NERSC Accomplishments and Plans

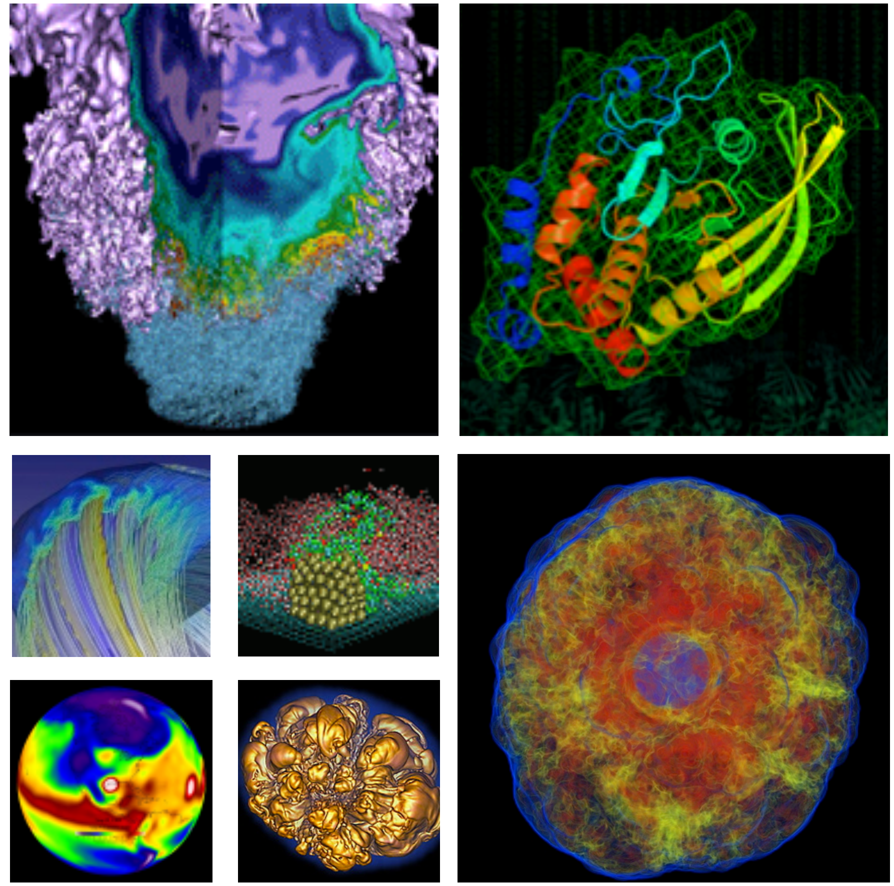

Francesca Verdier

NUG Business Meeting
February 12, 2013

February 11, 2013

NERSC Presented a Ten Year Strategic Plan to DOE ASCR

The following slides are excerpts.

