Workshop Goals & Process

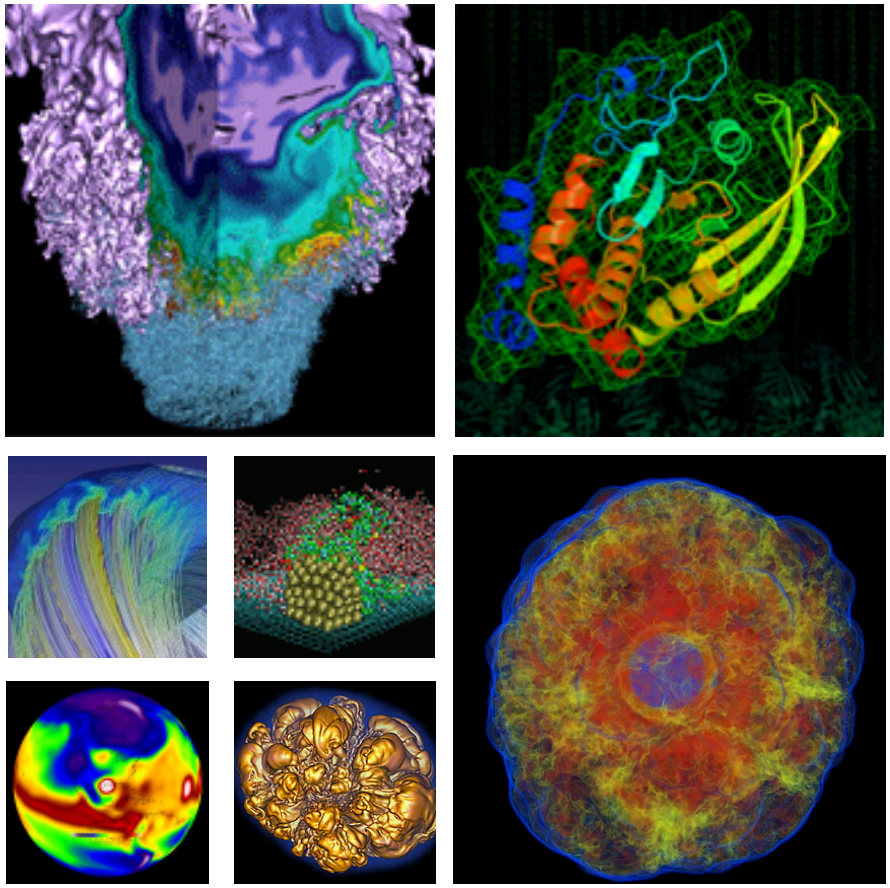

Richard Gerber
NERSC User Services

March 19, 2013

Overview

- **We're holding this review to ensure that**
  - you have the HPC resources you need to be successful in your research
  - NERSC can fulfill its mission to accelerate scientific discovery within the Office of Science
- **Your input helps NERSC**
  - create science-based justification for acquiring needed resources
  - focus on delivering the services that are important to you
  - make technology decisions
- **Result: NERSC can better provide what you need for your work**
- **This exercise benefits the Office of Science, FES, ASCR, NERSC, & you**

Process

- **Collect and refine requirements for 2017**
  - Case study worksheets
  - Discussions at this meeting
  - Post-meeting refinement of case studies
- **NERSC editors (Richard & Harvey)**
  - Check case studies for internal consistency and compare against historical trends
  - Aggregate requirements and summarize
  - Create draft report for you & FES to review
- **Send final draft to DOE FES office for final approval**
- **Publish final report**

Key Strategy

- **Key is to tie computational, storage, and services needs to achievement of scientific goals – as specifically as possible.**
  - Science -> codes & algorithms -> computation parameters -> resources needed

