


Astronomy 250







Exascale and Extreme Data Science at NERSC



Peter Nugent
Division Deputy for Scientific Engagement
Department Head for Computational Science

January 11, 2017

NERSC Overview



Background



- **NERSC is the National Energy Research Scientific Computing Center**
- **Founded in 1974**
- **Focused on open science**
- **Located at Lawrence Berkeley National Laboratory**
- **Operated for the U.S. Department of Energy (DOE) Office of Science**


NERSC Provides HPC and Data Resources for DOE Office of Science Research




U.S. DEPARTMENT OF
ENERGY

Office of
Science


Largest funder of physical
science research in U.S.




Biology, Environment




Computing




Materials, Chemistry,
Geophysics



Particle Physics,
Cosmology



Nuclear Physics, Nuclear
Astrophysics



Fusion Energy,
Plasma Physics


NERSC directly supports DOE's science mission



- DOE SC offices allocate 80% of the computing and storage resources at NERSC
- ALCC 10%
- NERSC Director's Reserve 10%


2014 NERSC Usage by DOE Program Office

(MPP Hours in Millions)




2014 NERSC Usage by Institution Type

(MPP Hours in Millions)




NERSC has a broad user base




642 international
users

5,000 users
from 47 U.S.
states




Strong focus on science


NERSC



Martin Karplus



Saul Perlmutter



George Smoot



Warren Washington



Nobel Prize in Physics 2015



Scientific Achievement


The discovery that neutrinos have mass and oscillate between different types

Significance and Impact

The discrepancy between predicted and observed solar neutrinos was a mystery for decades. This discovery overturned the Standard Model interpretation of neutrinos as massless particles and resolved the “solar neutrino problem”

Research Details

The Sudbury Neutrino Observatory (SNO) detected all three types (flavors) of neutrinos and showed that when all three were considered, the total flux was in line with predictions. This, together with results from the Super Kamiokande experiment, was proof that neutrinos were oscillating between flavors and therefore had mass



A SNO construction photo shows the spherical vessel that would later be filled with water.

NERSC helped the SNO team use PDSF for critical analysis contributing to their seminal PRL paper. HPSS serves as a repository for the entire 26 TB data set.


Q. R. Ahmad et al. (SNO Collaboration). Phys. Rev. Lett. 87, 071301 (2001)

Nobel Recipients: Arthur B. McDonald, Queen's University (SNO)
Takaaki Kajita, Tokyo University (Super Kamiokande)

We have a dual mission to advance the state-of-the-art in supercomputing




- We collaborate with computer companies years before a system's delivery to deploy advanced systems with new capabilities at large scale
- We provide a highly customized software and programming environment for science applications
- We are tightly coupled with the workflows of DOE's experimental and observational facilities – ingesting tens of terabytes of data each day
- Our staff provide advanced application and system performance expertise to users



Over 600 codes run at NERSC



NERSC 2015 Code Usage



- 10 codes make up 50% of the workload
- 25 codes make up 66% of the workload

NERSC's Edison System




- Edison is the HPCS* demo system (serial #1)
- First Cray Petascale system with Intel processors (Ivy Bridge), Aries interconnect and Dragonfly topology
- Very high memory bandwidth (100 GB/s per node), interconnect bandwidth and bisection bandwidth
- 5,576 nodes, 133K cores, 64 GB/node
- Exceptional application performance



A significant portion of the NERSC workload runs at scale



2016 Edison Hours Breakdown by Job Size




High concurrency jobs are important across all science domains



Concurrency within science categories on Edison

Cores Used: ■ >16K ■ 1K - 16K



- Some fraction of every domain's workload runs with more than 16K cores.
- In almost all domains, more than half the workload uses more than 1K cores.

