BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

LBNL Site Report
SLURM User Group 2022

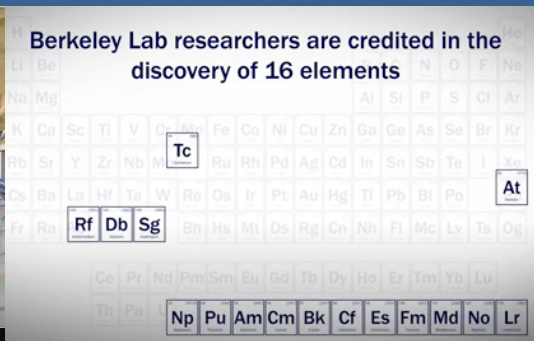

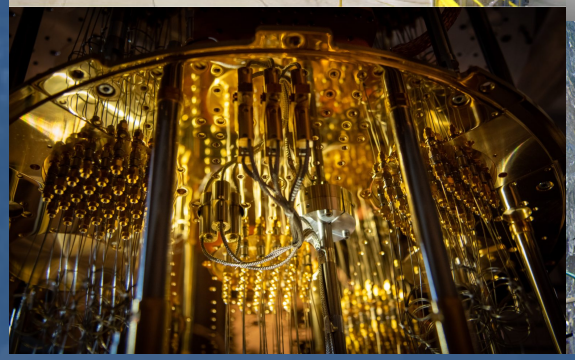

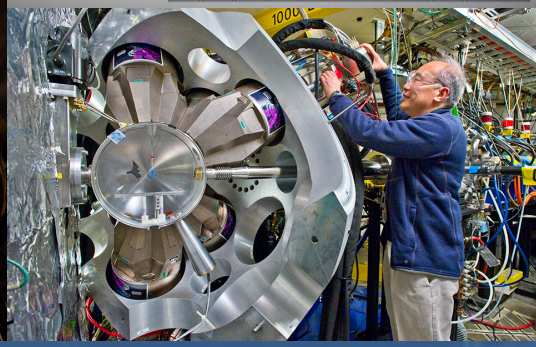

Wei Feinstein
High Performance Computing
Scientific Computing Group

Acknowledgment

Nicolas Chan
Shawfeng Dong
SchedMD support team

BERKELEY LAB

Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

Founded in 1931 by Nobel Prize winning physicist Ernest Lawrence
Sixteen elements are discovered at the Lab
Nine scientists have been awarded the Nobel Prize
Operates major National User Facilities for the Office of Science DOE

BERKELEY LAB

Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

How Molecular Foundry scientists model at the nanoscale

"Nanoscale" refers to particles and processes at a certain size. **The Molecular Foundry is one of five national user facilities for nanoscale science research around the country that provides state-of-the-art instruments and expertise to users from all over the world.** These centers bring together people of various specialties to work and interact together in this multidisciplinary field.

The Foundry's Theory Facility works to develop models of nanoscale materials or phenomena that can be tested using computational experiments run on computers. If the model is a good enough simulation of reality, then these computational experiments can sometimes serve as less labor-intensive and more detailed studies than actual experiments.

"Much like an chemist needs a lab...we need computing infrastructure. We need lots of fast processors. We need to connect them with high-speed networks. We need efficient compilers and libraries to make the most of that hardware," said David Prendergast, director of the Foundry's Theory of Nanostructured Materials Facility

READ THE FULL ARTICLE

How the Materials Project connects computational and experimental materials

"You can think of the computer as a virtual lab where I can do the same processes, but at a more rapid pace. I can get results very quickly. I can fail fast so I can try over and over and over again until I find that one combination that is going to work well," said Shyam Dwaraknath, a materials research scientist in the Applied Energy Materials Group under Materials Project co-founder and director Kristin Persson.

Since launching in 2011, the **Materials Project**, a Department of Energy program based at Berkeley Lab, has been on the forefront of materials discovery with a giant, searchable repository of data available to the whole materials science community.

Using computational materials science, which integrates supercomputers, advanced mathematics, and quantum mechanics, researchers can virtually simulate thousands of compounds every day to find the best candidates to test in the laboratory.

