


NERSC Users Group Monthly Meeting




August 25, 2016

Agenda



- Cori Phase II Update
- Data Day debrief
- NESAP & resources for porting to KNL
- Edison Scratch Filesystem Updates
- AY 2017 ERCAP Allocation Requests

Cori Phase II Update



Tina Declerck
Computational Systems Group

August 25, 2016

Agenda

- **Prep for Cori Phase 2**
- **Cori Phase 2 Installation**
 - System Arrival & Installation
 - Current Status
- **Projected Timeline**
- **NERSC pre-merge testing**
- **Merge plan**
- **Post Merge**
- **Acceptance Testing**

Preparation for Cori Phase 2

- **Additional Pipe work**
 - Project change early in the process
 - Completed in time for ph2 arrival
- **SMW 8.0 / CLE 6.0 upgrade**
- **SLURM**
 - Prepping software for KNL
- **Gerty**
 - Test platform with KNL
 - Allows early testing of KNL functionality

Delivery

- **System Arrival**
 - 1 day per week over 3 weeks
 - 12, 12, and 28 cabinets
 - 11 trucks



Physical Installation



- Cabinets Placed
- Plumbing & Electrical
 - Cray Site Engineering connected the system plumbing
 - Hard wire connections completed on August 4th

Physical Installation



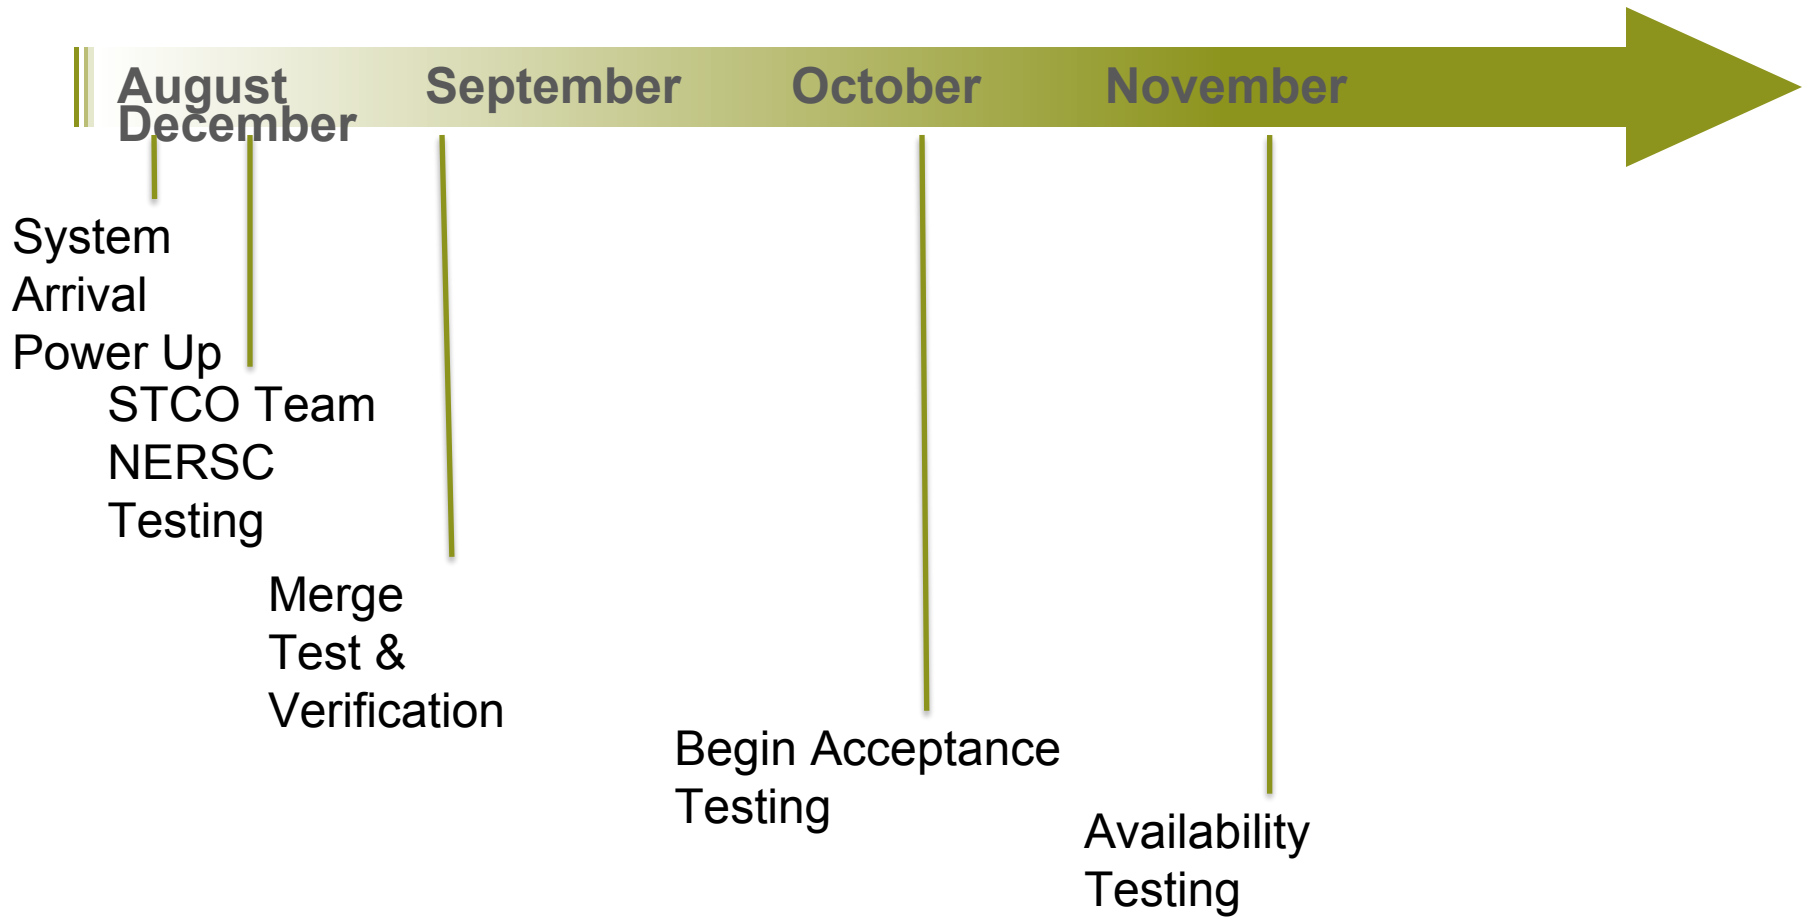
- **HSN Cabling**

- 7 days of HSN cabling; completed on August 1st.
- Optical HSN cables are in ready to be connected to Cori Phase 1 when we merge

System Test & Check Out

- **Cray Diagnostics**
 - ~2 weeks of testing to shake out the new racks
 - Check cable connections
 - Check for bad links and replace cables
 - Check for memory and processor fallout
- **Allowed NERSC to check cooling system**
 - Ensure pumps could handle the additional load
 - Cooling sufficient for new cabinets
- **Verify Power**
 - Ensure power requirements of the KNL racks
 - HPL and mem tests put high load on the system

Projected Timeline



NERSC Pre-Merge Testing

Goal: Ensure KNL is operating well

- Run full scale KNL job
- Ensure fallout rate is within expectations
- Ability to boot nodes into different modes
- Benchmark SSP runs
- Verify UME rates
- Testing the variable blower speeds


Goal: Keep the Haswell partition operating

- Verify ability to upgrade KNL
- Verify ability to update KNL bios
- Test affect on system with hardware actions