NERSC has moved to the Computational Research and Theory (CRT) Facility




- **Four story, 300 offices, 20Ksf HPC floor, 12.5->40 MW**
- **Located for collaboration**
 - LBNL, CRD, ESnet, UCB
- **Exceptional energy efficiency**
 - Natural air and water cooling
 - Heat recovery
 - PUE < 1.1 (total facility Eng/IT Eng)



NERSC Systems Timeline



2007	NERSC-5	Franklin	Cray XT4	102/352 TF
2010	NERSC-6	Hopper	Cray XE6	1.28 PF
2014	NERSC-7	Edison	Cray XC30	2.57 PF
2016	NERSC-8	Cori	Cray XC	30 PF
2020	NERSC-9			100PF-300PF
2024	NERSC-10			1EF
2028	NERSC-11			5-10EF



NERSC Timeline



**NRP
complete
12.5 MW**

2015

**Staff
move in**

**NERSC-8
Cori
Phase I**

2016

**Edison
Move
Complete**

**NERSC-8
Cori
Phase II**

**CRT
25MW
upgrade**

2016-18

2020

**NERSC-9
150-300
Petaflops**

**CRT
35+ MW
upgrade**

2021







**NERSC-10
Capable
Exascale
for broad
Science**

2024

2028

**NERSC-11
5-10
Exaflops**

Exascale at NERSC



Cori has an energy-efficient processor



- **We didn't deploy accelerators or GPUs in Edison**
- **Disruptions in programming models are a challenge for NERSC**
- **What's changed for Cori?**
 - **Energy-efficient architectures are needed to meet the science needs of our users**
 - **Heightened awareness about Exascale among application teams**
 - **Many codes are being adapted for next generation systems**
 - **Technology changes (e.g., self-hosted many core chips, tighter CPU/GPU integration) are making the transition easier**

Cori Supports the Office of Science HPC Workload and Data-Intensive Science



- **Cray system with 9,300 Intel Knights Landing compute nodes – deliver in mid 2016**
 - Self-hosted, (not an accelerator) manycore processor 68 cores per node
 - On-package high-bandwidth memory at ~490GB/sec
- **Robust Application Readiness Plan**
 - Outreach and training for user community
 - Application deep dives with Intel and Cray
 - 8 post-docs integrated with key application teams
- **Data Intensive Science Support**
 - 10 Haswell processor cabinets (Phase 1) to support data intensive applications
 - NVRAM Burst Buffer with 1.5PB of disk and 1.5TB/sec
 - 28 PB of disk, >700 GB/sec I/O bandwidth in Lustre bandwidth



Intel “Knights Landing” Processor



- **Single socket processor - Self-hosted, not a co-processor, not an accelerator**
- **68 cores per processor with support for four hardware threads each; more cores than current generation Intel Xeon Phi™**
- **512b vector units (32 flops/clock – AVX 512)**
- **3X single-thread performance over current generation Xeon Phi co-processor (KNC)**
- **High bandwidth on-package memory, up to 16GB capacity with memory bandwidth of ~490 GB/sec**
- **Presents an application porting challenge to efficiently exploit KNL performance features**

To run effectively on Cori users will have to:



- **Manage Domain Parallelism**

- independent program units; explicit

- **Increase Thread Parallelism**


- independent execution units within the program; generally explicit

- **Exploit Data Parallelism**

- Same operation on multiple elements

- **Improve data locality**

- Cache blocking;
Use on-package memory




```
| --> DO I = 1, N  
|      R(I) = B(I) + A(I)  
| --> ENDDO
```

Transitioning the broad SC workload to energy-efficient architectures




- Provide platforms for codes that are not transitioned
- Work with vendors to ensure that MPI+OpenMP provides reasonable performance in the near term
- Focus on 20 codes that are proxy's for much of the workload (NESAP)
- Provide libraries and training to the broader community
- Establish broader forums for sharing information (e.g., the Intel Xeon Phi User Group)




NESAP Codes

NERSC




Advanced Scientific Computing Research

Almgren (LBNL) **BoxLib**
Trebotich (LBNL) **Chombo-crunch**






High Energy Physics

Vay (LBNL) **WARP & IMPACT**
Toussaint(Arizona) **MILC**
Habib (ANL) **HACC**




Nuclear Physics

Maris (Iowa St.) **MFDn**
Joo (JLAB) **Chroma**
Christ/Karsch
(Columbia/BNL) **DWF/HISQ**