SCIENTISTS CAN EXPLORE THOUSANDS OF COMBINATIONS OF A MATERIAL IN THE COURSE OF A DAY USING HIGH-PERFORMANCE COMPUTING. Looking ahead, the Materials Project is integrating machine learning to train computers to "see" (model) materials and molecules the way a human does.

READ THE FULL ARTICLE

Using computation to uncover how Arctic soil microbes respond to the changing climate

"Soil has amazing amounts of microbial cells," said **Neslihan Taş**, a research scientist in Berkeley Lab's **E**arth and Environmental Sciences Area. The microbiomes of different soils around the world vary widely. Taş' focus is Arctic soils — specifically permafrost. Her research consists of three parts: fieldwork, sequencing, and analysis.

For her sequence analysis, Taş partnered with the ScienceIT to leverage Lawrencium, the Lab's institutional cluster, and her own dedicated computational server, also managed by ScienceIT.

On the testing side, ScienceIT sets the groundwork by helping Taş troubleshoot her scripts so that her group can conduct trials of the data analysis. The group works in constant communication with her to develop frameworks for streamlining computational pipelines. TAŞ CALLED SCIENCEIT HER "LIFELINE."

READ THE FULL ARTICLE

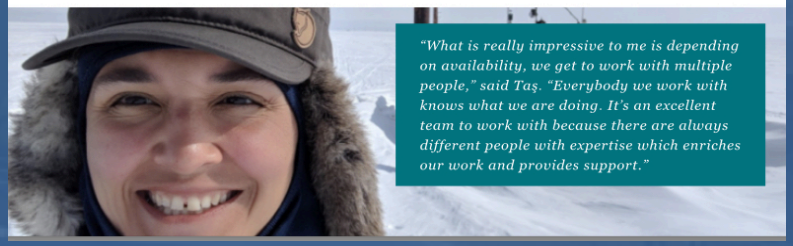

"What is really impressive to me is depending on availability, we get to work with multiple people," said Taş. "Everybody we work with knows what we are doing. It's an excellent team to work with because there are always different people with expertise which enriches our work and provides support."

Autofocus for X-ray Crystallography: How AutoML Targets Samples at the ALS

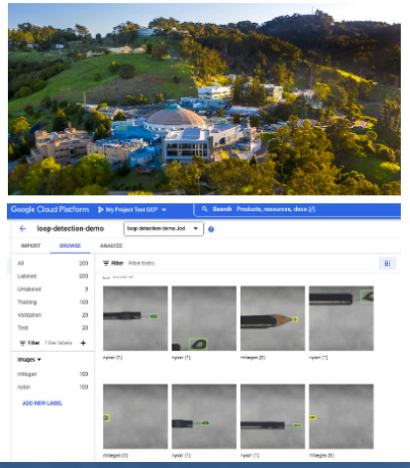

AUTOMATED MACHINE LEARNING (AUTOML) FRAMEWORKS AIM TO AUTOMATE TASKS SO NON-EXPERTS CAN TAKE ADVANTAGE OF MACHINE LEARNING ON A LARGE SCALE. Macromolecular Crystallography is the use of X-rays to study the structures of proteins, DNA, and RNA at an atomic resolution. The technique relies on alignment of crystal samples within high energy X-rays beams which are produced at the synchrotrons, such as Berkeley Lab's **Advanced Light Source (ALS).**

This step has historically been performed through a manual click-to-center strategy or traditional machine vision strategies based on edge detection and other rigid procedural techniques. **Dr. Scott Classen,** a Biophysicist Research Scientist in the Molecular Biophysics and Integrated Bioimaging Division, devised a method to automate this process of X-raying samples using LoopDHS, a custom machine learning (ML) program .