Merge Plan

- **Prepare for the merge**
 - Pre-lay cables
 - Prepare system software
- **Unconfigure current DataWarp**
- **Shutdown Cori Ph1 & Ph2**
- **Complete hardware actions (1.5 weeks)**
 - Move blades
 - Cable work
 - Connect Aries
- **Software Configuration/Updates**
 - System configuration
 - DataWarp configuration
 - SLURM configuration
 - Monitoring updates
 - Ensure latest KNL updates are installed
- **Boot and Test**

Data Day Debrief



Annette Greiner
Data & Analytics Services Group

August 25, 2016

Data Day 2016



72 in-person attendees

30 remote attendees

102 Total Participants



Tutorials and talks on

Machine learning

Python

Spark

Visualization

Burst buffer

Data management


Hack day on workflows with

Spark

Burst buffer

A big thank-you to  for sponsoring Data Day!

NESAP & KNL Resources



Brian Friesen
Application Performance Group

August 25, 2016

Cori Phase 2 is here!



- **All 52 cabinets of Cori Phase 2 (“Knights Landing”) are installed and powered on in Wang Hall!**
- **Cray is testing the machine right now**
 - They will turn it over to NERSC staff when they’re done testing
- **After NERSC does our own tests, we will open it to a subset of users (NESAP) for a period of “early science” for about 12 weeks**

Cori Phase 2 is here!



- **Additional users with codes that have demonstrated readiness for KNL will get access after the NESAP teams.**
- **We anticipate that all users will have access to the system for code development and testing in early 2017.**
- **See the Cori timeline here:**
<http://www.nersc.gov/users/computational-systems/cori/cori-phase-ii-schedule/>

Application readiness at NERSC




- **NESAP: NERSC Exascale Science Application (Readiness) Program**
- **Multi-pronged effort to prepare existing NERSC codes for energy-efficient, manycore architectures that we will likely see on exascale computers**
- **Consists of collaborations among NERSC staff, NERSC postdocs, Cray/Intel staff, and application developers (you!)**
- **Next step on the road to exascale: Cori Phase 2**

- **Components of application readiness:**
 - 8 postdocs dedicated 100% to different Tier-1 applications
 - consultants spending 20% on > 2-3 Tier-1/2 applications
 - Technical deep-dives with Intel/Cray every ~2 months
 - Training sessions with Intel/Cray
 - Early access to Knights Landing hardware
 - Knowledge transfer among Tier-1/2/3 codes and the rest of NERSC users

- **Resources to help you prepare your codes for Cori Phase 2:**
 - General info about Intel Xeon Phi (“Knights Landing”) and Cori Phase 2:
<http://www.nersc.gov/users/computational-systems/cori/cori-phase-ii/>
 - Application porting and performance (lots of info here!):
<http://www.nersc.gov/users/computational-systems/cori/application-porting-and-performance/>

- **Resources to help you prepare your codes for Cori Phase 2:**
 - Application case studies:
 - <http://www.nersc.gov/users/computational-systems/cori/application-porting-and-performance/application-case-studies/>
 - Using the Burst Buffer:
 - <http://www.nersc.gov/users/computational-systems/cori/burst-buffer/>

Edison Scratch Filesystem Updates



Douglas Jacobsen
Computational Systems

August 25, 2016


- **Update filesystem software to latest stable version from Cray (Neo2.0update19)**
 - Performance benefits from GRIDRAID
 - Large numbers of bug fixes for performance, scalability and reliability
 - Improved manageability (owing to bug fixes)
 - Expectation of reduced load in certain areas, reducing likelihood of contention
- **Will destroy all data on /scratch1 and /scratch2**
 - Reformat to gridraid requires this