Quantitative Method

- **Quantitative requirements are very important**
  - Hours needed
  - Archival data storage needed
  - Disk storage needed
- **For hours and archival storage**
  - Requirements from this review are summed
  - Scaled to full FES need by the fraction of 2012 FES usage represented by case studies
  - Important: Associate each case study with 2012 NERSC repo or repos
  - New projects' requirements added in separately
- **Like to do the same for Scratch and Permanent Disk**
  - Please state 2012 usage and 2017 need so we can create a ratio

Hours Required

- **The unit of "Hour" is defined as 1 Hopper core hour.**
- **Please state your requirements in these units**
  - How much computing will you need in multiples of a Hopper hour?
  - For this exercise, ignore the architecture – we will normalize this when future systems arrive, based on average application performance
- **Give your best estimate for 2017 specifically**
  - Remember that each year's usage has historically been 2X the previous year's

Computational Hours

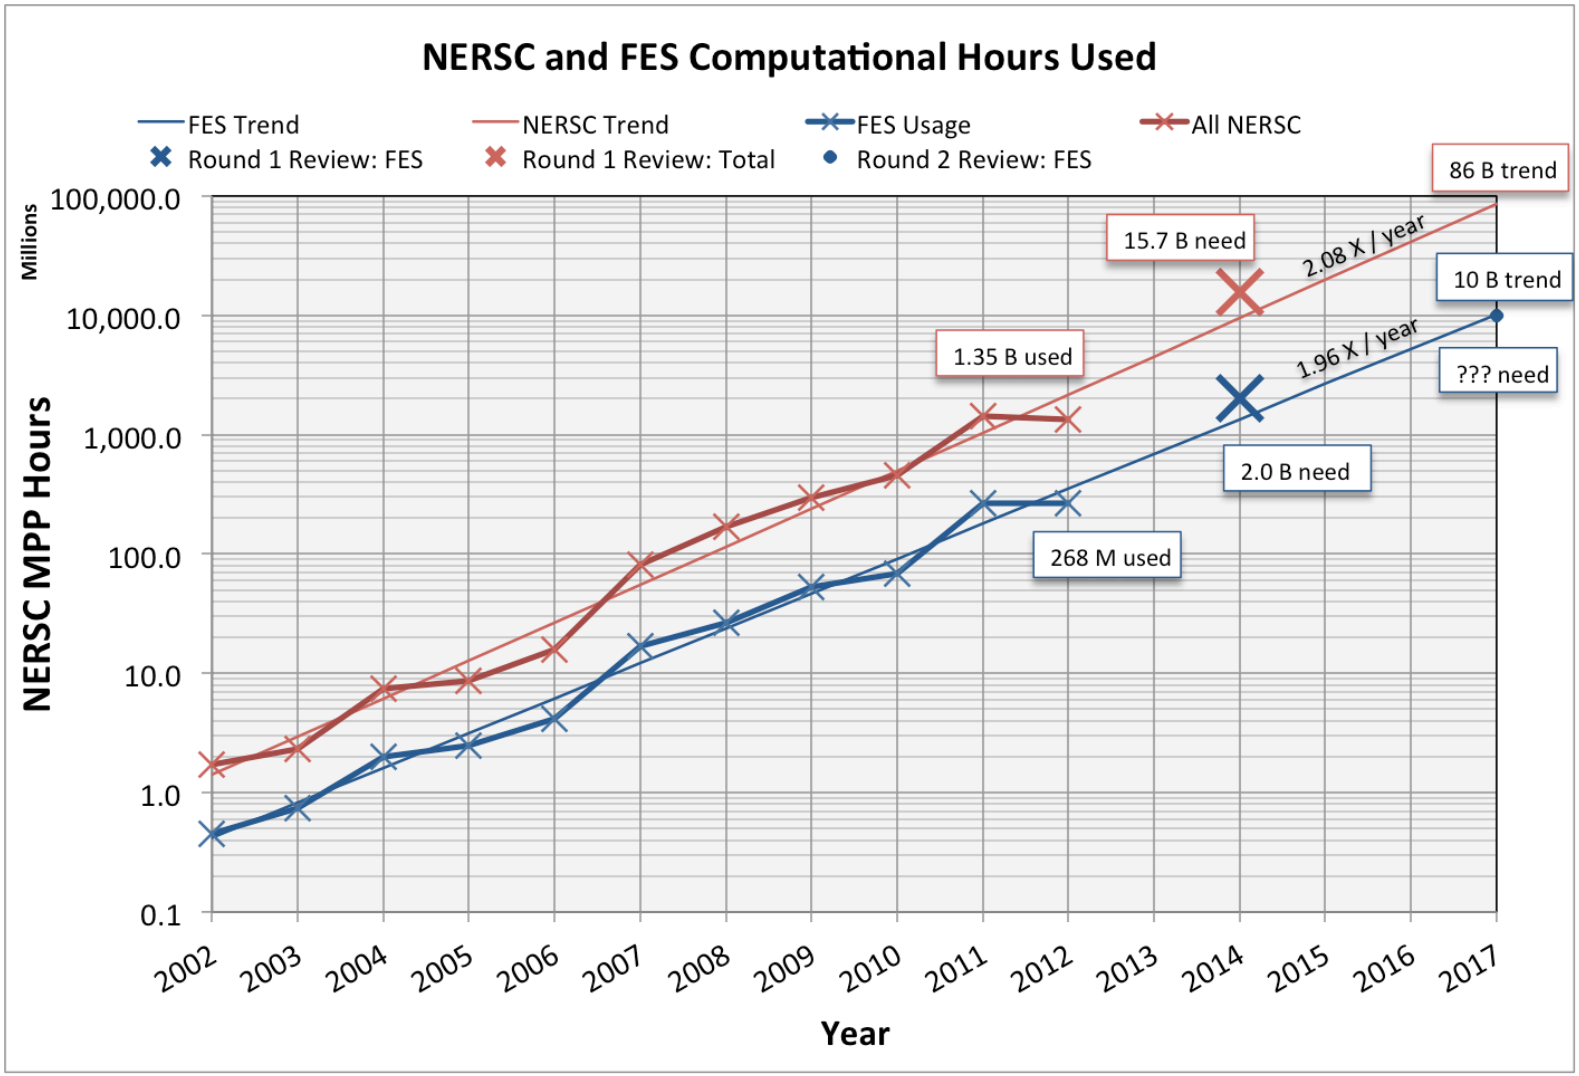

Data Storage Requirements

- **Archival storage estimate for 2017**
  - This is an aggregate number: Σ all years
  - Historical trend: 1.5-1.7 X / year
- **Scratch (temporary)**
  - What is the maximum you will need at any given time during 2017?
  - Not just what you will need for a single run
- **Permanent disk space**
  - What will you need for source code, data files or executables that will be constantly accessed and/or shared, etc.