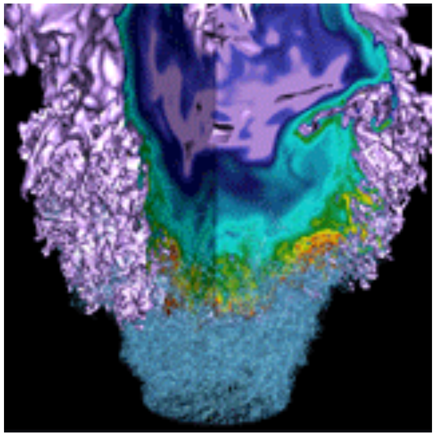

NERSC

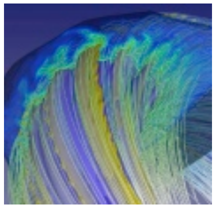

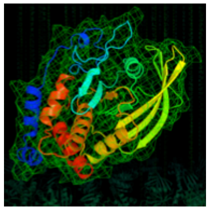

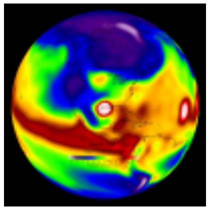

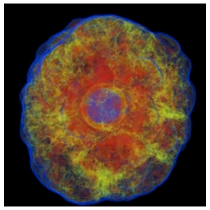

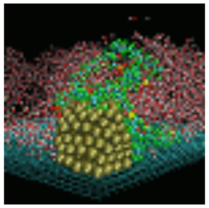

NERSC Today

1,500+ publications per year

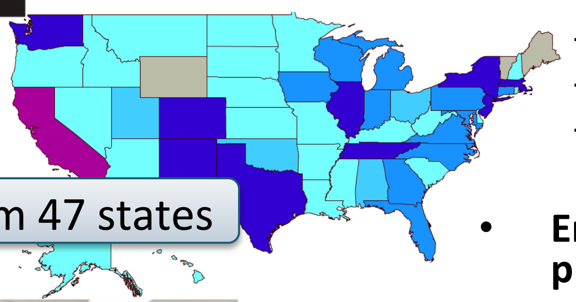

4,500 users from 47 states

Petaflop systems for science

- **Only HPC Center that supports production computing at all scales for the entire Office of Science**
  - Average 2X yearly increase in compute capacity
  - Average 1.7X yearly increase in HPSS data stored
  - Acquires systems that support production science: highly available, useable, with a science friendly security model
- **Already a significant data center**
  - Main repository of Office of Science scientific data
  - Has been a net importer of data since 2007
  - Supporting data intensive science for ASCR (extreme I/O), BER (JGI), BES (ALS), HEP/NP (PDSF cluster)
- **Enormous scientific output: about 1,500 refereed publications per year**
- **Ease of use and user support focus has resulted in increasing user base: about 500 net new users per year**
  - In 2012 we support 4500 users, 700 science projects, and 600 codes

NERSC's Newly Proposed Mission Statement

Accelerate scientific discovery at the DOE Office of Science through high performance computing and extreme data analysis.

NERSC collaborates with computer companies to deploy advanced HPC and data resources

- Hopper (N6) and Cielo (ACES) were the first Cray petascale systems with a Gemini interconnect
- Edison (N7) will be the first Cray petascale system with Intel processors, Aries interconnect and Dragonfly topology (serial #1)
- N8 and Trinity (ACES) are being jointly designed as on-ramps to Exascale
- Architected and deployed data platforms including the largest DOE system focused on genomics
- One of the first facility wide filesystems

We focus on the scientific impact of our users

NERSC

- **1,500 journal publications per year**
- **10 journal cover stories per year on average**
- **Simulations at NERSC were key to 2 Nobel Prizes (2007 and 2011)**
- **Supernova 2011fe was caught within hours of its explosion in 2011 and telescopes from around the world were redirected to it the same night**
- **Data resources and services at NERSC played important roles in two of Science Magazine's Top Ten Breakthroughs of 2012 — the discovery of the Higgs boson and the measurement of the Θ_{13} neutrino weak mixing angle**
- **MIT researchers developed a new approach for desalinating sea water using sheets of graphene, a one-atom-thick form of the element carbon. Smithsonian Magazine's fifth "Surprising Scientific Milestone of 2012."**
- **Four of Science Magazine's insights of the last decade (3 in genomics, 1 related to cosmic microwave background)**

13 Journal Covers in 2012

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

We support a diverse workload

- **Many codes (600+) and algorithms**
- **Computing at Scale and at High Volume**