Edison \$SCRATCH Update



- **Begins Monday, 08/29/2016 08:00 Pacific**
- **Will redefine \$SCRATCH**
 - Will point to “/invalid/path/scratch/maintenance”
 - Jobs requesting scratch1, scratch2, or SCRATCH licenses will be held and then cancelled
 - Will attempt to identify jobs referring to \$SCRATCH or \${SCRATCH} and provide an error upon submission
- **Maintenance scheduled through 9/12/2016 Noon**
- **At conclusion will resume normal definition of \$SCRATCH**
- Your directory has a 50% likelihood of changing filesystems when we return to service
- **You can use cori scratch (\$CSCRATCH) in the meantime**

AY 2017 ERCAP Renewals



Clayton Bagwell

User Engagement Group - Account
& Allocation support

August 25, 2016

ERCAP

“Energy Research Computing Allocations Process”

All projects that wish to continue to use NERSC resources in AY 2017, must submit an ERCAP renewal request. Including:

- **DOE Production (including SciDac & CSGF)**
- **Startup**
- **Education**
- **Director's Reserve**

Submission Deadline: September 19, 2016 23:59 PDT

Note: ALCC is exempt since they run on a different allocation schedule.

**Award decisions will be emailed to PIs/PI Proxies
on December 9th.**

ERCAP Application Deadlines

<http://www.nersc.gov/users/accounts/allocations/deadlines/>



Powering **Scientific Discovery** Since 1974

search...

Login

Site Map | My NERSC | Share

HOME ABOUT SCIENCE AT NERSC SYSTEMS **FOR USERS** NEWS & PUBLICATIONS R & D EVENTS LIVE STATUS

FOR USERS

- » Live Status
- » User Announcements
- » My NERSC
- » Getting Started
- » Connecting to NERSC
- » Accounts & Allocations

Glossary

User Accounts

Allocations

Overview and Eligibility

Apply for Your First Allocation

Allocation Request Form (ERCAP)

2017 Call for Proposals

Application Deadlines

Request More Time

Allocation Reductions

Managing Your Allocation

DOE Allocation Managers

NESAP Application Process

Past Calls for Proposals

NIM (NERSC Information Management portal)

Home » For Users » Accounts & Allocations » Allocations » Application Deadlines

ALLOCATION PROPOSAL (ERCAP) APPLICATION DEADLINES

Award Type	ERCAP Submission Open	ERCAP Submission Deadline	Award Decision	Allocation Period
Startup and Education	Ongoing	November 10, 2016	Within 3 weeks of applying	18 months from award (but must be renewed for the next allocation year)
Production (DOE awarded) for the rest of allocation year (AY) 2016	Ongoing	November 10, 2016	Within 4 weeks of applying (pending sufficient DOE reserves of time)	through January 9, 2017
All continuing projects must be renewed for 2017. New projects may apply after the deadline.	August 9, 2016	September 19, 2016	December 9, 2016	January 10, 2017 - January 8, 2018



U.S. DEPARTMENT OF **ENERGY**

Office of Science



Submitting the ERCAP renewal request:

<http://www.nersc.gov/users/accounts/allocations/request-form/>



The screenshot shows the NERSC website interface. At the top, there is a search bar and navigation links for Site Map, My NERSC, and Share. The main navigation menu includes Home, About, Science at NERSC, Systems, For Users, News & Publications, R & D, Events, and Live Status. The 'For Users' section is expanded, showing a list of links including Live Status, User Announcements, My NERSC, Getting Started, Connecting to NERSC, Accounts & Allocations, Glossary, User Accounts, Allocations, Overview and Eligibility, Apply for Your First Allocation, Allocation Request Form (ERCAP), 2017 Call for Proposals, Application Deadlines, Request More Time, Allocation Reductions, Managing Your Allocation, DOE Allocation Managers, NESAP Application Process, Past Calls for Proposals, NIM (NERSC Information Management portal), Awarded projects, ALCC, and Computational Systems. The main content area is titled 'THE NERSC ALLOCATION REQUEST FORM (ERCAP)' and contains the following text:

Requests to use NERSC resources are submitted annually via a web form known as the **ERCAP** (Energy Research Computing Allocations Process) Request Form. ERCAP is accessed through the **NERSC Information Management (NIM)** web interface and is available year-round.