Archival Storage

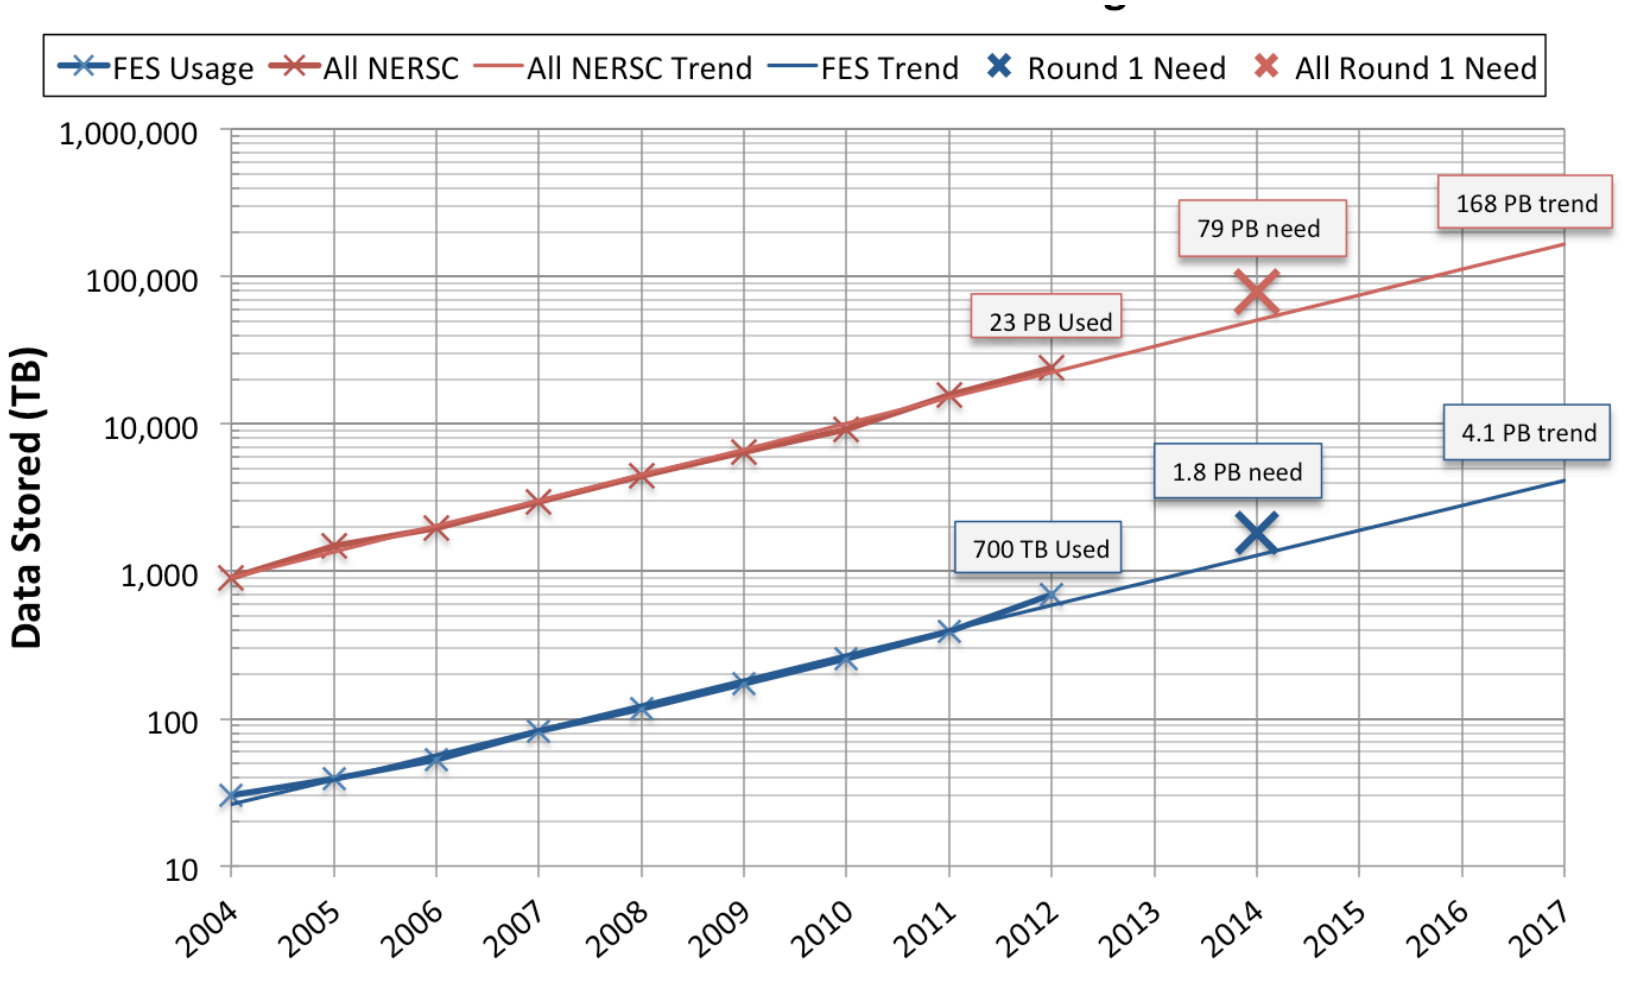

Logistics: Schedule

- **Agenda on workshop web page**
  - http://www.nersc.gov/science/requirements/FES
- **Mid-morning / afternoon break, lunch**
- **Today: Case study presentations & discussions**
- **Self-organization for dinner**
- **Wednesday: overview, review, and reach agreement on key findings**
- **Report: FES Intro + PI case studies + NERSC summary**
  - Final Case Studies due May 1
  - Richard / Harvey review
  - PI/DOE draft review June 15ß
  - Final: August 1 (?)
- **Final reports from 2009-2011 workshops (Target: 2014) on web**
  - http://www.nersc.gov/science/requirements

Logistics: Presentation to Remote Participants

- **We need your view graphs in advance**
  - Email
  - Web download
  - USB stick
- **The laptop at the front is sharing its screen with remote participants**
  - We'll load your presentations onto it

Questions?

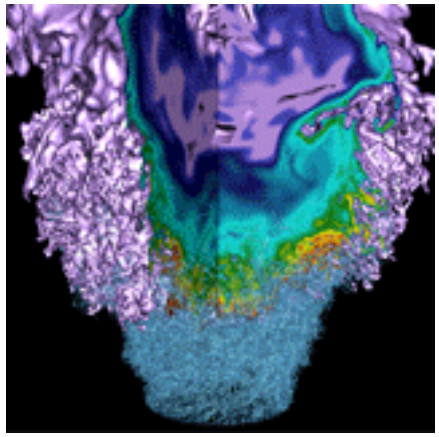

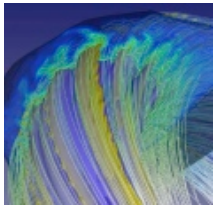

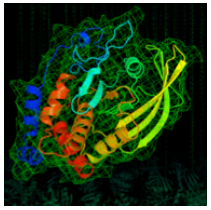

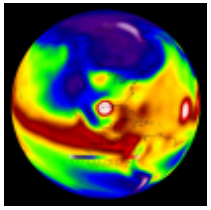

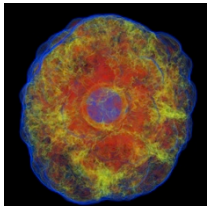

