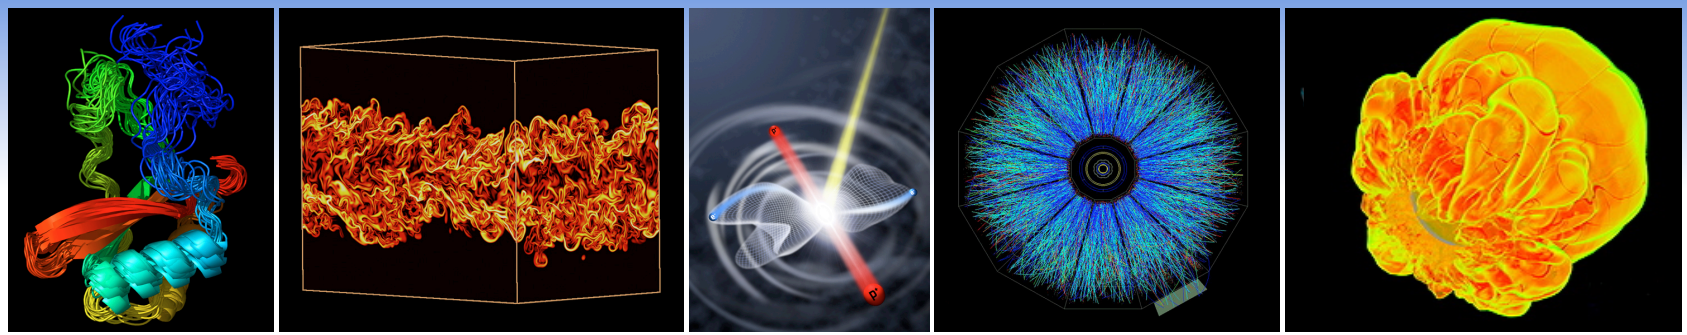

HEP/NERSC/ASCR Requirements Workshop
Large Scale Computing and Storage Requirements for High Energy Physics

NERSC
Lawrence Berkeley National Laboratory

Workshop Held Nov 12-13, 2009, Rockville, MD

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC

BERKELEY LAB

NERSC Requirements Workshops

- **Goal**: Ensure that NERSC continues to provide the world-class facilities and services needed to support DOE Office of Science Research
- **Method**: Hold workshops to derive and document each DOE SC Office's HPC requirements for NERSC in 2013-14
- **Deliverable**: Reports that includes both the HPC requirements and supporting narratives, illustrated by specific science-based case studies
- **Use**: Guide NERSC procurements and service offerings; help NERSC, ASCR, Program Offices advocate for the HPC resources needed to support your science

