Data Day Intro

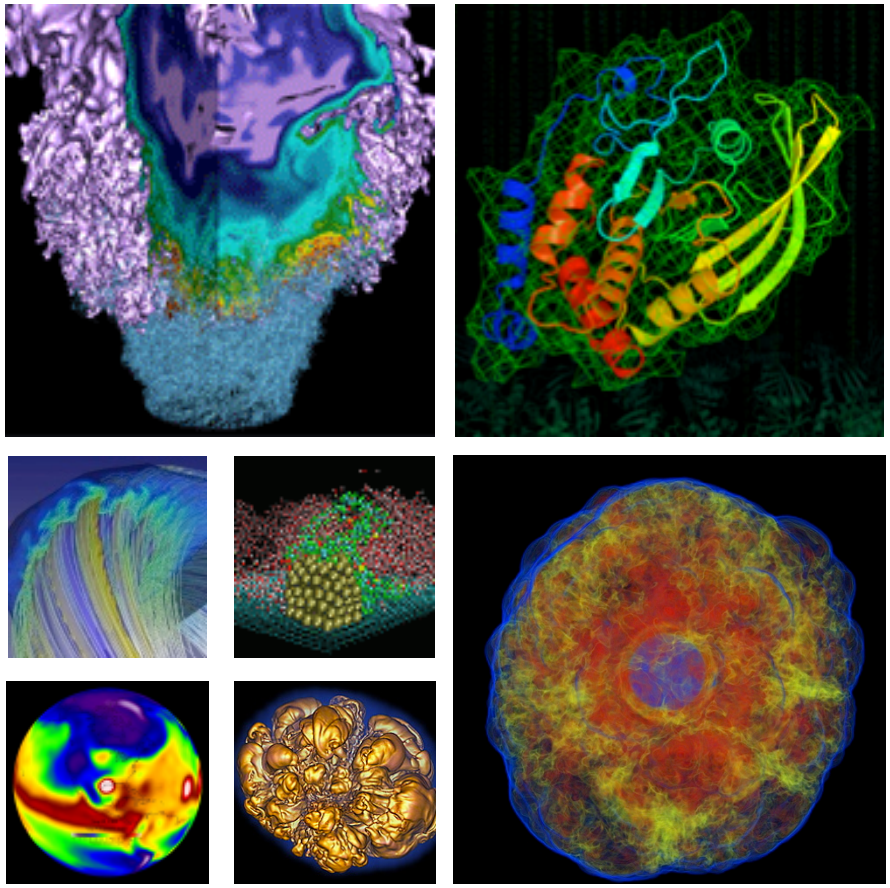

Katie Antypas
Data Department Head

August 22, 2016

NERSC has been supporting data intensive science for a long time

Palomar Transient Factory
Supernova

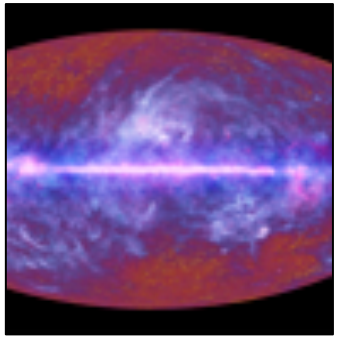

Planck Satellite
Cosmic Microwave Background Radiation

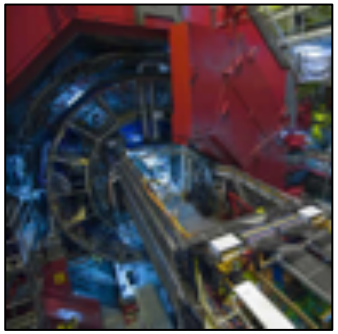

Alice
Large Hadron Collider

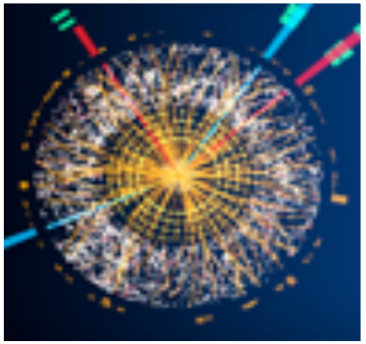

Atlas
Large Hadron Collider

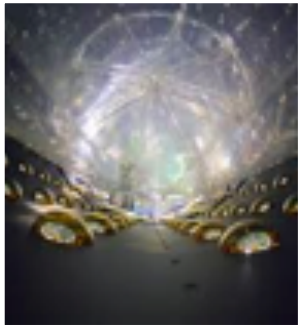

Dayabay
Neutrinos

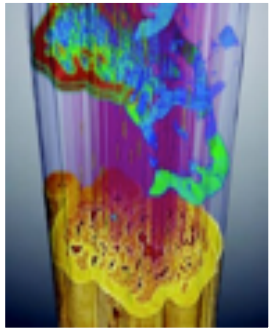

ALS
Light Source

LCLS
Light Source

Joint Genome Institute
Bioinformatics

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Growth of 'data' projects

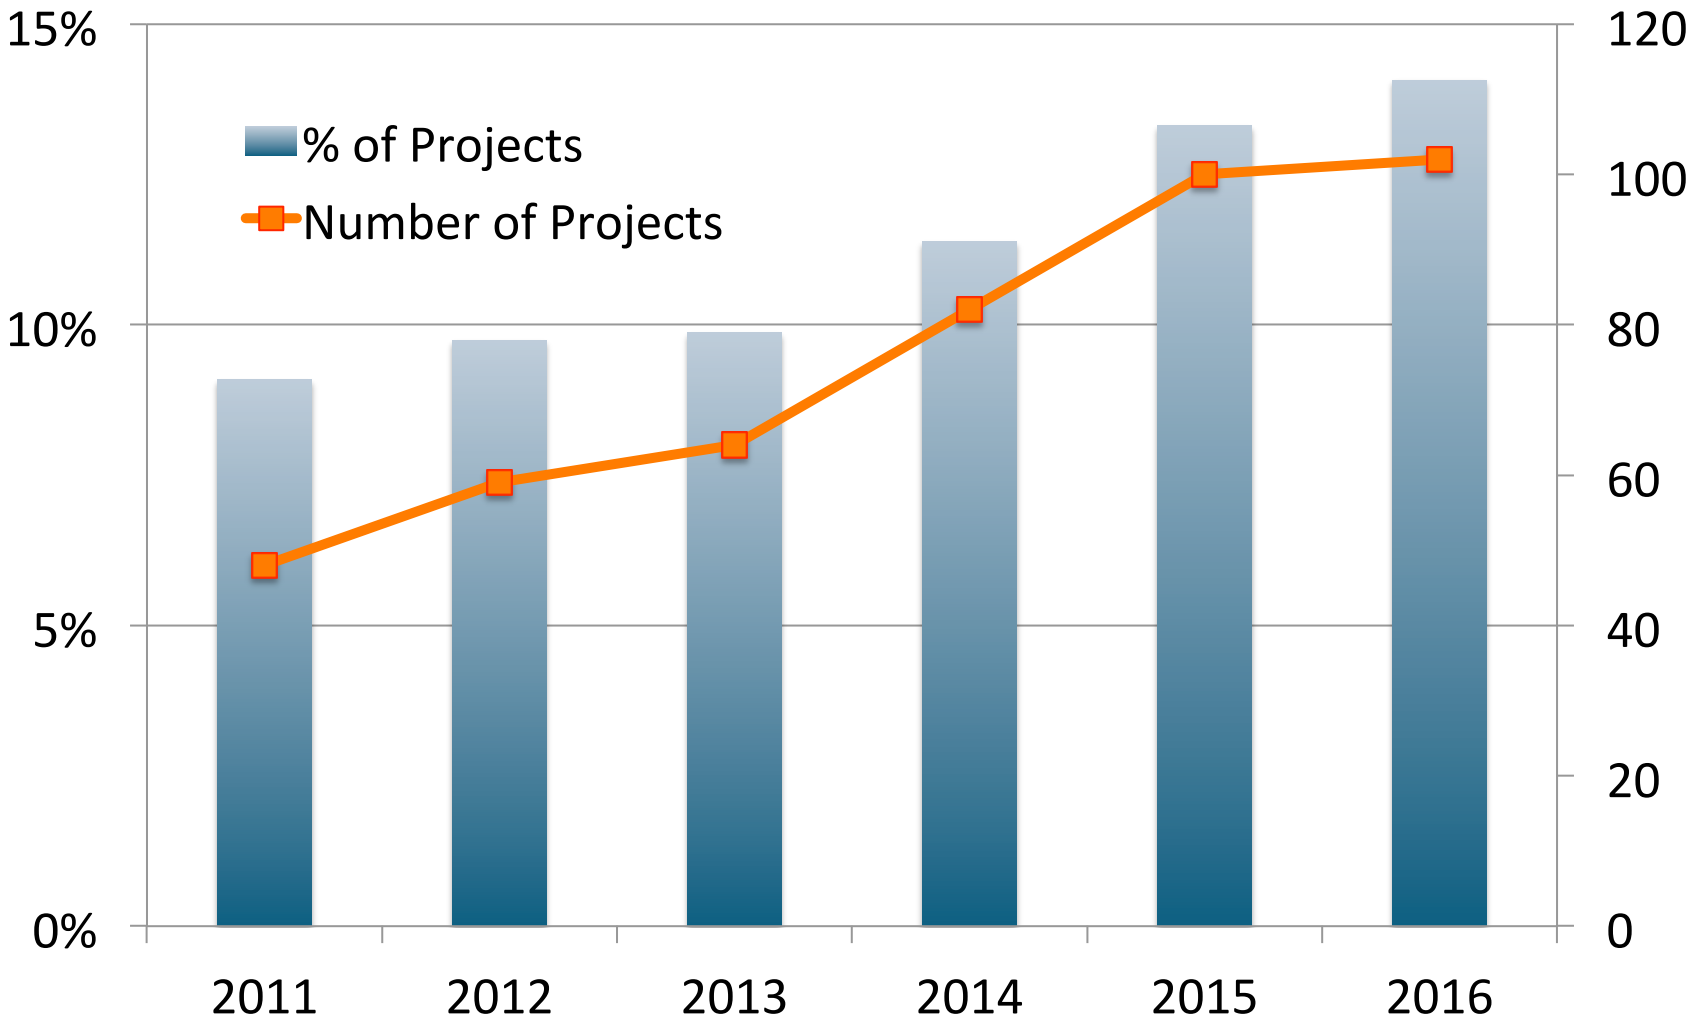

What do we mean by data-intensive projects?

- Projects analyzing data from experimental or observational facilities
