SubMIT News: Status, & Updates

David Walter

11th February 2025

Physics Basic Computing Services (subMIT)

Users Meeting

Introduction

Our mission

- Provide basic computing services in the MIT physics department
- Enable easy access for newcomers to start their physics analysis
- Support advanced customization for experiences users
- Ensure sufficient and efficient computing resources through fair sharing

We do the maintenance such that you don’t have to care about

- System configuration, upgrades, security
- Software installation and management
- Integration with external resources and services

So you can focus on doing great physics

Project organization

Steering committee

- Oversight
- Funding

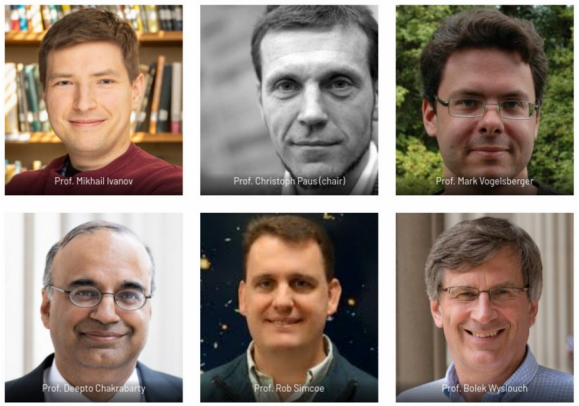

Project team

- Implementation
- Operation
- Maintenance
- Support

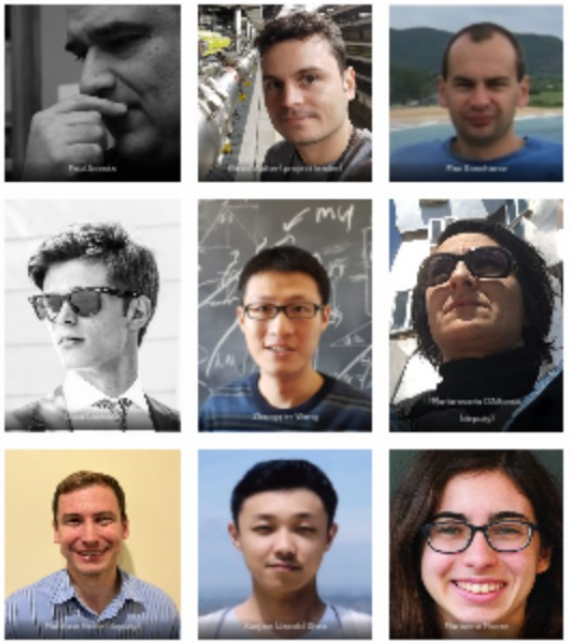

Users group

- Information flow between user community and project team
- Feedback
- Requests

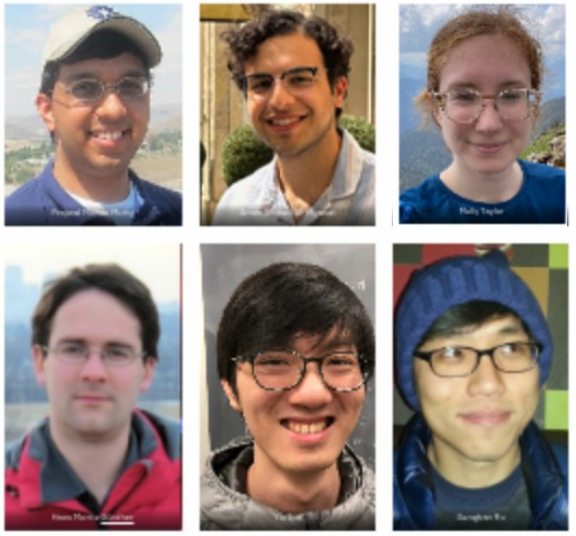

Resources

Local batch system

- O(1000 cores); ~30GPUs
- Interactive SSH login pool
- Jupyterhub access

Convenient software environment

- Alma Linux 9 native
- Python, C++, Java, ...
- Containers (singularity/ podman)
- Virtual environments (Conda)

Access to larger external resources

- Open Science Grid (OSG)
- CMS Tier-2 and Tier-3
- LQCD Cluster
- Earth, Atmospheric & Planetary Sciences (EAPS)

