Support for Astronomy & Astrophysics at NERSC

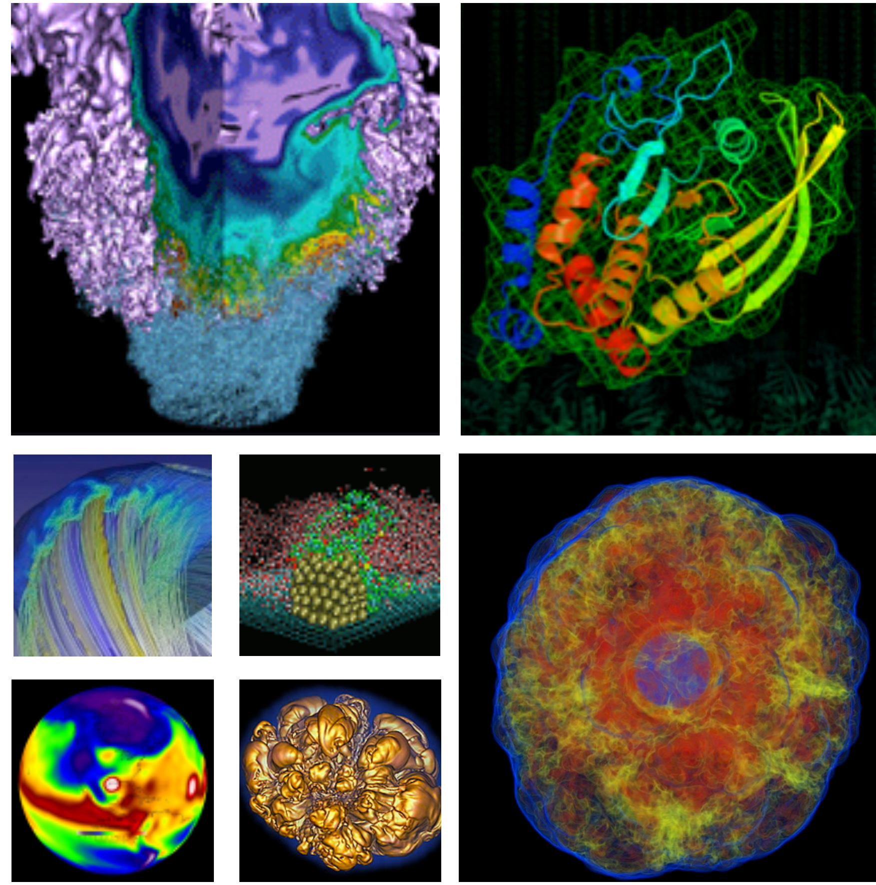

NERSC 40 YEARS at the FOREFRONT

Richard Gerber
Senior Science Advisor
User Services Group Lead
NERSC @ Berkeley Lab

April 9, 2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Outline

- What is NERSC
- Who uses NERSC?
- How to get access to NERSC Resources
- NERSC Requirements Gathering

NERSC is the Production HPC & Data Facility for DOE Office of Science Research

NERSC 40 YEARS at the FOREFRONT

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in U.S.

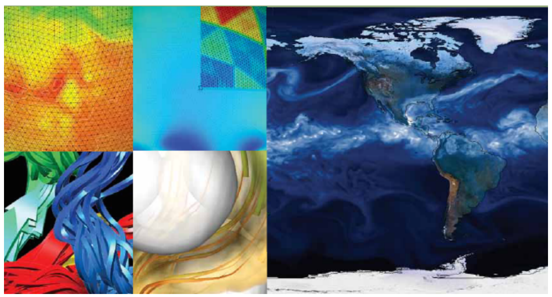

Bio Energy, Environment

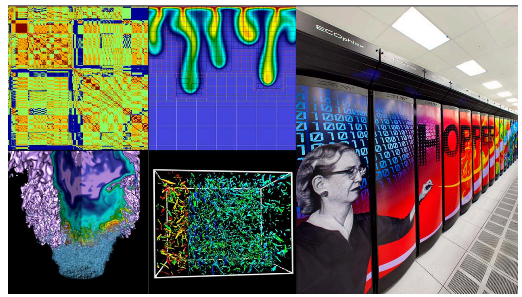

Computing

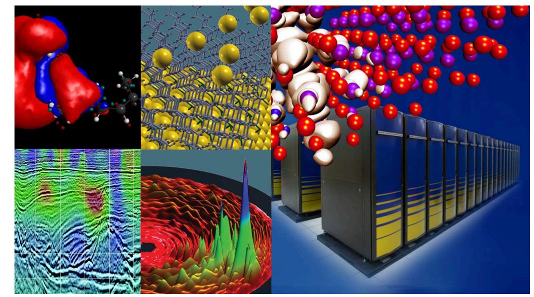

Materials, Chemistry, Geophysics

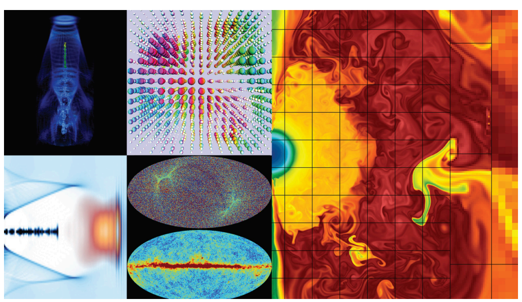

Particle Physics, Astrophysics

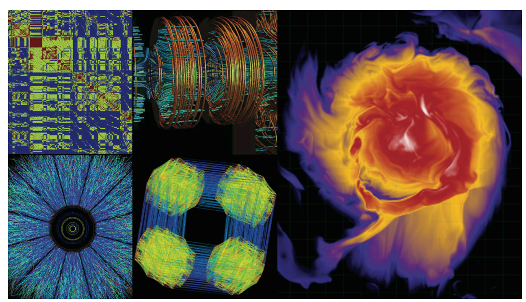

Nuclear Physics

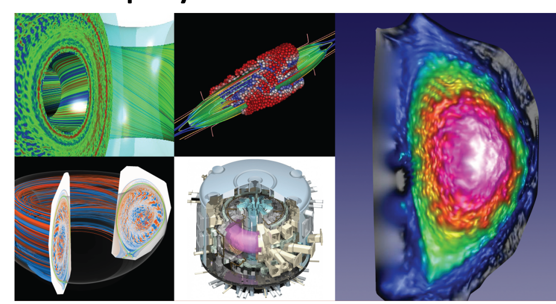

Fusion Energy, Plasma Physics

U.S. DEPARTMENT OF ENERGY | Office of Science

Some Benefits of NERSC

- Large, state-of-the-art computing and data systems
- Consulting, system admin, 24x7 operations support
- Well maintained software environment, prebuilt optimized applications & libraries
- Designed for massive parallelism, but supports all scales
- Easily share data and codes
- Easy to use account management
- Large permanent archival data storage
- Ongoing technology refreshes
- Word-class cybersecurity
- Open science environment
- Web-based science /data gateways

NERSC

Focus on Science

- A word-class resource to support world-class science.
- 1,500 refereed journal publications per year
- Supports Nobel-prize winning projects: Chemistry 2013, Physics 2011, Peace 2007

Large diverse user community

- 5,000 users, 700 projects
- From 48 states, 65% from universities
- Many large international collaborations

Science-driven systems and services

- Designed to support science
- Optimized for scientific productivity at cutting-edge scale

17 Journal Covers in 2013

Demographics

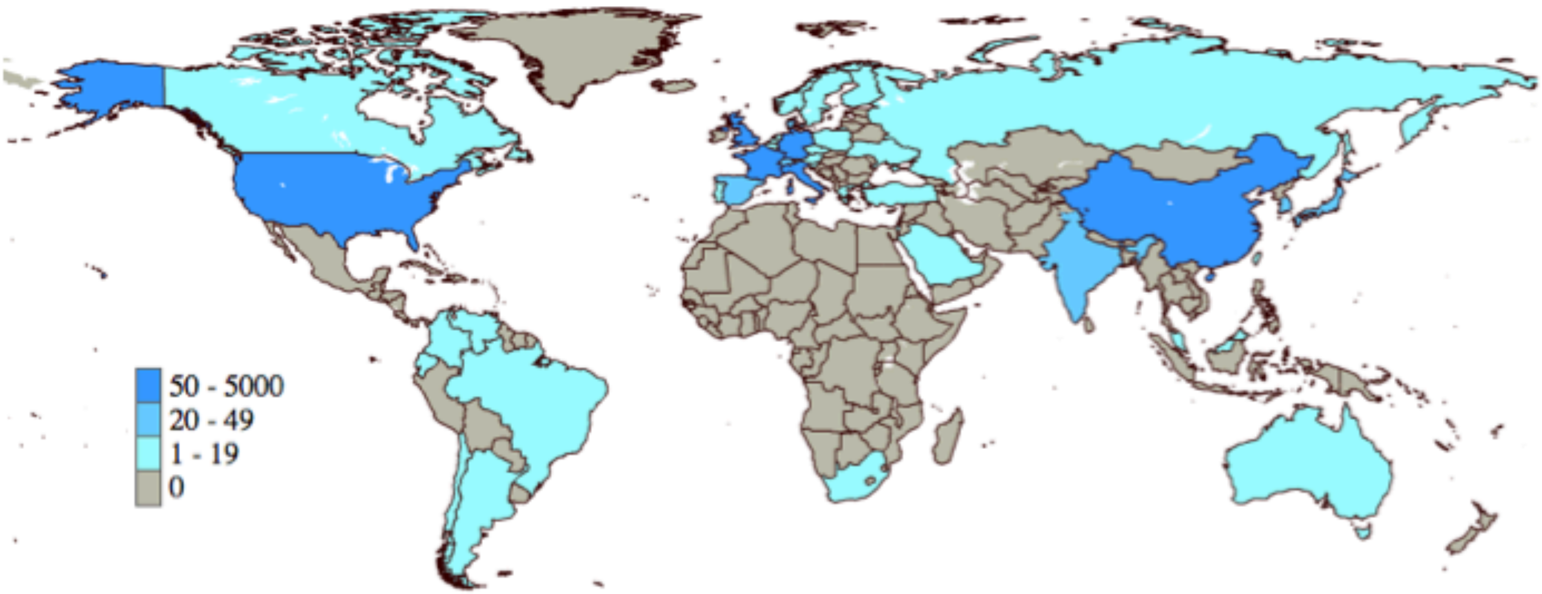

642 users from outside the U.S.

