



Lawrence Berkeley
National Laboratory



HPCOR 2014

Supporting Data-Producing Facilities and
Instruments

Co-Chairs: David Skinner & Stuart Campbell

Contributors

- Kaki Kelly, LANL
- David Smith, LLNL
- Jeff Cunningham, LBNL
- Yao Zhang, ANL
- John Harney, ORNL
- Stuart Campbell, ORNL
- Ilana Stern, NCAR
- David Skinner, NERSC
- Craig Ulmer, Sandia
- Dino Pavlakos, Sandia
- Rudy Garcia, Sandia
- Jack Deslippe, NERSC

Some Questions ?

- How does a Data-Producing Facility or Instrument plug in to HPC facilities ?
- Do we support facilities under the existing user model or is a new method required (e.g. facility to facility)
- Facilities can be categorized into 3 main types:
 - 'Bursty' - e.g. LCLS, Telescopes, BER Model generation (simulation)
 - Constant stream of data - e.g. SNS, ALS, APS, etc., LHC, JGI, Observation (ARM)
 - Distributed Aggregation – e.g. sensor networks, fluxnet, CDIAC

What are your major strategies and initiatives over the next 5 & 10 years? How do they affect staffing levels?

- Strategies are in flux for new areas of computing for data-intensive workloads.
- Build on existing interfacility partnerships toward generalized services for data intensive science.
- Discuss and plan early with next-gen facilities, on-board early
- Computational stewardship of traditional HPC alongside emerging data-intensive workloads. Balance the entire portfolio of computing, data and network needs.
- Draw the best and brightest staff from various program offices involved. Data Intensive science will require domain specialists as well as CS savvy staff. Consider diverse teams drawn from across programs.
- Will require investment from both sides.

What are your current efforts and/or site configuration in this area?

ORNL

- ADARA – Streaming data from SNS to HPC resources
- CAMM – Integration of simulation with experiment
- CADES – Cross cutting effort to provide lab-wide data services and domain expertise
- ACME – end-to-end testbed for production runs of the Community Earth System Model)
- ASCR/BES Data Pilot Project (APS/SNS)

LLNL

- ALICE
- PCMDI climate archive

What are your current efforts and/or site configuration in this area?

LBL

- LCLS-II pilot work (Photon Science Speedway),
- ASCR/BES pilot
- SpotSuite (ALS/SSRL/LCLS Simulation and Analysis Framework, Portal)
- Daya Bay Neutrino Experiment
- Palomar Transient Factory
- NEWT, FireWorks
- PDSF, JGI

NCAR

- GLADE – PB scale disk
- ESG – publication of Community Earth System Model (CESM) runs

ANL

- Integrating Simulation and Observation: Discovery Engines for Big Data
- ASCR/BES pilot (SNS/APS)

What are your mandates and constraints?

- Security constraints - In some leadership computing facilities (ALCF and OLCF), HPC resources and experimental facilities (APS and SNS/HFIR) reside in different security enclaves. Communication/messaging between enclaves is restricted in one direction
- Provide queuing for both traditional batch and real-time computing
- Users requesting open source solutions

How to do you forecast future needs and requirements?

- Computational facilities need to be engaged with the DOE experimental facilities future plans/ upgrades etc...
- Attention to ASCR Requirements Reviews (ongoing between program offices)

What are the biggest challenges and gaps between what you can do today and what will be required in 5 - 10 years?

- We are missing a single large scale data analysis paradigm suitable for HPC. No MPI.
- Need standardization across facilities and centers in terms of both data formats and workflow engines, environments to allow software written for one facility to work everywhere.
- Many data producing facilities codes get their start on the workstation
- Uncertainty around computer architecture for emerging Data Intensive spaces. Will Exascale do DI?

What opportunities exist for productive collaborations among DOE HPC centers?

- Federated Identity Management
 - OpenID, Shib, etc...
- Need standardization across facilities and centers in terms of both data formats and workflow engines, environments to allow software written for one facility to work everywhere.
- Facilities to enable users to co-analyze data from multiple experimental facilities (e.g. SNS and APS)

Describe some practices that you think are effective as well as lessons learned that would be helpful to other centers?

Measure

- Benchmark co-development (3 types of facilities)
- Performance data fusion with facility data (DTWG++, gridftp + PBS logs)

Educate

- WAN data methods for users
- Help staff then users pick between workflow tools
- Facility DMP roles and responsibilities (and timeframes, curation)
- Menu of science data services from HPC facilities

Coordinate

- More communication and meeting between facilities
- Proposals and resource co-scheduling (beamtime + compute-time)
- Federate ID's to minimize number of auth hops for users

Innovate

- Workflow testbeds (mock up and beta test proposed WF's out of production)
- Upstreaming analysis, triggering, veto systems
- Co-design with next-gen detectors and instruments (get in early)