Top Codes by Algorithm

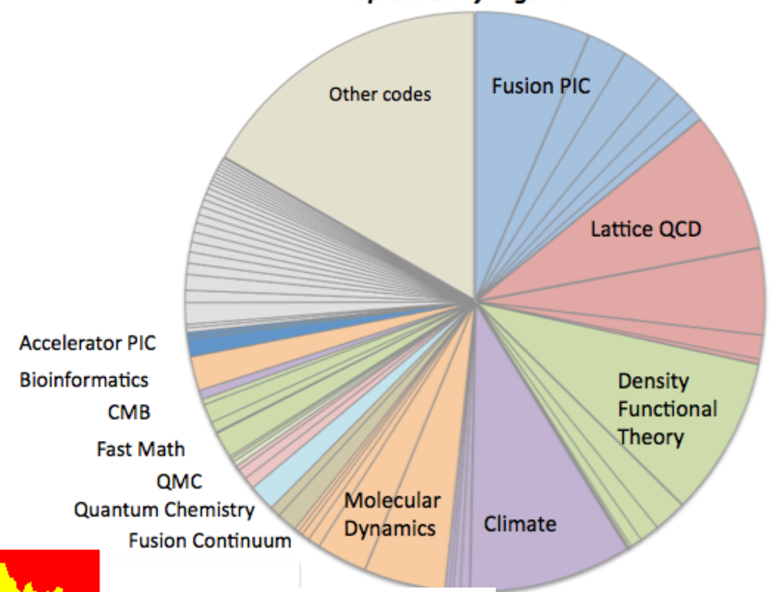

2012 Job Size Breakdown on Hopper

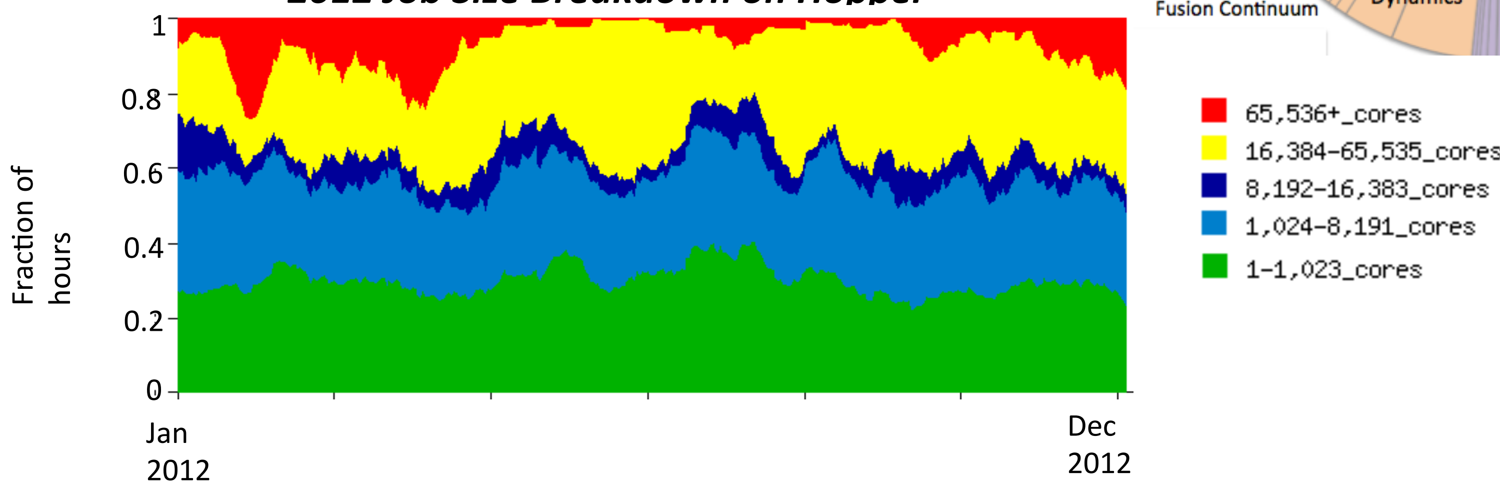

Application Code Distribution

Top Application Codes

Jan – Nov 2012

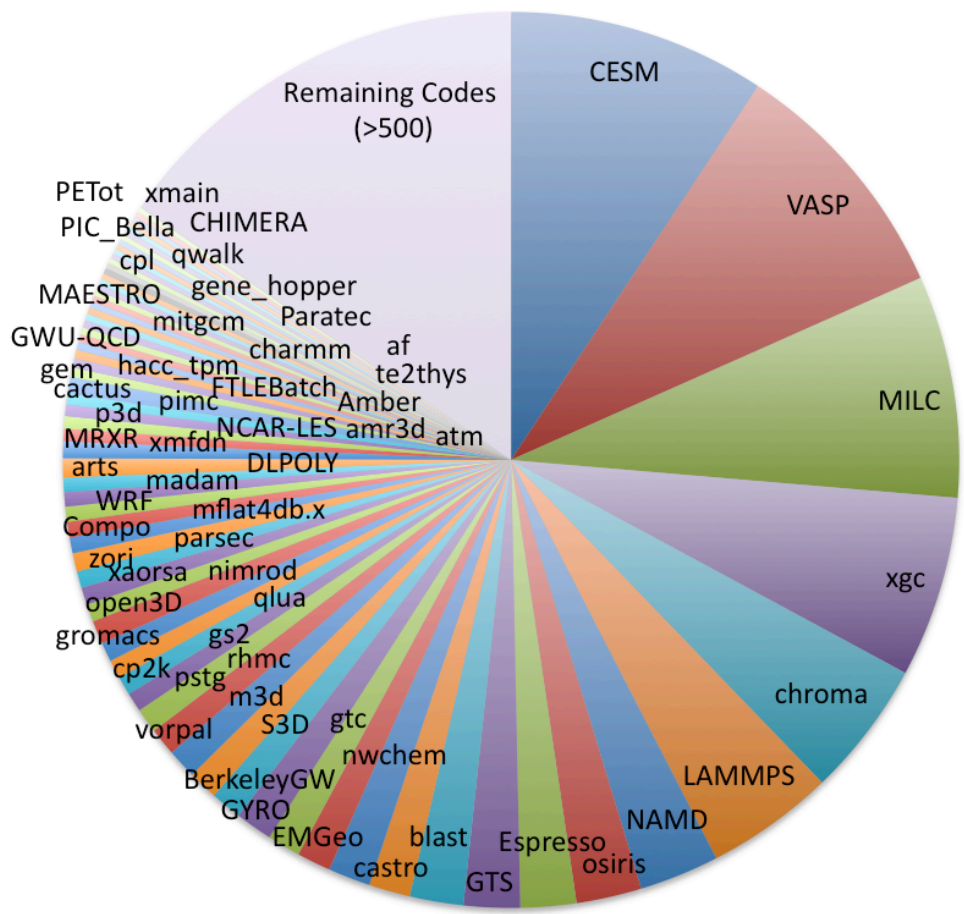

- **10 codes make up 50% of workload**
- **25 codes make up 66% of workload**
- **75 codes make up 85% of workload**
- **remaining codes make up bottom 15% of workload**

Approximately 80% of the workload needs to transfer to NERSC-8
(20% can remain on Edison for the next few years)

Our operational priority is providing highly available HPC resources backed by exceptional user support

- **We maintain a very high availability of resources (>90%)**
  – One large HPC system is available at all times to run large-scale simulations and solve high throughput problems
- **Our goal is to maximize the productivity of our users**
  – One-on-one consulting
  – Training (e.g., webinars)
  – Extensive use of web pages
  – We solve or have a path to solve 80% of user tickets within 3 business days