Basic Energy Sciences

Kent (ORNL) **Quantum Espresso**
Deslippe (NERSC) **BerkeleyGW**
Chelikowsky (UT) **PARSEC**
Bylaska (PNNL) **NWChem**
Newman (LBNL) **EMGeo**





Biological and Environmental Research

Smith (ORNL) **Gromacs**
Yelick (LBNL) **Meraculous**
Ringler (LANL) **MPAS-O**
Johansen (LBNL) **ACME**
Dennis (NCAR) **CESM**



Fusion Energy Sciences

Jardin (PPPL) **M3D**
Chang (PPPL) **XGC1**









Goals and Objectives for the NERSC-9 Project




- 1. Provide a significant increase in computational capabilities over the Edison system, 16-25x on a set of representative DOE benchmarks.**
- 2. Meet the needs of extreme computing and data users by accelerating workflow performance.**
- 3. Provide a vehicle for the demonstration and development of exascale-era technologies**
- 4. Delivery in the 2020 time frame**


Data Intensive Science




NERSC has been supporting data intensive science for a long time




Palomar Transient
Factory
Supernova




Planck Satellite
Cosmic Microwave
Background
Radiation




Alice
Large Hadron Collider



Atlas
Large Hadron Collider



Dayabay
Neutrinos



ALS
Light Source




LCLS
Light Source



Joint Genome
Institute
Bioinformatics

NERSC users import more data than they export!



Traditionally we had deployed separate Compute Intensive and Data Intensive Systems



Compute Intensive




Data Intensive



Carver



Genepool




PDSF

What has changed? Increased data rates and new sensing capabilities




LCLS
Light Source



Advanced Lightsource Upgrade



Environmental
sensors




Next generation
electron microscope



Sequencers that fit into
the palm of your hand

- In the next 5 years, data rates will be approaching Tb/sec for many instruments
- Infeasible to put a supercomputer at the site of every data generator


What has changed? Coupling of experiments with large scale simulations



*Nyx simulation of
Lyman alpha
forest*



*Kitt Peak National
Observatory's Mayall 4-meter
telescope, planned site of the
DESI experiment*



*New climate
modeling methods,
produce new
understanding of ice*



*Genomes to
watersheds*

But how different really are the compute and data intensive platforms?



Policies

- Fast-turn around time. Jobs start shortly after submitted
- Can run large numbers of throughput jobs

Software/Configuration

- Support for complex workflows
- Communication and streaming data from external databases and data sources
- Easy to customize user environment

Hardware


- Local disk for fast I/O
- Some systems (not all) have larger memory nodes
- Support for advanced workflows (DB, web, etc)

Differences are primarily software and policy issues with some hardware differences in the ratio of I/O, memory and compute

NERSC is making significant investments on Cori to support data intensive science



- High bandwidth external connectivity to experimental facilities from compute nodes (Software Defined Networking)
- NVRAM Flash Burst Buffer as I/O accelerator
 - 1.5PB, 1.5 TB/sec
 - User can request I/O bandwidth and capacity at job launch time
 - Use cases include, out-of-core simulations, image processing, shared library applications, heavy read/write I/O applications
- Virtualization capabilities (Docker)
- More login nodes for managing advanced workflows
- Support for real time and high-throughput queues




NERSC is exploring Burst Buffer Use Cases beyond checkpoint-restart



- **Accelerate I/O**
 - Checkpoint/restart or other high bandwidth reads/writes
 - Apps with high IOP/s e.g. non-sequential table lookup
 - Out-of-core applications
 - Fast reads for image analysis
- **Advanced Workflows**
 - Coupling applications, using the Burst Buffer as interim storage
 - Streaming data from experimental facilities
- **Analysis and Visualization**
 - In-situ/ in-transit
 - Interactive visualization




