


NUG Teleconference April 2013



Richard Gerber
NERSC User Services

April 11, 2013

Connection Info



Connection Info

Topic: NUG Web Conference

Date and Time:

Thursday, April 11, 2013 11:00 am, Pacific Daylight Time (San Francisco, GMT-07:00)

Event number: 291 825 57

Event password: edison

<https://nersc-training.webex.com/> and chose from the list of events.

Teleconference information

1-866-740-1260


PIN: 4866820

- **NUG 2014 Planning: NERSC 40th Anniversary**
- **NERSC Achievement Awards 2013**
- **Edison-forum mailing list**
- **Batch wait time statistics**
- **Edison update: Timeline for Phase 2**
- **Miscellaneous**
- **Special Discussion Topic: Fair Share Scheduling on Edison**

NUG 2014 & NERSC's 40th Anniversary



- We'd like to do something special for NERSC users
- Advice and help from NUG
- Proposal
 - Committee of NUGEX and NERSC staff to plan and execute NUG 2014
 - Volunteers?



- **Propose: Awards Committee**







- Made up of ~ 6 NERSC Staff and ~6 NUGEX members (one representing each office)
- Review and modify nomination procedures
- Manage reviews and make recommendations to NERSC
- NERSC Director will make final decision based on committee recommendations
- NUGEX Chair and Co-Chair select NUGEX committee members



- **109 users signed up**
- **I had hoped users would exchange info, tips, make observations, suggest changes/improvements to Edison**
- **Not much activity**
 - 5 staff postings
 - 1 user posting
- **What do users want?**
 - 105 thought it might be useful
 - What would be?

- **Phase 2 system scheduled to arrive late June 2013**
- **Two weeks of system downtime needed to convert Phase 1 and integrate into additional 24 racks**
- **User access to Phase 2 around August 1**
 - Possibly, a very short early user access period
- **Scratch file systems 2 & 3 have arrived and awaiting installation**
 - One additional 1.6 PB @ 35 GB/sec
 - One 3.2 PB @ 70 GB/sec : for I/O intensive applications


Miscellaneous



Euclid Retirement Fallout?



- **Euclid has been retired.**
 - Is it missed?
 - Is there a need for a visualization / analytics machine separate from Edison/Hopper/Carver?









- **NUGEX has been helping NERSC evaluate requests for startup accounts**
 - Feedback?
 - Is this an interesting way to be more involved with NERSC?
 - Is there a better way to do this?
- **Scheduled Outages Calendar**
 - You asked for it, now you have it!
 - <http://www.nersc.gov/events/nersc-scheduled-system-outages/>
 - You can subscribe to this Google Calendar by clicking a button at the bottom of the page

- **We published instructions at**

<https://www.nersc.gov/users/software/nersc-user-environment/modules/#toc-anchor-3>

- **Install your software in /project or \$HOME**
- **Install the TCL module file (example on web)**
 - Doesn't have to be in same place as software
- **Load your additional module directory tree with "module use ..."**

Fair Share Scheduling



- **What is fair share?**
 - An attempt to schedule jobs so that users/groups get a certain “share” of a machine averaged over a specified time period
 - Your batch job priority depends on your “share” percentage and your recent usage
- **Why fair share on Edison now?**
 - We are obligated to give DARPA Mission Partner users a 25% share of Edison (up to 20M hours)
 - We want to gain experience and evaluate the possibility of adding a fair share component to production scheduling

Fair Share Scheduling Experiment on Edison



- **What are current usage goals?**

- DARPA Mission Partners: 25%
- Each DOE Office gets a share of the remaining 75% based on its relative ERCAP allocation

- **How does it work?**

- These are the major components considered in calculating your job's priority (which is constantly being updated)

Approx.
equal
contributors



- The queue in which you are running
- The job's wait time in the queue (aging)
- **Your office's share and your office's recent (~2 days) usage**
- **We plan to add a user component to keep one user from using up an office's share**

This is the only new part



- **What is the status?**

- Fair share has been in effect for over a month
- But we are still learning and experimenting
- The system described here has been in place for a few days

- **What are the possible implications of fair share from a user perspective?**
 - A more “fair” scheduler from the perspective of hours used over the long term
 - More unpredictable job start times
 - Lower job priority if you have significant recent usage
 - Higher job priority if you haven’t run lately
- **What do you want?**
 - What do you think about fair share?
 - How would you like it to work?
 - And you can’t say “I want my jobs to run when I want them to run!”

Next NUG Teleconference



- **Next scheduled: Thu. May 2, 2013**
- **Send suggested topics and comments to ragerber@lbl.gov**
- **Special Discussion Topic**
 - Case study: Best practices for building applications on Edison based on experiences with material sciences codes



National Energy Research Scientific Computing Center