Review Goals & Process

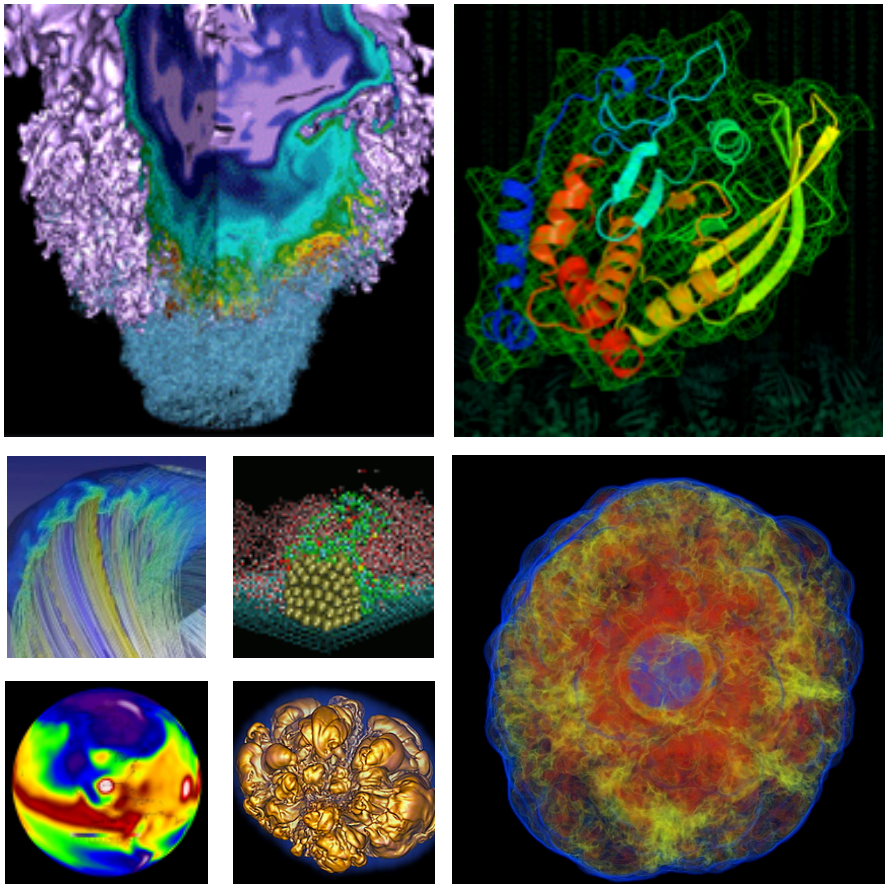

Richard Gerber

NERSC Senior Science Advisor to the Director

User Services Group Lead

January 15, 2014

Accelerating Scientific Discovery

17 journal covers in 2013

We love science!

NERSC holds Program Requirements Reviews so we can provide facilities and services that best enable scientific discovery.

We're hear to listen to your needs and help guide you to express them as requirements.

ASCR's Mission

The mission of the Advanced Scientific Computing Research (ASCR) program discover, develop, and deploy computational and networking capabilities is to to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE). ***A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures***, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

In addition to gathering your requirements, we're asking you to help us determine what NERSC can do to help meet this challenge.

Overview

- **Goals**
  - Gather computing, storage, and HPC services required to support ASCR research through 2017
  - Gather your input on how NERSC can get its users ready for exascale
  - Collect a set of project-based "case studies" with your research goals and how your HPC requirements support achieving those goals
  - Ultimately, a written report for DOE
- **Review findings allow NERSC/ASCR to**
  - create a science-based justification for acquiring resources
  - provide services that are important to you
  - select the appropriate technologies for the user base
- **This exercise benefits the Office of Science, ASCR, NERSC, & you**