Palomar Transient
Factory Pipeline:
Use Burst Buffer as cache
for fast reads



VPIC – in situ visualization
of a trillion particles

Burst Buffer Software Development Efforts



Create Software to enhance usability and to meet the needs of all NERSC users


- Scheduler enhancements
 - Automatic migration of data to/from flash
 - Dedicated provisioning of flash resources
 - Persistent reservations of flash storage
- Caching mode – data transparently captured by the BB nodes
 - Transparent to user -> no code modifications required
- Enable In-transit analysis
 - Data processing or filtering on the BB nodes – model for exascale

Enhanced Networking for Cori



Progress

- HW and SW installed and configured
- Simple outbound BW testing shows 4X improvement in bandwidth 5.5 Gb/s vs 20Gb/s



Initial Science Uses Cases

- General Atomics – 5x improvement talking to an external database used in a real-time workflow
- Globus-url-copy to CERN test point – 100x faster!
- LCLS to Cori BB now 100x faster!

Next Steps


- Scale Testing 160 Nodes to 1 GW
- Multi-stream In-bound transfers
- Med Term: scheduler integration
- Long Term: software defined networking circuit testing and integration

Shifter: Containers for HPC



Challenge and Opportunity

- Data Intensive computing often require large, complex software stacks
- Docker becoming standard package to run applications.



Innovation


- Shifter is a NERSC R&D effort, in collaboration with Cray, to support User-created Application images.
- Shifter provides “Docker-like” functionality for HPC

Impact and Early Successes


- Shifter has enabled multiple projects to quickly make use of NERSC (e.g. LCLS, LHC, DES)
- Shifter can improve job-startup times and application performance (e.g. Python)
- Shifter will be supported by Cray and is under evaluation by other HPC centers




Observational Cosmology Data Sets



1-2 PB




2-3 PB



~500 PB

- Imaging surveys of the sky
 - Digital images cover a fraction of a square degree (41,000 square degrees on the sky)
 - Typically 2,000 * 4,000 pixels with most detections occupying only 9 pixels
 - 20,000 detections with a signal-to-noise ratio > 5 per image
 - Done in several filter passbands to provide not only shape and brightness, but “color” information on each detection (used to determine its approximate type and distance)
- Spectroscopic surveys of the sky
 - For each object, precisely measure the flux as a function of wavelength
 - Use the imaging surveys to provide targets for spectroscopy

Observational Data Pipelines




SN Subtraction



The Dark Energy Survey

Lensing Analysis

- Imaging Pipeline
 - Runs in real-time (to provide immediate quality assessment) as well as annually and at the end of the survey.
 - Handles image de-trending: removal of electronic biases, pixel-to-pixel response, removal of artifacts, object identification, etc.
 - Aligns images wrt each other and with an overall sky solution
 - Determines the depth of each image (were there clouds, and how much?)
 - Measures the location, shape of the object (stars, galaxies, etc.), and their brightnesses
 - Often employs several codes, patched together with perl, python, bash, etc. with calls to databases in between.
 - Can have different ones within the same survey depending on the science




Observational Data Pipelines




DES Pipeline at NCSA

- Job manager is Condor
- Uses scores of codes to process the data and identify objects
- Communicates with an Oracle db at every major point in the pipeline
- Transfers data from one system to another with wget
- Pushes final images and catalogs back to main storage
- Very much tuned to the hardware and unix environment at NCSA




Major problems occur when a bug is found in the pipeline ar reprocessed. How do you get back to real-time with a limite


Solution: Docker + HPC Resources + ESnet



Docker is a shipping container system for code



Docker + HPC + ESnet



DECam Processing at NERSC



- Project began, from scratch at NCSA, on June 2nd
 - NCSA was able to move their code onto a set of docker containers in 2 weeks and run it locally and determine it produced the same results
 - Biggest issue was getting HTCondor to play nicely through docker with their server at NCSA.
 - Tested bandwidth for image processing and determined that a raw image (1GB) and processed image and catalogs (4+GB) could be shipped via ESnet from NCSA to NERSC and back in 1/50th the time it took to process (negligible I/O)
 - Pushed the container to docker.io so it could be pulled to NERSC
- **Mid-June NERSC was able to securely allow docker to see the Scratch directories on the NERSC Global Filesystem from Jessup - faster I/O**
 - Allows processing to occur via docker and post-processing analysis from any of the other HPC systems at NERSC.
 - First parts of DECam pipeline began running July 1 on a single node
- **By August 7th there was a request for an increase to 10 nodes**
- **By September 29th running on up to 20 nodes (80 node docker-enabled pool)**

In September we began running *both* LSST and DECam pipeline work at ORNL, ANL & BNL through docker as well.