READ THE FULL ARTICLE

- Reduce computing barriers in their breakthrough research
- Science-centric and comprehensive approach

OPEN OnDemand

MyLRC - Laboratory Research Computing Access Management System

Google Cloud Platform

aws

globus
a uchicago non-profit service

DDN STORAGE

ceph

DELL Compellent

WEKA

VAST

High Performance Computing (HPC)

HPC Systems

- Lawrencium supercluster: ~1300 compute nodes
- Over 10 division clusters: ~1700 compute nodes
- Multiple generations of CPUs, GPUs, large memory and AMD nodes
- Interconnected with EDR, FDR InfiniBand
- High performance NFS storage for home and projects
- 2.5 PT Lustre parallel filesystem for scratch space

BERKELEY LAB

Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

High Performance Computing (HPC)

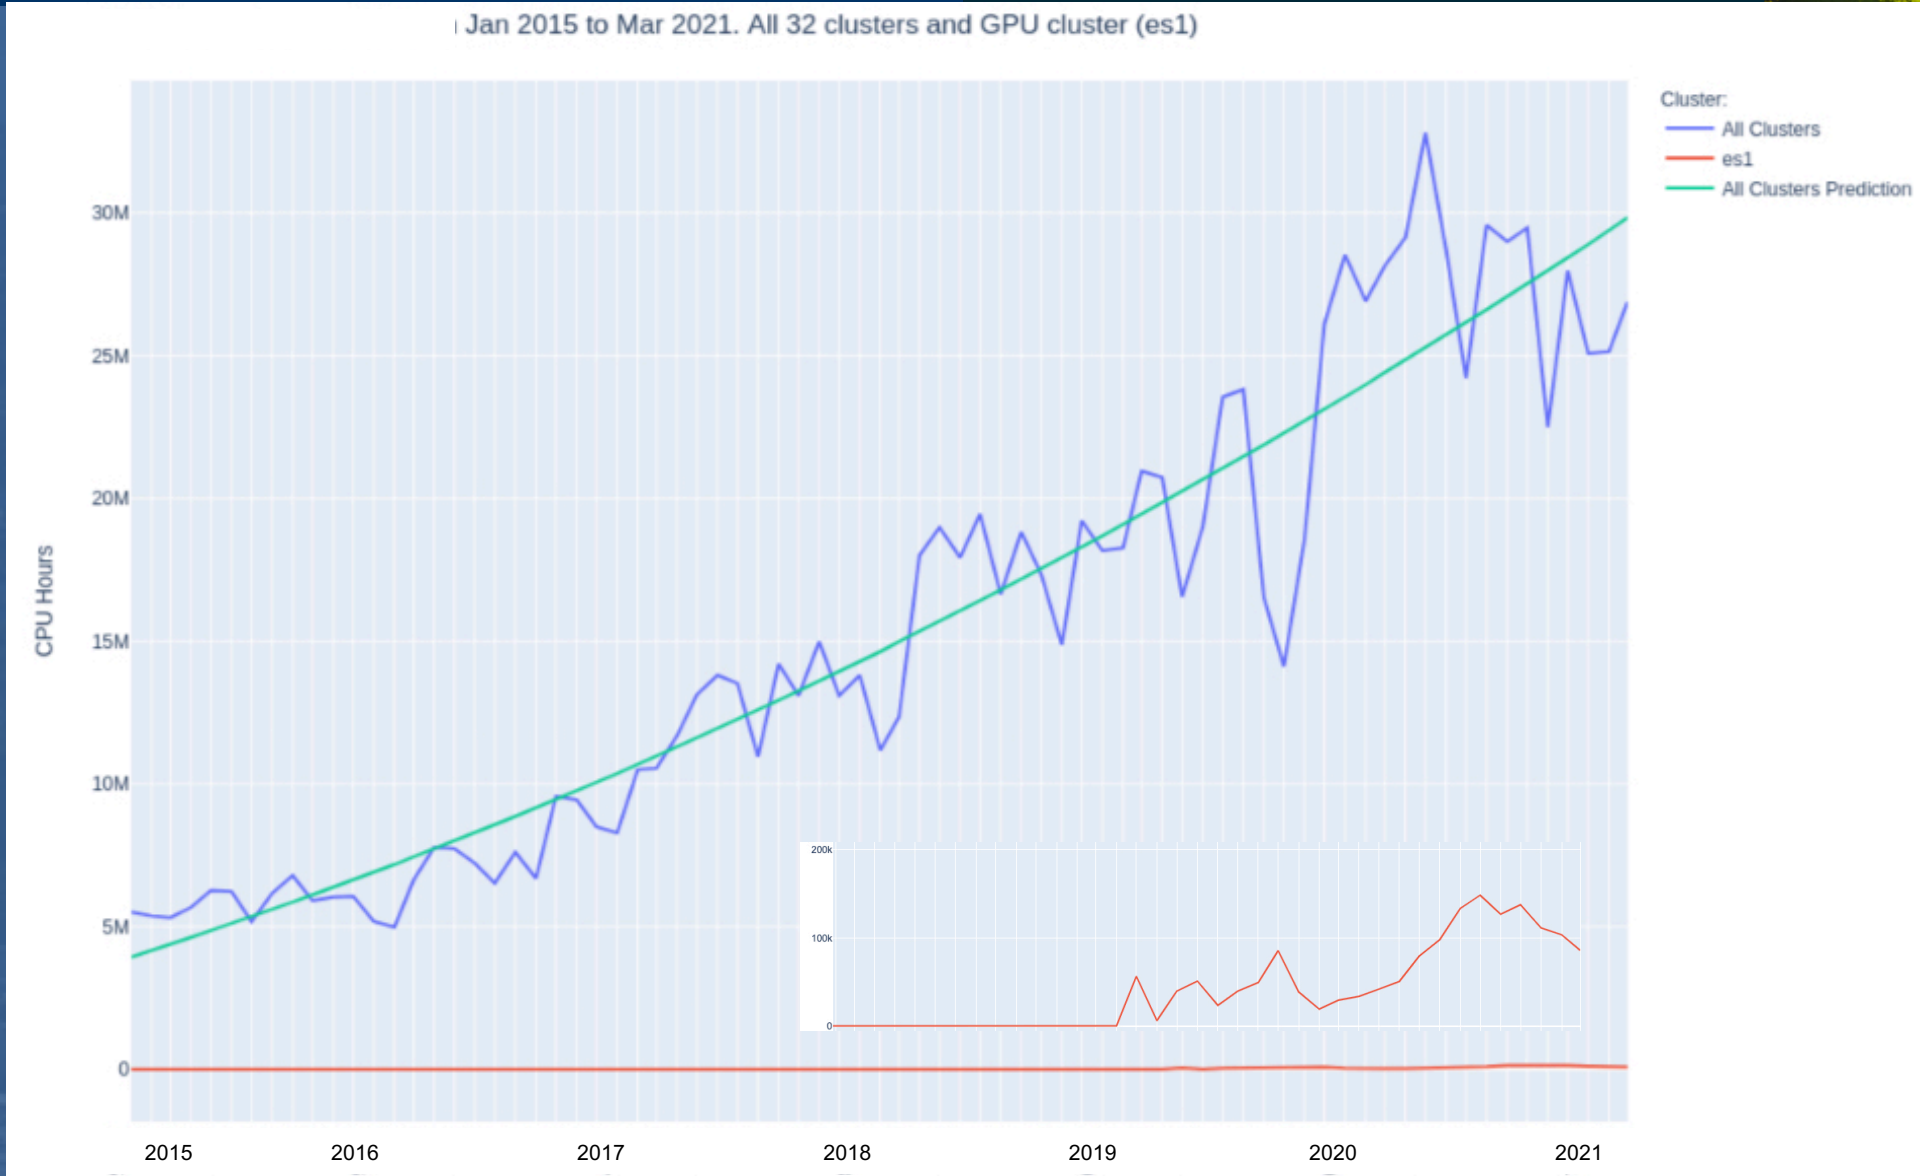

