NERSC Users Group 2013

Warren Focke

2013/02/20

What is NERSC?

- “The primary scientific computing facility for the Office of Science”
- Main facility is a Cray XE6 with ~150k cores
- Cray XC30 under construction
- IBM dataplex, ~10k cores
  - This is where they've stuck me
  - And I'm fine with that
- Also becoming a data repository

- Supercomputers are very different from batch clusters
  - It's all about the network
  - Scheduling is hard – 2-d packing problem
- Focused on MPI
  - Highly parallel
  - They expect you to run on ~1e4 cores
  - For short times ~1hr
- It's common for long jobs to checkpoint & restart

- Do I/O from a single node, distribute via MPI calls
- Monte-Carlo or event reconstruction could run well on this sort of setup
  - It would require some adjustment in thinking
  - And would probably underutilise the network

- They also have some more batch-like systems
  - But you have to pay for those
  - Staff is aware that some classes of users would like to see more of this
- They are will work with customers to customize solutions
- Planck got custom-managed hardware by pledging 2.5 FTEs to use it

- NEWT – RESTful API allowing access of NERSC resources through the web
  - Staff has used it to implement a lot of their own management toools
- Users are doing some really cool stuff with it
- Science gateways
  - QCD
  - Astronomy
  - Materials science
  - Climate

