NERSC Requirements Reviews

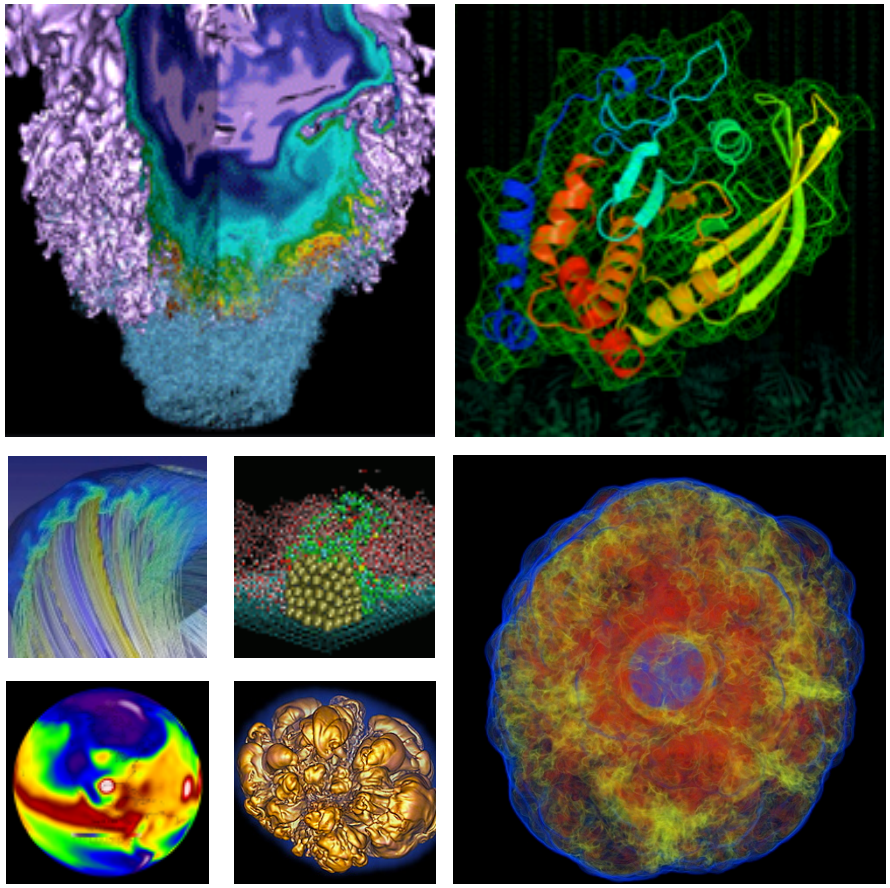

Richard Gerber
Harvey Wasserman
Requirements Reviews Organizers

February 19, 2014

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

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)

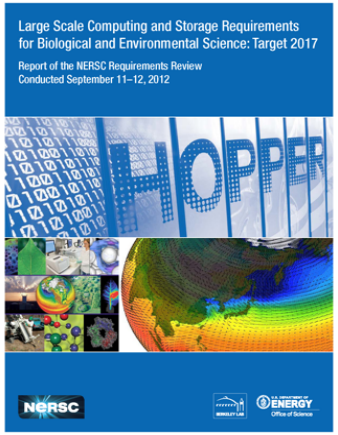

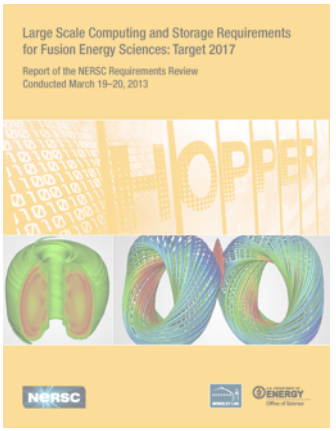

http://www.nersc.gov/science/hpc-requirements-reviews/reports/

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Impact

- **Highly regarded within DOE**
- **Scientific justification for ASCR budget requests**
  - Quantitative requirements
  - Documented needs from science teams
- **Basis for NERSC 7 and NERSC 8 Mission Need documents**
- **Influences NERSC services directions**
  - e.g. application readiness, support for high-throughput computing, planning for NERSC data services
- **HEP report formed much of the Distributed Computing and Facility Infrastructures portion of HEP community's "Snowmass" Report**