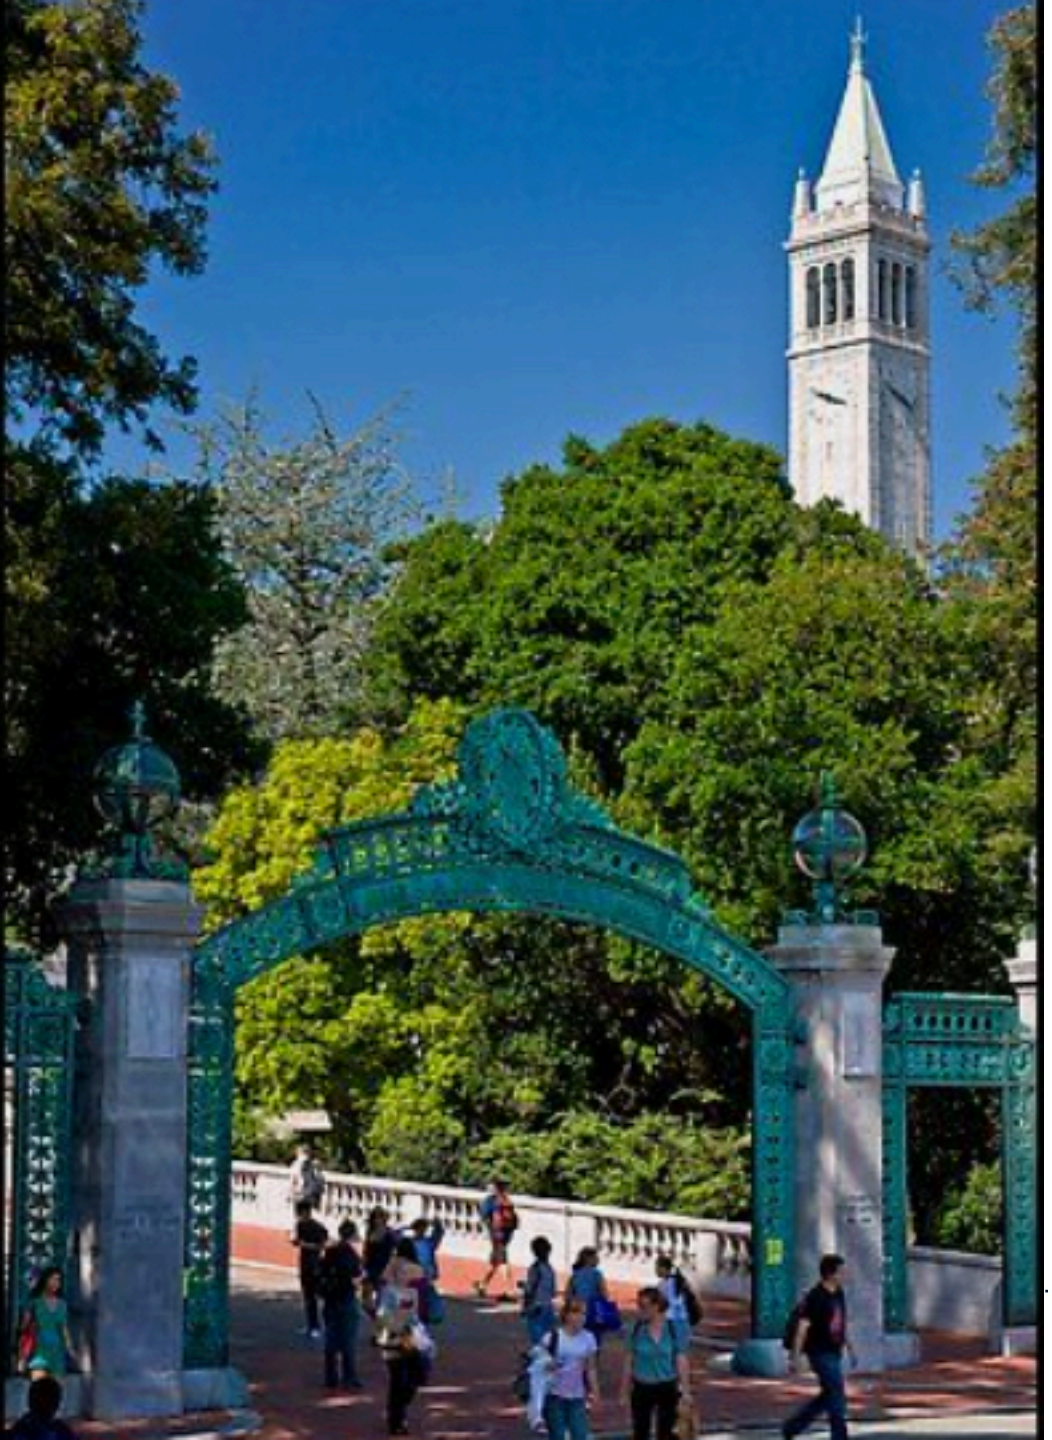

- Manage HPC service for UC Berkeley Research Computing program
- Provide Berkeley faculty, researcher, students with state-of-the-art computing systems of Savio and SRDC (secure computing)
- Auxiliary benefits:
  - Improved grant competitiveness
  - Incentive for recruitment and retention
  - Achieving significant economies of scale with centralized data center