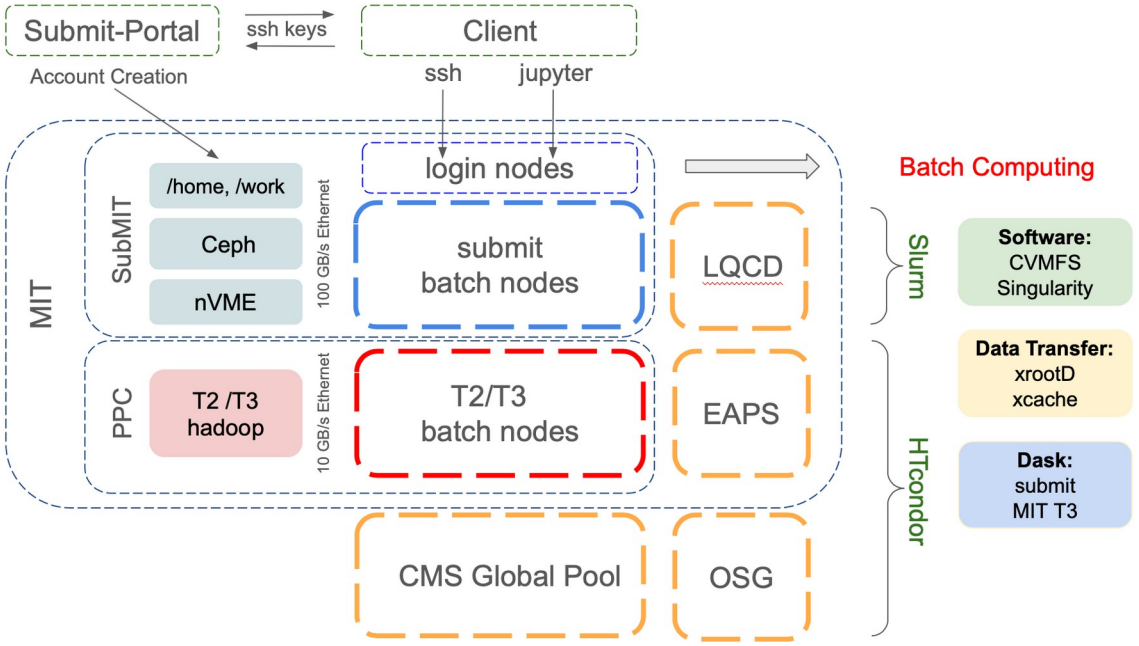

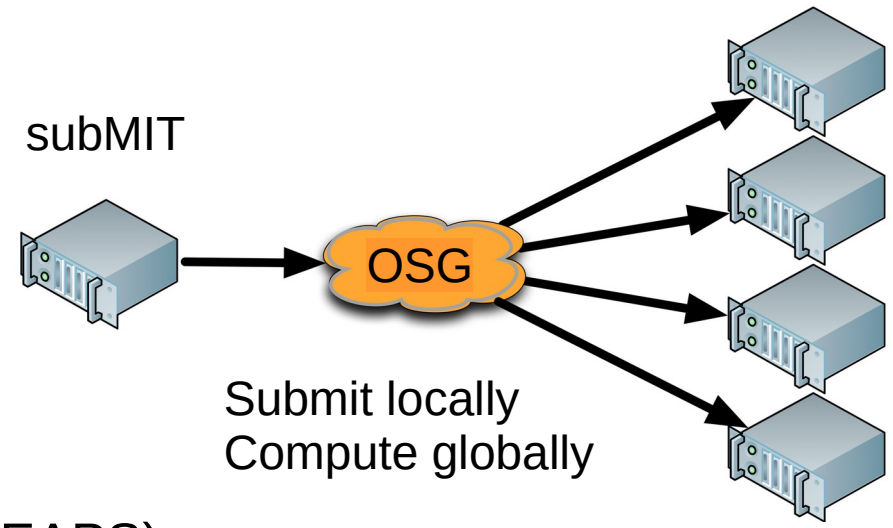

SubMIT overview

Website

- Overview and general information
- Account creation
- Documentation: User guide
- A2rchi (chatbot)
- Monitoring systems
- Direct JupyterHub access