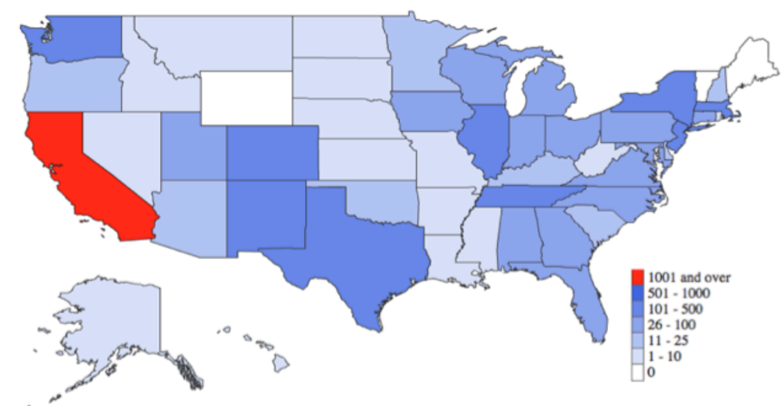

NERSC Systems Today

NERSC 40 YEARS at the FOREFRONT

Edison: 2.57PF, 357 TB RAM

Cray XC30 5,576 nodes, 134K Cores

Hopper: 1.3PF, 212 TB RAM

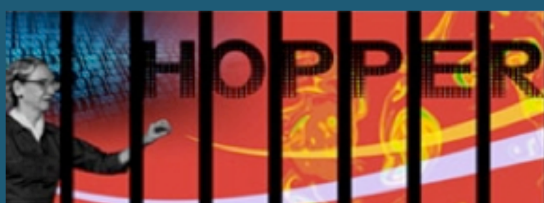

Cray XE6 6,384 nodes 150K Cores

.6 PB Local Scratch 163 GB/s

16 x FDR IB

2.2 PB Local Scratch 70 GB/s

5 x QDR IB

80 G

50 G

5 G

12 G

Global Scratch — 3.6 PB, 5 x SFA12KE

/project — 5 PB, DDN9900 & NexSAN

/home — 250 TB, NetApp 5460

HPSS — 50 PB stored, 240 PB capacity, 20 years of community data

Production Clusters
Carver, PDSF, JGI,KBASE,HEP
14x QDR

Vis & Analytics **Data Transfer Nodes**
Adv. Arch. Testbeds **Science Gateways**

Ethernet &
IB Fabric
Science Friendly Security
Production Monitoring
Power Efficiency
WAN

2 x 10 Gb

1 x 100 Gb

Software Defined Networking

ESnet
Energy Sciences Network

Edison Quick Facts

First Petaflop system with Intel "Ivy Bridge" processors & Cray Aires High Speed Network	
Nodes	5,576 dual-socket with 64 GB memory
Processors	Intel "Ivy Bridge" 12-core, 2.4 GHz
Network	Cray "Aires" Dragonfly Topology
Scratch Disk	7.6PB with >165 GB/sec bandwidth
Peak / Sustained	2.67 PF / 260 TF
Global Network Bandwidth	> 11 TB/sec
Node Memory Bandwidth	90 GB/s

High-Impact Results on Day One

NERSC's users started running production codes immediately on Edison.

Top projects: carbon sequestration, artificial photosynthesis, complex novel materials, cosmic background radiation analysis

Edison is very similar to Hopper, but with 2-5 times the performance per core on most codes.

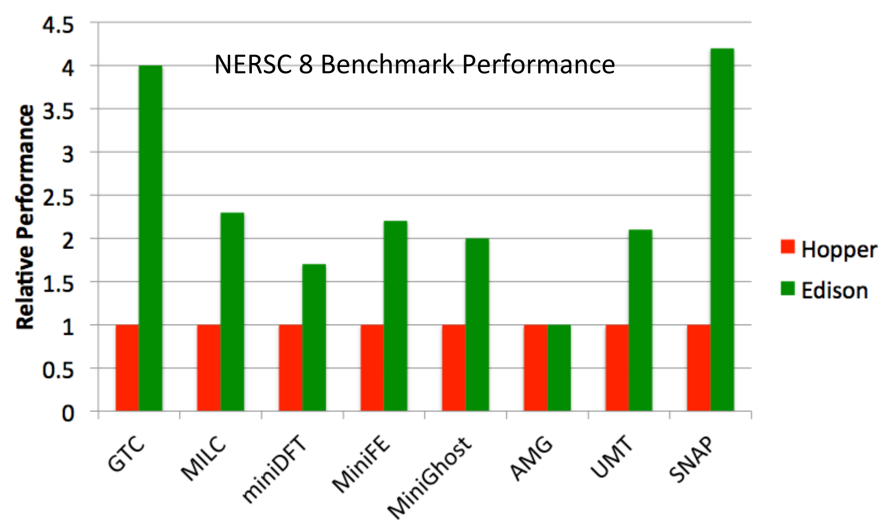

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC Supports Jobs of all Kinds and Sizes

High Throughput: Statistics, Systematics, Analysis, UQ

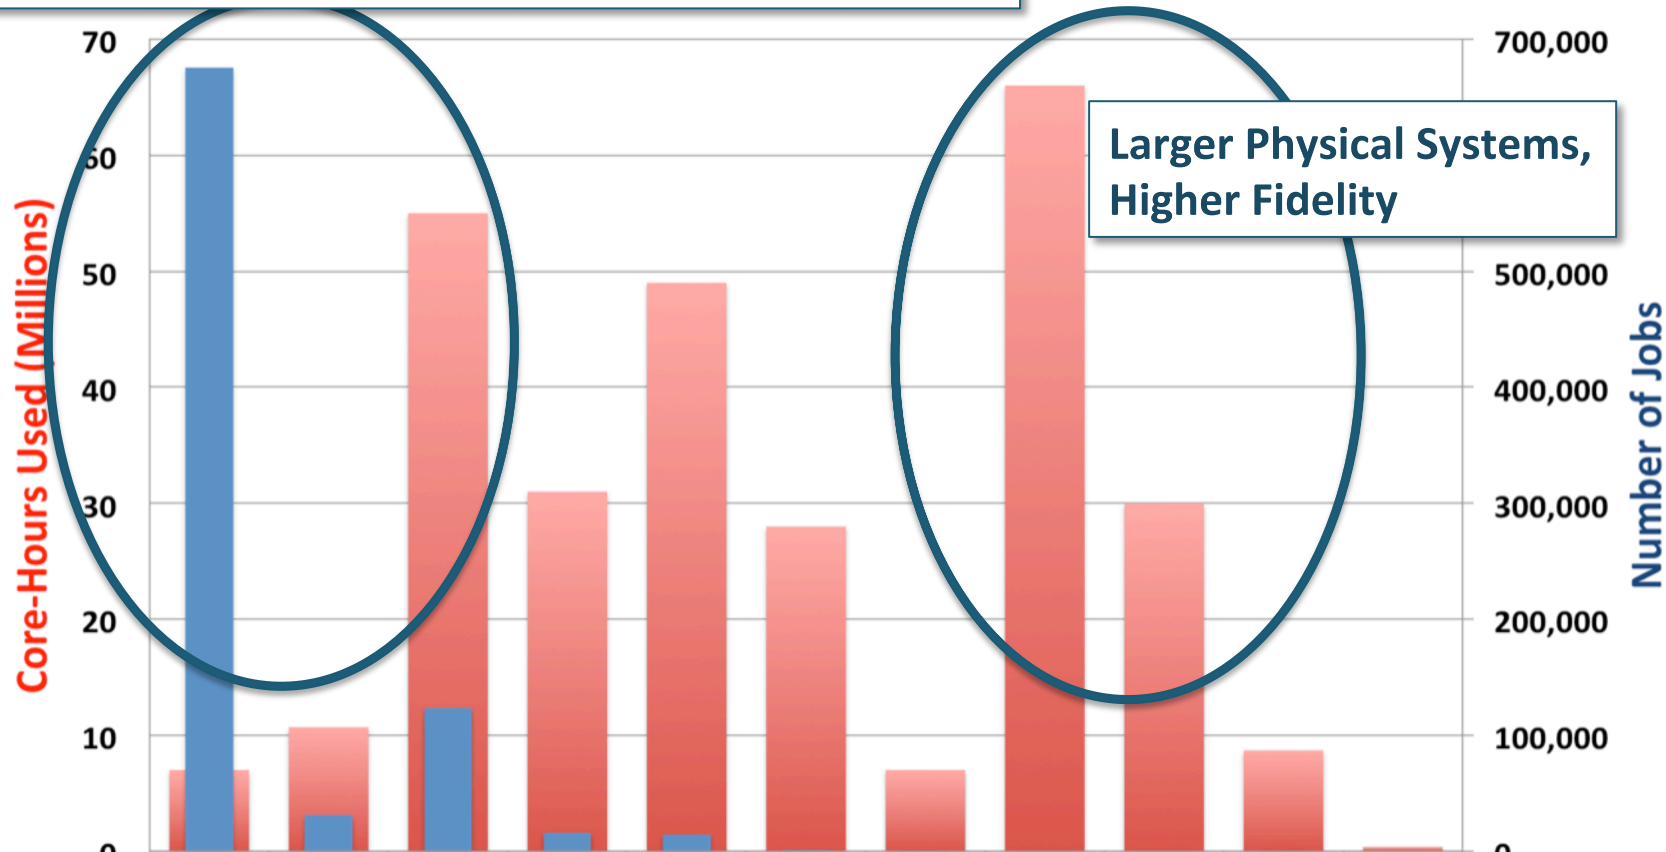

	1-31	32-63	64-511	512-1K	1K-4K	4K-8K	8K-16K	16K-32K	32K-64K	65-128K	128K+
Hours	7	10.7	55	31	49	28	7	66	30	8.7	0.382
Number	675,233	30,991	123,776	15,672	14,618	1,604	826	667	202	56	15

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC-8 Mission Need

The Department of Energy Office of Science requires an HPC system to support the rapidly increasing computational demands of the entire spectrum of DOE SC computational research.

