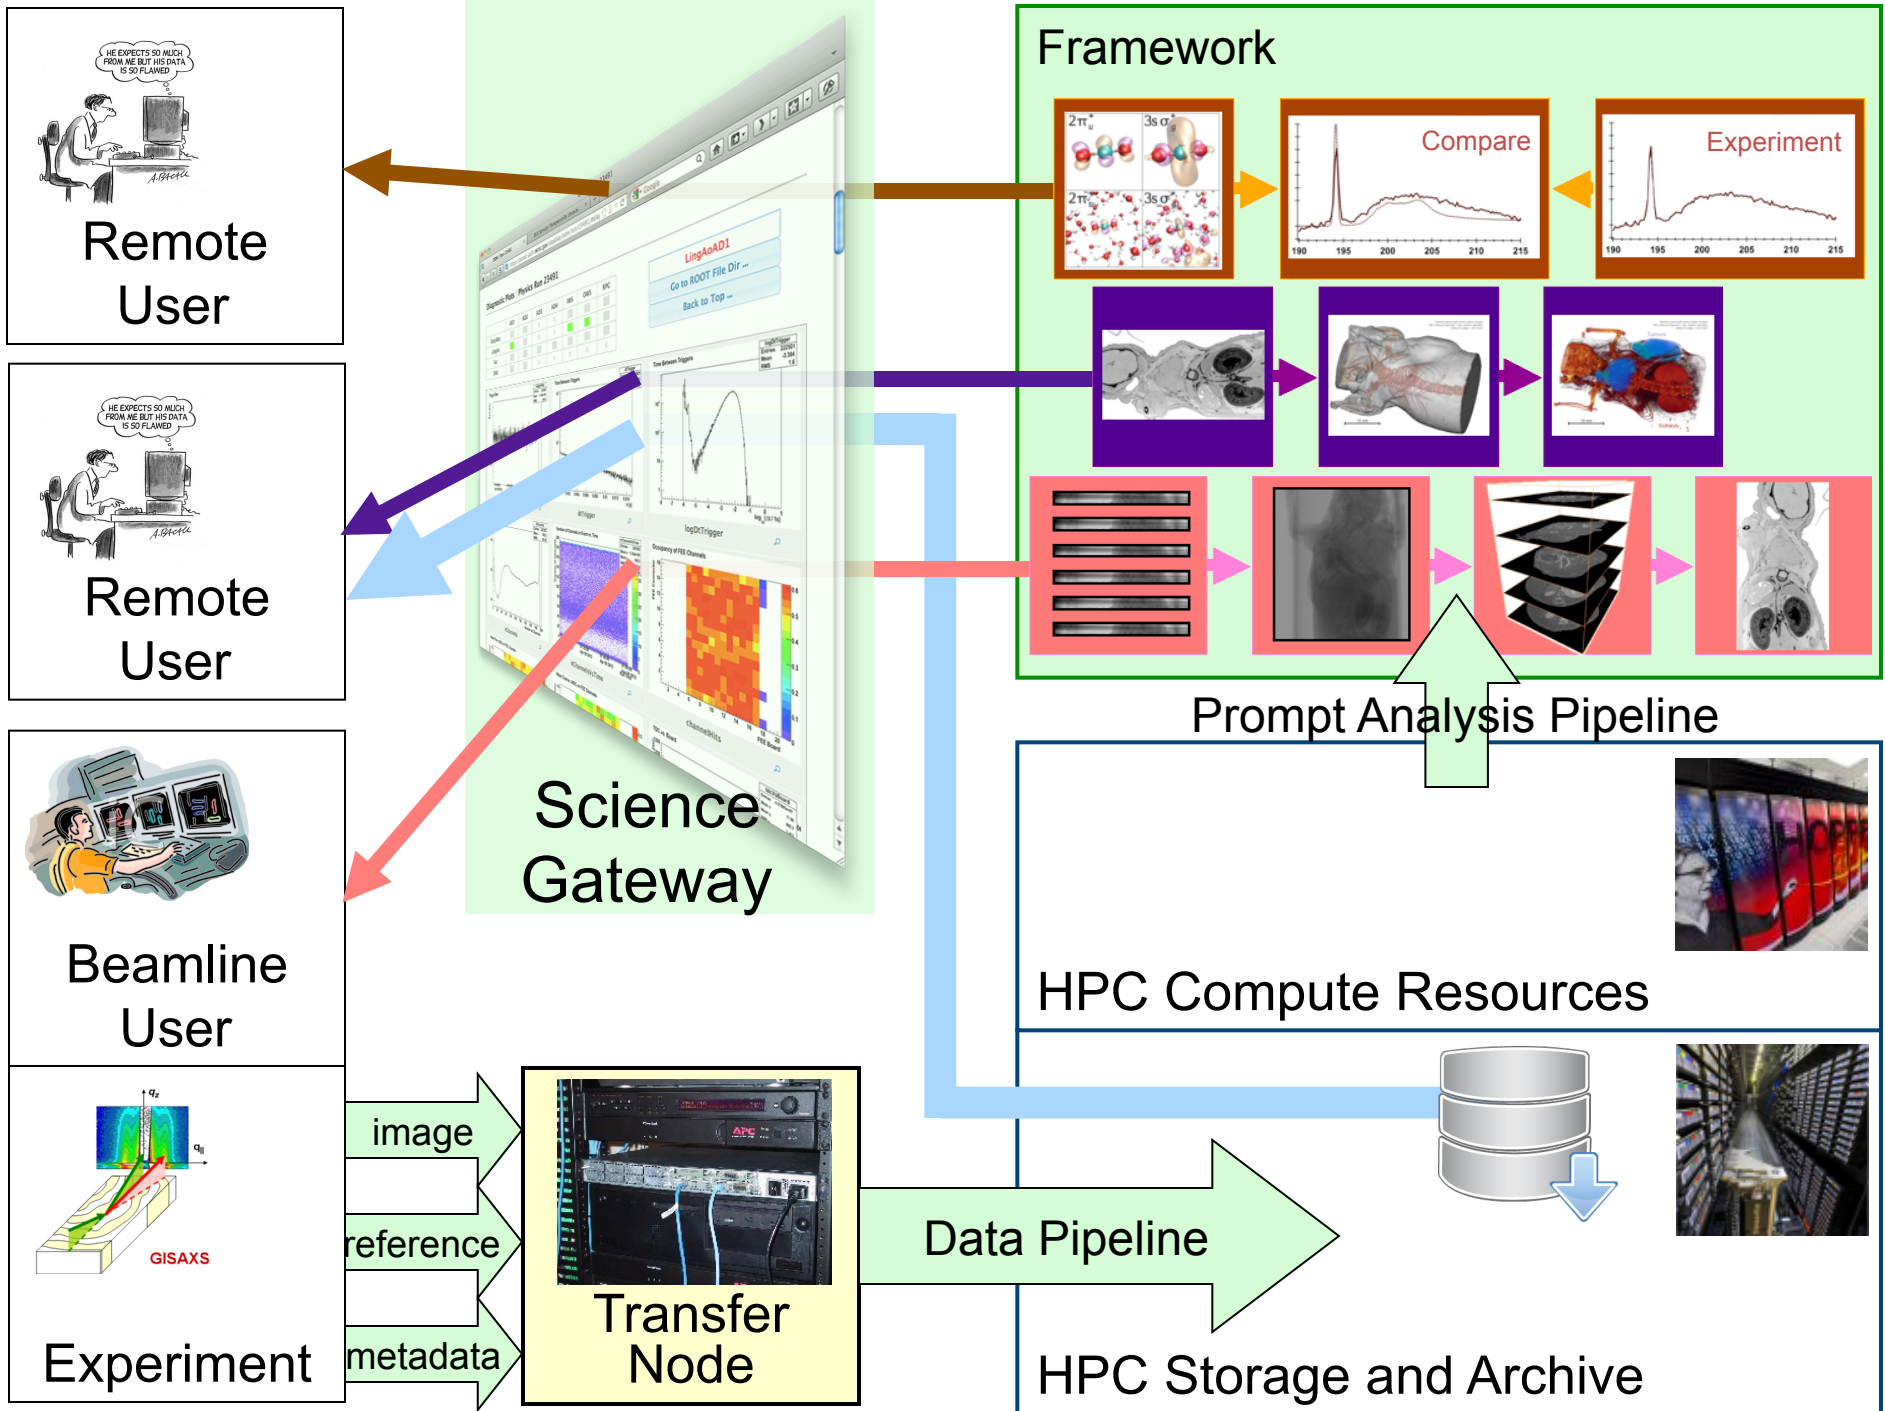
ALS Data Portal


POWERED BY



- Users are asking for faster scans, but are not prepared for the consequences: they are overwhelmed by data rates/volumes
- Most don't have the background to use high performance computers
- SPOT Suite allows users to take advantage of high performance computers, to overcome their data problems

A screenshot of a web browser displaying the SPOT Suite Portal. The browser's address bar shows "spot.nersc.gov/index.html". The page has a dark navigation bar with the ALS logo and menu items: Home, About, Status, Portal Tools, Data Browser, and Simulation Tools. On the right of the navigation bar, it says "Welcome dula" and has a "Logout" button. The main content area is a light gray box with the heading "Welcome to the SPOT Suite Portal for ALS Data and Simulations." Below the heading, it says "Log in with your NERSC account at the top right to gain access to NERSC compute resources (required for ALS simulation and analysis tools)." and "N.B. This site is very much a work in progress. Please report any problems with the page to