


Allocations



Richard Gerber & Clayton Bagwell


NUG Business Meeting
March 24, 2016

Allocations Summary

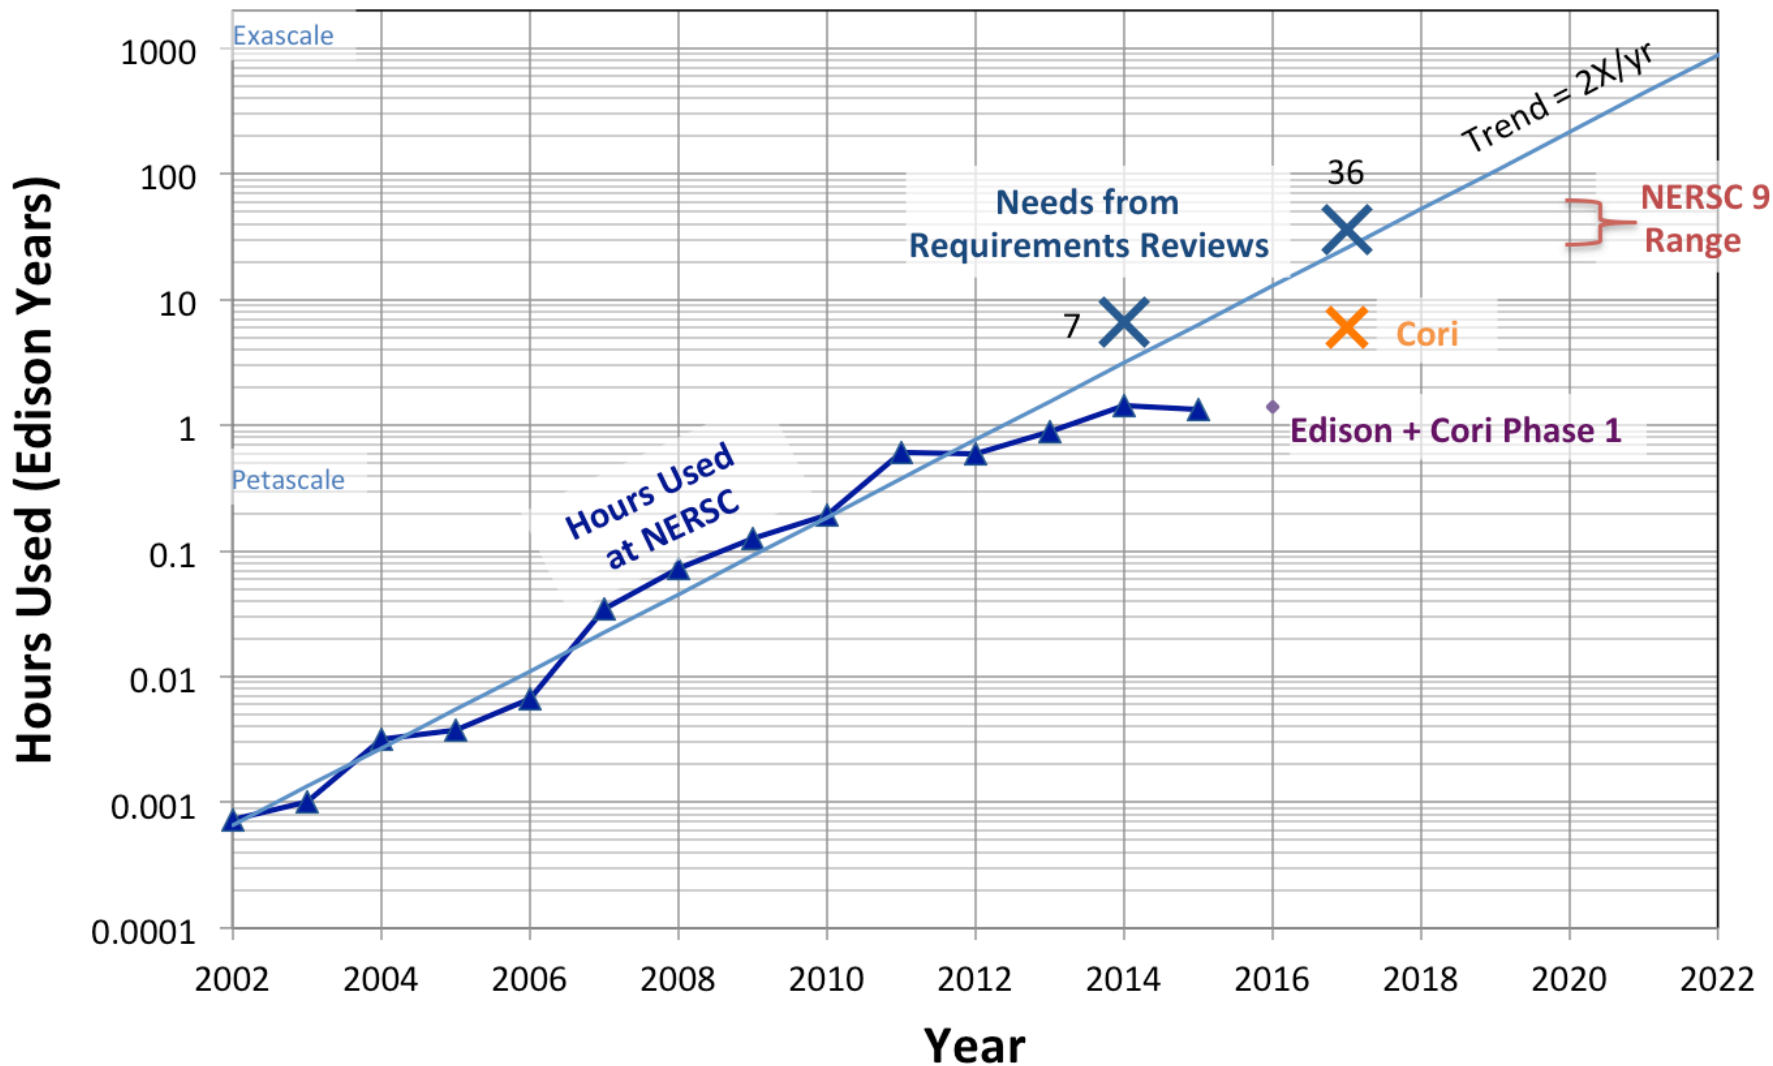


- **How the NERSC pie is distributed**
 - A little history
- **DOE Offices & Programs**
- **NERSC reserves**
- **The ERCAP process**
- **How user accounts and allocations work**
- **What happens when user/repo run out of time**
- **Q & A**

Allocations History



Compute Hours Used at NERSC




- **Cori Phase 2 will arrive this summer**
- **Full Cori system (Phase 1+2) will be in production starting sometime in early 2017**
- **We anticipate a DOE Production pool of between 4.8 and 6.0 billion MPP hours equivalent.**
 - 2.4 billion in 2016
- **We will likely refactor the base unit (MPP hours)**
 - Now: Hopper equivalent core hours
 - Possible new: Cori Phase 2 equivalent core or node hours or maybe something else

- **NERSC commitment to DOE in AY2016: 3 Billion MPP Hours**
 - 2,400 M for DOE Production (mission computing): DOE Managers
 - 300 M for ALCC (ASCR Leadership Computing Challenge): ASCR competition
 - 300 M for NERSC Director's Reserve: NERSC
- **Additional time set aside for NERSC miscellaneous: ~72 M hours**
 - NERSC Overhead – 65 M
 - Startup projects – 5 M
 - Education – 1.5 M
 - Guests – 500 K
- **Additional time is available if system downtime is less than estimated or new resources become available (e.g. preproduction systems)**