- Provide a significant increase in computational capabilities, at least 10 times the sustained performance of the Hopper system on a set of representative DOE benchmarks
- Delivery in the 2015/2016 time frame
- Provide high bandwidth access to existing data stored by continuing research projects.
- Platform needs to begin to transition users to more energy-efficient many-core architectures.

NERSC Data Resources

- **Global shared filesystems (aka NGF)**
  - Connected to all NERSC computational systems
  - Large, fast, permanent data storage
  - Intended for data sharing within and among projects
  - Many PBs
  - Default quotas ~ 5-10 TB, but often increased
- **Hopper and Edison have dedicated "local" scratch systems**
  - 2 PB & 7.6 PB, respectively
- **Archival storage system**
  - HPSS tape-backed storage
  - Permanent, many 10s of PB
  - No quotas per se, current 240 PB capacity
- **Grid enabled for fast and easy transfers**
- **Dedicated data transfer nodes**
- **Science Data Gateways**

Solving the Puzzle of the Neutrino

- **HPC and ESnet vital in the measurement of the important "θ_{13}" neutrino parameter.**
  - Last and most elusive piece of a longstanding puzzle: why neutrinos appear to vanish as they travel
  - The result affords new understanding of fundamental physics; may eventually help solve the riddle of matter-antimatter asymmetry in the universe.
- **HPC for simulation / analysis; HPSS and data transfer capabilities; NGF and Science Gateways for distributing results**
  - All the raw, simulated, and derived data are analyzed and archived at a single site
  - => Investment in experimental physics requires investment in HPC.
- **One of Science Magazine's Top-Ten Breakthroughs of 2012**

The Daya Bay experiment counts antineutrinos at three detectors (shown in yellow) near the nuclear reactors and calculates how many would reach the detectors if there were no oscillation. transformation.

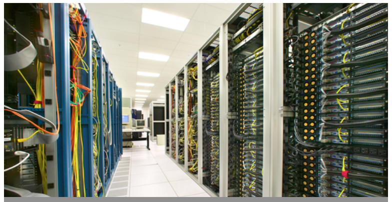

NERSC's PDSF cluster

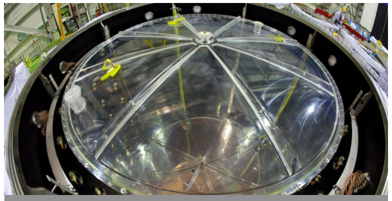

Daya Bay detectors

U.S. DEPARTMENT OF ENERGY | Office of Science

HEP

Science AAAS

PI: Kam-Biu Luk (LBNL)

BERKELEY LAB

The Planck Mission

- **A European Space Agency (+NASA) satellite mission to measure the temperature and polarization of the Cosmic Microwave Background.**
  - The echo of the Big Bang: primordial photons have seen it all.
  - Fluctuations encode all of fundamental physics & cosmology.
  - Planck results assumed by all Dark Energy experiments.

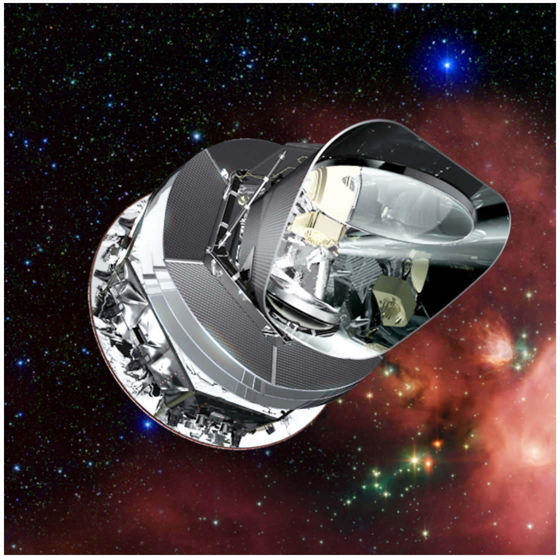

- **Realizing the full scientific potential of Planck requires very significant computing resources**
  - Tiny signal (μK - nK) requires huge data volume for sufficient S/N
  - 72 detectors sampling at 30-180Hz for 2.5 years => 10^{12} samples.
  - Analysis depends critically on Monte Carlo methods
    - Simulate and analyze 10^4 realizations of the entire mission!
- **One of Physics World's Top 10 Breakthroughs of 2013**

CXIDB
Coherent X-ray Imaging Data Bank

Daya Bay
θ13

STAR

NERSC 40 YEARS at the FOREFRONT

DEEP SKY

OpenMSI
Mass spectrometry imaging

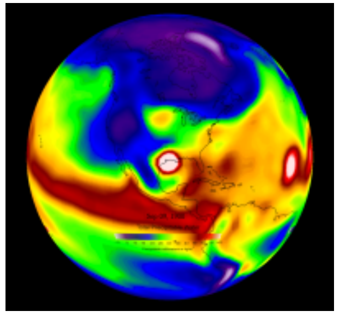

Earth System Grid Federation
ESGF

PALOMAR TRANSIENT FACTORY

esa
PLANCK

MATERIALS PROJECT

JGI
DOE JOINT GENOME INSTITUTE
US DEPARTMENT OF ENERGY
OFFICE OF SCIENCE

IceCube

LUX

KamLAND

CUORE

ALS

ALICE

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Who Uses NERSC? : Compute Hours