- Projects combining modeling/simulation with experimental/observational data
- Projects with complex workflows that require large amounts of data movement
- Projects using analytics in new ways to gain insights into scientific domains

Some exemplars

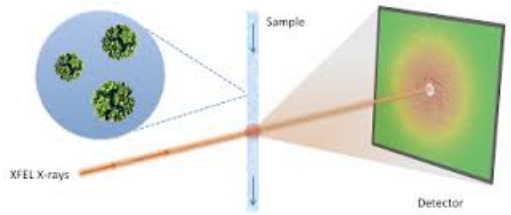

ASCR: Algorithms for next generation light sources
PI: Sethian

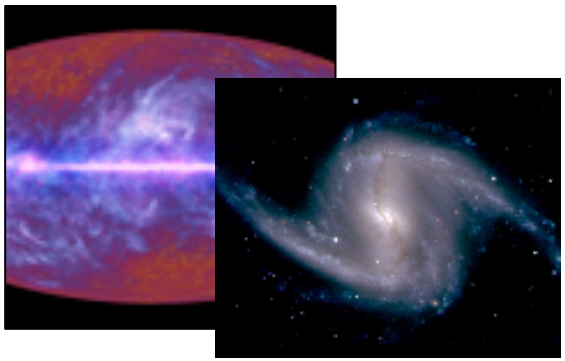

HEP: CMB Data Analysis for Planck Satellite
PI: Borrill

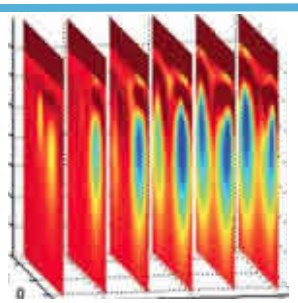

BES: Large Scale 3D Geophysical Inversion & Imaging
PI: Newman

ALS

BES: Advanced Light Source
PI: Banda

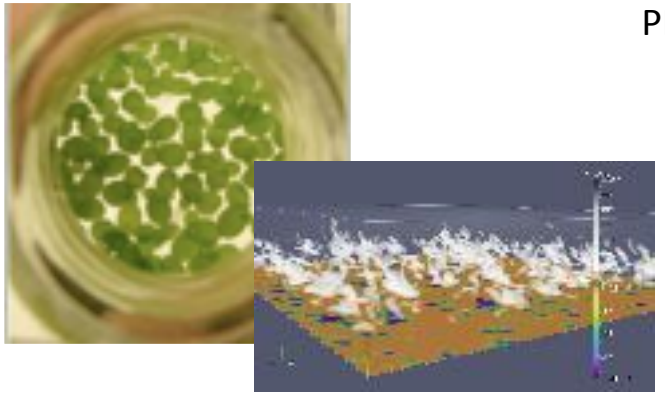

BER: Joint Genome Institute, Production Sequencing
PI: Ruben/Acting

BER: Development of the LES ARM Symbiotic Simulation and Observation Workflow
PI: Vogelmann