Distribution Among Offices for 2016



2016 Allocation By DOE Office (Millions)




2,400
Million
Total

Distribution Among Reserves for 2016



2016 NERSC Allocations by Reserve (Millions)




2,400
Million
Total

Request vs. Allocated



NERSC 2016 Allocations

■ Request ■ Available ■ Allocated



Purpose of Startups



- Many new NERSC projects begin with a Startup allocation
 - New to HPC
 - Not yet ready to run production codes on thousands of CPU cores
- Limited to 50,000 MPP hours (Cray Hopper equivalent hours)
- Request for a Startup project can be made at any time during the year
- Good for up to 18 months
 - Must then convert to DOE Production with allocation coming from DOE Office/Program reserve
 - Production projects are expected to have codes that can effectively use hundreds to thousands of processors.

Division Director's Reserve Supports Strategic Programs and High-Impact Science



Focus Areas

- Industrial partners
- Application readiness
- Data infrastructure development
- Applied energy offices

In 2015 and 2016 also used to augment high-impact science teams and buffer uncertain, limited hours (move, Cori Phase 1 and 2 arrival, Hopper retirement). In 2017 will likely have competitive program (a la NISE, Data Challenge)

Category	Number of Projects	Allocated from DDR (Millions of MPP Hours)
Director's Reserve Projects	45	161
Additional Support for High-Impact Mission Science	221	139
TOTAL		300

Description of ERCAP Process



- **Users (PIs or their proxies) apply for time through the ERCAP process (via NIM)**
 - Open to anyone whose project has relevance to DOE Office of Science
 - Success more likely for Office of Science funded research
 - ERCAP cycle has September deadline for application renewals, but accept throughout the year
 - Success only possible if DOE program managers hold a reserve from their initial allocation pool
- **DOE “Allocations Managers” from each Office evaluate requests and make awards**
 - Allocation year runs from 2nd Tuesday in January through 2nd Monday in January the following year

ERCAP Information We Collect – Suggestions Welcome



- **Project Overview**

- PI's information
- Other Authorized Preparers – PI Proxy(ies)
- Senior Investigators
- Project Title
- URL for relevant, project related web site
- Project Name
- Project Class
 - DOE Base Funding, ALCC, CSGF, SBIR, SciDac, Startup, Education
- Sponsoring Site
- Science Category
- DOE Office & Program

- **DOE Funding**
 - Yes – DOE Manager & Grant Info
 - No – Where does funding come from & what is the relevance of your research to the mission of one of the DOE Office of Science Office/Programs
- **Project Description**
 - Summary – non-technical
 - Long – more technical PDF
- **Accomplishments**
- **Publications**

- **Resources Requested and Justification**
 - MPP hours, Storage SRUs, Scratch Space, Real Time needs
- **Code Description(s)**
 - Code name and description, # nodes used, Languages, Parallel Programming Model
- **Data & Analytics**
 - I/O formats/libraries, workflow tools,
 - Imaging/Visualization tools, Analytics Frameworks
 - How much data moving to NERSC, Web access?
- **Other HPC support**

- **Users must be part of one or more repositories**
 - They spend against repos (bank account)
 - PIs/Managers can limit by percentage or absolute hours
 - Users have default repo but can choose which to charge against if they have more than one
- **User Management**
 - Repo PI/Managers manage their users via NIM
 - Add via the NIM interface
 - Set user limits via NIM interface

What Happens When a User/Repo Runs Out of Time



- **Option 1: If balance zero or less, they can still submit and run jobs**
 - Jobs run at low priority in scavenger queue
 - First year for this policy; in previous years users were prohibited from running at NERSC
- **Option 2: They can ask DOE NERSC Allocations Manager for their area for more time**
 - This lets them run at normal priority
 - Few program managers have time left in reserve

Questions & Answers





National Energy Research Scientific Computing Center
Thank you