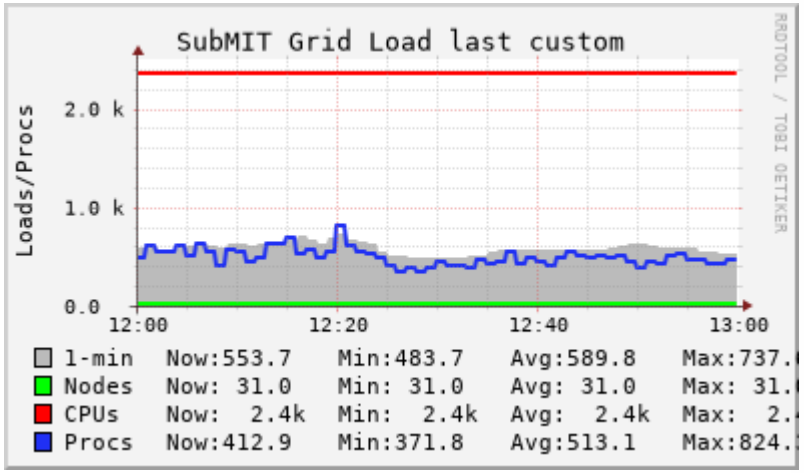

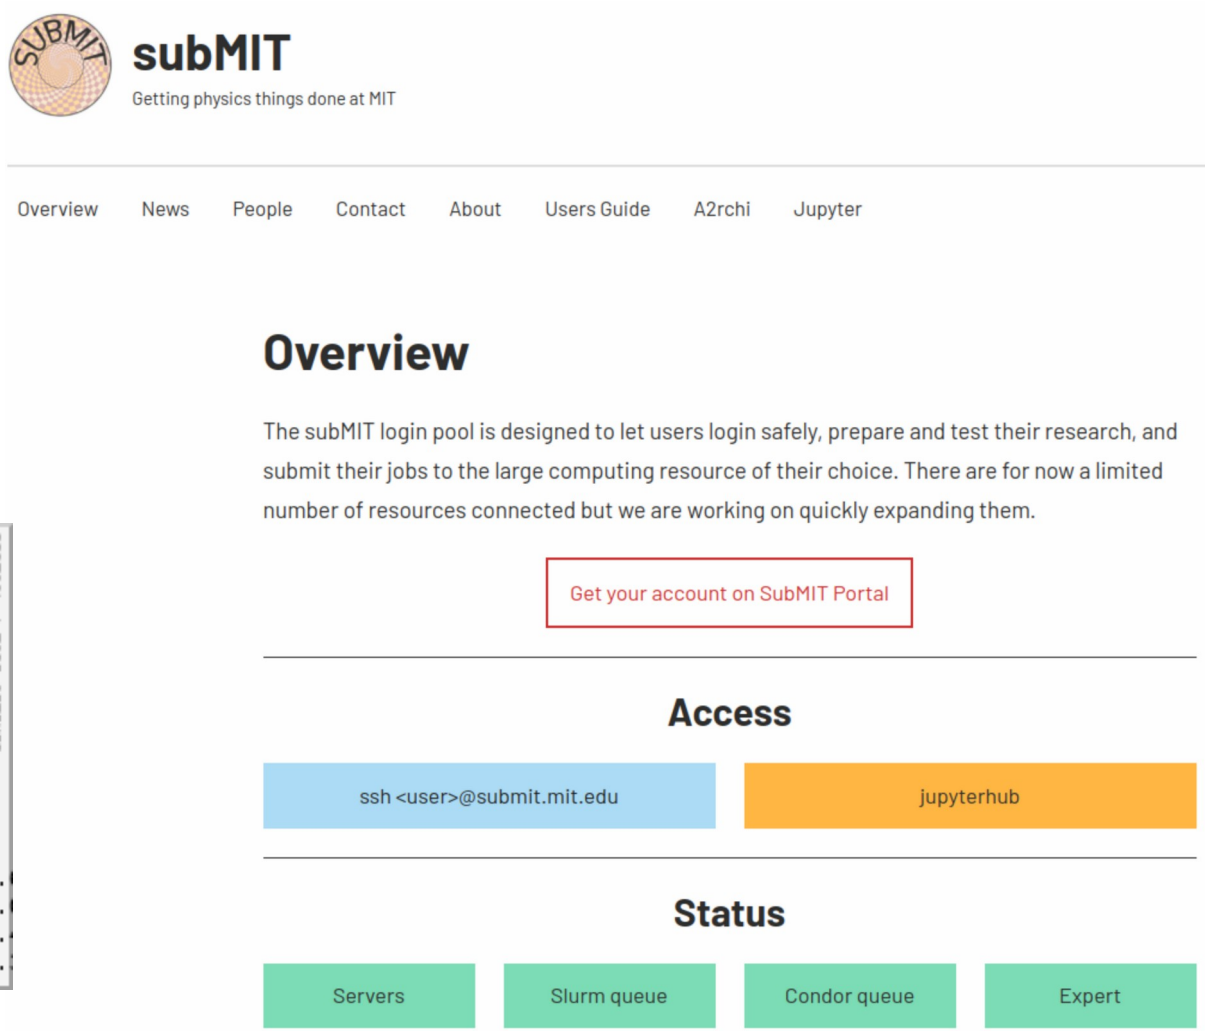

User support

User support is a key feature of the system

- Contact: submit-help@mit.edu
- Slack workspace: https://mit-submit.slack.com/
  - “help-desk” channel

Beyond basic troubleshooting

- Help users make optimal use of the available resources
- Expert advice on designing/improving workflows
- Customize and evolve system configuration, accommodate user needs as appropriate