Workshop Schedule

BER
May 7-8, 2009
Report Published

FES
Aug. 3-4, 2010
Report in progress

HEP
Nov. 12-13, 2009
Report Published

ASCR
Dec. 2010 or Jan 2011

BES
Feb. 9-10, 2010
Report in progress

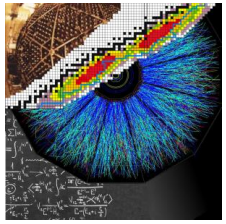

NP
2011

HEP Workshop

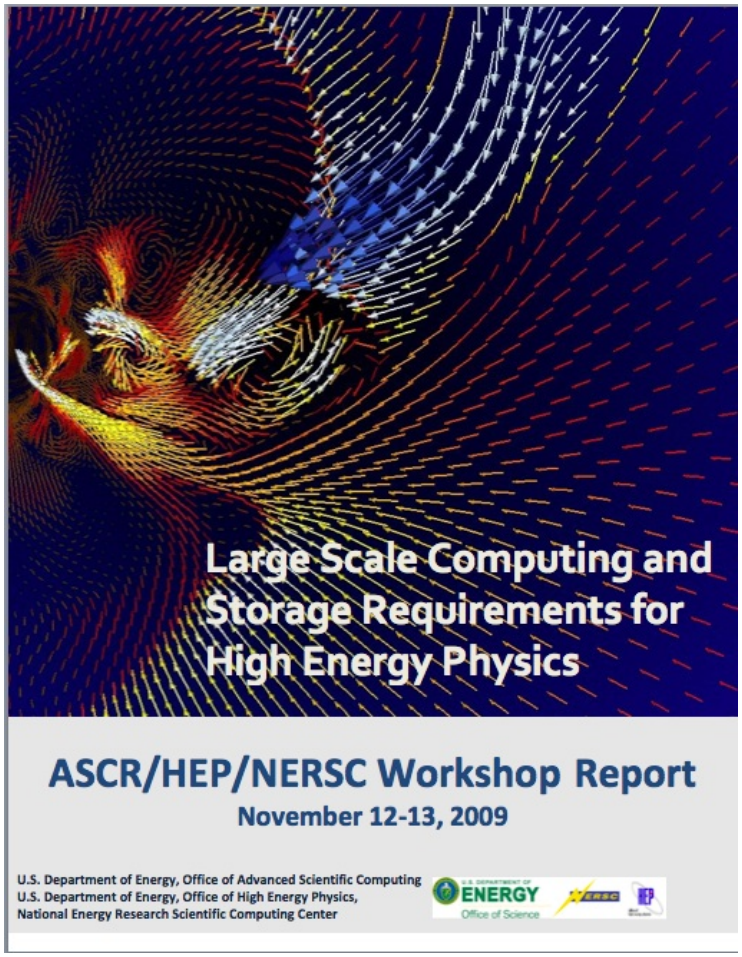

Nov. 12-13, 2009

Washington D.C. (Rockville, MD)

Report approved by DOE

Workshop Participants

DOE Office of Science

Yukiko Sekine, ASCR
Amber Boehnlein, HEP
Alan Stone, HEP

NERSC / LBNL

Kathy Yelick, NERSC Director
J. Shalf, Advanced Technologies
J. Bell, Computational Research
H. Wasserman, User Services
R. Gerber, User Services

Workshop Participants

HEP Scientists

Accelerator Physics & Design

David Bruhwiler, Tech-X
Cameron Geddes, LBNL
Rich Lee, SLAC
Panagiotis Spentzouris, FNAL
Chengkun Huang, UCLA

Theoretical Astrophysics

Mike Norman, UCSD
Stan Woosley, UCSC

Astrophysical Data Analysis

Julian Borrill, LBNL
Peter Nugent, LBNL
Alex Szalay, JHU

Lattice QCD

Paul McKenzie, FNAL
Doug Toussaint, U Arizona

Data Analysis & Detector Design

Craig Tull, LBNL

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Case Studies

Accelerator Physics

Community Petascale Project for Accelerator Science and Simulation
Advanced Modeling for Particle Accelerators
Electromagnetic Modeling of Accelerator Structures
Simulation of Laser Plasma Wakefield Particle Accelerators
Petascale Particle-in-Cell Simulations of Plasma Based Accelerators

Astrophysics Modeling and Simulation

The Cosmic Frontier – Structure Formation
Type Ia Supernovae
Core-Collapse Supernovae

Astrophysical Data Analysis

Palomar Transient Factory & La Silla Supernova Search
Cosmic Microwave Background Data Analysis for the Planck Satellite

Lattice QCD

MIMD Lattice Computation (MILC) Collaboration

HEP Data Analysis and Detector Simulation

85% of NERSC HEP Usage

HEP Workshop Requirements

1. **Resources**

   Science teams need access to a significant increase in computational resources to meet their research goals.

2. **Data**

   Science teams need to be able to read, write, transfer, store online, archive, analyze, and share huge volumes of data.

HEP Workshop Requirements

3. **Emerging Architectures**

   Science teams need guidance and support to implement their codes on future architectures.

4. **Throughput**

   Projects need predictable, rapid turnaround of their computational jobs to meet mission-critical time constraints

BER & HEP Requirements Comparison

HEP and BER basic requirements are the same.

