


NERSC Users Group Monthly Meeting









July 20, 2017

Agenda



- **Plans for Data Day/NUG 2017**
- **NERSC Achievement Awards**
- **Scratch Purging**
- **Edison Update**
- **Cori Allocations, Accounting, & Usage**
- **MyNERSC Data Dashboard**

Data Day/NUG 2017



<https://www.nersc.gov/users/NUG/annual-meetings/nersc-data-day-and-nug2017>

- **Data Day Topics**
 - Machine learning
 - Burst buffer
 - Shifter
 - Data management
 - Python
 - Esnet
- **Data Data Schedule**
 - Demo and Science Talk, Sep 19
 - Hackathon, Sep 20 AM

Data Competition









- **Wed AM, Sept 20**
- **A combined Data Day and NUG event**
- **Teams will be given data to analyze and present meaning results for an award**
- **Details to come**

- **NERSC Status and Updates**
 - NERSC Status
 - Cori and Edison Update
 - N9 Project Update
 - Storage 2020
 - Data HPC Update
 - NIM and ERCAP Update
 - Security and MFA Update
 - User Survey Results

- **NERSC Science and Technology Day**
 - NERSC High Impact Science
 - Preparing Users for Cori KNL
 - NESAP Update
 - ECP talk for Lunch Plenary
 - Talks from NERSC Achievement Awards
 - Lightning talks from Large Scale Cori runs
 - UPC++ with the new V1.0 code

- **Data Day and NUG2017 info and registration at**
 - <https://www.nersc.gov/users/NUG/annual-meetings/ner-sc-data-day-and-nug2017>

NERSC Achievement Awards



NERSC HPC Achievement Awards



- **Annual awards to recognize extraordinary users of NERSC HPC resources.**
 - NERSC Award for Innovative Use of High Performance Computing
 - NERSC Award for High Impact Scientific Achievement
 - Each award has Open Category and Early Career Category
 - Total of 4 awards
- **Awards to be presented at annual NUG meeting in Sept**
- **Help us with nominations please! Details and nominations forms are available at**
 - <http://www.nersc.gov/science/nersc-hpc-achievement-awards/>

Early Career Award Alumni



David Cohen-Tanugi

Cofounder at EMBR labs (MIT)

Tanmoy Das

Assistant Professor of Physics, Indian Institute of Science (IISc)

Edgar Solomonik

Assistant Professor of Computer Science, University of Illinois (UC Berkeley)

Victor Ovchinnikov

Research Fellow of Computational Chemistry and Biophysics, Harvard University (Harvard)

Anubav Jain

Research Scientist in Chemistry, Lawrence Berkeley National Lab (LBNL)

Ken Chen

Research Fellow, National Astronomical Observatory of Japan (UC Santa Cruz)

Taylor Barnes

Grace Hopper Postdoctoral Fellow, Lawrence Berkeley National Laboratory (Caltech)







Nathan Howard

Research Scientist, MIT Plasma Science and Fusion Center (MIT)

Scott French

Google, Inc. (UC Berkeley)

Scratch Filesystem Purging



Scratch Filesystem Purging









- **Purging has started on Cori \$SCRATCH (/global/cscratch1)!**
- Edison /scratch{1..3} has been ongoing
- As a reminder, inactive files will be purged:
 - **Cori** purged after 12 weeks (84 days)
 - **Edison** purged after 8 weeks (56 days)
- Remember to backup your important files
 - **HPSS** is a great resource for long term preservation
 - If you do archive something, and you no longer need it on scratch, please make scratch cleanup a part of your workflow

Scratch Filesystem Purging









- **We're here to help!**
 - If you are having trouble working around the purge schedule, please reach out to us
- TIP: If you see files missing, check for a stub we leave behind. It will be called `.purged.<date>` and will list removed files.
- We're always grateful if you can proactively clean up your unneeded files
- As a reminder, **attempting to circumvent purges is not allowed**. Unless you are actively using your data, please be considerate of other scratch users. The scratch filesystems are a shared resource, and the purge policy is there to ensure we can all have space to do our work.

Edison Update



- **All jobs in the queue from before upgrade will be deleted**
- **All codes must be recompiled**
 - You can compile on Cori now (and during Edison downtime) to get ready for Edison. To compile for Edison from Cori, do
 1. `module swap craype-haswell craype-ivybridge`
 2. `compile as usual`
- **Most job scripts need to be updated**
 - The process/thread/memory affinity control will be changed to be the same as Cori
 - Check Edison running jobs page for updated job scripts
<http://www.nersc.gov/users/computational-systems/edison/running-jobs/>

Cori Allocations, Accounting, & Usage









Cori KNL (Phase 2) Charging



- **Cori KNL production (charging) began on July 1**
- **2,400 M additional NERSC Hours distributed to repos by program managers**
 - <https://www.nersc.gov/users/accounts/allocations/2017-cori-knl/awards/>
 - Check your allocation in NIM
- **“NERSC Hours” can be used anywhere**
 - Cori KNL, Cori Haswell, Edison
 - Cori KNL wait times 3X shorter than Edison, 6X shorter than Cori Haswell
- **Big Job Discount**
 - Cori KNL jobs $\geq 1,024$ nodes get 20% discount
 - Edison discount ends July 31

Data Dashboard



Data Lifecycle Project



- **NERSC project to develop better tools to help users**
 - Find data
 - Manage and share data
 - Archive and stage data
- **First phase deployed**
 - Project filesystem
 - Focused on the most common problems
 - Quota
 - Permission drift

Data Dashboard Demo



- **Data Dashboard**
 - Expand functionality: step through file list, sort by access time, provide sums by directory
 - Working with storage group and Cray to harvest information from scratch file systems
- **Data Management and Sharing**
 - Beta program to enforce group readability in several repo's project directories
- **Data Archiving**
 - Working with vendors to deploy and evaluate tools
- **Something missing? Let us know**

NERSC