HEP: Dark Energy Survey
PI: Habib

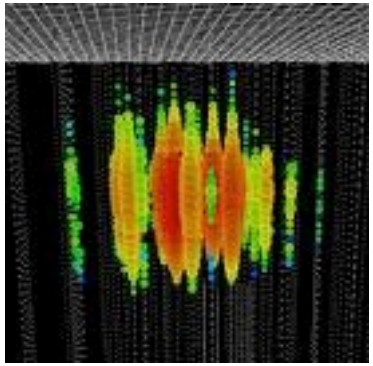

NP: Simulations and Analysis for IceCube
PI: Palczewski

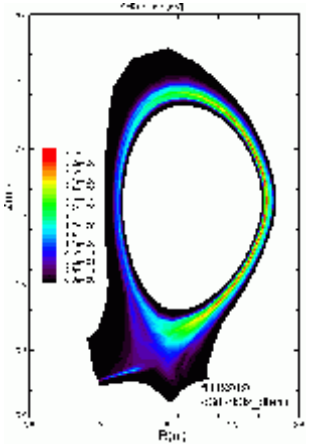

FES: LLNL MFE Supercomputing
PI: Maxim

NERSC's Org Chart

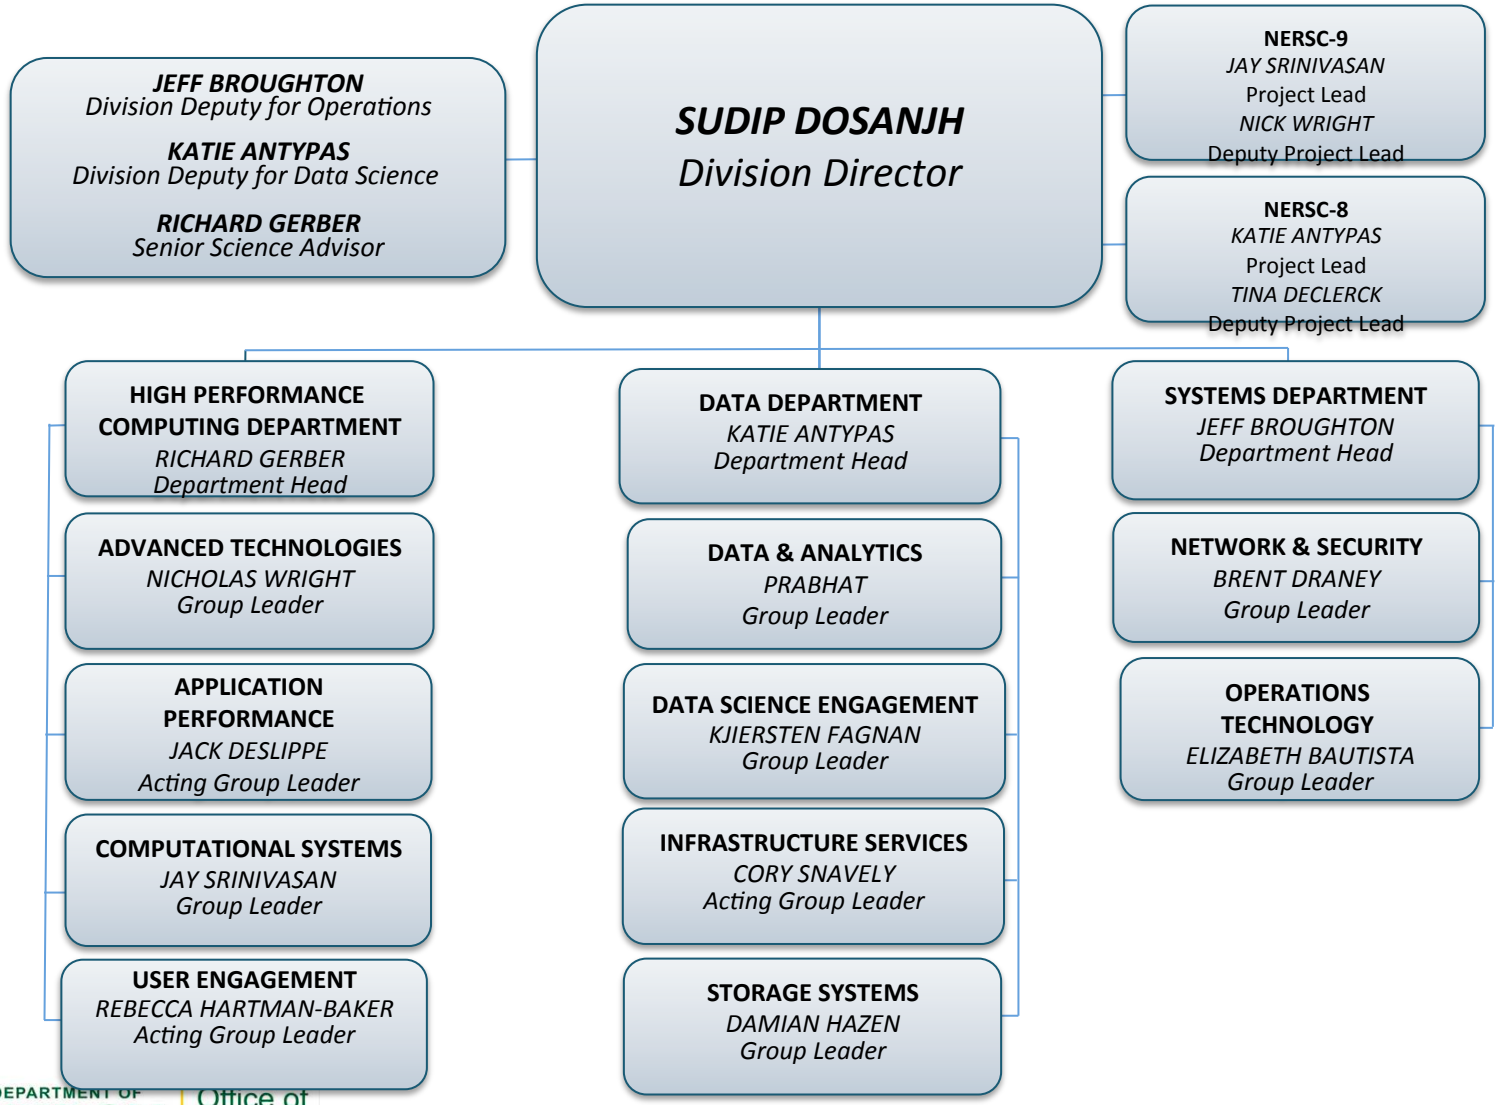

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

DAS Goal:
"Enable Data-Intensive Science at Scale"