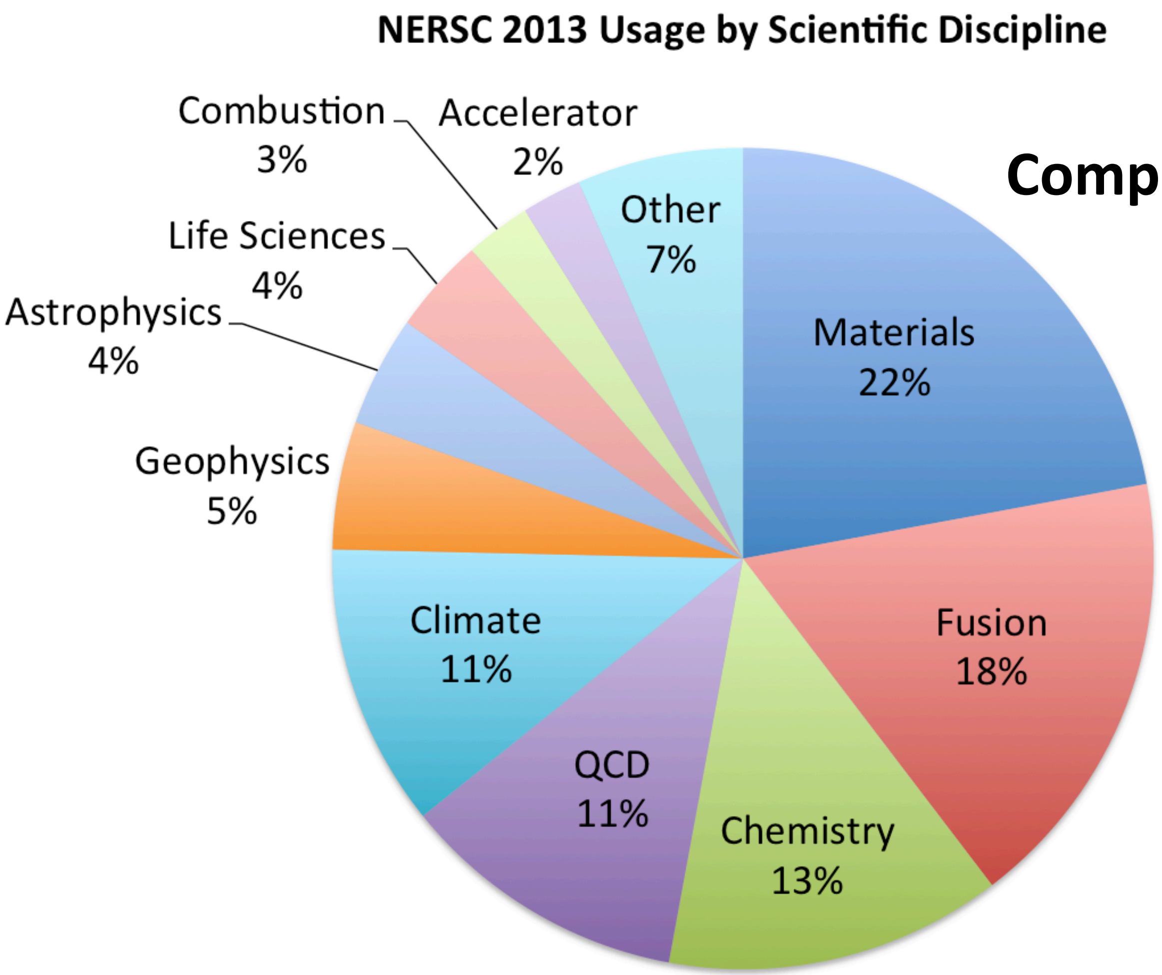

Compute Core Hours

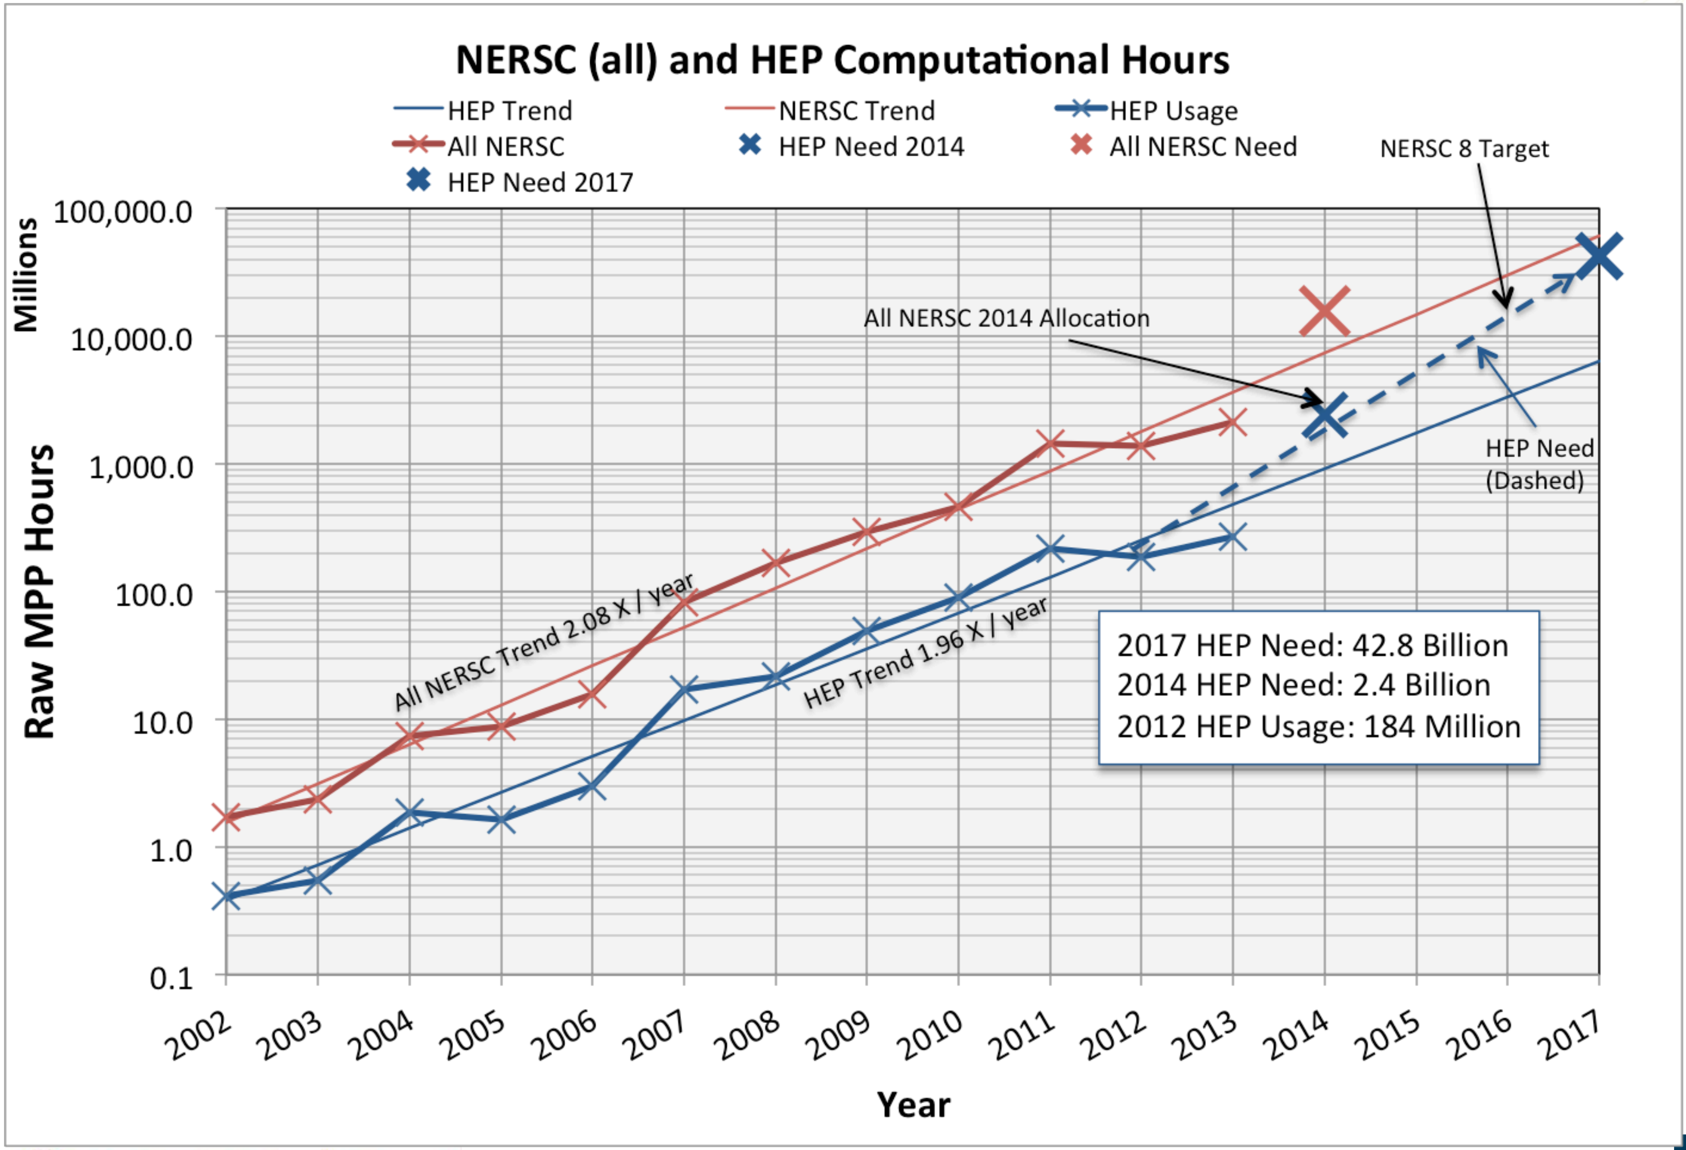
NERSC (all) and HEP Computational Hours
HEP Trend
NERSC Trend
HEP Usage
All NERSC
HEP Need 2014
All NERSC Need
HEP Need 2017
NERSC 8 Target
All NERSC 2014 Allocation
HEP Need (Dashed)
All NERSC Trend 2.08 X / year
HEP Trend 1.96 X / year
2017 HEP Need: 42.8 Billion
2014 HEP Need: 2.4 Billion
2012 HEP Usage: 184 Million
Raw MPP Hours
Millions
100,000.0
10,000.0
1,000.0
100.0
10.0
1.0
0.1
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
Year