Voluntary (anonymous) survey

- Please tell us what you like or don’t like

Status

Diverse user base

- From different departments
- From undergrad to professors

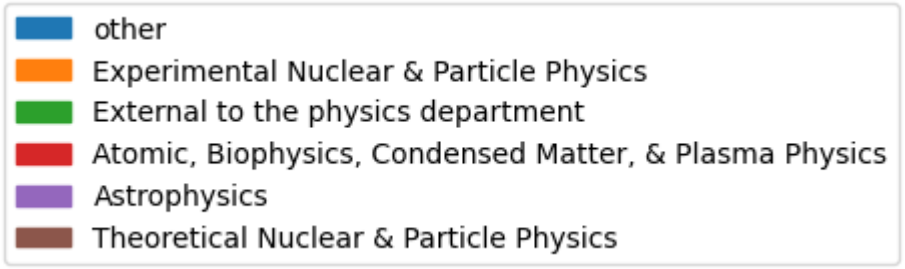

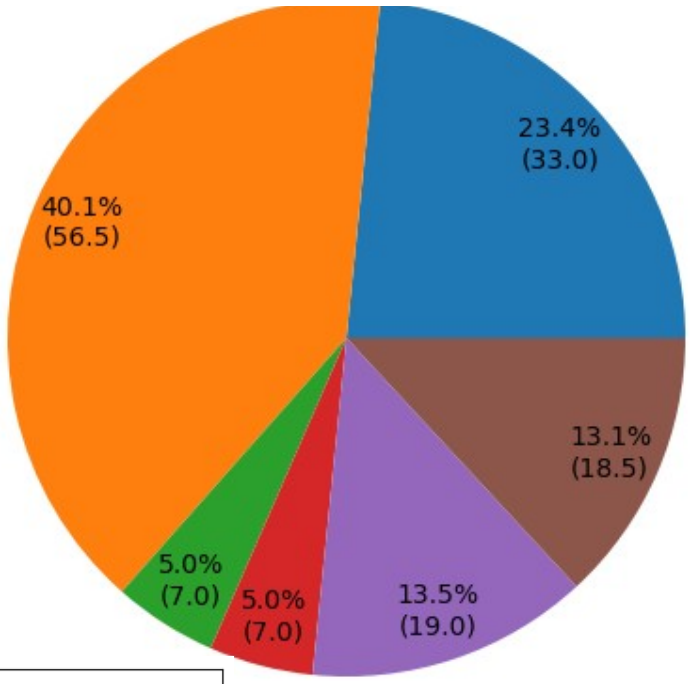