PRRs Influenced the Selection of Edison

- **Findings from first round of PRRs (Program Requirements Reviews)**
  - The NERSC community would not be ready to effectively use accelerators in production by 2014
  - There is a need for improved I/O rates and disk storage
  - Many codes benefit from more memory per node, faster single-processor performance, and a high-bandwidth, low-latency interconnect
  - Productivity is more important than "feeds and speeds"
- **PRRs findings formed basis of NERSC 7 Mission Need Statement**
  - Edison has fast commodity Intel x86 processors, 64 GB/node memory, 6+ PB of /scratch, and novel high bandwidth, low-latency Aries interconnect
  - Adoption by NERSC community was immediate, with little porting effort
  - Performance is running 2X-4X that of Hopper on a per-core basis

Review Process

- **Collect and refine requirements for 2017**
  - Case study worksheets
  - Discussions at this meeting lead to high-level findings
  - 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 & ASCR to review
- **Send final draft to DOE ASCR office for final approval**
- **Publish final report**

Key Strategy

Tell us what you need to support your research – as specifically as possible

- computational and data resources
- HPC services
- software

Tell us how you are preparing your software for exascale architectures

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 ASCE need by the fraction of 2013 ASCR usage represented by case studies
  - Important: Associate each case study with a 2013 NERSC repo or repos
  - New projects' requirements added in separately
- **Like to do the same for Scratch and Project shared disk**
  - Please state 2013 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

Data Requirements

- **Archival storage estimate for 2017**
  - This is an total stored, not data added in 2017
  - 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 for a single run)
  - What bandwidth do you need?
- **Project shared disk space (permanent)**
  - What will you need for source code, data files or executables that will be constantly accessed and/or shared, etc.
- **What data services and software do you need?**

Logistics: Schedule

- **Agenda on workshop web page**
  - http://www.nersc.gov/science/requirements/ASCR
- **Mid-morning / afternoon break, lunch**
- **Case study presentations & discussions**
- **Report: ASCR Intro + PI case studies + NERSC summary**
  - Final Case Studies due Feb. 21
  - Richard / Harvey review and iterate with you
  - PI/DOE draft review April 1
  - Final: May 1
- **Final reports from previous 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
- **We will stay on time**
  - Descriptive and concise science justification
  - Please emphasize requirements

Questions?

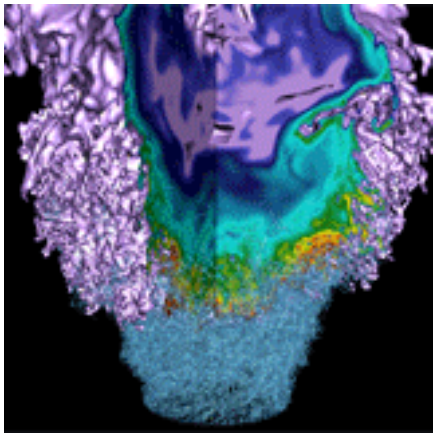

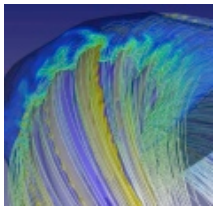

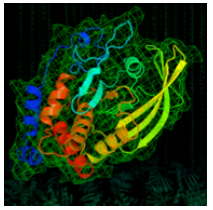

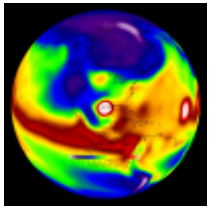

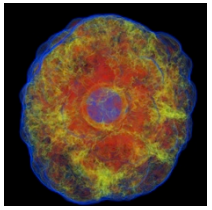

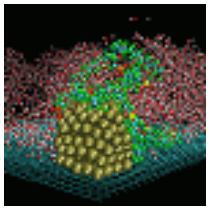

National Energy Research Scientific Computing Center

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

Accelerating Scientific Discovery

We love science!

NERSC holds Program Requirements Reviews so we can provide facilities and services that best enable scientific discovery.

We're hear to listen to your needs and help guide you to express them as requirements.