Who Uses NERSC? – Archival Storage

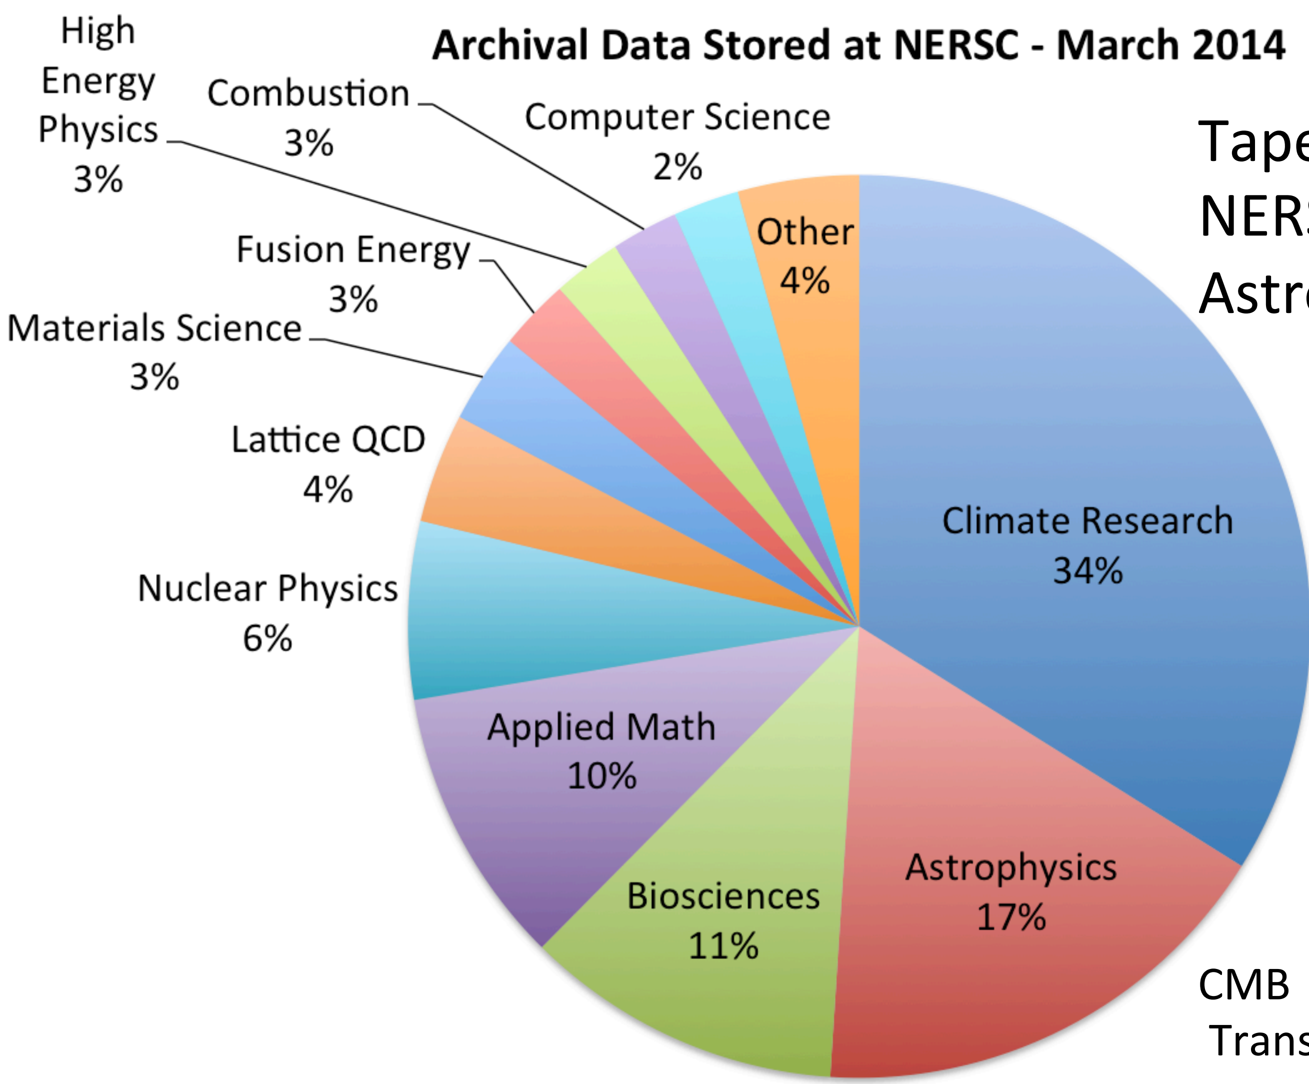

Tape-Backed Storage

NERSC Total: 38 PB

Astro: 6.5 PB

CMB

Transient/SN searches

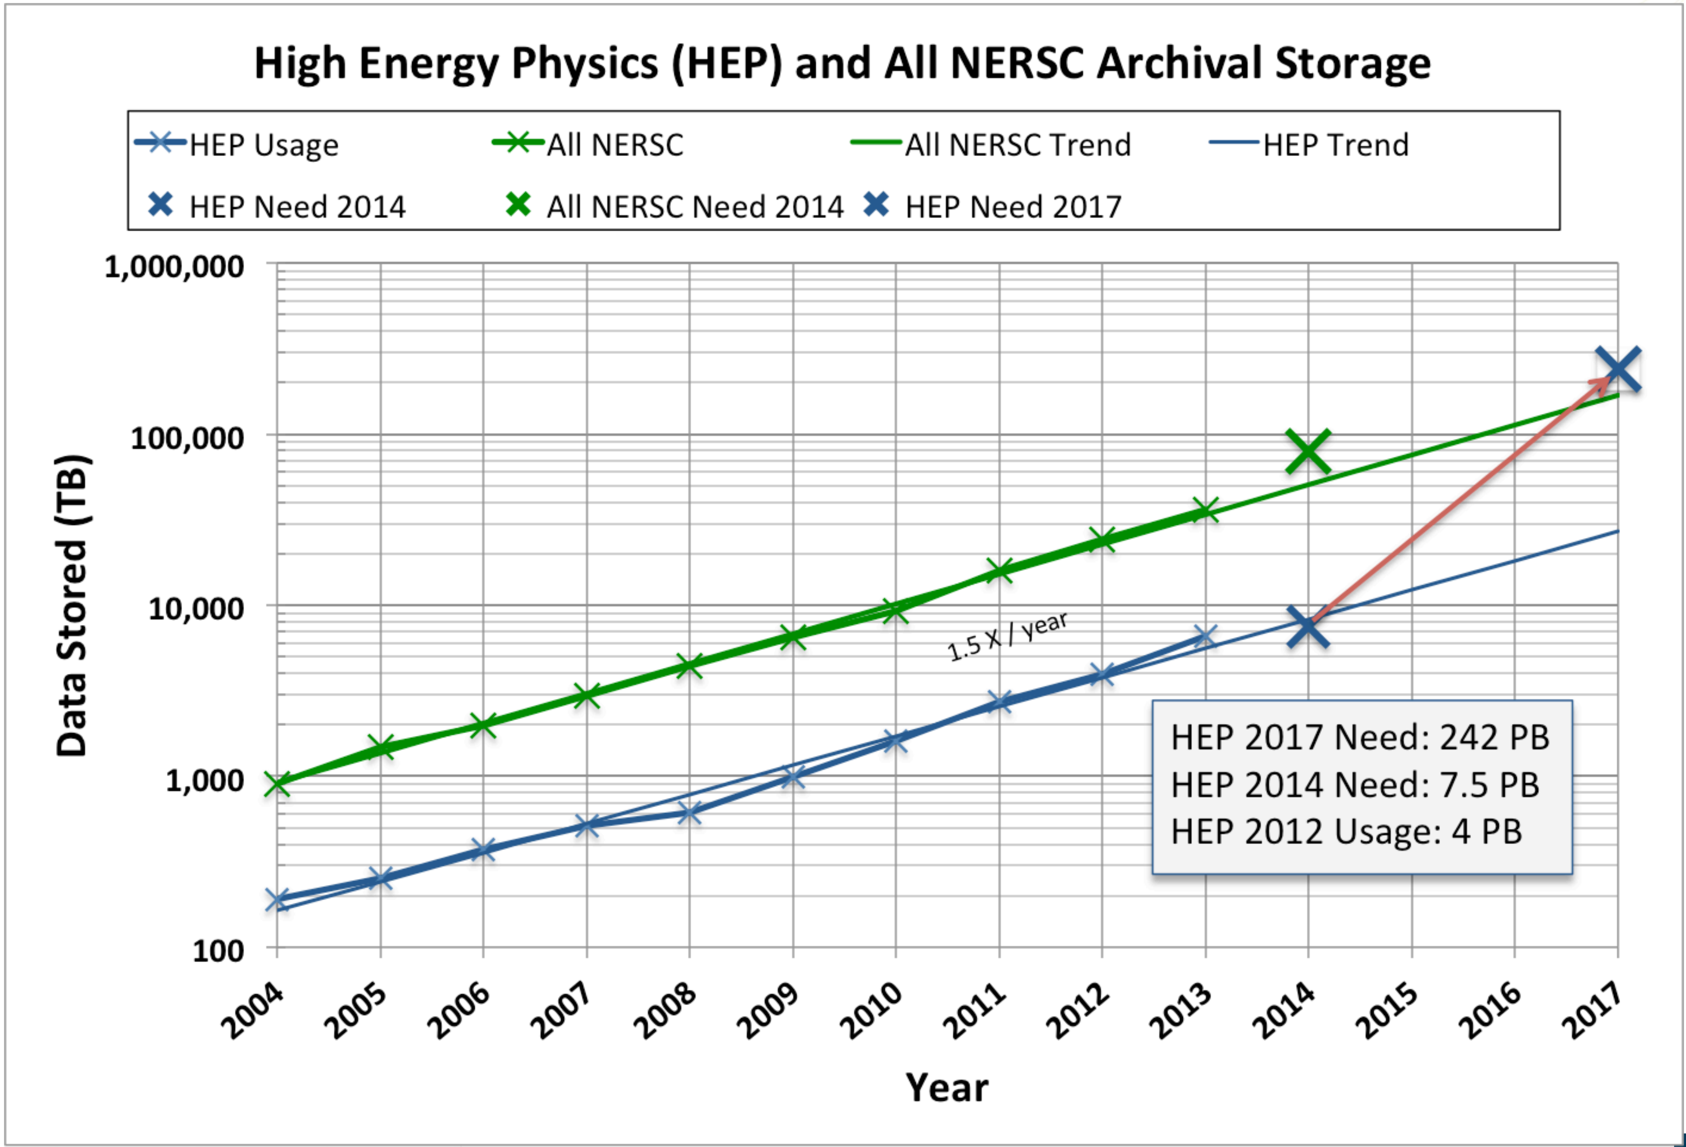
High Energy Physics (HEP) and All NERSC Archival Storage
HEP Usage
All NERSC
All NERSC Trend
HEP Trend
HEP Need 2014
All NERSC Need 2014
HEP Need 2017
Data Stored (TB)
1,000,000
100,000
10,000
1,000
100
1.5 X / year
HEP 2017 Need: 242 PB
HEP 2014 Need: 7.5 PB
HEP 2012 Usage: 4 PB
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
Year