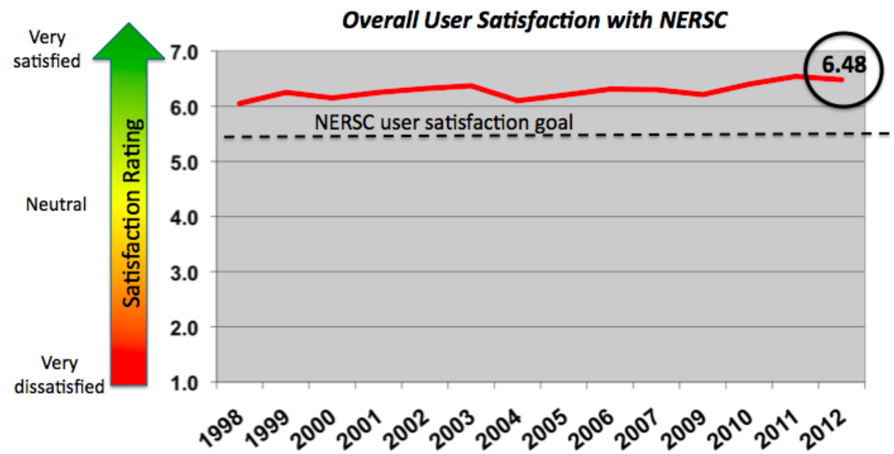

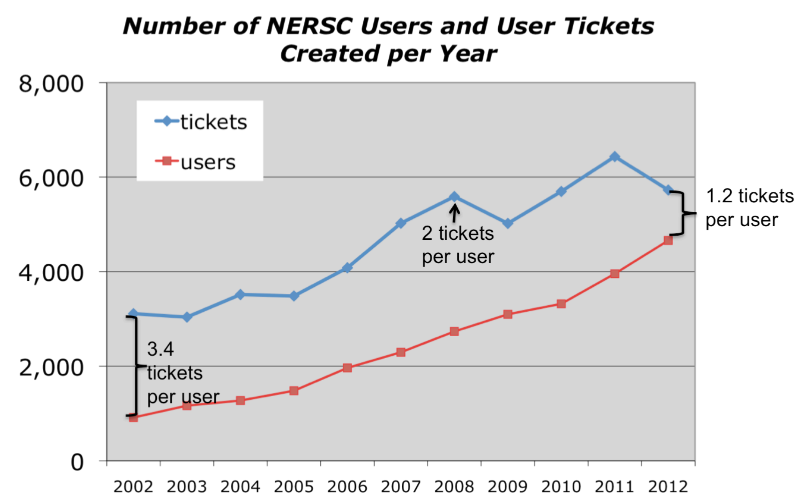

3

Keeping up with user needs will be a challenge

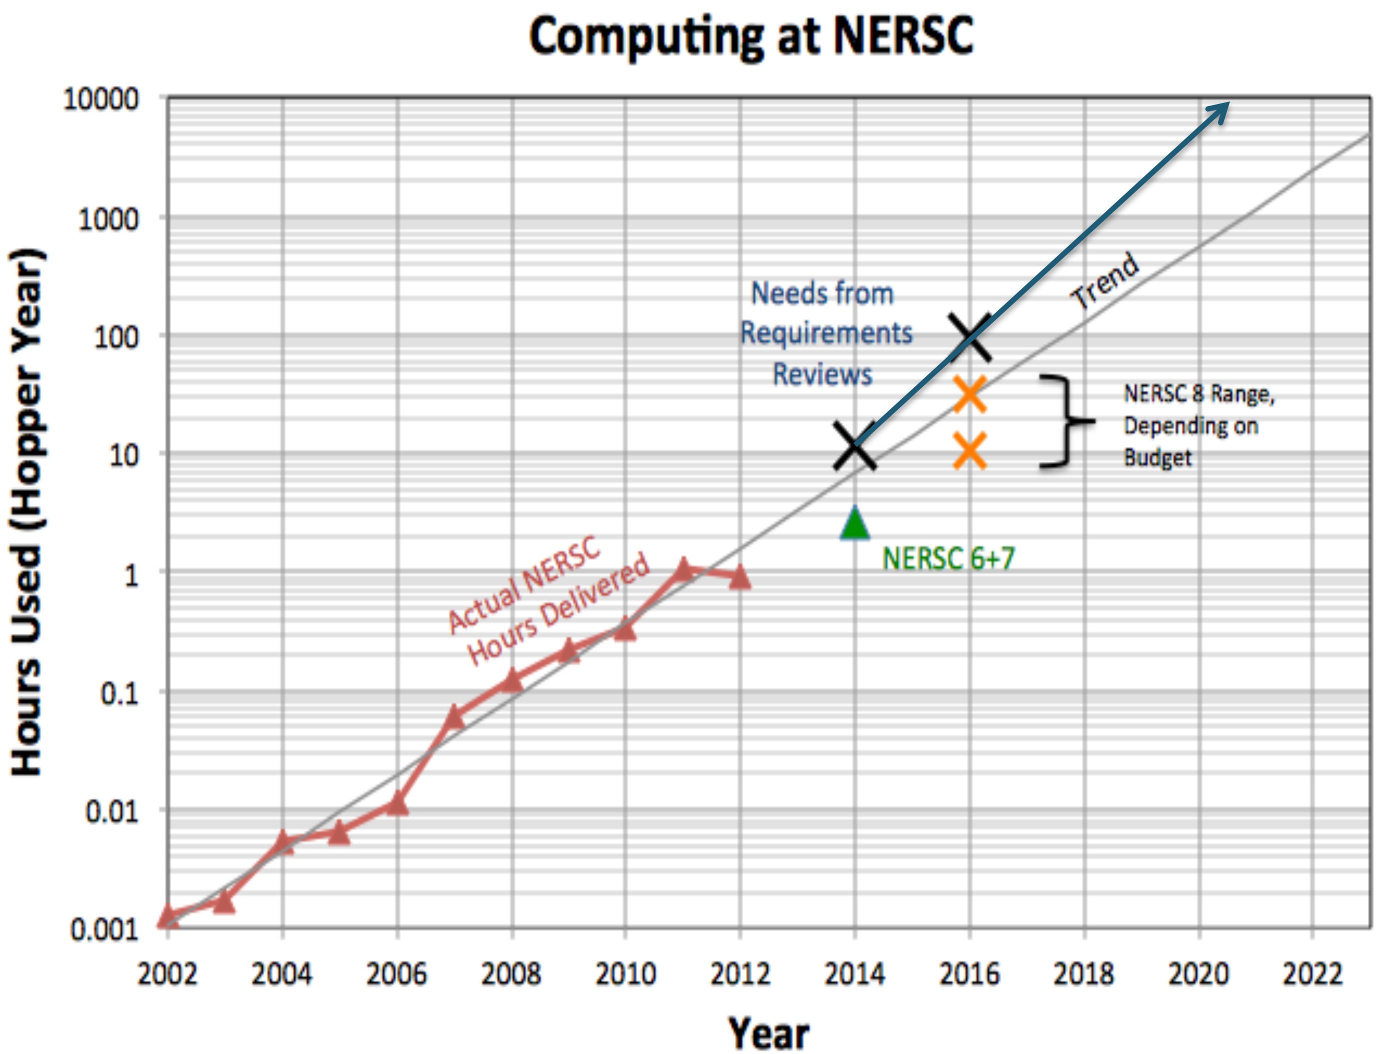

Future archival storage needs

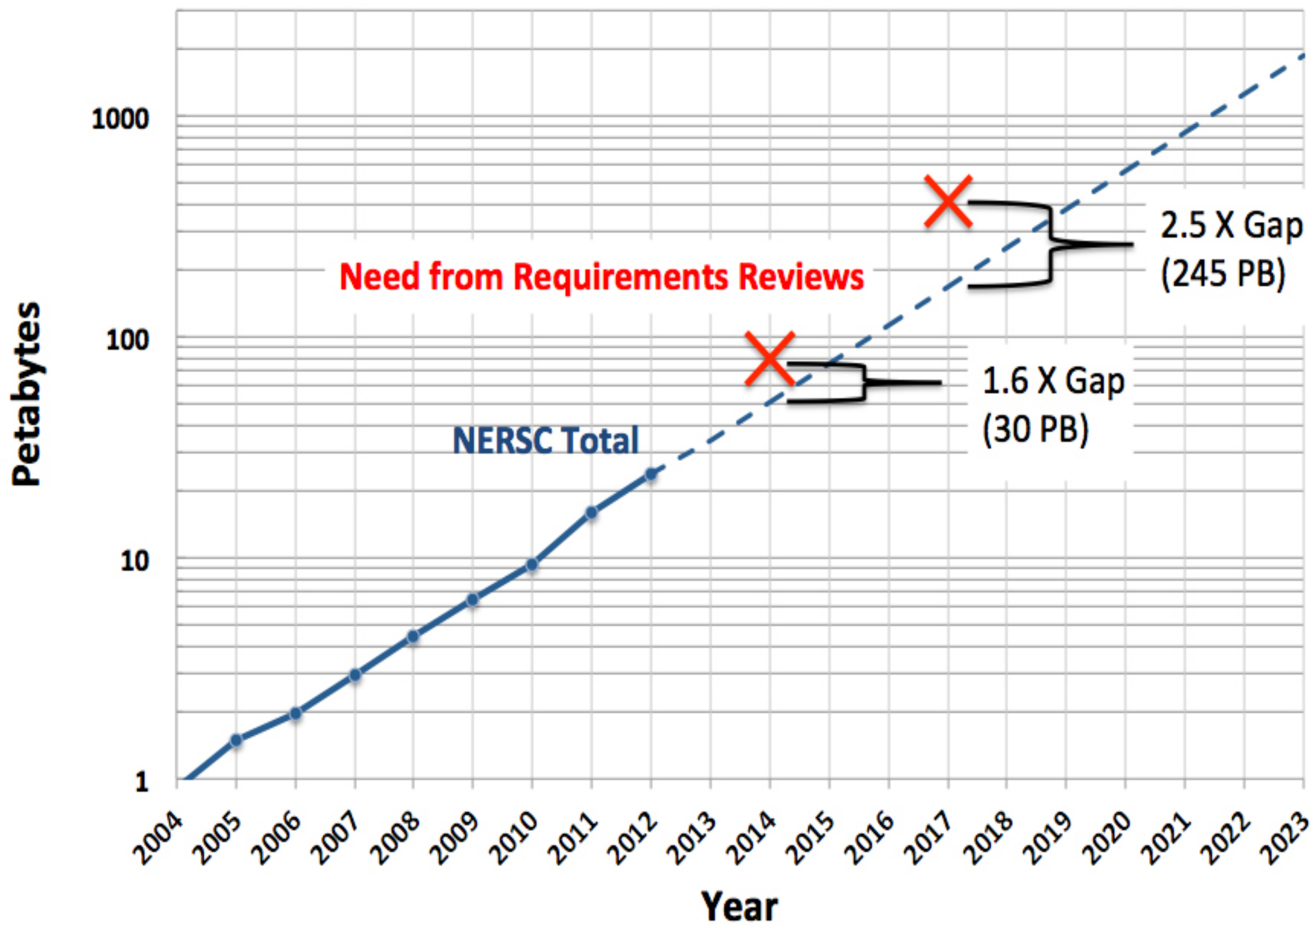

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Strategic Objectives

- **Meet the ever growing computing and data needs of our users by**
  - providing usable exascale computing and storage systems
  - transitioning SC codes to execute effectively on many core architectures
  - influencing the computer industry to ensure that future systems meet the mission needs of SC
- **Increase the productivity, usability, and impact of DOE's user facilities by providing comprehensive data systems and services to store, analyze, manage, and share data from those facilities**

Providing usable Exascale computing and storage systems

- **We made NERSC-7 an x86-based system because our broad user base wasn't ready in 2013 for GPUs, accelerators or greatly increased threading**
- **We will deploy pre-Exascale systems in 2015 (NERSC-8) and 2019 (NERSC-9), and an Exascale system in 2023. Our strategy is:**
  - Open competition for best solutions
  - Focus on the performance of a broad range of applications, not synthetic benchmarks