If you are new to NERSC and want to apply for an allocation as a new Principal Investigator please read [Applying for your First NERSC Allocation](#).

In general existing projects should be renewed for the next allocation year - you should not start a new request if you have an existing repository. Renewing carries forward most of the responses to the request questions and preserves the project's user list.

Renewing a Current Project

If you have existing repositories you should usually renew them; this preserves your user lists. Point your browser to nim.nersc.gov, and login to NIM. Click the My Stuff menu and select **My ERCAP Requests**. You will see a list of your ERCAP requests.

To renew an existing project for the next allocation year click its Renew link. Your renewed ERCAP Project Overview will load. After you click the **Save Project Overview** button the project PI name and title will appear. You can then click on any ERACP tab in any order, but before switching tabs you must click a tab's Save button to save any edits you have made to that section.

On the right side of the page, there is a 'TABLE OF CONTENTS' section with the following items:

1. Renewing a Current Project
2. Starting a New Request
3. Project Overview Tab
4. DOE Funding Tab
5. Project Description Tab
6. Resources Requested Tab
7. PDSF Tab
8. Code Description(s) Tab
9. Data & Analytics Tab
10. Validating Current User List (For Renewal Requests)
11. Other Info
12. Submitting your request
13. Printing and saving to disk
14. Advice on answering long text questions
15. ERCAP error messages
16. Deleting a request

At the bottom right, there is a link for 'How to Submit a 2016 ERCAP Request'.

What's New for AY 2017?



The currency at NERSC is the “NERSC Hour”