While a small test, roughly 200khrs spent processing to date and another 100khrs post-processing, it was able to nicely increase the pool of resources available to the DES Data Management team to plow through the rest of DES Year-1 reprocessing and still handle the current season's real-time processing. Roughly 15TB of raw data were processed leading to nearly 100TB of processed imaging data, co-adds and catalogs.

It also provided an excellent, and natural, way to share the LSST pipeline. Currently, LSST's OpSim (the Operations Simulator - an astronomical survey simulator developed to understand how to optimize the survey strategy by simulating sequences of observations or visits) is up and running and phoSim (the Photon Simulator - a set of fast photon Monte Carlo codes used to calculate the physics of the atmosphere, telescope, and detector) will be available soon.

Go to: <https://registry.hub.docker.com/u/lsst/opsim> and try it out!


...and now we have to set up a way to share the data.

Real-time Queue at NERSC



- **NERSC has made a small pool of nodes available for immediate turnaround / “Realtime” computing**
 - Up to 32 nodes in realtime queue (1024 cores)
 - Realtime nodes have higher priority than other queues
 - Pool can shrink or grow as needed based on demand
- **Approved projects have a small number of nodes available on-demand without queue wait times**

Raw Machine Hours by Science Area (in millions)



- Prototype queue used by a handful of projects at NERSC
- Real-time queue accounts for <1% of time at NERSC
- NERSC is tracking usage and use cases closely

PTF Discovery and Follow-up



**Instrumentation, system
design, first results**

Law, Kulkarni, Dekany et al. 2009 PASP 121 1395L

Science plans

Rau, Kulkarni, Law et al. 2009 PASP 121 1334R

2010 survey status

Law et al. 2010 SPIE 7735




U.S. DEPARTMENT OF
ENERGY

Office of
Science



ModSim 2016



HPWREN Network



PTF Camera









92 Mpixels, 1" resolution, R=21 in 60s

PTF Key Projects	
Various SNe	Dwarf novae
Transients in nearby galaxies	Core collapse SNe
RR Lyrae	Solar system objects
CVs	AGN
AM CVn	Blazars
Galactic dynamics	LIGO & Neutrino transients
Flare stars	Hostless transients
Nearby star kinematics	Orphan GRB afterglows
Type Ia Supernovae	Eclipsing stars and planets
Tidal events	H-alpha 1/2 sky survey



The power of PTF resides in its diverse science goals and follow-up.

PTF Science

▼► Detected transients will be followed up using a wide variety of optical and IR, photometric and spectroscopic followup facilities.




The power of PTF resides in its diverse science goals and follow-up.



Real or Bogus – Machine Learning Analysis

NERSC



4096 X 2048 CCD images - over 3000 per night – producing 1.5M bogus detections, 50k known astrophysical objects and only 1-2 new astrophysical transients of interest every night. Machine learning is used to wade through this sea of garbage.

PTF Database



	R-band	g-band
images	1.82M	305k
subtractions	1.52M	146k
references	29.2k	6.3k
Candidates	890M	197M
Transients	42945	3120

All in 851 nights from 2009-2012.

An image is an individual chip (~0.7 sq. deg.)

The database is 1 TB.


Now doubled in size from iPTF from 2013-2016.

Pipeline...




Observatory

2.5 MB/s



Processing/db




12 MB/s

Subtractions



PTF
Classification

**4 MB/s
(crude)**




1 GB/s

**0.5 MB/s
(full)**

NERSC GLOBAL FILESYSTEM
250TB of purchased space

One of 4 such pipelines
running at NERSC...

PTF Turn-around




What does “real-time” subtractions really mean?