Lawrencium - Institutional Condo Program

Three types of accounts

- Annual allocation account: free for lab PIs
- Condo account: hardware purchase by PIs
- Recharge account: SU purchase

Lawrencium SLURM Configuration

- LRC partitions: lr[3, 4 5, 6], lr_bigmem, es1 (GPU), cm (AMD) …
- Condo QoS: customized based on PIs' preference
- General QoS: lr_normal, lr_debug, low_prio
- QoS access
  - Condo users: condo_qos to the partition they buy in, low_prio to all partitions.
  - Free 300K SU users: lr_normal to all partitions
  - Recharge: lr_normal and low_prio to all partitions

BERKELEY LAB

Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY

High Performance Computing (HPC)

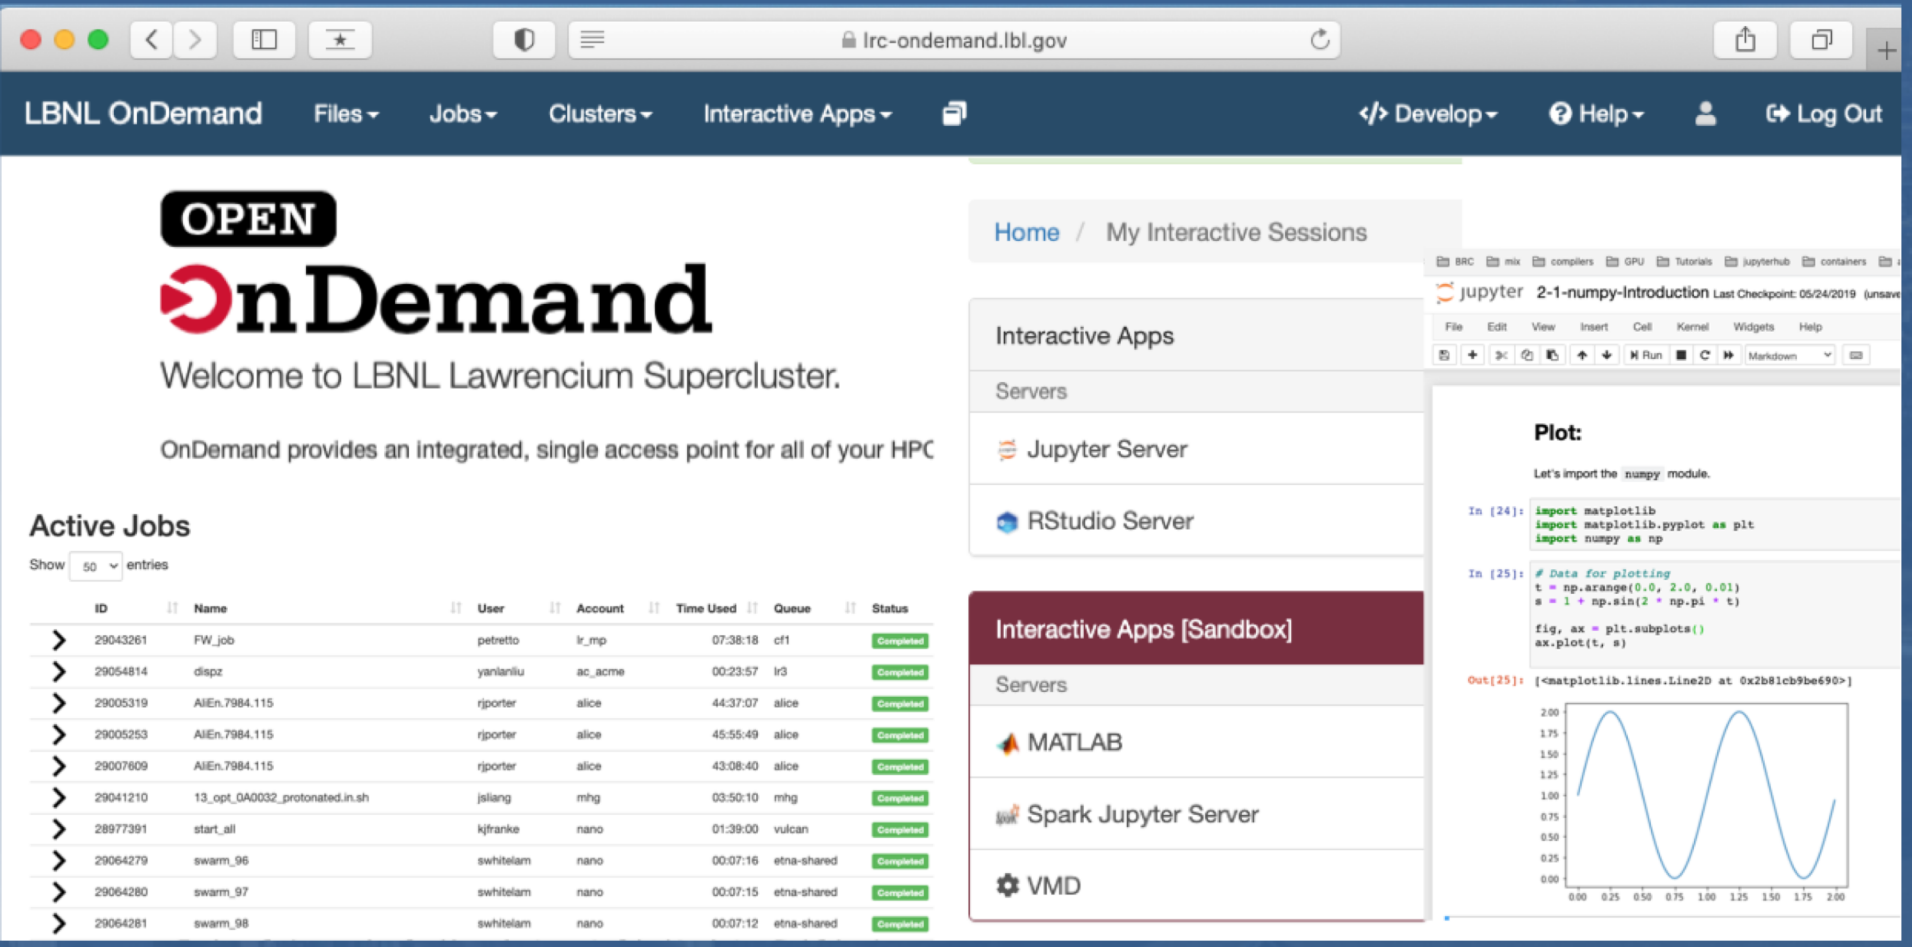

BERKELEY LAB

Bringing Science Solutions to the World

MyLRC - Laboratory Research Computing Access Management System

BERKELEY LAB
Bringing Science Solutions to the World

Home Center Summary Help Log In

Welcome to MyLRC

MyLRC is a user portal for managing access to the clusters and other resources provided by the Laboratory Research Computing (LRC) program.

- Join projects and gain access to Lawrencium and other clusters.
- Create new projects and manage project users.
- Request or purchase computing allowances.
- View details of current and past jobs, and allowance usages.
- And more!

For more information, refer to our documentation.

Berkeley Lab: Log In | Other: Log In

Science IT
High Performance Computing

Powered by ColdFront Version 1.0.1 | GitHub

LRC HPC Resources Scientific Impact

Active Allocations and Users by UCB Campus Divison or Department

Show 10 entries Search:

UCB Campus Division or Department	Active Allocation Count	User Count
Other	295	2396

Showing 1 to 1 of 1 entries Previous 1 Next

Total Active Users: 1268
Total Principal Investigators: 203

Resources and Allocations Summary

Show 10 entries Search:

Resource Name (Type)	Active Allocation Count
LAWRENCIUM Compute **(Cluster)**	283
ALICE Compute **(Cluster)**	1
ALSACC Compute **(Cluster)**	1

Job List

Viewing only jobs belonging to you and belonging to projects in which you are a PI or manager. To view all jobs select "Show All Jobs" in the search form below and search, or click here.