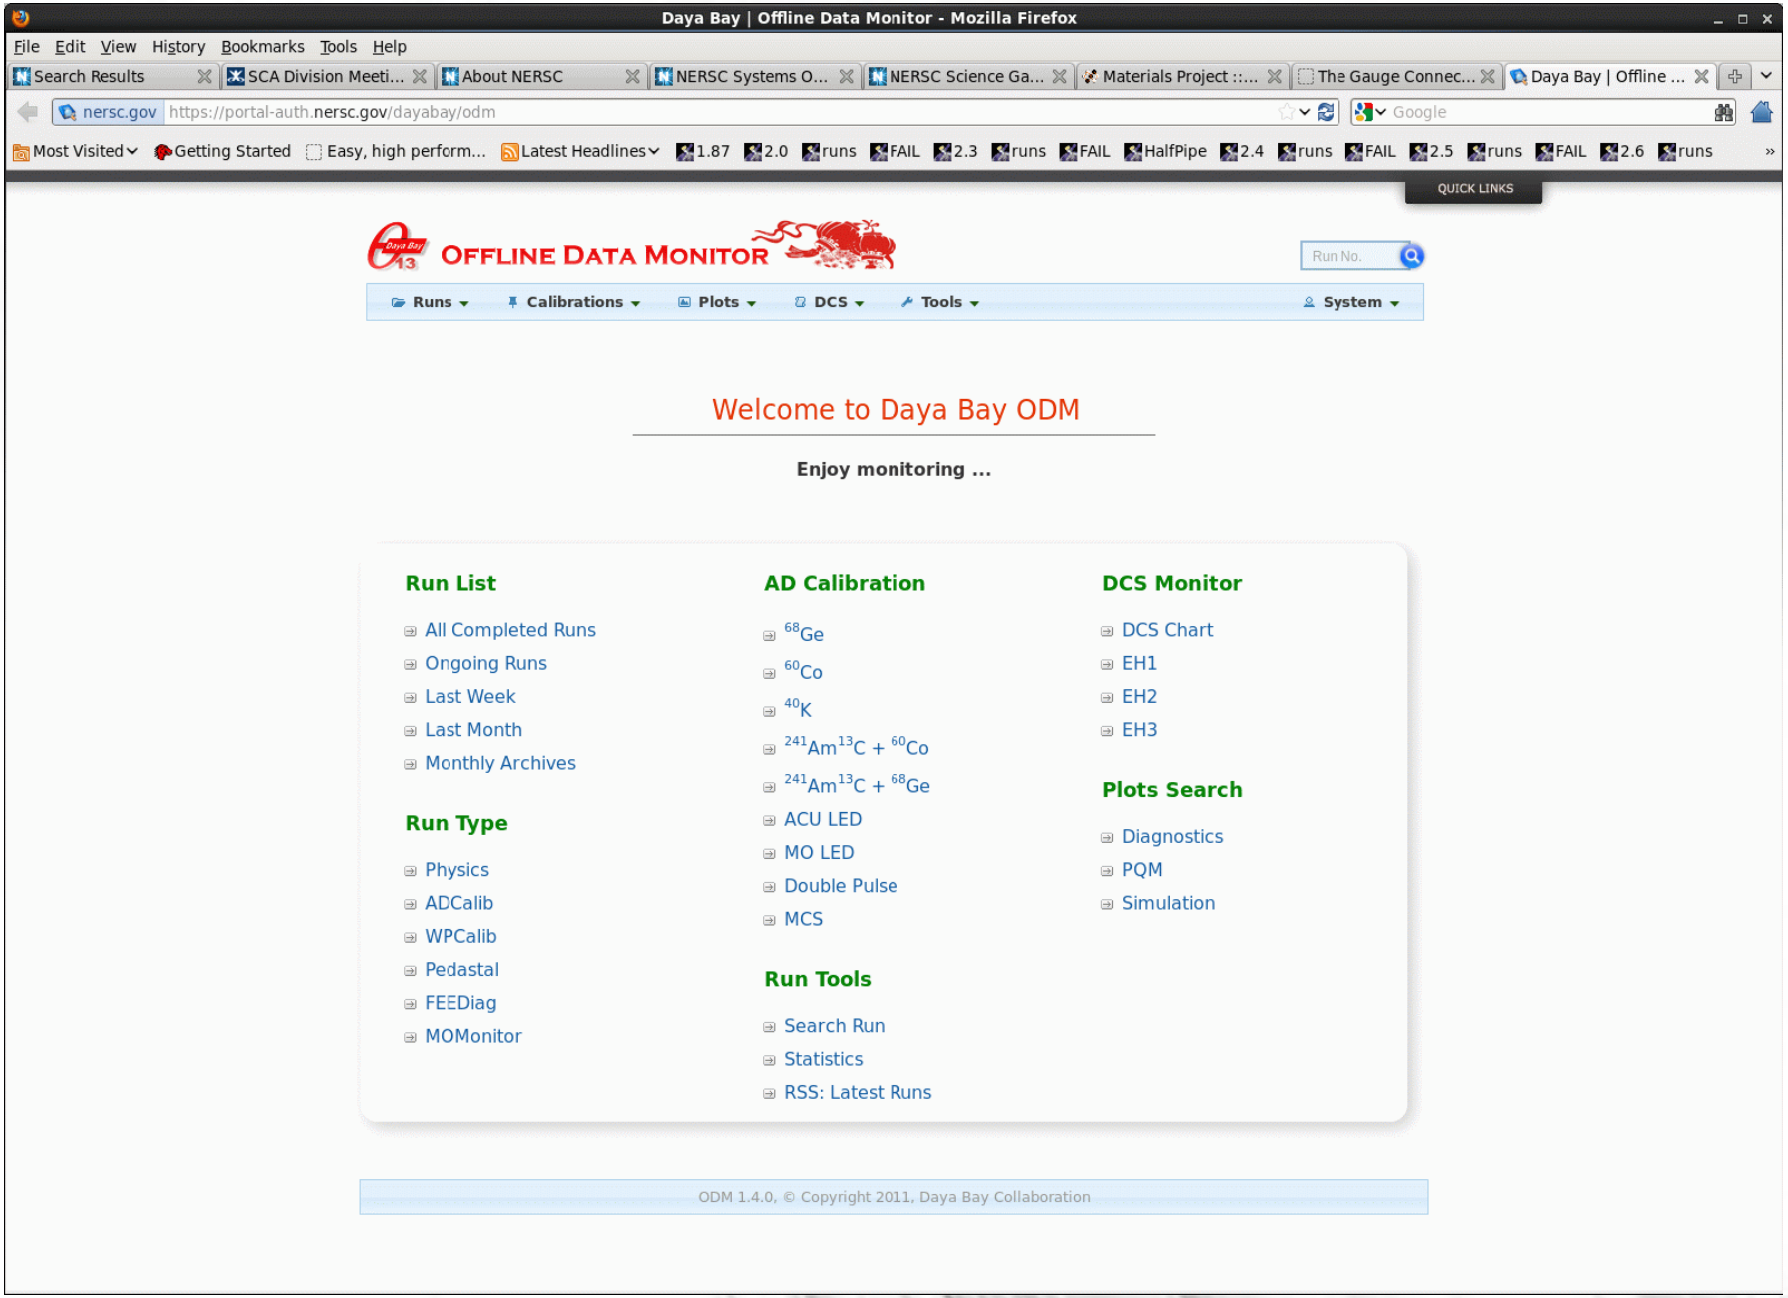
Daya Bay | Offline Data Monitor - Mozilla Firefox
File Edit View History Bookmarks Tools Help
Search Results
SCA Division Meeti...
About NERSC
NERSC Systems O...
NERSC Science Ga...
Materials Project ::...
The Gauge Connec...
Daya Bay | Offline ...
nersc.gov
https://portal-auth.nersc.gov/dayabay/odm
Google
Most Visited
Getting Started
Easy, high perform...
Latest Headlines
1.87 2.0 runs FAIL 2.3 runs FAIL HalfPipe 2.4 runs FAIL 2.5 runs FAIL 2.6 runs
QUICK LINKS
OFFLINE DATA MONITOR
Run No.
Runs
Calibrations
Plots
DCS
Tools
System
Welcome to Daya Bay ODM
Enjoy monitoring ...
Run List
All Completed Runs
Ongoing Runs
Last Week
Last Month
Monthly Archives
Run Type
Physics
ADCalib
WPCalib
Pedastal
FEEDiag
MOMonitor
AD Calibration
68Ge
60Co
40K
241Am13C + 60Co
241Am13C + 68Ge
ACU LED
MO LED
Double Pulse
MCS
Run Tools
Search Run
Statistics
RSS: Latest Runs
DCS Monitor
DCS Chart
EH1
EH2
EH3
Plots Search
Diagnostics
PQM
Simulation
ODM 1.4.0, © Copyright 2011, Daya Bay Collaboration