Who Uses NERSC? – Permanent Disk

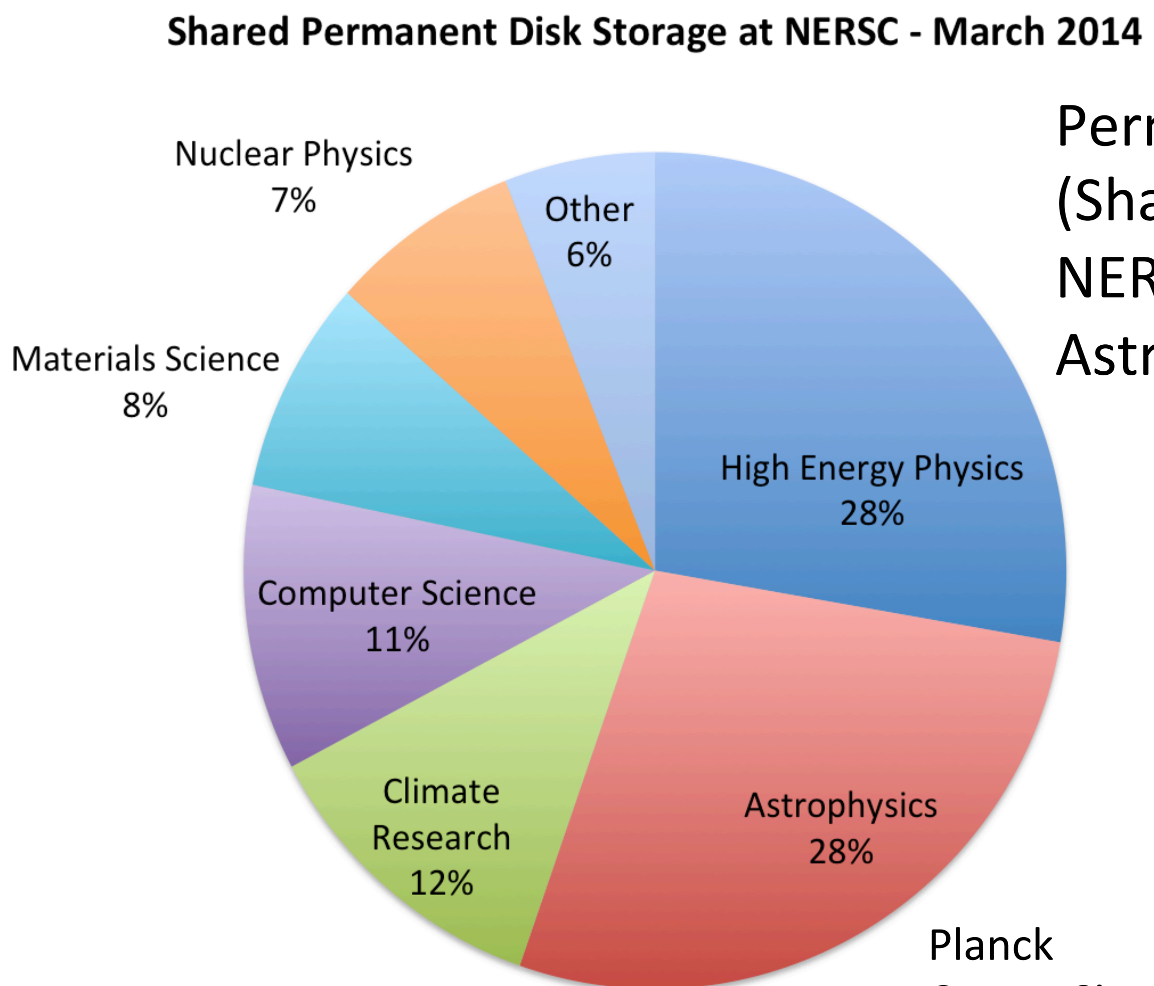

Permanent Spinning Disk (Shared)

NERSC Total: 3.25 PB

Astro: 0.9 PB

Planck
Cosmo Simulation Data

DOE Usage by Office

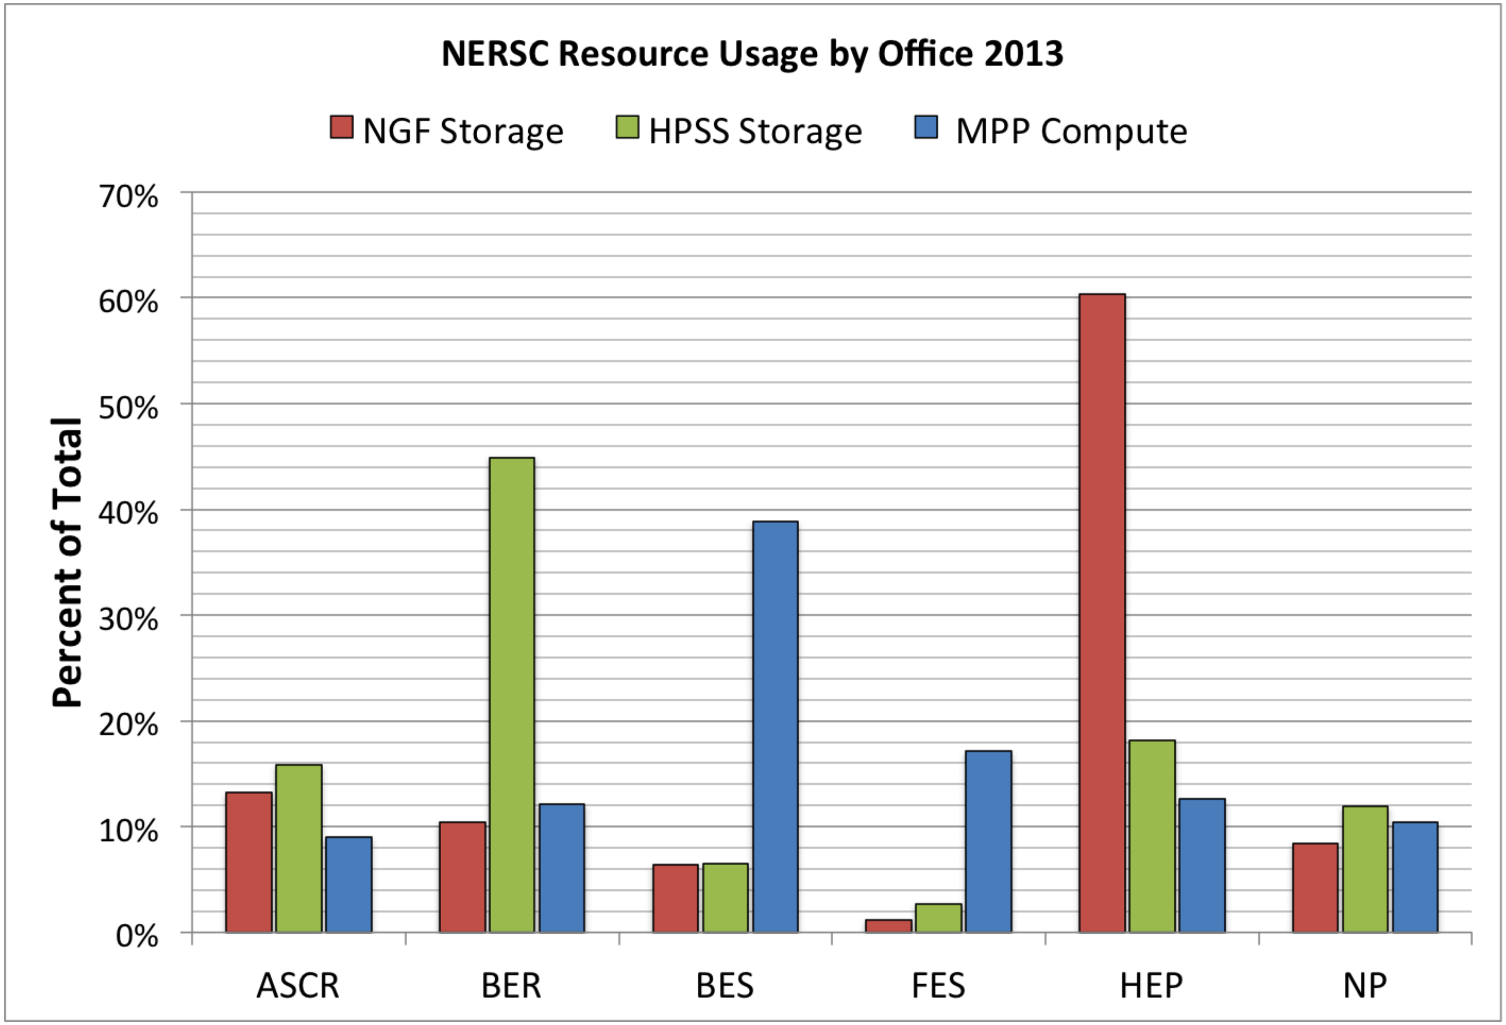

Astronomy & Astrophysics Projects at NERSC

- **55 Projects in 2014**
  - 250 Million hours of compute time allocated
  - 6.5 PB of archival data currently stored
  - 1 PB on permanent spinning disk shared among project members (/project)
- **Science Emphasis**
  - Planck data analysis and synthetic observations/maps
  - Supernova searches & transients
  - Cosmological simulations
  - Supernova simulations
  - Other: Neutrino astrophysics, radio astronomy data analysis, galaxy formation, X-ray bursts, MHD, …

HPC Services at NERSC

User Services

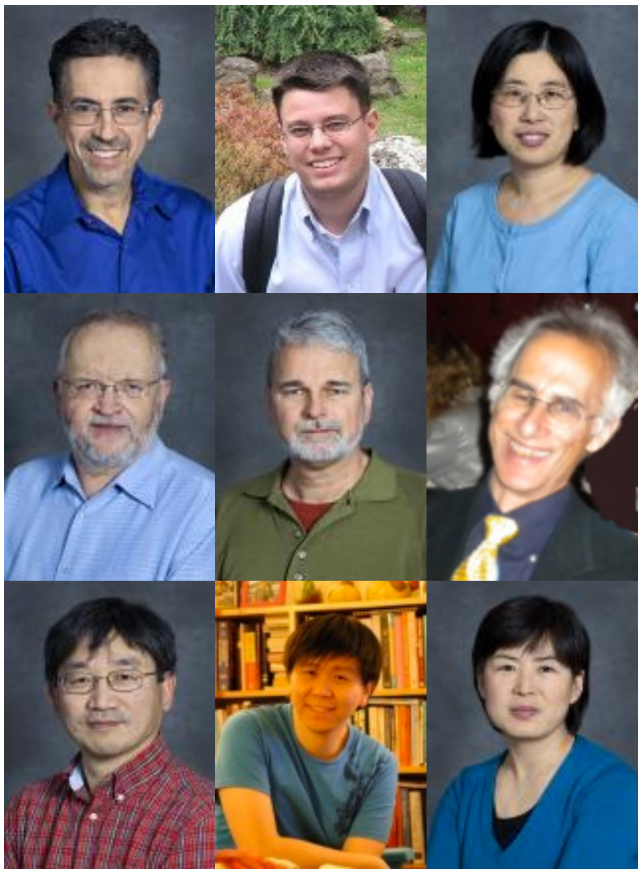

HPC Consultants (1 open position)