BER	HEP
Need more allocation	Need more allocation
Data storage and I/O	Data storage and I/O
Programming future architectures	Programming future architectures
Job turnaround, throughput	Job turnaround, throughput

HEP Workshop Findings: Expanded

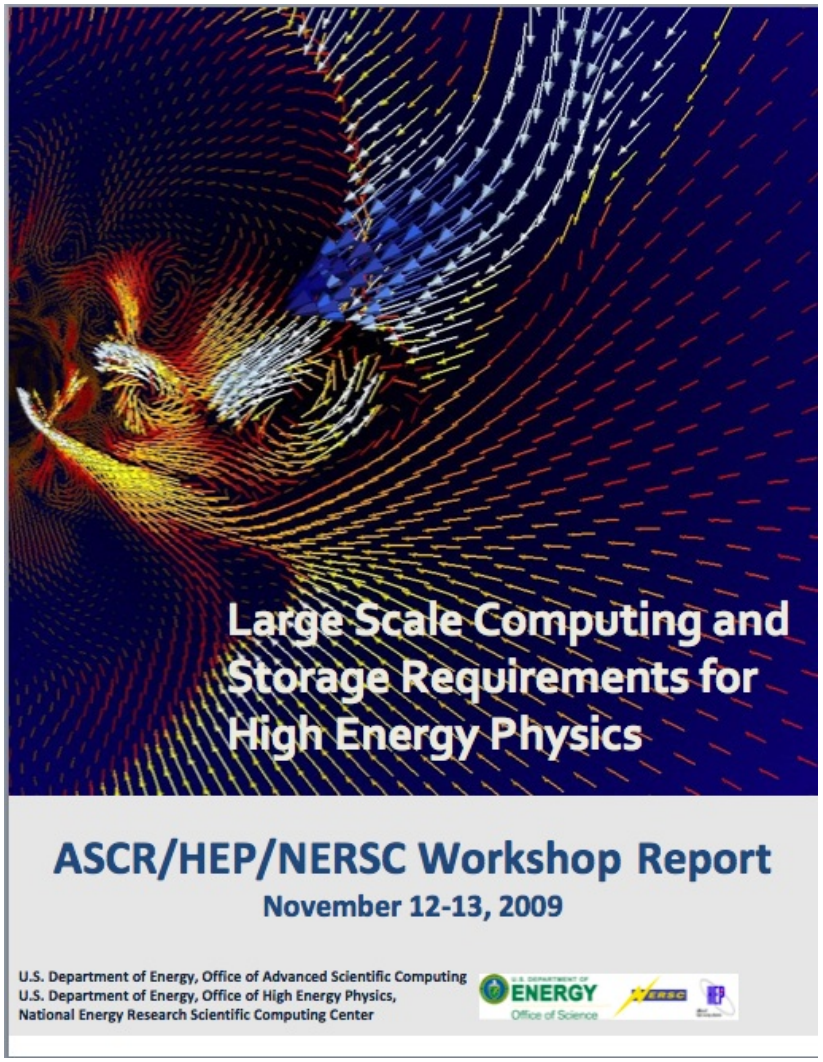

1. Allocations Needs 2009-13

- Need more now: current allocations are limiting
- 35-fold increase in hours
- Adequate time for
  - production
  - development
  - validation and verification

