


NERSC Requirements Reviews



Richard Gerber
Harvey Wasserman
Requirements Reviews Organizers

February 11, 2013

Requirements Reviews



- 1½-day reviews with each Program Office
- Computing and storage requirements for next 5 years
- Participants
 - DOE ADs & Program Managers
 - Leading scientists using NERSC & key potential users
 - NERSC staff



High Energy Physics




Fusion Research

Reports From 6 Requirements Reviews Have Been Published



- Computing and storage requirements for 2013/2014
- Executive Summary of requirements
- Case studies
- Second round, for 2017 requirements, are underway – BER & HEP done, FES in March

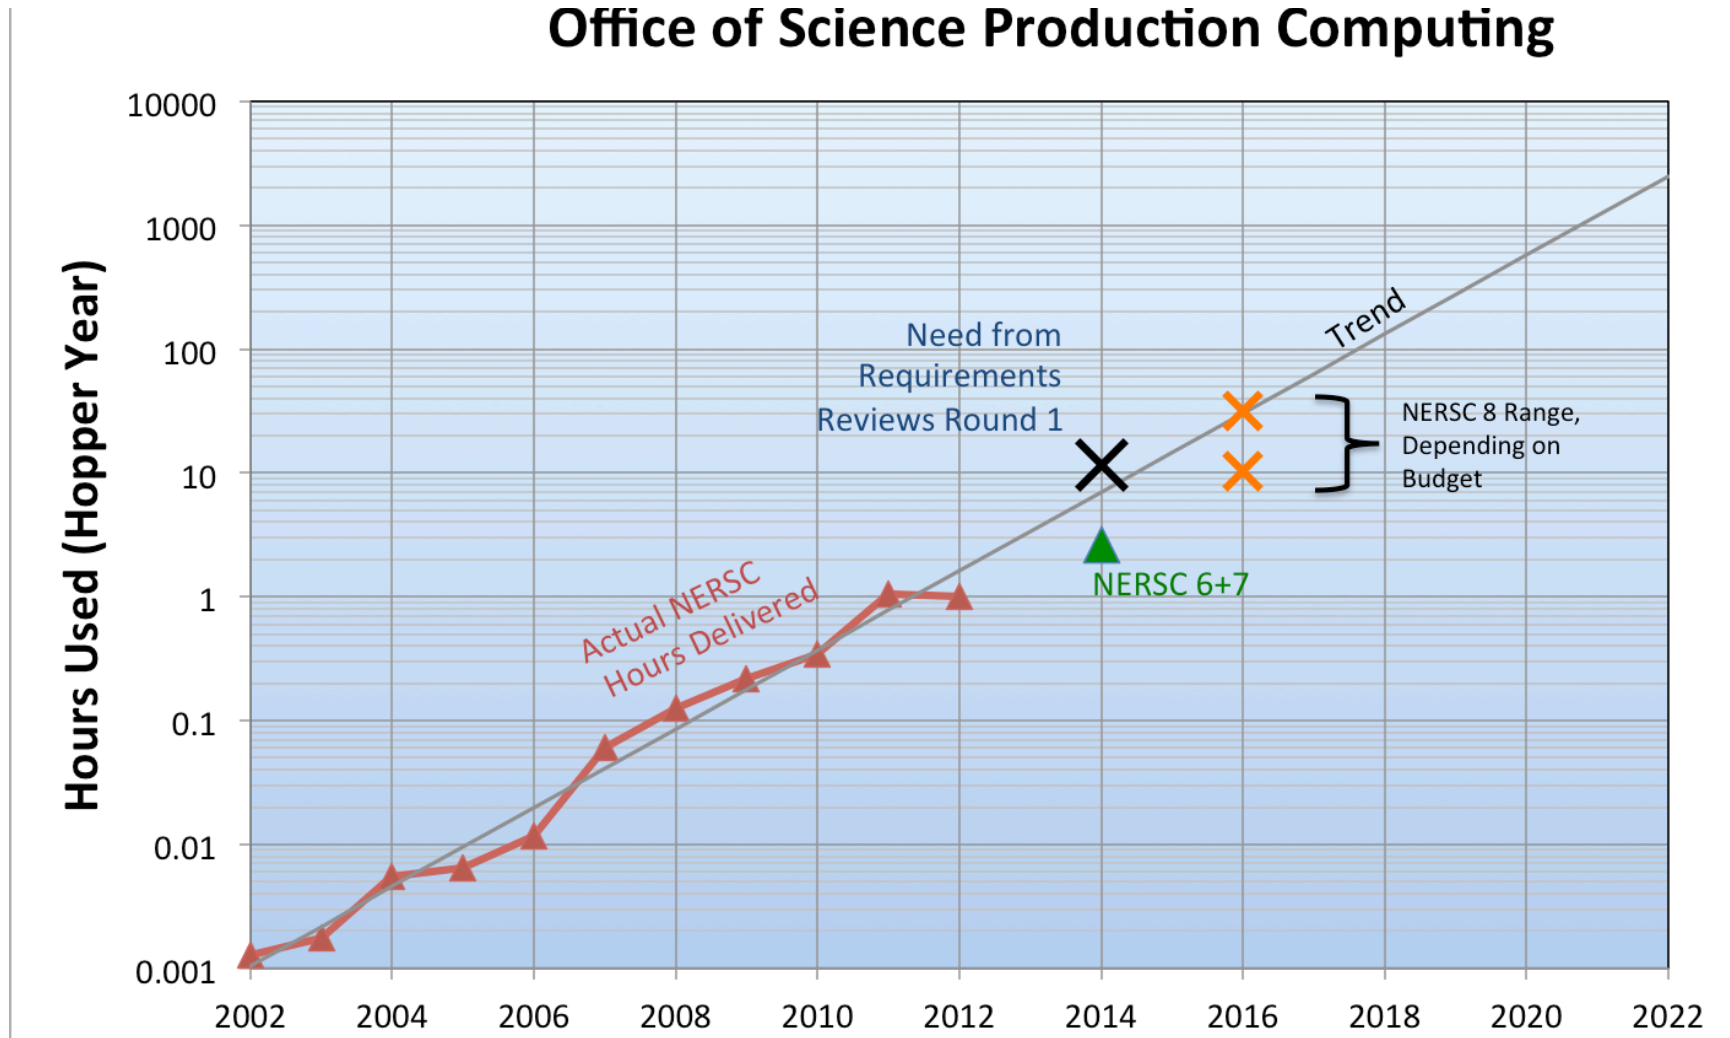


- **Highly regarded within DOE**
 - Quantitative requirements
 - Documented needs from science teams
- **Formed the foundation for NERSC 7 and NERSC 8 Mission Need documents**
- **Influencing NERSC services directions**
 - e.g. application readiness, support for high-throughput computing, planning for NERSC data services
- **Model for DOE Data Requirements reviews**
 - “*Data Requirements from NERSC Requirements Reviews*,” (Yelick & Gerber) document distributed at ASCR “Data Summit” in January 2013
- **Next HEP report likely to be basis for Distributed Computing and Facility Infrastructures portion of HEP community’s “Snowmass” Report**