Production Computing Trend

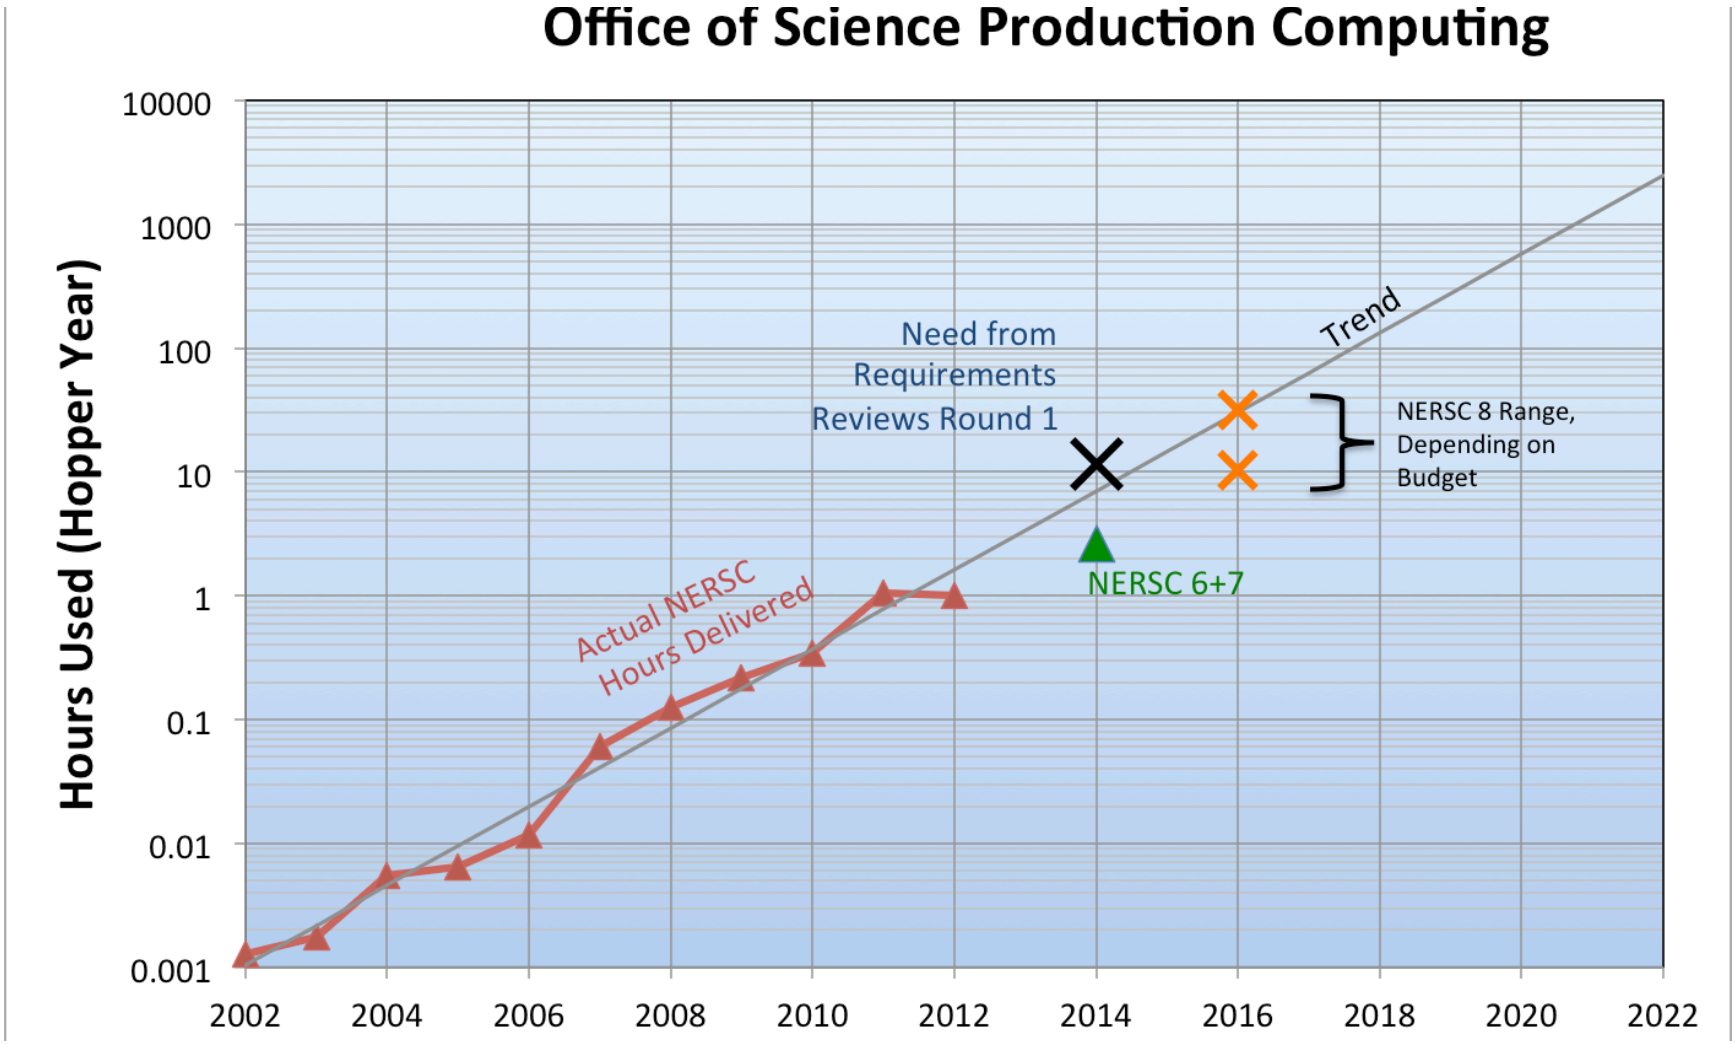

Keeping up with user needs will be a challenge

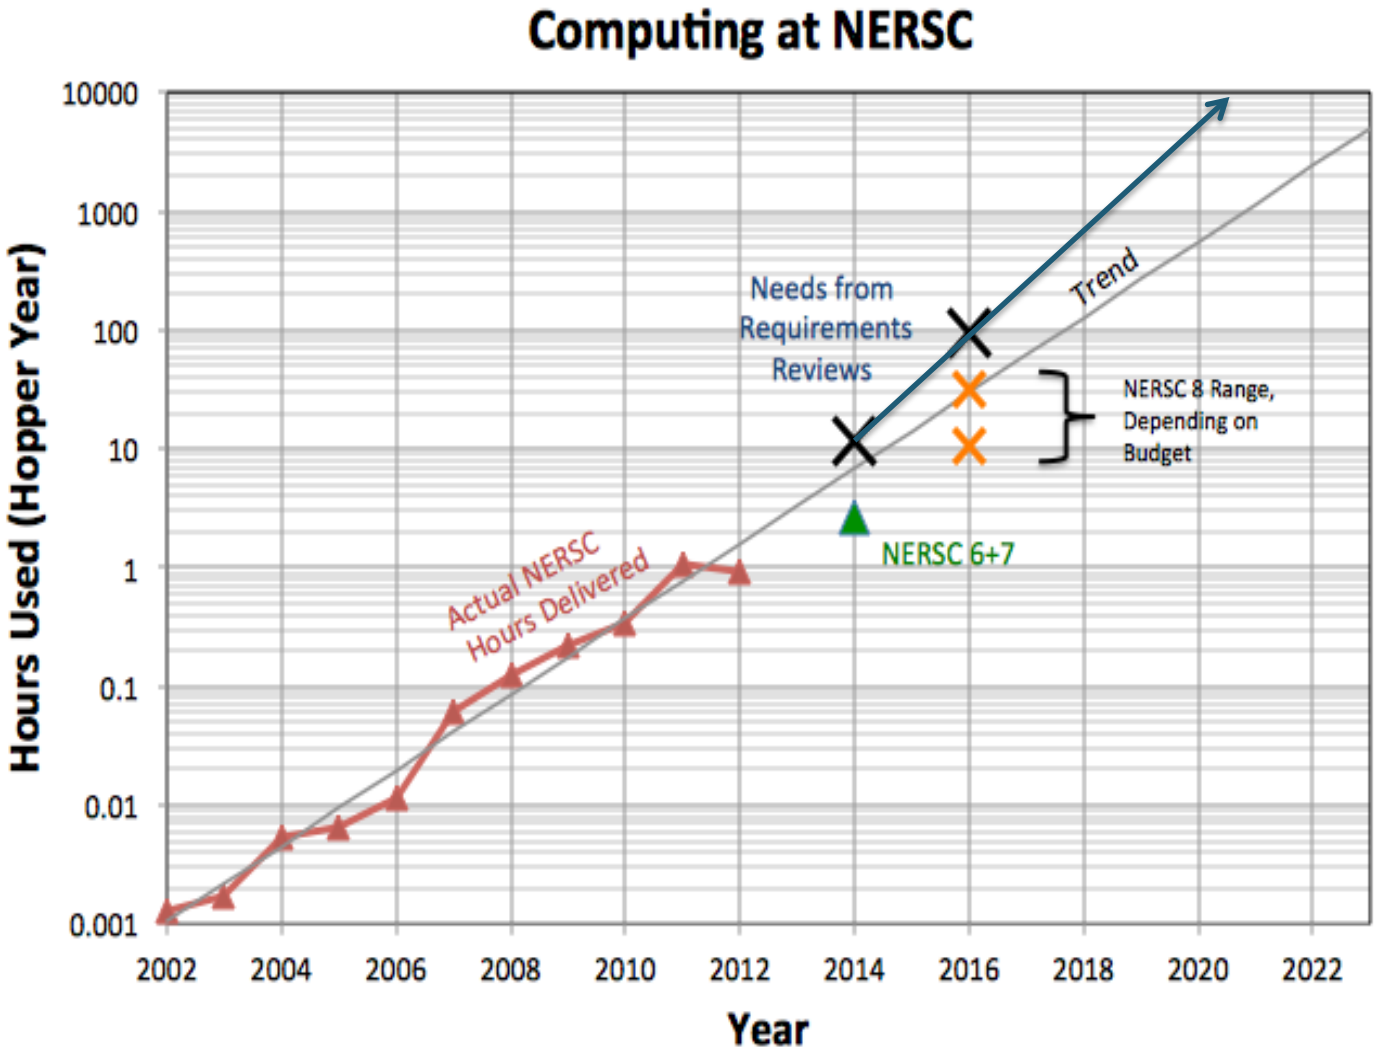

Future archival storage needs

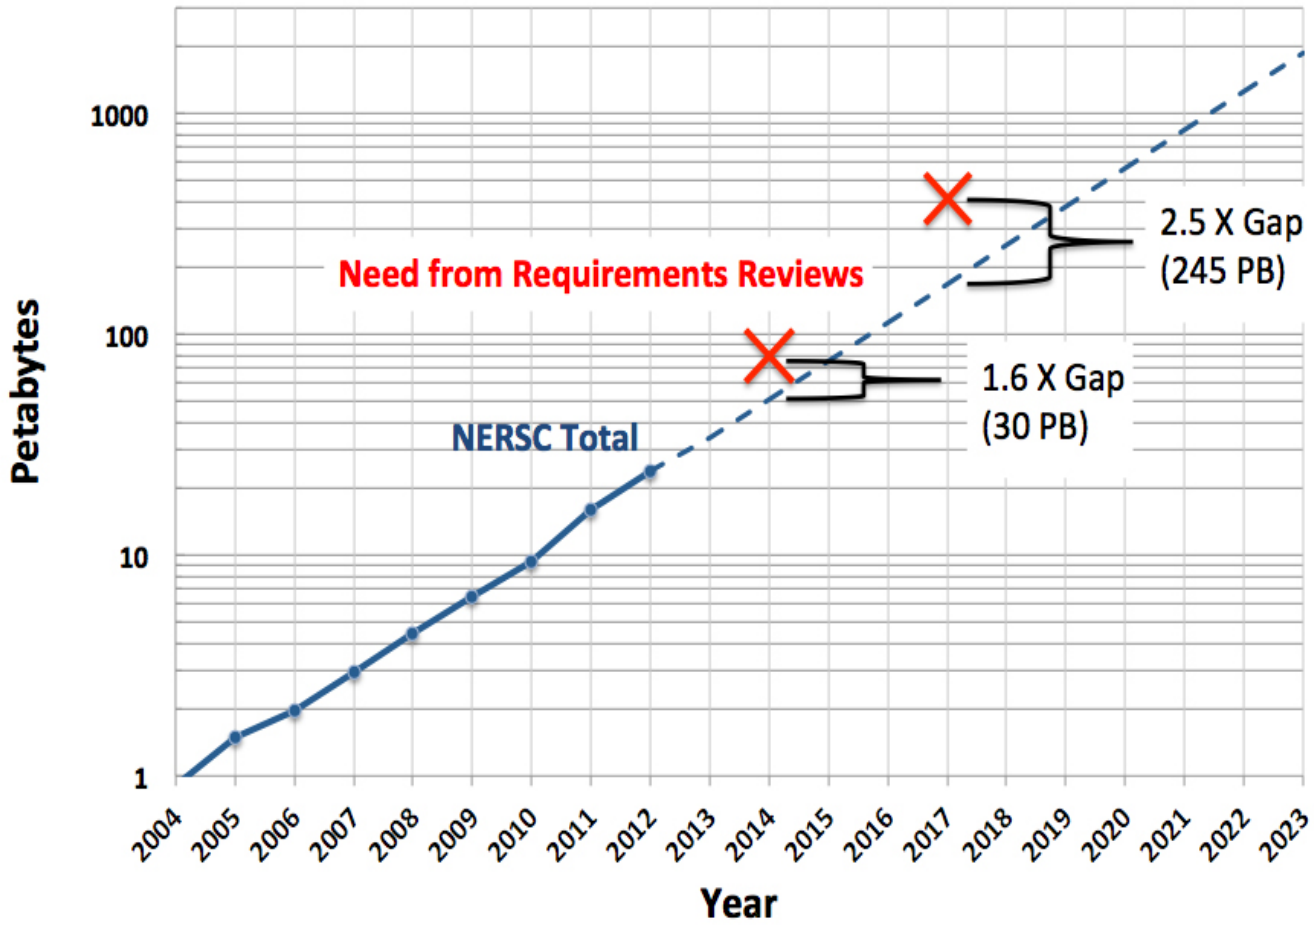

Priority Needs Across all Offices

- **More hours**
  - Progress is already limited by allocations
- **Science at Scale**
  - Requirements for science at scale
- **Science through Volume**
  - High throughput workflows (e.g., for data analysis)
  - Ensemble runs for V&V, statistics, & exploration
- **Science in Data**
  - Data storage, I/O bandwidth, data management tools

Priority Needs Across all Offices

- **Standard applications, libraries & tools**
  - Essential for productivity on next-generation systems
  - Scientists heavily invested in established HPC software
- **High performance, available, stable systems**
  - "Leading but not bleeding"
  - Availability and stability crucial for throughput
  - Expensive to deal with job failures & workflow interruptions