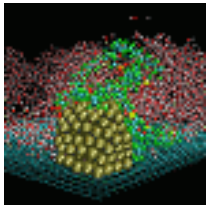

National Energy Research Scientific Computing Center

Terms

- **"Memory"**
  - Volatile or "RAM"
  - Each "node" has a pool of RAM shared among all cores on the node
  - "Global memory requirement" means the sum of all the RAM on the nodes on which your job is running
- **"Many Core"**
  - "Processors" with 100s+ of "light-weight" cores
  - Slower clock speeds (energy efficient)
  - Not self-hosted; need a master CPU (today)
  - Special ways needed to write programs
  - GPUs and Intel Phi

Storage Terms

- **"Scratch storage"**
  - Temporary, purged after ~6 weeks
  - Fast: 10s – 100s of GB/sec
  - Not backed up
  - Access from a single system (at least at high performance)
  - Default quotas: ~ 10s TB + today

Burst Buffers

- **"Permanent storage"**
  - Not purged
  - Usually backed up (feasible into the future?)
  - Somewhat less performant
  - Maybe sharable
  - Center-wide access
  - Default quotas: ~10s GB (Home) to ~10-100 TB (Project) today
- **"Archival Storage"**
  - Permanent & long term
  - Much slower access time
  - No quotas: up to 10 PB today