- **Provide world-class, production quality software services for all major Data capabilities:**
  - Analytics, Management, Workflows, Transfer, Access, Visualization
- **Pioneer evaluation, research and deployment of Big Data technologies**
  - Focusing on productivity and performance
- **Engage with stakeholders to enable scientific discovery in a data-driven world**
  - Users, Computing Sciences Staff, Vendors, Researchers (Industry, Academia)

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**
- **More login nodes for managing advanced workflows**
- **Support for real time and high-throughput queues with SLURM**
- **Virtualization capabilities with Shifter (docker containers)**

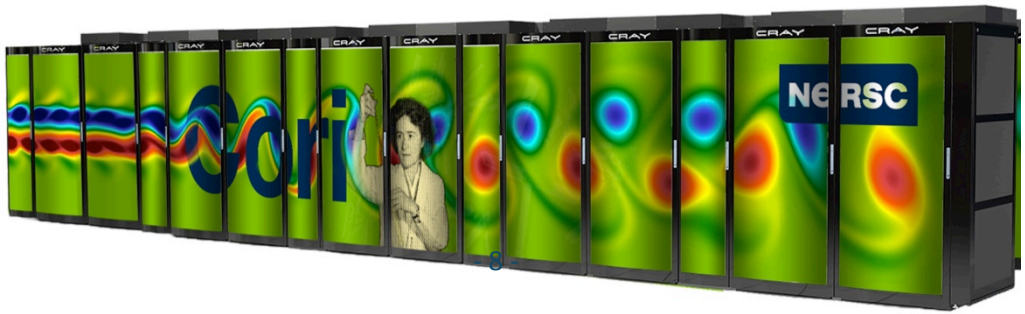

...through Innovative Technologies

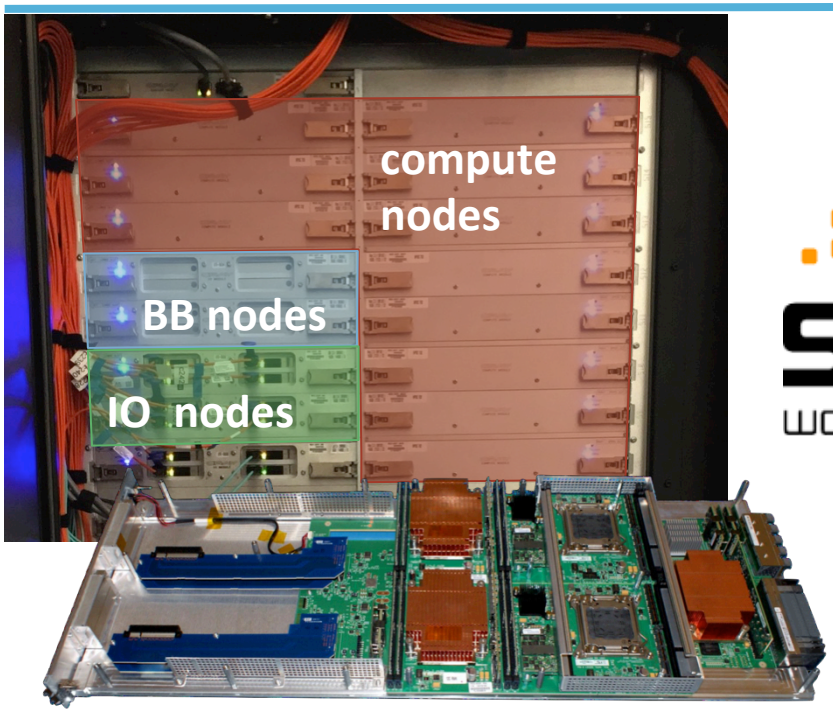

slurm workload manager

SHIFTER

docker

```
In [10]: # overlaying the small H&E and MS images

registered_ms_image = ird.transform_img_dict(my_images[2], result)
big_registered_ms_image = imresize(registered_ms_image, optical_image.shape, interp='bicubic')

# cut out low intensity region of MS image for easy viewing of underlying H&E
masked_big_ms_image = np.ma.masked_where(big_registered_ms_image < 100, big_registered_ms_image)

# plot the two images overlayed
f = plt.figure(1, figsize=(20, 20))
plt.imshow(optical_monochrome, alpha=0.7, cmap=cm.Greys_r)
plt.imshow(masked_big_ms_image, alpha=0.3, cmap=cm.jet)
plt.axes().set_axis_off()
```