- **Preparation for emerging technologies**
  - Access to testbed or prototype systems
  - Assistance with application readiness

Round 2 In Progress: Target 2017

- **Reviews with BER, HEP, FES, BES, ASCR completed**
  - BER, HEP published; FES in final draft
- **NP planned for April 2014 (Just in! April 29-30)**
- **Target 2017 preliminary results**
  - Continued need for computation hours at or beyond historical trend
  - Increasing focus on data needs and capabilities
  - Application readiness is a major concern
    - Needs for porting help and robust and ubiquitous software libraries

Target 2017 Focus Items

- **Improved I/O capability and efficient I/O libraries**
- **Ability to run at largest scale, but also support ensemble runs and HTC**
- **Math libraries and solvers optimized for next-generation architectures**
- **Help getting application codes ready for next-generation systems**
  - Early access to testbeds and prototype systems
- **Data analysis and collaboration facilities and software**

Increasing Emphasis on Data

- **BER (2017):** "*Access to more computational and storage resources … and the ability to access, read, and write data at a rate far beyond that available today*"
- **HEP (2017):** "*Need for more computing cycles and fast-access storage; support for data-intensive science, including*
  - *Improvements to archival storage*
  - *Analytics (parallel, DBs, services, gateways etc.)*
  - *Sharing, curation, provenance of data*
- **ASCR (2014):** "*Applications will need to be able to read, write, and store 100s of terabytes of data for each simulation run. Many petabytes of long-term storage will be required to store and share data with the scientific community.*"
- **BES (2014):** "*[There is a need to support] … huge volumes of data from the ramp-up of the SLAC LINAC Coherent Light Source (LCLS) [and other experimental facilities in BES].*"
- **FES (2017):** "**I***/O libraries that will run efficiently on next generation architectures*"
- **NP (2014):** *Needs include*
  - *"Useable methods for cross-correlating across large databases …"*
  - *"[…] grid infrastructure, including the Open Science Grid (OSG) interface […]. "*
  - *[…] The increased capacity afforded by GPUs has resulted in […] a significant increase in IO demands in both intermediate and long term storage. "*

Requirements Reviews Methodology

- **Invite representative set of ~10 case studies from each program office**
  - Selected from current large NERSC users
  - Guided by PMs insight into future directions
- **Together, arrive at estimate of requirements ~5 years hence for each case study**
- **Forecast aggregate 5-year need for each office by considering total sum of case study needs as representative of entire office requirements**
- **Needs for "opportunity communities" are quoted separately**
  - Thus reported requirements may be underestimates

National Energy Research Scientific Computing Center