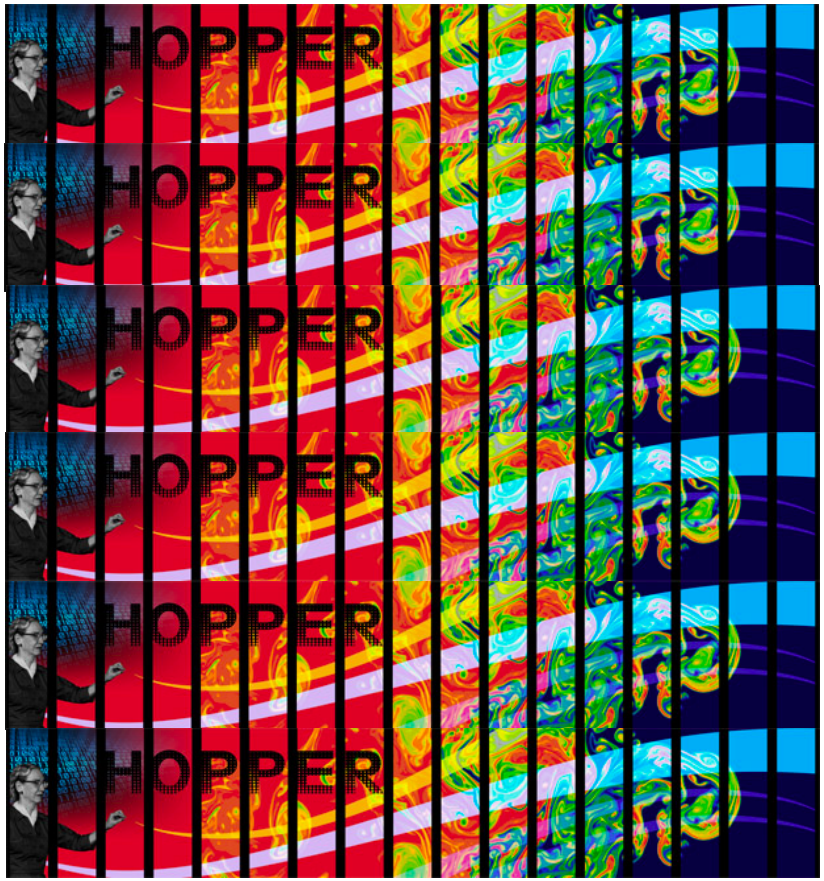

2. Data Needs 2009-13

- 10X increase in online storage
- I/O libs like HDF5 and netCDF
- Data analysis and vis tools
- Tools to manage PBs
- Data sharing
- 128 GB memory for vis apps
- Parallel databases

3. Future Architecture Needs

- Help choosing programming models
- Testbed machines as prototypes of possible future NERSC systems
- Development of, and support for, existing libraries
- Support for UPC and Co-Array Fortran
- Full-featured OS for some groups

4. Throughput Needs

- Stable, available, reliable systems. Dealing with system failures is expensive.
- Good throughput for many ensemble runs at modest concurrency
- Ready access and turnaround to support code development
- Real-time data processing on a fixed schedule
- Frameworks for sophisticated job management

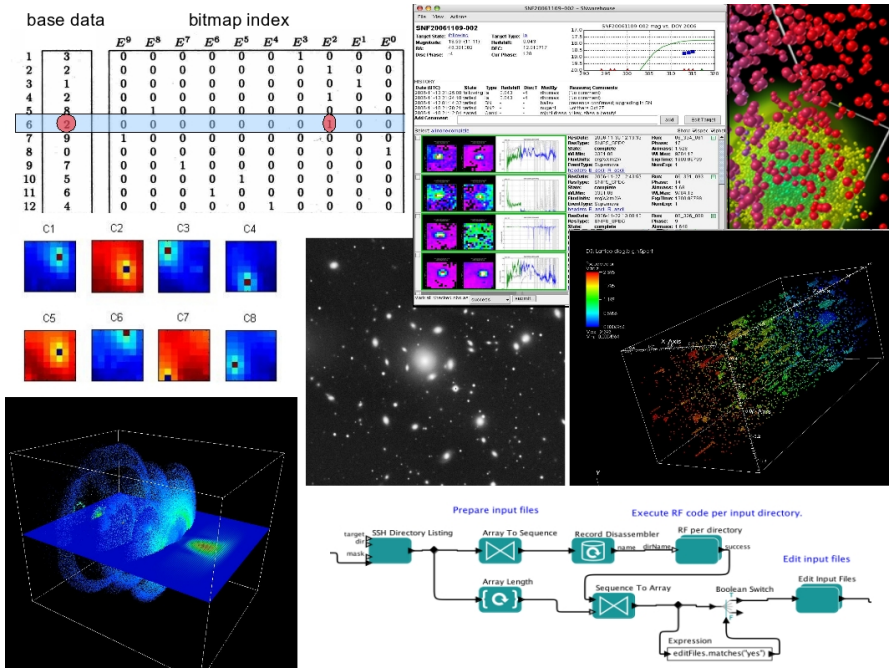

HEP Workshop Findings

1. Resources
2. Data
3. Emerging Architectures
4. Throughput

Questions

- How easy is it for you to translate what you need to meet your scientific objectives into quantitative computing and storage requirements?
- How can we help you do so?