spot-help@lists.lbl.gov". There is a blue button labeled "Learn more. »". Below this box are three images: the Berkeley Lab logo, an aerial view of the Lawrence Berkeley National Laboratory building, and a close-up of a large red and black data visualization or display wall.





Tomography Pipeline

Prep Sample



X-ray Image



Store/ Pack Data



Move Data



Process/ Simulate



SPOT FRAMEWORK

1. Go to SPOT.NERSC.GOV

2. Log in with a NERSC username.

3. Search for data by name, date...

4. View 2D AND 3D data

5. Launch jobs on NERSC

SPOT Suite transfers, processes, and presents data as it is collected, giving "real-time"

The Project:

In 2012, ALS was granted its first [Facility Allocation](#).

There is not a single scientific objective.

By 2017 we anticipate:


- Deploying a web-based portal at NERSC that seamlessly provides data and computational tools of scale to all ALS Users who require them
- Experiment with extending our portal to other BES storage ring facilities
- Enhancing the suite of HPC analysis applications available to users
- Solving the real-time feedback requirements

Present Focus:

- Deploy a pilot data portal to a tomography beamline
- Enroll users (controlled by the beamline scientist) in that portal
- Make the portal “user proof”
 - The users should focus on their science, not on the the challenges of moving large data sets and computing them.
- Begin deployment of the portal to multiple scattering and microdiffraction beamlines
- Continue development of HPC applications for use in the portal environment
 - Moving the large data sets to NERSC is not the end-game. A suite of applications tools must be made available.

ALS Data Portal

POWERED BY



- Micro-Tomography (8.3.2) portal is in “beta”, with 18 users now signed up with NERSC accounts
- Work is in progress for SAXS-WAXS (7.3.3), Micro-Diffraction (12.3.2) and Ptychography (COSMIC)
- Many questions still to be answered, for example:
 - Data retention/sharing policies
 - Allocation of NERSC resources to ALS users (storage and compute)

ALS Data and Simulation | spot.nsls.gov/index.html




ALS Home About Status Portal Tools Data Browser Simulation Tools Welcome dula Logout

Welcome to the SPOT Suite Portal for ALS Data and Simulations.