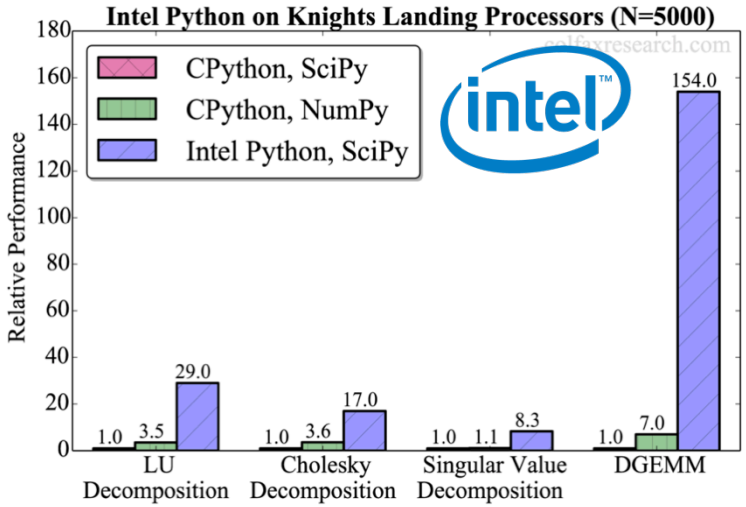

Spark

ANACONDA

jupyter

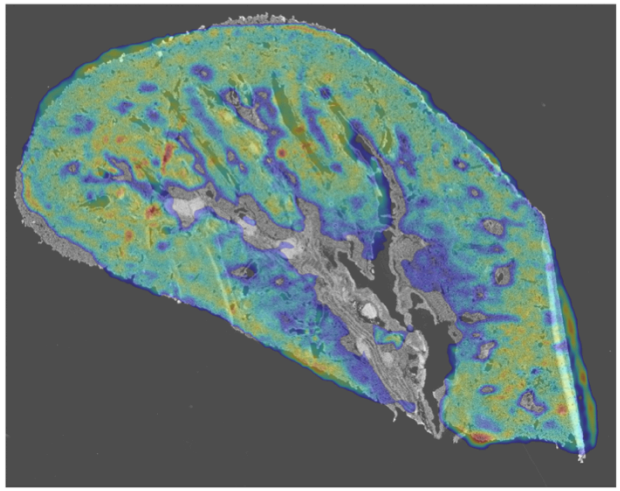

NERSC's Data and Analytics Software Portfolio

Capabilities	Technologies
Data Transfer + Access	globus online, GridFTP, php, python django, newt
Workflows	FireWorks, swift
Data Management	HDF, netCDF, ROOT Data Analysis Framework, PostgreSQL, SciDB, mongoDB, MySQL
Data Analytics	python, jupyter, IPython, Spark, R, TensorFlow, theano, Caffe, Wolfram Mathematica 10
Data Visualization	VisIt, ParaView

Have a great day at Data Day!!!

- We want to hear your feedback on how the day went
- Reminder that NERSC's call for proposals (ERCAP) is open now to apply for time in 2017

Extra slides

What has changed? Increased data rates and new sensing capabilities

LCLS
Light Source

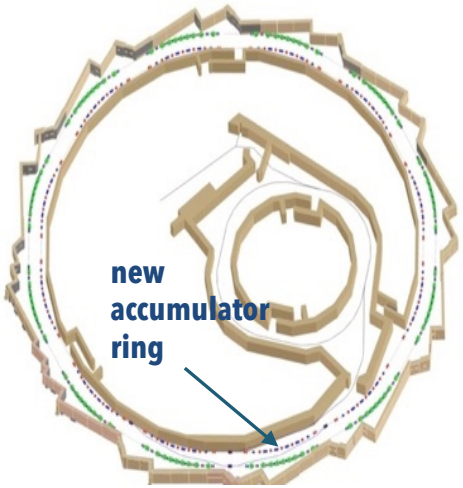

Advanced Lightsource Upgrade

Environmental
sensors

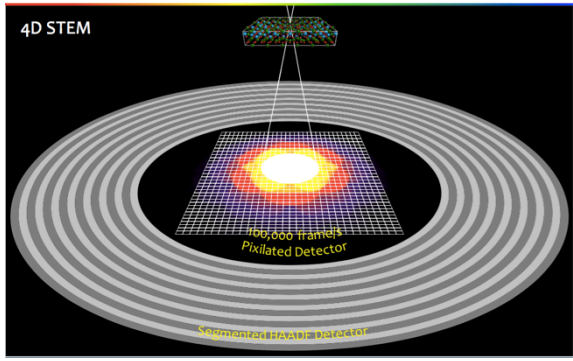

Next generation
electron microscope

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

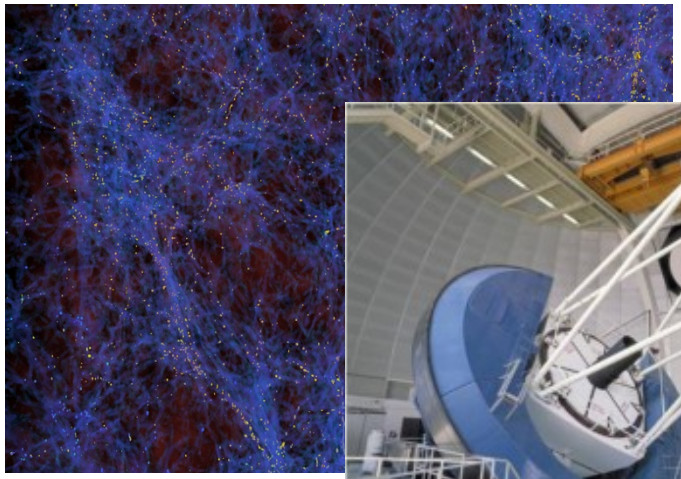

Nyx simulation of Lyman alpha forest

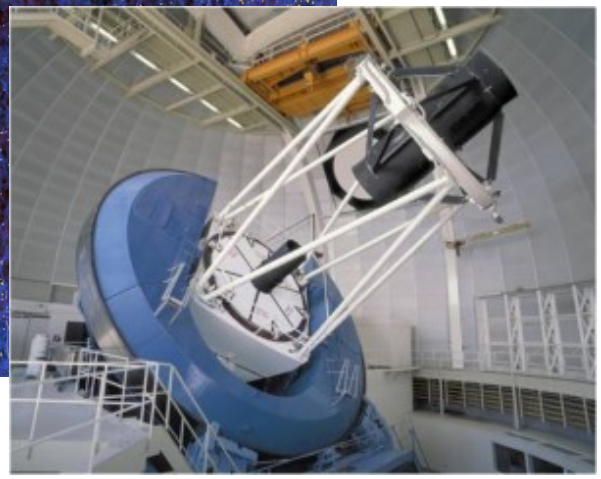

Kitt Peak National Observatory's Mayall 4-meter telescope, planned site of the DESI experiment

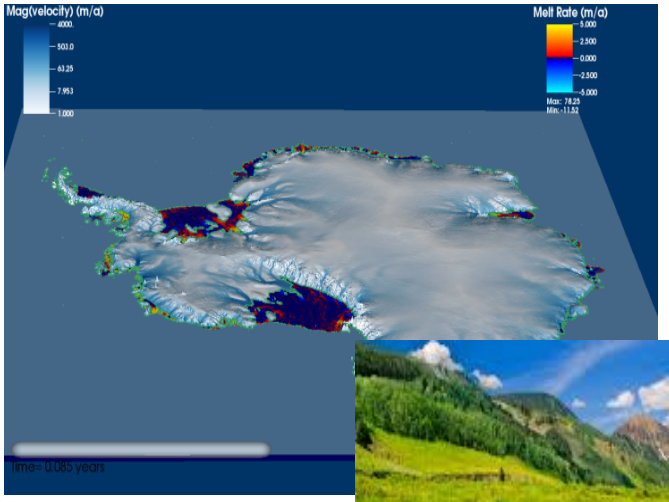

New climate modeling methods, produce new understanding of ice

Genomes to watersheds

Burst Buffer: Non-volatile storage in HPC system for application I/O acceleration