Filter +

Export Job List to CSV

Slurm ID	Username	Project	Job Status	Partition	Submit Date	Service Units
1001450	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.44
1001449	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.02
1001448	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.02
1001447	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.02
1001446	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.00
1001445	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.06
1001442	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.01
1001441	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.18
1001440	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	0.74
1001439	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 06, 2022	1.92
1001438	wfeinstein	ac_scsguest	COMPLETED	slmtest	Sep. 05, 2022	8.02

Home Center Summary Project **Requests** Jobs Admin Help wfeinste

- MyLRC Administration
- All Allocations
- All Jobs
- All Projects
- All Users
- Project Reviews
- User Search
- Requests
  - All Requests
  - Cluster Access Requests
  - New Project Requests
  - Project Join Requests
  - Project Removal Requests
  - Project Renewal Requests

Request Hub

Below are all the requests in MyLRC. Click on a request s…t type. To perform actions on a specific request, click the butto… go to the request's main page and perform the actions th…ur requests.

Collapse All − Expand All +

Cluster Access Requests 4 pending reques

Pending Cluster Access Requests

#	Request Time	User Email	Cluster Username			User	Billing ID	Status
4	Aug. 16, 2022	fengchenliu@berkeley.edu	No cluster account.				N/A	Pending
6	Aug. 16, 2022	fengchenliu@lbl.gov	fengchenliu	ac_scsguest	445	N/A	N/A	Pending
8	Aug. 23, 2022	meli@lbl.gov	meli	pc_newone	622	N/A	108276-002	Pending
9	Aug. 23, 2022	meli@lbl.gov	meli	ac_mp	417	N/A	108276-002	Pending

Page 1 of 1 Previous Next

Completed Cluster Access Requests

#	Request Time	User Email	Cluster Username	Project	Allocation	Host User	Billing ID	Status
1	Aug. 05, 2022	wfeinstein@lbl.gov	wfeinstein	ac_scsguest	445	N/A	N/A	Complete

Page 1 of 1 Previous Next

Job Detail: 1001450

Job Information

Slurm ID:	1001450
Username:	wfeinstein
Project:	ac_scsguest
Job Status:	COMPLETED
Submit Date:	Sept. 6, 2022, 12:18 p.m.
Start Date:	Sept. 6, 2022, 12:18 p.m.
End Date:	Sept. 6, 2022, 12:21 p.m.
Partition:	slmtest
Nodes:	n0001.slmtest0
Service Units :	0.44
Quality of Service:	normal
Number of CPUs:	8
Number of Required Nodes:	1
Number of Allocated Nodes:	1
Raw Time (seconds):	0.054444443
CPU Time (seconds):	0.43555555

User Portal SLURM Banking Plugins

job_submit_plugin (job submission): Estimate maximum job cost based on submission parameters, and reject job if the API reports that the user/account has insufficient service units available.

spank_plugin (job running): Report job and estimated cost to the API.

job_completion_plugin (job completing): Modify job in API to reflect actual usage.

Portal Plugins Written in Rust

To ensure safety memory management

OpenSSL to make HTTPS connections to the portal API

Compiled with the SLURM source code

Caveat: potential plugin code revision with SLURM upgrade

Plugin Configuration

Enable submit andcompletion plugins
/etc/slurm/slurm.conf
...
JobSubmitPlugins=job_submit/slurm_banking
JobCompType=jobcomp/slurm_banking

Enable spank plugin
/etc/slurm/plugstack.conf
…
optional /etc/slurm/spank/spank_slurm_banking.so

Configure our plugins using /etc/slurm/bank-config.toml
To enable or disable the plugins
Partition names and pricing
URL of the user portal API, API token

More SLURM Plugins

cpu_gpu_ratio_plugin (job submission): ensure the ratio (gres=1, --ntasks=2) is set appropriately

spank_private_tmpshm: create per-job tmp (/tmp, /var/tmp and shm /dev/shm)

spank_collect_script: collect job submission scripts

Resources

Portal Slurm plugins: https://github.com/ucb-rit/slurm-banking-plugins

More information, documents, tips of how to use Lawrencium supercluster http://scs.lbl.gov/

Contact us at hpcshelp@lbl.gov scienceIT@lbl.gov