Production Computing Trend




Office of Science Production Computing



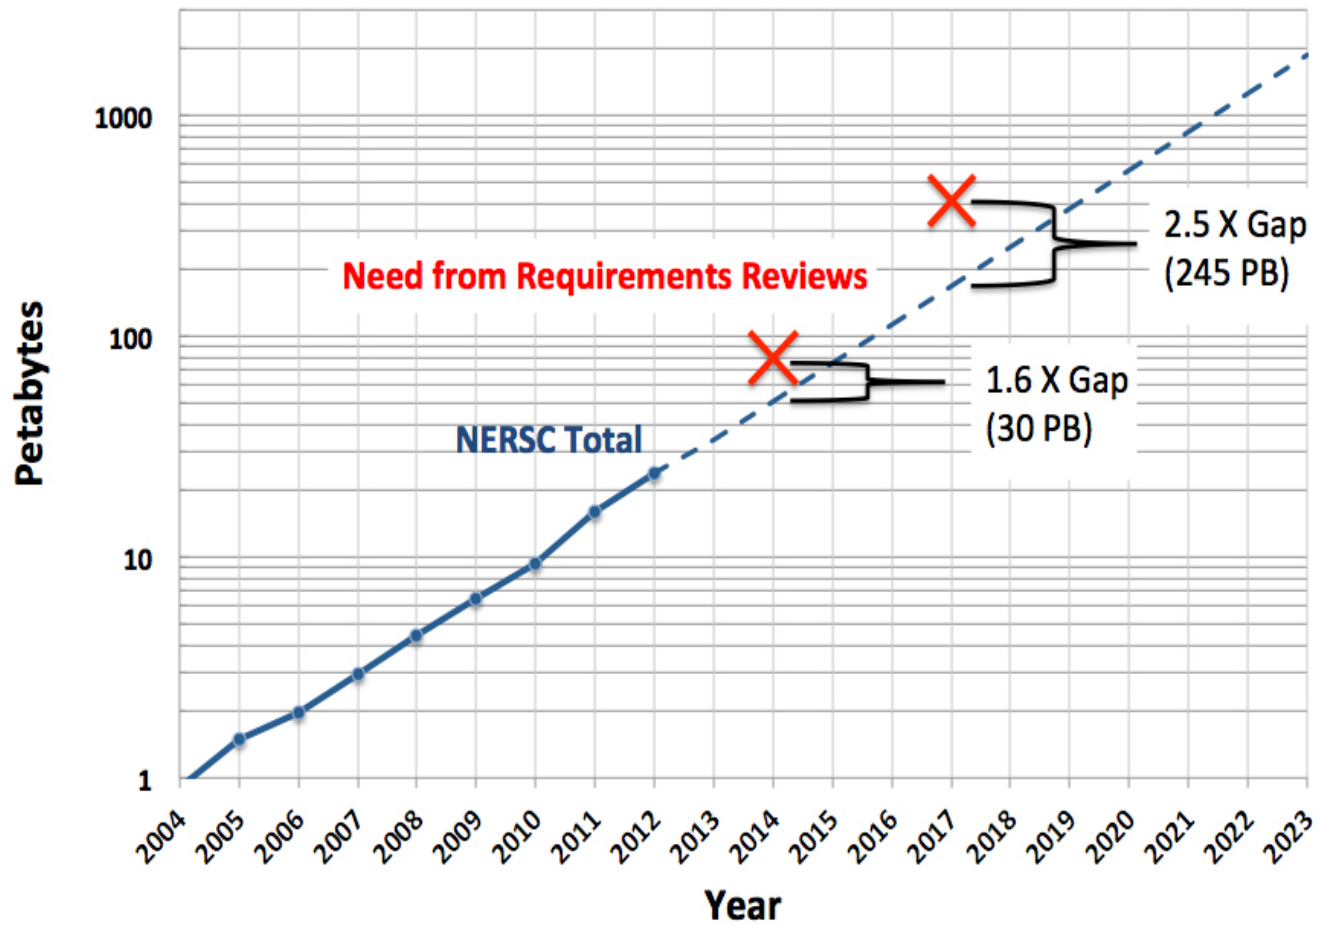
Keeping up with user needs will be a challenge



Computing at NERSC



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
 - Scientists heavily invested in standard HPC software
- **Highly available, stable systems**
 - Necessary for throughput
 - Expensive to deal with job failures & workflow interruptions
- **Preparation for emerging technologies**
 - Access to testbed or prototype systems
 - Assistance with application readiness

Increased data emphasis in requirements reviews



- **BER (2017 draft):** *“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 pre-draft):** *“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 (2014):** *“[Researchers need] data storage systems that can support high-volume/high-throughput I/O.”*
- **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.”*

Round 2 In Progress: Target 2017




- **Reviews with BER and HEP completed**
 - Reports in progress
- **FES planned for March 2013**
- **BES targeted for 2013**
- **Early results**
 - Continued need for computation hours at or beyond historical trend
 - Increasing focus on data needs and capabilities
 - Application readiness is a major concern
 - Early access to testbeds and prototypes requested
 - Needs for porting help and robust and ubiquitous software libraries

Requirements with 6 program offices



- Reviews with 6 program offices every 3 years
- Program managers invite representative set of users (typically represent >50% of usage)
- Identify science goals and representative use cases
- Based on use cases, work with users to estimate requirements
- Re-scale estimates to account for users not at the meeting (based on current usage)
- Aggregate results across the 6 offices
- Validate against information from in-depth collaborations, NERSC User Group meetings, user surveys



Tends to underestimate need because we are missing future users



National Energy Research Scientific Computing Center

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