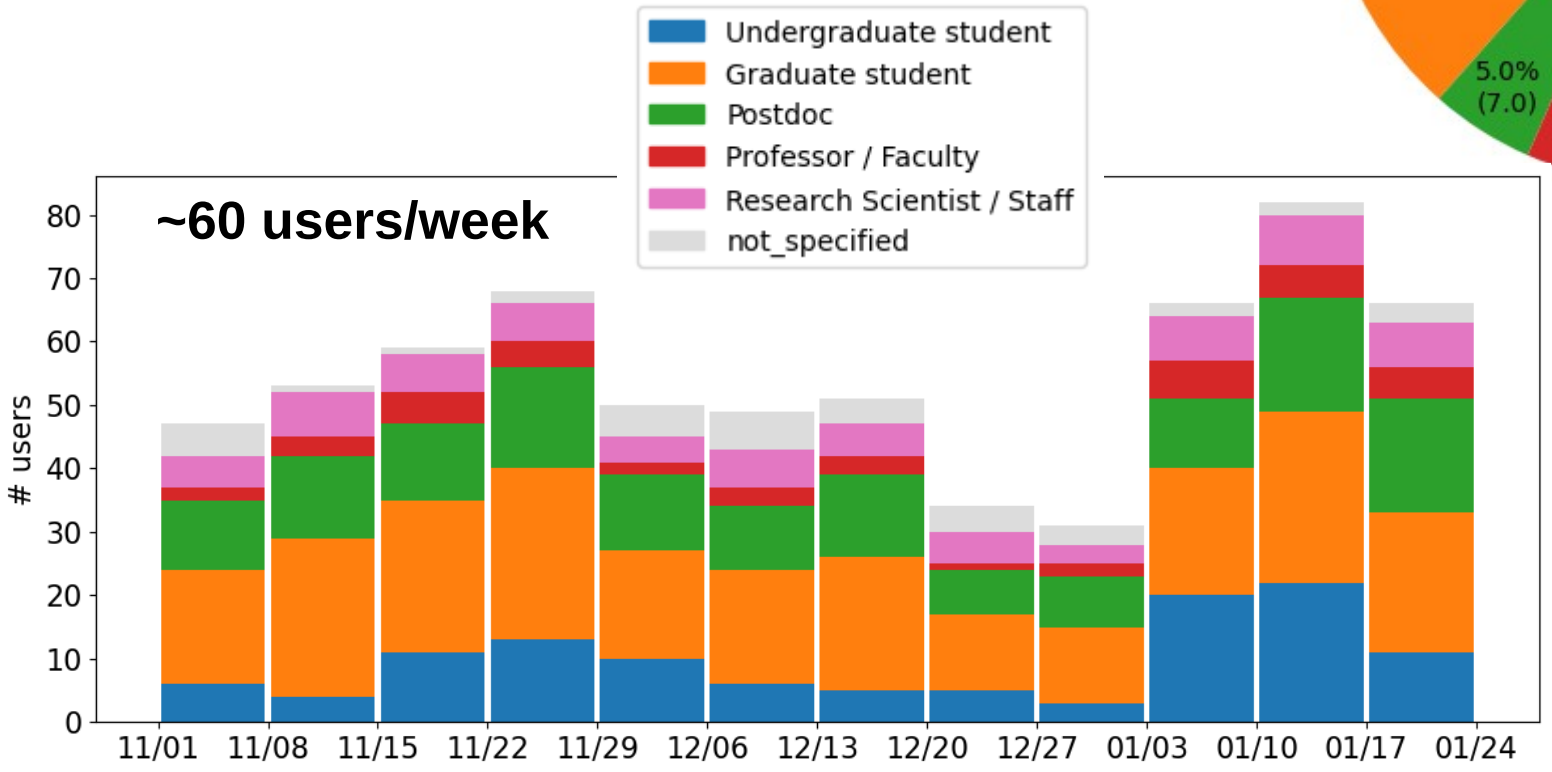

News: Deactivating login for old users

Established practice to deactivate and remove old users

- Periodically deactivate login for old users in January each year
- Remove SSH keys for users w/o MIT touchstone
- First time this year: 28 users