- **Opportunities/Challenge**
  - First facility to deploy a configurable Burst Buffer
  - Diverse NERSC workload pushes Burst Buffer functionality beyond checkpoint/restart use case
- **NERSC Contribution**
  - NERSC-funded NRE with Cray to develop DataWarp software
  - Early User Program enabled 30 projects to get up and running using the Burst Buffer, with NERSC staff assistance
- **Impact**
  - *CUG 2016 Best Paper Award*
  - Multiple SC16 paper submissions from early users
  - Early success for scientists accelerating their science workflows using the Burst Buffer

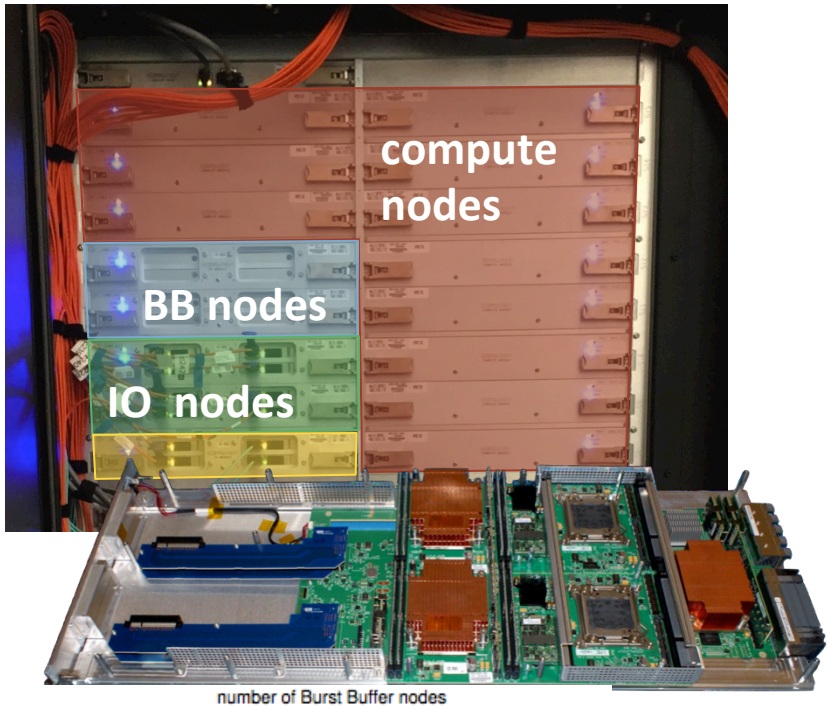

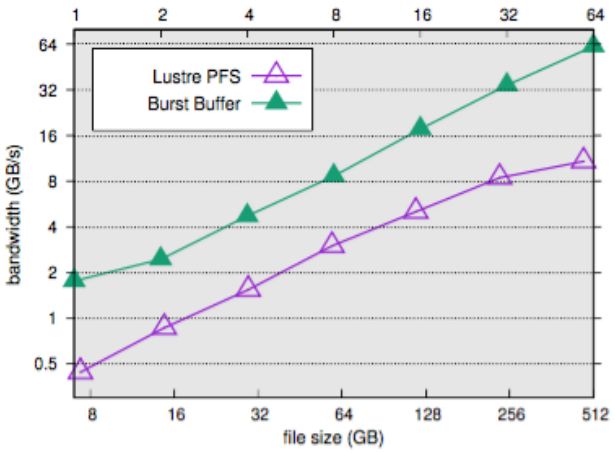

Enhanced Networking for Cori

Progress

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

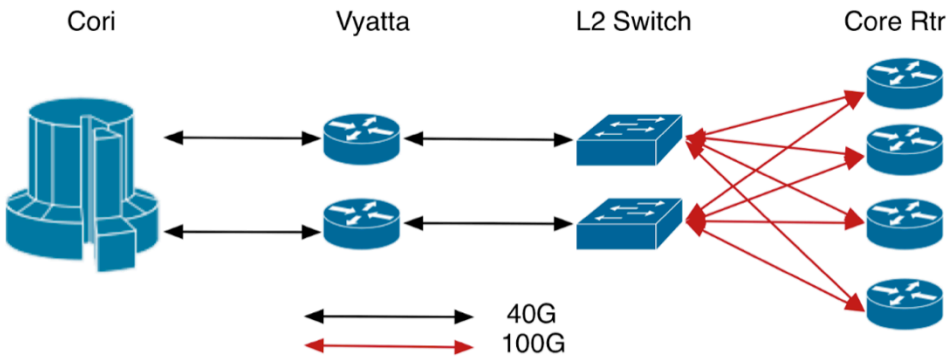

Initial Science Uses Cases

- 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)
- 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

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**

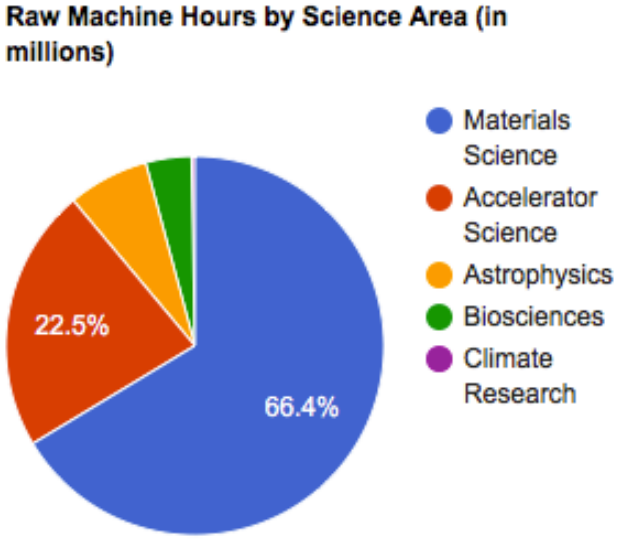

- Prototype queue used by a handful of projects at NERSC
- Real-time queue accounts for <1% of time at NERSC
- NERSC is tracking usage and use cases closely

Science Use Case: iPTF

- Nightly images transferred
- Subtractions performed
- Candidates inserted in database
- Typical turn-around time < 5 minutes