JGI Consultants

PDSF Consultant

Postdoc
Program
Coming Soon

Data & Analytics Services

+ Many partial FTEs

Year to Year comparison: JGI overall satisfaction with NERSC

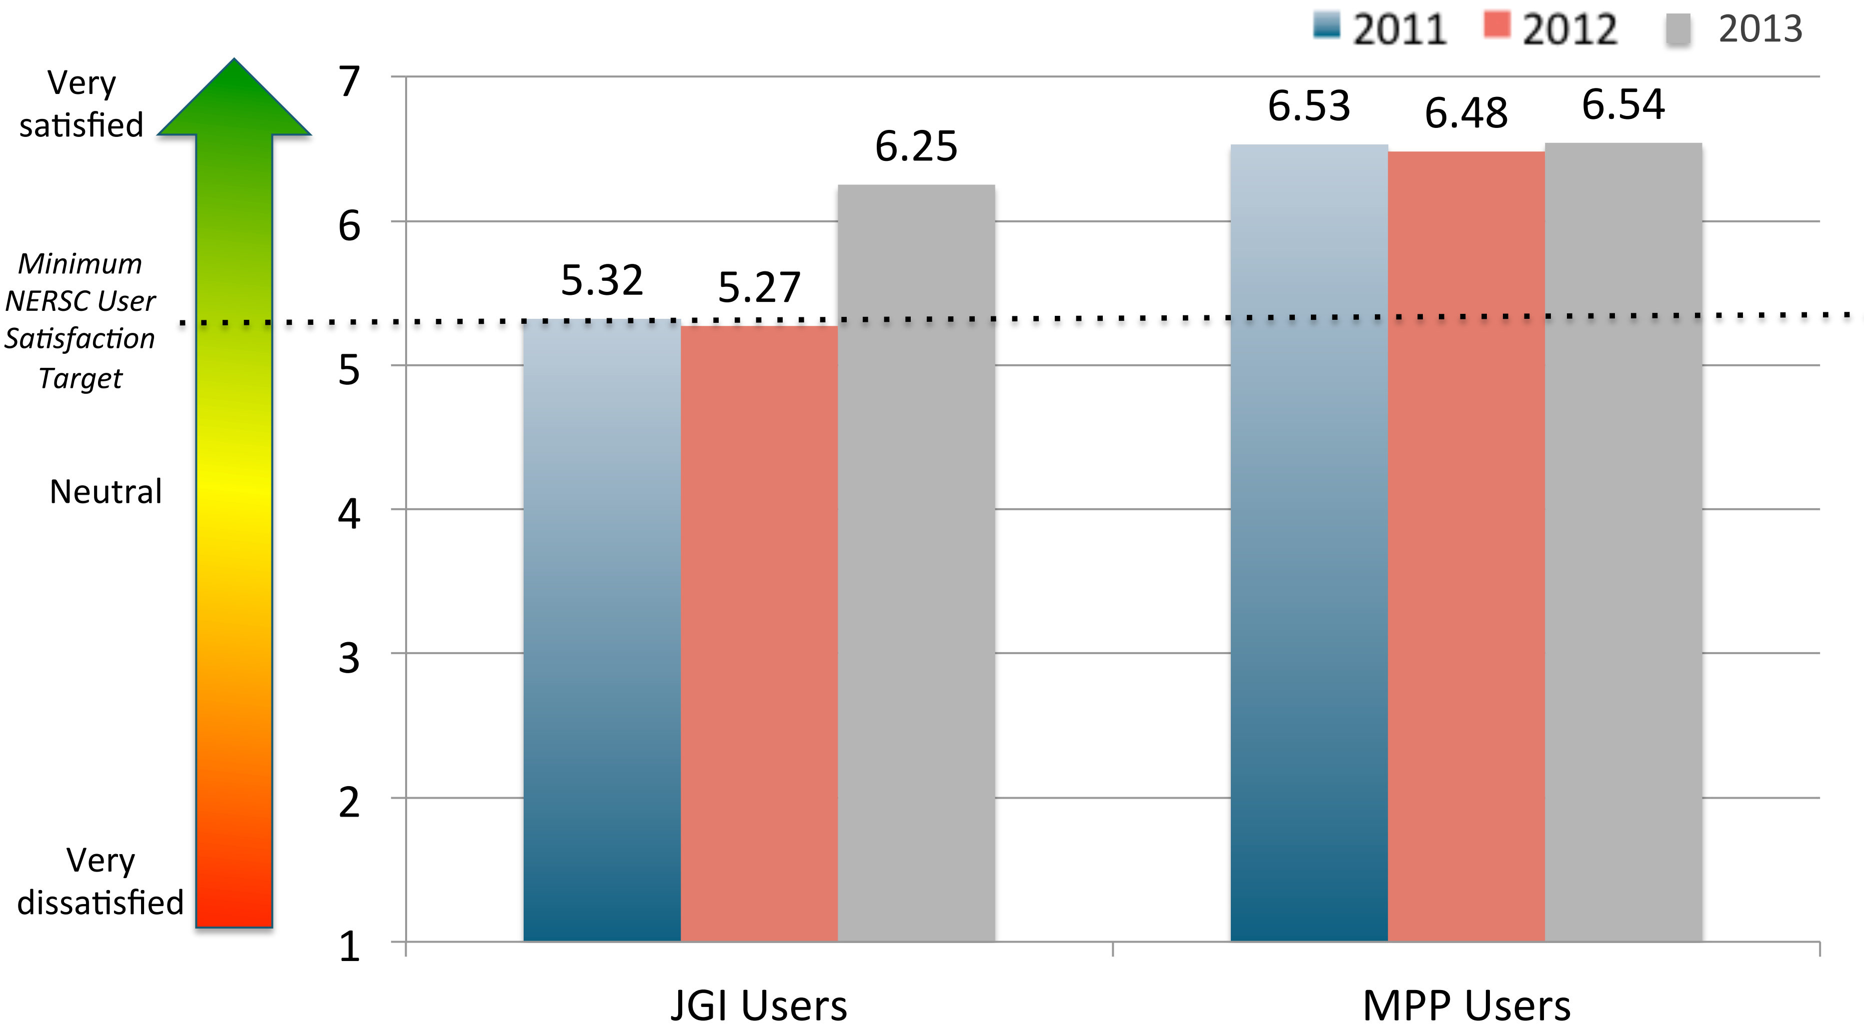

How to Get Access to NERSC Resources

- **“ERCAP” allocations process**
  - 80% of compute hours allocated by DOE program managers to projects doing research within the DOE mission
  - 10% allocated through ALCC (high-risk, high-payoff)
  - Archival storage (tape) also allocated
  - Project funding from DOE not required; at discretion of program managers
- **NERSC Director’s Reserve for strategic projects**
  - 10% of computer time (250 M hours)
  - NISE and Data Initiative exploratory programs
- **Startup Projects**
  - At NERSC’s discretion
  - Up to 50 K hours for 18 months

How to Get Access to NERSC II

- **Buy-in model for hardware and support**
  - PDSF cluster: Nuclear and High Energy Physics
  - Genepool cluster & file systems: Joint Genome Institute
- **A La Carte resources run by NERSC**
  - Planck bought a rack of a compute cluster
  - Fixed cost for 5 years of shared spinning disk (coming soon)
  - The Materials Project has dedicated nodes
  - Science Gateways

High Performance Computing is …

… the application of supercomputers, data systems, networking, and advanced algorithms & workflows to scientific problems that are either too large for standard computers or would take too long on them.