  - General-purpose architectures are needed in order to support a wide range of applications, both large-scale simulations and high volumes of smaller simulations
  - Earlier procurements to influence designs
  - Leverage Fast Forward and Design Forward
  - Engage co-design efforts
  - Transition users to a new programming model

NEW

Strategy for transitioning the SC Workload to Energy Efficient Architectures

- **We will deploy testbeds to gain experience with new technologies and to better understand emerging programming models and potential tradeoffs.**
  - In particular, we will deploy a testbed representative of the NERSC-8 architecture as soon as it is determined.
  - Already have GPU testbed (Dirac) and an Intel MIC cluster is being installed.
- **We will have in-depth collaborations with selected users and application teams to begin transitioning their codes to our testbeds and to NERSC-8**
  - We will choose partnerships based on level of interest, expected application usage, and algorithmic diversity
- **We will develop training and online resources to help the rest of our users based on our in-depth collaborations, as well as on results from co-design centers and ASCR research**
  - We will leverage our existing training vehicles, such as online webinars, as much as possible.
- **We would like to add consultants with an algorithms background who can help users when they have questions about improving the performance of key code kernels**

Extreme Data Strategy

- **Partner with DOE experimental facilities to identify requirements and create early success**
  - NERSC pilot projects have shown automated data pipelining, indexing, search, and tape archive to be a good fit with current and future ALS needs.
- **Develop and deploy new data resources and capabilities**
  - We would like to accelerate NERSC's traditional storage growth rate to meet rapidly increasing user requirements for capacity and bandwidth.
  - We would like to enhance the data processing capabilities of NERSC-7 in 2014 by adding large memory visualization/analysis nodes, adding a flash-based burst buffer or node local storage, and deploying a throughput partition for fast turnaround of jobs.
  - We project deploying new follow-on data-centric systems in 2017 (NERSC Data-1) and 2021 (NERSC Data-2).
- **We would like to provide new classes of HPC expertise required for data-intensive workloads**
  - Database-driven workflows and storage
  - Scalable structured and unstructured object stores
  - Application software solutions to traverse massive data for search or analysis
  - Sophisticated web-based gateways to interact with and leverage data
  - Comprehensive scientific data curation beyond simple archiving
- **Leverage ESnet and ASCR research to create end-to-end solutions**