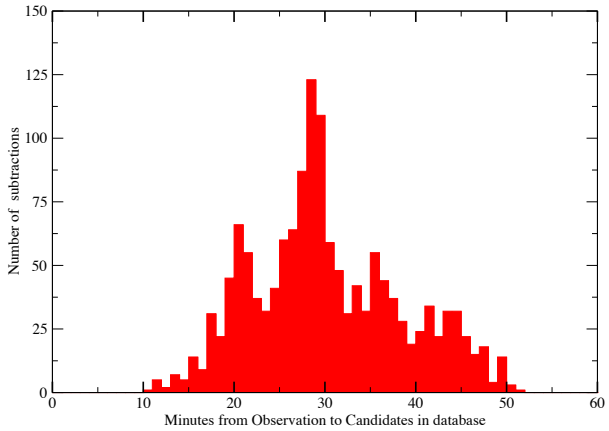

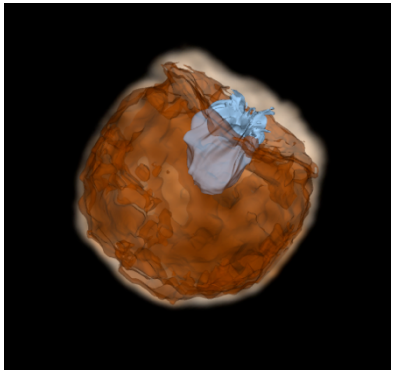

PI: Kasliwal, Nugent, Cao

DISCOVERIES
Yi Cao, *et al.* (2015) *Nature*,
"A strong ultraviolet pulse from a newborn Type Ia supernova"

Science Use Case: Advanced Light Source

- **Image reconstruction algorithms run on Cori**
- **3D volume rendered on SPOT web portal**
- **ALS beamline users receive instant feedback**

Production running at ALS beamlines:

- 24x7 Operation
- 176,293 Datasets
- 155 Beamline Users
- 1,050 TB Data Stored
- 2,379,754 Jobs at NERSC

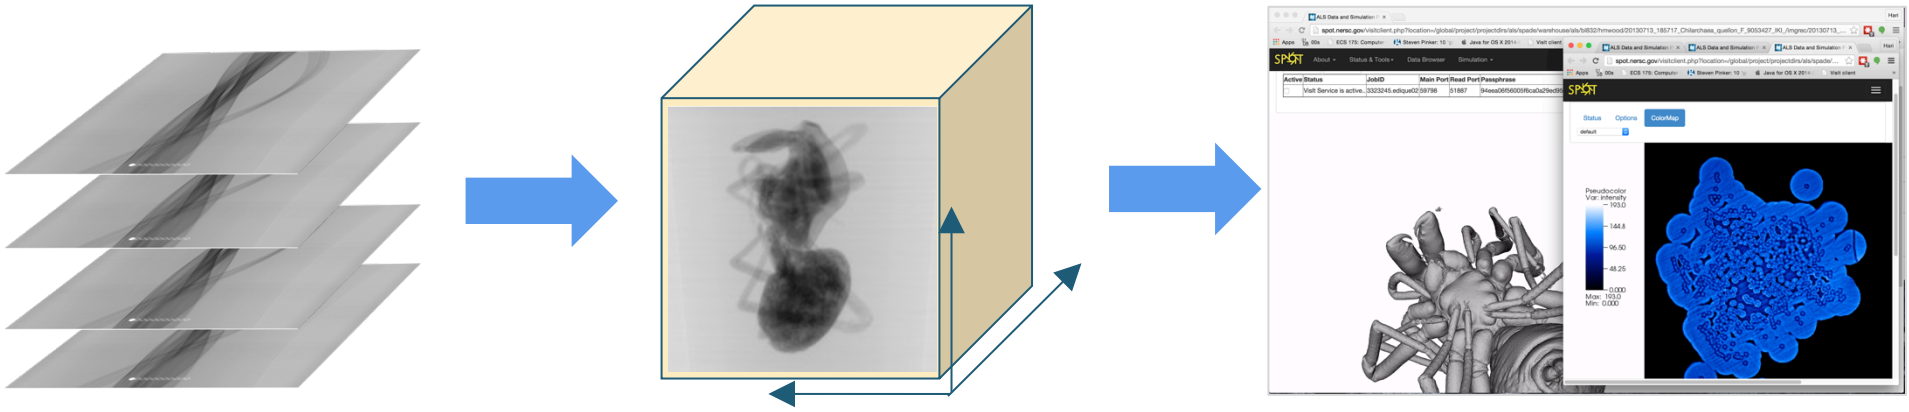

NERSC users import more data than they export!