Log in with your NERSC account at the top right to gain access to NERSC compute resources (required for ALS simulation and analysis tools).

N.B. This site is very much a work in progress. Please report any problems with the page to spot-help@lists.lbl.gov

[Learn more. »](#)

© LBL LDRD 2013

Current HPC Usage

Hours used in 2012-2013:

- 2012 (first year) 1M hours requested, 2M hours used
- 2013 2M hours requested, over 4M hours used.

Machines currently used:

- Edison
- Hopper
- Carver

Necessary software, services or infrastructure:

- Data analytics and visualization, collaboration tools, web interfaces, federated authentication services, and gateway support will be needed.

Current HPC Usage (continued)

Typical parallel concurrency, number of runs per year?

- Varies significantly from beamline to beamline:
- Tomographic reconstruction codes use ~100 cores concurrently (split across multiple jobs)
- Monte-Carlo based GISAXS codes (to be deployed in production in the next calendar year) can scale to 10's of thousands of cores.
- We run multiple jobs concurrently when a new data set arrives before the completion of the analysis on a previous set. This is compounded as datasets arrive from multiple beamlines.

Scaling:

- In general weak scaling (by increasing the amount of data to be analysed) is more important to us than strong scaling.
- This may vary from beamline to beamline; in particular, the GISAXS codes are able to strong scale effectively.

HPC Requirements for 2017

Compute hours needed;

- 45,000,000 (some dedicated?)

Changes to parallel concurrency, run time, number of runs per year

- Number of concurrent users could increase by a factor of 5x to 10x, increasing the number of concurrent runs.

Changes to necessary software, services or infrastructure

- Real-time feedback is needed by the beamlines to determine the quality of the data and whether the beamline instrumentation is properly set.
- Data simulations may need to be run before taking data. While not real-time, the simulations need to be timely
- Workflow demands innovation in batch queuing (LDRD, NEWT)
- HPC consultancy and applied math expertise

Strategies for New Architectures

- The GISAXS codes in particular have been optimized for GPUs
- Tomographic codes capable of efficiently utilizing a GPU are in development.
- We believe GPU's at NERSC, whether in the form of a rack, midsize cluster or large cluster would benefit the ALS workload.
- Porting to other many-core architectures is likely possible (particularly after the work in identifying areas for on-node parallelism has already been completed during the GPU porting phase).
- Work closely with J. Sethian on algorithm development for image processing.

Summary

- Large and numerous data sets at ALS have become a reality.
- The Data Pilot is very promising
 - Compute and storage requirements are increasing through 2017 and beyond.
- ALS Facility allocation will extend the the concept of user support at ALS.
- At NERSC, the ALS Facility allocation expands the traditional operating parameters at NERSC
 - Users without ERCAP: Security, length of access, and other policy issues
 - Real-time and HTC on HPC : Beamline feedback needs fast, scalable analysis
- Advent of the “Data User” for ALS data
 - Will data have to be held indefinitely? Who/what is responsible?
 - How will the Data User access be provided and controlled? Who/what is responsible?
- How will all of this be supported?

DOE SC Data Policy Directions Derived from COMPETES

- **Principles**
 - Use Data to Drive Discovery
 - Share & Preserve for Society
 - Think Costs & Benefits
- **New Requirements**
 - All proposals must have a Data Management Plan (DMP)
 - DMPs must make research data accessible @ publication
 - Facility data policies control resources & may need approval
- **SC Policy Approach**
 - Support DOE/SC Goals
 - Accept Community Input
 - Max(Program Flexibility)
 - Min(Science Burden)

*Draft given at
NUFO&BERAC 2013
(Laura Biven, DOE/SC)*

- **Resources (TBD)**
 - Learn from NERSC Data Pilot