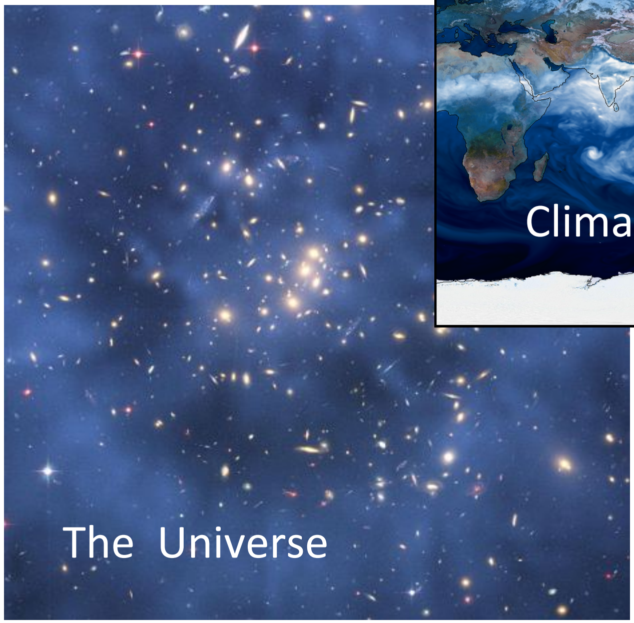

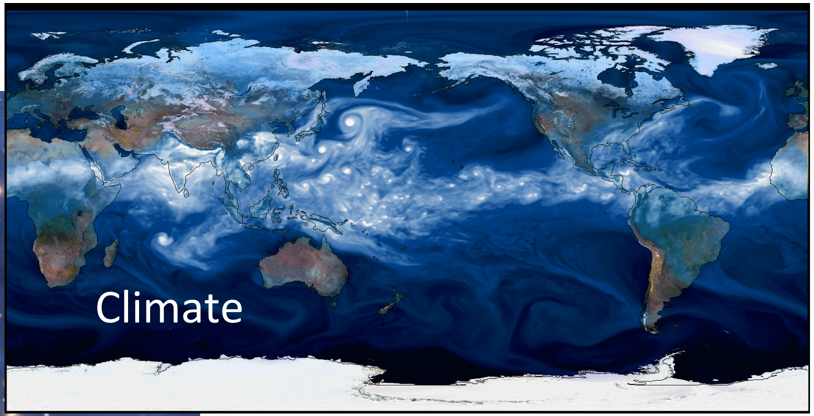

Understanding How Proteins Work

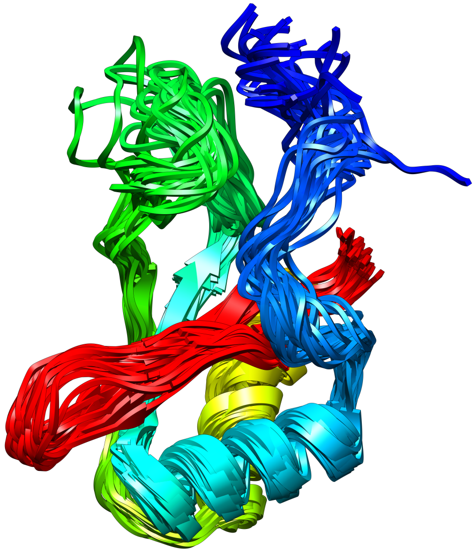

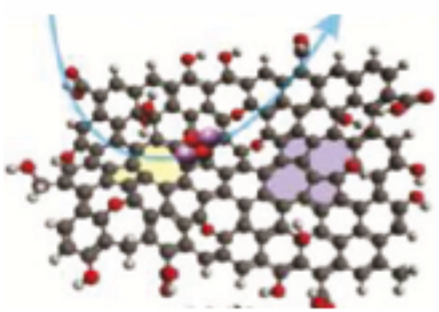

Designing Better Batteries

HPC is a Tool for Discovery

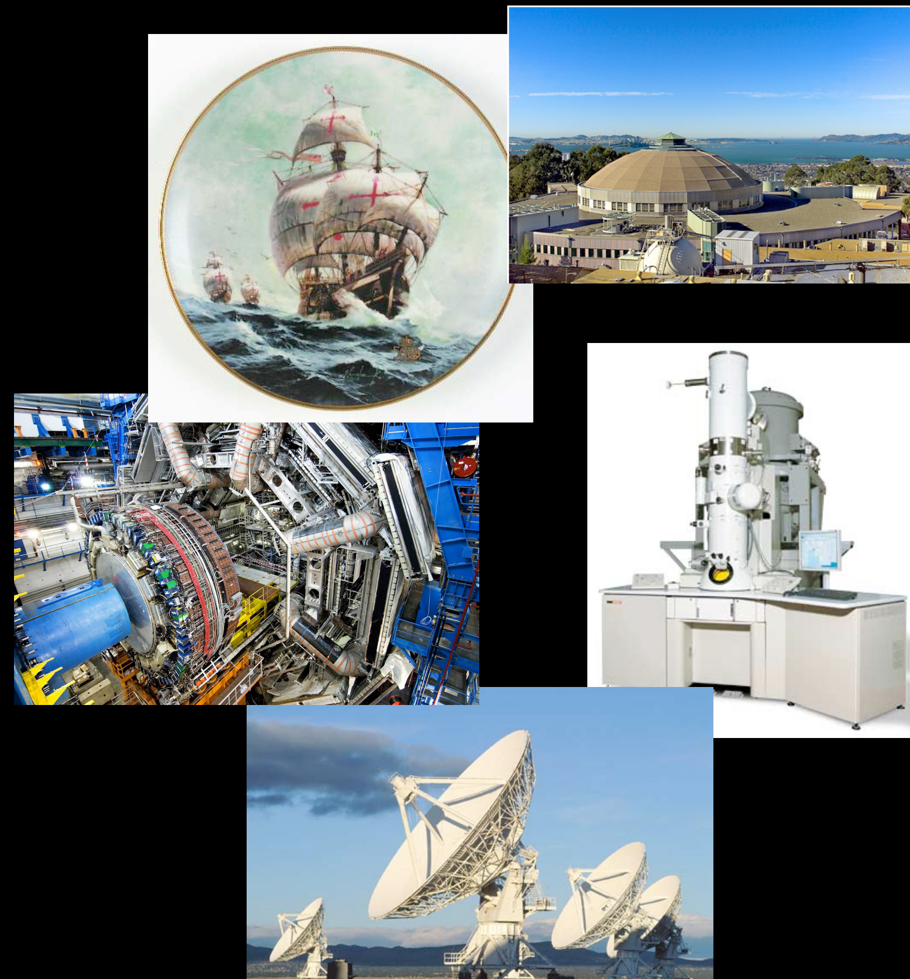

Additional Slides

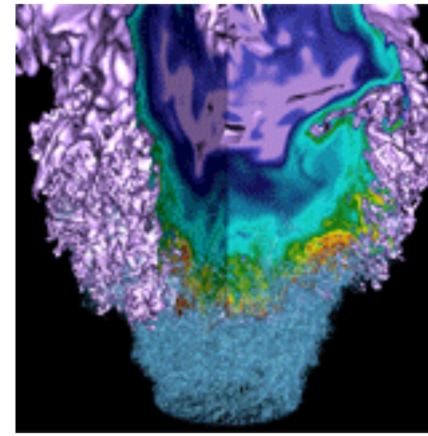

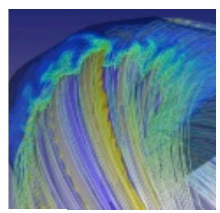

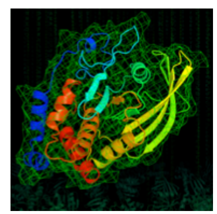

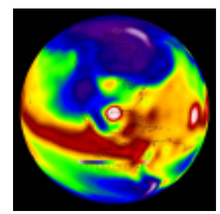

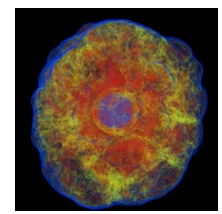

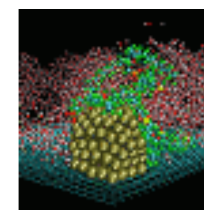

Requirements Reviews

1½-day reviews with each Program Office

Computing and storage requirements for next 5 years

- Participants
  - DOE ADs & Program Managers
  - Leading NERSC users & key potential users
  - NERSC staff & CS Experts

Scientific Objectives

Computing, Storage, Software, Services Requirements

Reports From 8 Requirements Reviews Have Been Published

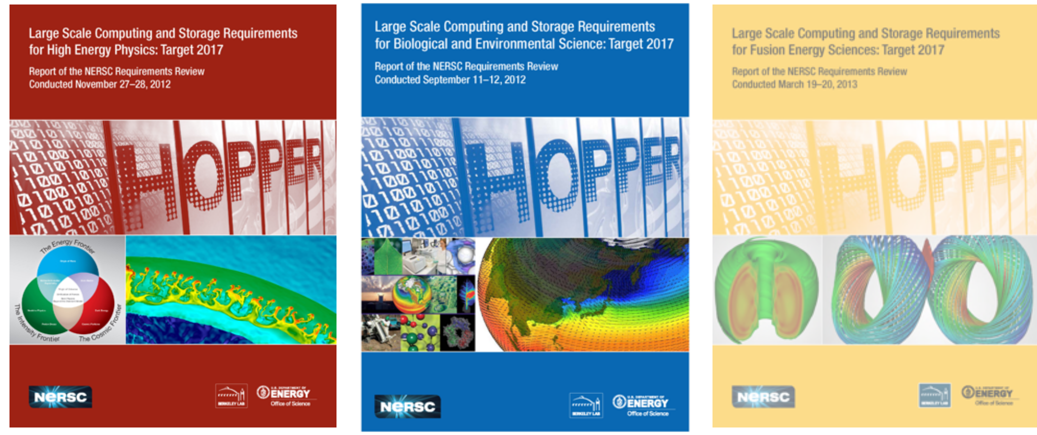

- Computing and storage requirements for 2014 & 2017
- Executive Summary of requirements
- Case studies
- Second round, for 2017 requirements, will be completed in April 2014 (NP)

http://www.nersc.gov/science/hpc-requirements-reviews/reports/

Impact

- **Highly regarded within DOE**
- **Scientific justification for ASCR budget requests**
  - Quantitative requirements
  - Documented science goals & needs from science teams
- **Basis for NERSC 7 and NERSC 8 Mission Need documents**
- **Influence on NERSC services**
  - e.g. application readiness, support for high-throughput computing, planning for NERSC data services
- **Derivative publications and reports:**
  - HEP community's "Snowmass" Report
  - DOE ASCR white paper on data needs

HEP Executive Summary

- More computing and data resources needed
- Vastly improved I/O capabilities and better facilities for data-intensive science
- Need to support both large-scale and ensemble runs
- Assistance needed to transition to next-generation processors
- There are communities within DOE HEP that are not traditional users of large HPC centers, yet have a profound need for additional computing, storage, and analysis facilities. (LHC, sky surveys, …)

HEP Cosmic Frontier Requirements - MPP

Project	Repos	2012 Usage	2017 Need	Factor Increase
Experimental Cosmology	LSST, Boss, BigBoss, DES, DESSN, PTF, DESI, COSMO	2 M	82 M	41
Cosmological Simulations for Sky Surveys	Cosmosim, cusp, hacc	24 M	10,000 M	417
CMB Analysis	Planck, usplanck, mp107	13 M	500 M	38
Supernova Studies	m1400	13 M	200 M	15

32X is “Normal”

HEP Cosmic Frontier Requirements - HPSS

Project	Repos	2012 Usage	2017 Need	Factor Increase
Experimental Cosmology	LSST, Boss, BigBoss, DES, DESSN, PTF, DESI, COSMO	40 TB	1,000 TB	25
Cosmological Simulations for Sky Surveys	Cosmosim, cusp, hacc	70 TB	10,000 TB	143
CMB Analysis	Planck, usplanck, mp107	550 TB	50,000 TB	91
Supernova Studies	m1400	100 TB	2,000 TB	20

7.6 X is "Normal"

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

HEP Cosmic Frontier Requirements - /project

Project	Repos	2012 Usage	2017 Need	Factor Increase
Experimental Cosmology	LSST, Boss, BigBoss, DES, DESSN, PTF, DESI, COSMO	20TB	500 TB	25
Cosmological Simulations for Sky Surveys	Cosmosim, cusp, hacc	120 TB	10,000 TB	83
CMB Analysis	Planck, usplanck, mp107	200TB	5,000 TB	25
Supernova Studies	m1400	3 TB	200 TB	67

7.6 X is "Normal"

Extreme Data Strategy

- Develop and deploy new data resources and capabilities
- Partner with DOE experimental facilities and projects to identify requirements and create early success
- Provide expertise and services for extreme data
- Leverage ESnet and ASCR research to create end-to-end solutions

Thank you.