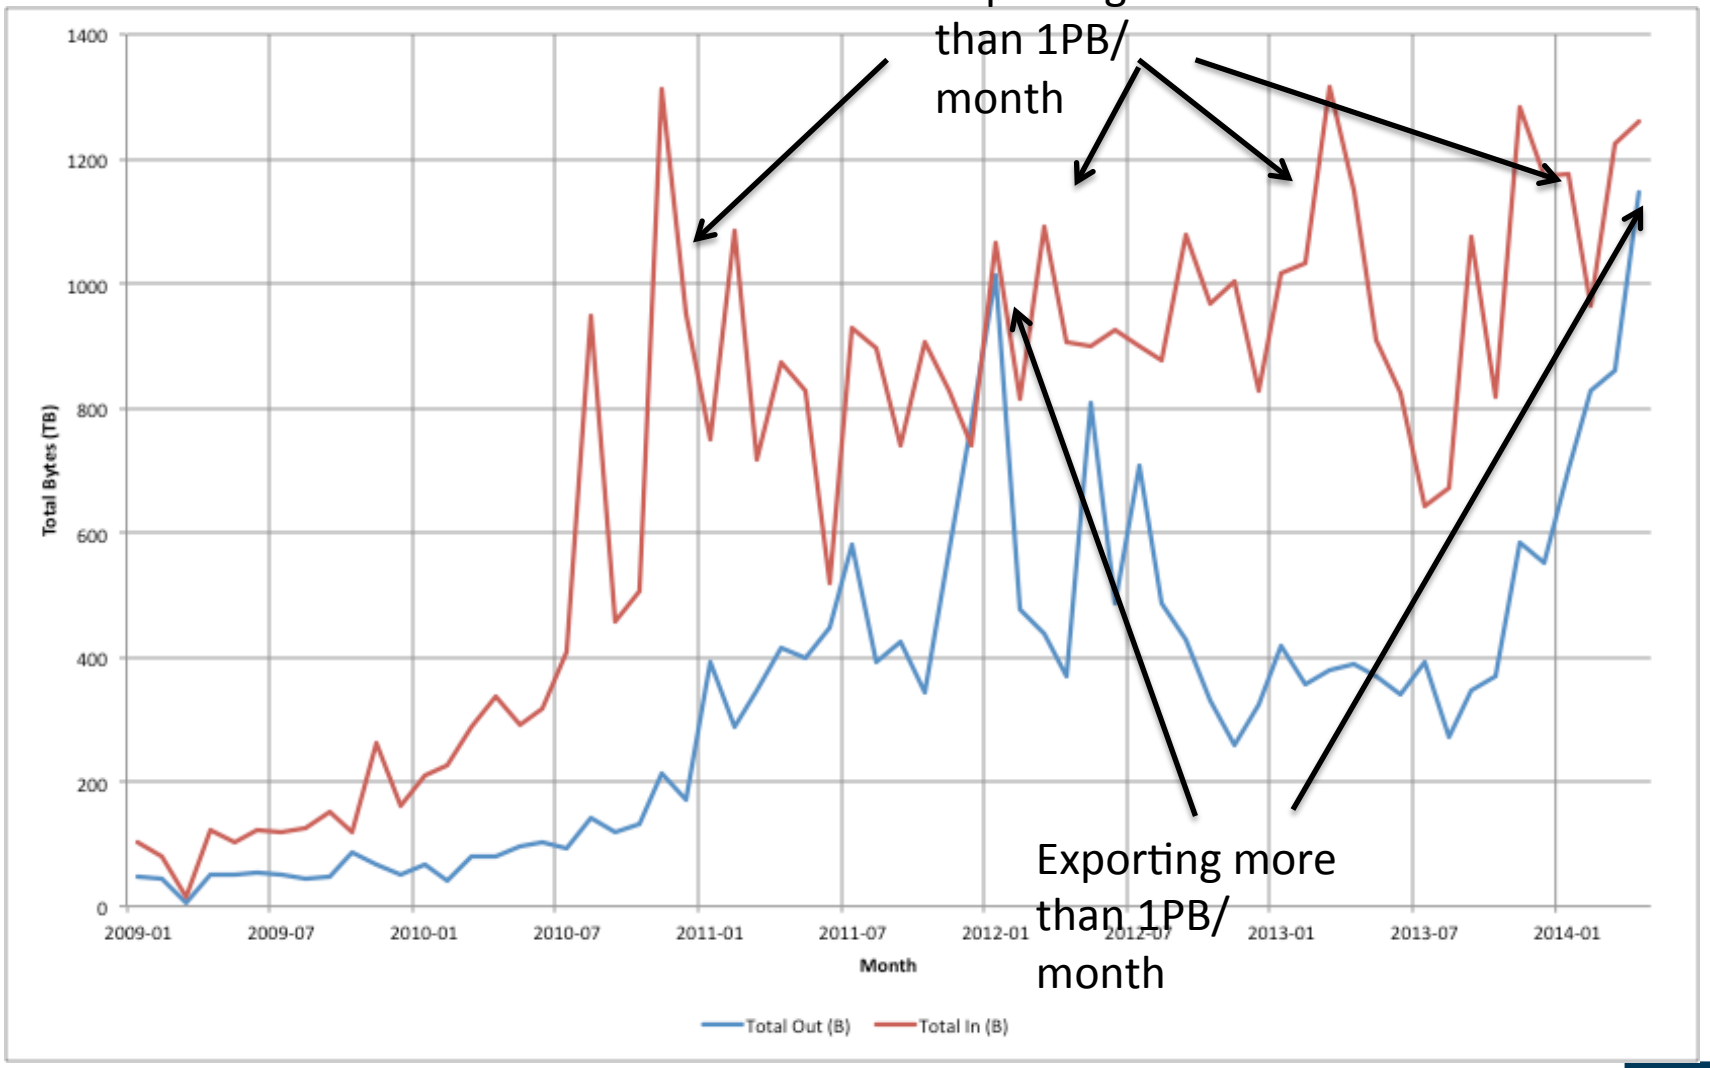

Cryo-Electron Microscopy of Transcription Initiation

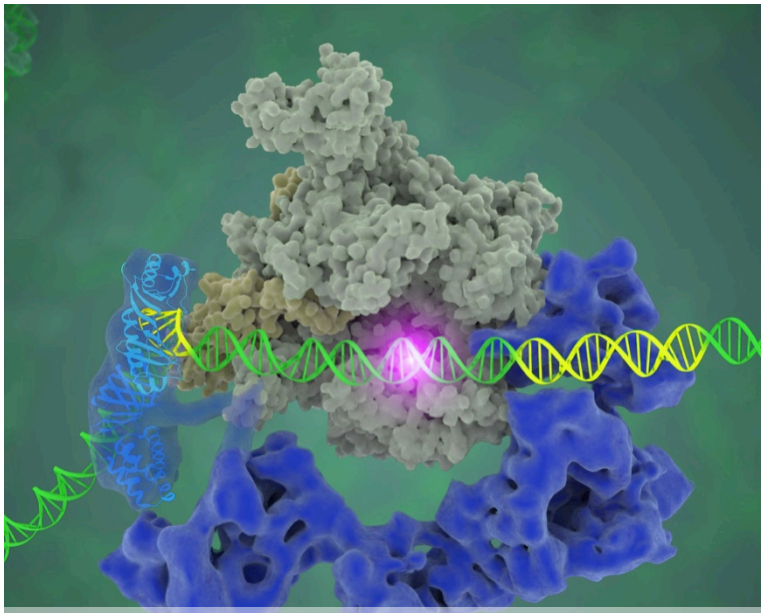

3D structure of a TFIID-containing complex with a key role in the regulation of gene expression

*Louder et al. (2016), Nature **531** (7596): 604-619*

Cryogenic electron microscopy (Cryo-EM) plays key role in allowing determination of 3D structures of highly complex molecular assemblies.

Challenge:

Perform extensive image classification and maximum likelihood optimization tasks for large CryoEM datasets 0.5M-1M images, total size O(100)GB

Result:

10X speed-up using NERSC systems, heavy value on time-to-knowledge for experiments, usage of real-time queue

Impact: Two nature publications, “Near-atomic resolution visualization of human transcription promoter opening (2016)” and “Structure of promoter-bound TFIID and model of pre-initiation complex assembly (2016)”