Near Term Plans: covered in the remaining presentations

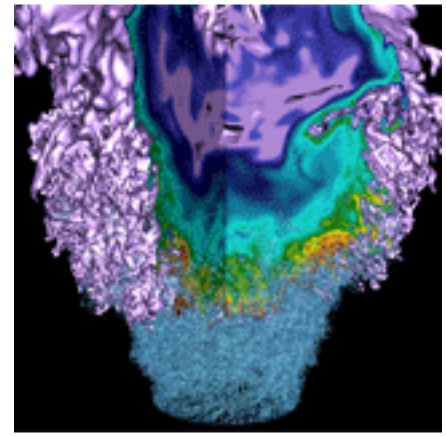

NERSC

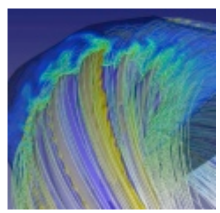

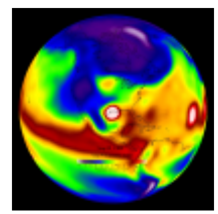
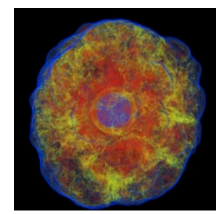
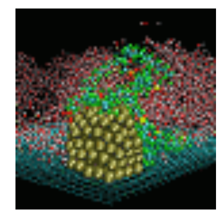

We are deploying the CRT facility to meet the ever growing computing and data needs of our users

- **Four story, 140,000 GSF**
  - Two 20Ksf office floors, 300 offices
  - 20K -> 29Ksf HPC floor
  - Mechanical floor
- **42MW to building**
  - 12.5MW initially provisioned
  - WAPA power: Green hydro
- **Energy efficient**
  - Year-round free air and water cooling
  - PUE < 1.1
  - LEED Gold
- **Occupancy Early 2015**

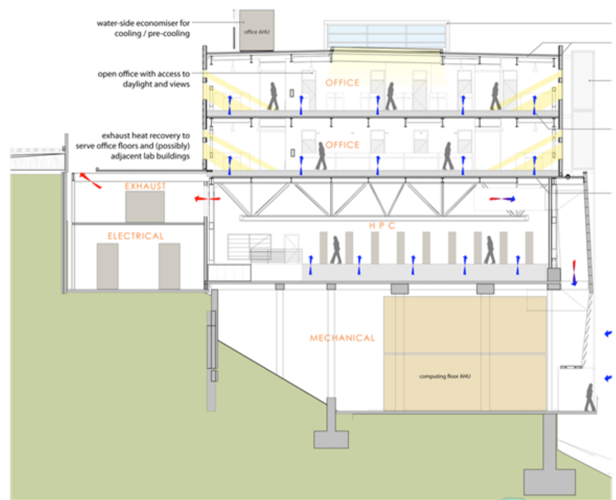

Edison Installation Timeline

- **Phase 1**
  - 6 login nodes
  - 4 cabinets (664 nodes, 10,624 cores) with Sandy Bridge processors
  - 1 file system (35GB/s, 1.6 PB)
  - 20 early users enabled Feb. 5
  - All users likely by end Feb
- **Phase 2 (Summer/Fall 2013)**
  - Add 24 cabinets with Ivy Bridge processors
  - Upgrade first 4 cabinets with Ivy Bridge
  - Complete all 3 file systems (140 GB/s, 6.4 PB)
  - All users likely by end of Fall 2013
- **Charging starts AY 2014**

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC-8 and partnership with Trinity project

- Procure an HPC system to support rapidly increasing computational demands of NERSC users
- Provide a significant increase, at least 10 times the sustained performance of Hopper on a set of representative benchmarks

Trinity/NERSC-8 Collaboration

- Delivery in the 2015/2016 time frame
- Platform needs to begin to transition users to more energy-efficient many-core architectures.
- Plans for joint Trinity/NERSC-8 RFP calling for two distinct systems of similar technology with the <u>intention</u> to award both systems to the same vendor.