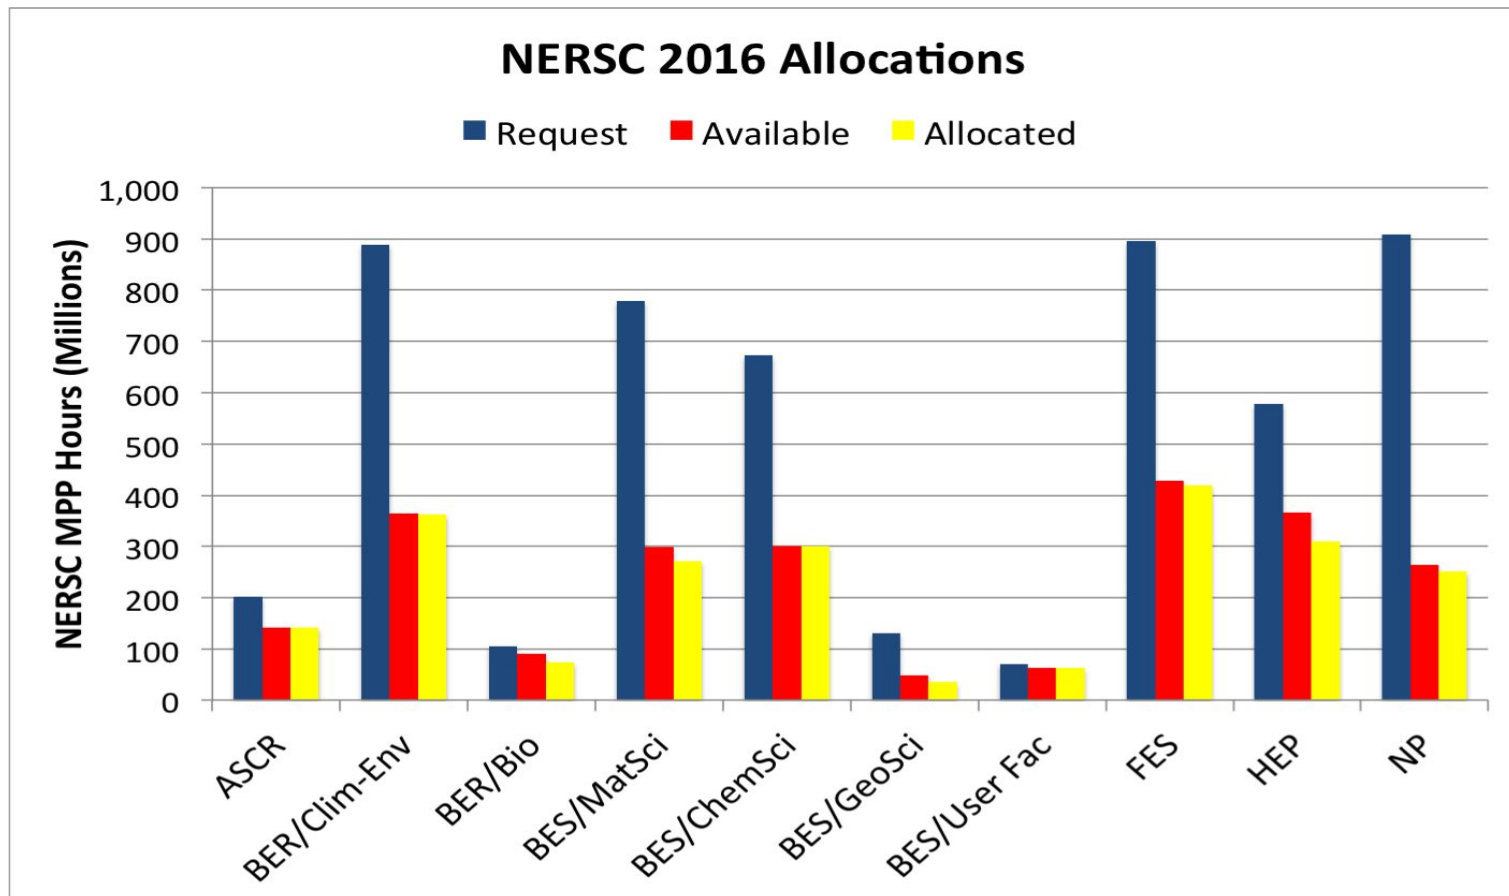
- Equivalent to “MPP Hour” in previous years
- Change in nomenclature in preparation for Cori Phase II and follow-on systems

- Edison: 48 NERSC Hours per hour per node used
- Cori Phase 1: 80 NERSC Hours per hour per node used
- Cori Phase 2: TBD based on measure application performance

Initial computing hours provided to DOE will initially be flat.

- **3 Billion NERSC hours to be distributed by DOE to the various offices.**
- **All Cori Phase 2 time will be allocated in the spring of 2017 for use in the second half of allocation year 2017.**

AY 2016 Requested vs Allocated



DOE Funding tab, Question # 10

If you have funding from sources other than the DOE Office of Science, we ask you to include information on all of your funding sources.

Resources Requested, added question 18.1

- Use of Experimental or Observational data
- We are asking you to identify if the primary role of this project is to:
 - Analyze data from experiments or observational facilities?
 - Create tools or algorithms for analyzing experimental/observational data?
 - Combine models and simulations with experimental/observational data?

Code Descriptions tab, question # 20

- **We reduced the amount of information requested**

Data & Analytics tab, questions 21 & 22

- **We removed this tab and the associated questions**

Helpful URLs:

<http://www.nersc.gov/users/accounts/allocations/deadlines>

<http://www.nersc.gov/users/accounts/allocations/overview/>

<http://www.nersc.gov/users/accounts/allocations/first-allocation>

<https://nim.nersc.gov>

Contact Information:

Accounts & Allocations:

accounts@nerisc.gov

allocations@nerisc.gov

Phone:

510-486-8612 or

1-800-66-NERSC opt 2

Questions?



NERSC