Home
Apps
Resources
About
References
Login or Register
MATERIALS PROJECT
A Materials Genome Approach
Accelerating materials discovery through advanced scientific computing and innovative design tools.
e.g. explore Fe2O3 or Li-Fe-O pd
Search
powered by MOOGLE
Database Statistics
30732 materials
3044 bandstructures
405 intercalation batteries
15175 conversion batteries
Register now for free, full access.
Unlimited access
Up to 500 search results
History of your searches and analyses
Or try the apps in demo mode
10 minute usage limit
Search results limited to 10 best matches
Just click an app to start
We're hiring! See our new opportunities for postdocs and web developers!
Materials Explorer
Search for materials information by chemistry, composition, or property.
Lithium Battery Explorer
Find candidate materials for lithium batteries. Get voltage profiles and oxygen evolution data.
Crystal Toolkit
Convert between CIF and VASP input files. Generate new crystals by substituting or removing species.
Phase Diagram App
Computational phase diagrams for closed and open systems. Find stable phases and study reaction pathways.
Reaction Calculator
Calculate the enthalpy of tens of thousands of reactions and compare with experimental values.
Structure Predictor
Predict new compounds using data-mined substitution algorithms.
Find out more about our open Materials API and pymatgen library for querying large amounts of data.
Tutorials
MATERIALS PROJECT Tube
Press Highlights
The New York Times
Beyond Fossil Fuels: Finding New Ways to Fill the Tank
The end of the petrolhead: Tomorrow's cars may just plug in
IMPROPER BOSTONIAN
The Big Picture - Material Good
ascr discovery
Building a better battery proves Material Genome Project's metal
all press
Latest News
pymatgen v2.5.2
by Shyue Ping Ong - 22 hours ago
Version 2.5.2 of pymatgen has been released! Check out the change log at https://pypi.python.org/pypi/pymatgen for information on the latest changes.
pymatgen v2.5.2
The Materials Project is hiring!
A new year, a new pymatgen
Journal article on the Python Materials Genomics library
all entries
subscribe
Materials Project is on Facebook and
Like
Send
437 people like this. Sign Up to see what your friends like.
+153 Recommend this on Google
How to cite the Materials Project

- LSST DESC was granted 2.5M hours
- Initial use is raytracing the telescope
  - Lots of variation in job runtimes makes it hard to schedule
- Hoping to use NEWT to run pipeline jobs at NERSC, controlled from SLAC
- Will need to set up parallel jobs that act as virtual batch farms
  - e.g. Reserve 32 nodes for 24 hours
  - Use a (somewhat experimental) queueing system taskfarmerMQ