For 95% of the nights all images are processed, subtractions are run, candidates are put into the database and the local universe script is run in < 1 hr after observation.

Median turn-around is 30m.

iPTF turn-around



Due to the X-SWAP project (Extreme-Scale Scientific Workflow Analysis and Prediction), funded through the ASCR LAB-1088 call (Analytical Modeling for Extreme-Scale Computing Environments), we have been able to understand and eliminate a lot of our inefficiencies and decrease the turn-around by an order of magnitude!


Better use of the Lustre filesystem (for everything), better use of OpenMP in all codes, reserved nodes, etc.

iPTF turn-around



We made major changes to the old pipeline.

- Pipeline completely instrumented for timings.
- Identified and fixed python load time on Edison (15min to 5 sec).
- Moved all I/O in processing to Lustre /scratch filesystem
- Now optimizing db access



Yi Cao's Caltech thesis May 3, 2016. This fall will become an eScience Postdoctoral Fellow at University of Washington.


Typical turnaround is now
< 5 minutes for 95% of the
data!

Instrumented Pipeline with 39 Checkpoints




Covers everything from:


- Pulling the data from the telescope
- I/O on scratch
- Subtraction software
- Running ML algorithms
- Loading the db with discoveries
- Performing difficult geometric queries to match with known stars, asteroids, previous discoveries, etc.
- Copying data from scratch to project



DB Access



I/O time on Cori




PTF Sky Coverage




To date:

- 2338 Spectroscopically typed supernovae
- 10^6 Galactic Transients
- 10^4 Transients in M31

178 publications, 7 in *Nature* and 3 in *Science* since late 2009



Overcoming wide & fast: iPTF13bxi in 71 deg²!




Singer et al. 2013

The second Fermi afterglow: iPTF13dsw at $z=1.87$!





Kasliwal *et al.* 2013b



Overcoming
Wide, Fast & Faint

Pinpointing the afterglow
amidst 30,000 candidates


Orphan Afterglow




iPTF14yb
Cenko et al. 2014

IPN found a GRB (localization ~ 200 - 300 sq. deg.)
~15 min before first detection.....

Flash Spectroscopy





SN2013cu (iPTF13ast)
 Gal-Yam et al. (2014) *Nature*






Flash Spectroscopy



NERSC



Flash Spectroscopy



Flash Spectroscopy



Gravitational Wave Trigger



GW150914



Signals in synchrony

When shifted by 0.007 seconds, the signal from LIGO's observatory in Washington (red) neatly matches the signal from the one in Louisiana (blue).

● LIGO Hanford data (shifted) ● LIGO Livingston data




Kasliwal *et al.* (2016)
(arXiv and ApJL)

GW150914



Going to have to be able to sift through a lot of stuff, and react quickly with follow-up, to get on the optical companion for a GW trigger.




ZTF will survey an order of magnitude faster than PTF.

	PTF	ZTF
Active Area	7.26 deg ²	47 deg ²
Overhead Time	46 sec	<15 sec
Optimal Exposure Time	60 sec	30 sec
Relative Areal Survey Rate	1x	15.0x
Relative Volumetric Survey Rate	1x	12.3x

3750 deg²/hour

⇒ 3 π survey in 8 hours







>250 observations/field/year
for uniform survey





Existing PTF camera
MOSAIC 12k

New ZTF camera:
16 6k x 6k e2v CCDs


Superfacility Concept



Experimental and observational science is at crossroads




- Data volumes are increasing faster than Moore's Law
- New algorithms and methods for analyzing data
- Infeasible to put a supercomputing center at every experimental facility



Superfacility Vision: A network of connected facilities, software and expertise to enable new modes of discovery

NERSC



- **Petigura: Earth-like Planets Around Sun-like Stars**
- **Wang: Gemini Planet Imager Data Reduction**
- **Liu: 21-cm Cosmology**
- **Shivvers: Optimized Spectrum Synthesis**
- **Stark: Lyman-alpha cosmology with Nyx**

...and lots of little projects covering everything from solar physics to cosmology.