Old users can restore access through sponsors

If we don’t get contacted, data of deactivated users is removed after 6 month

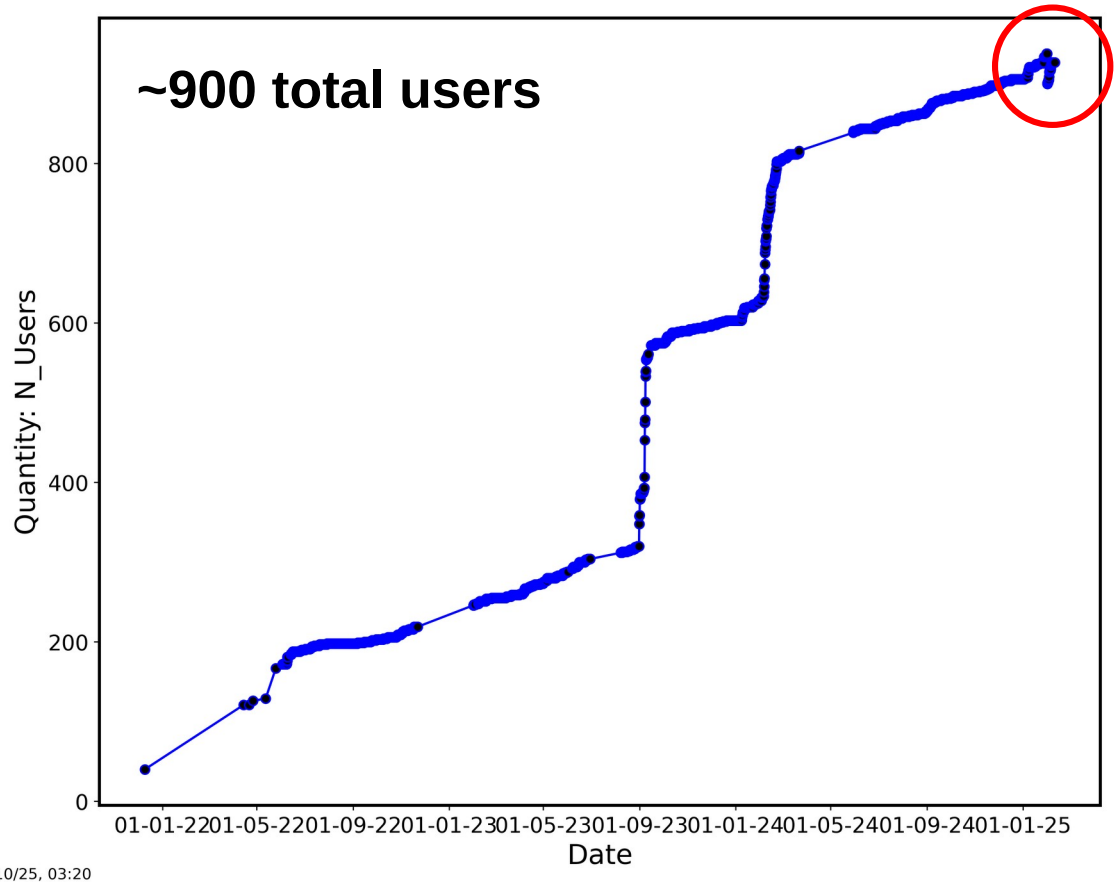

News: Website

Added list of publications that used significant resources for subMIT

- submit.mit.edu

Please let us know if your publication is missing

subMIT

Getting physics things done at MIT

Overview | News | People | Contact | About | Publications | Users Guide | A2rchi | Jupyter

Publications

Using significant subMIT resources

2024

- **A. Belley, J. Pitcher, T. Miyagi et al.**,
  Correlation of neutrinoless double-beta decay nuclear matrix elements with nucleon-nucleon phase shifts,
  2024. arXiv.
  Division: Theoretical Nuclear & Particle Physics | **Center:** LNS
- **S. Abe, T. Araki, K. Chiba et al.**,
  Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset,
  2024. arXiv.
  Division: Experimental Nuclear & Particle Physics | **Center:** LNS
- **J. Du Plessis, Z. Janelidze, B. Wessels**,
  A Primer on Chainmails: Structures for Point-free Connectivity,
  2024. arXiv.
  Division: Theoretical Nuclear & Particle Physics | **Center:** CTP
- **B. Binks, H. Guenther**,
  TESSILATOR: a one-stop shop for measuring TESS rotation periods,
  MNRAS, 533, 2024. DOI.
  Division: Astrophysics | **Center:** MKI
- **H. Guenther, P. Cheimetz, C. DeRoo et al.**,
  Arcus X-ray telescope performance predictions and alignment requirements,
  JATIS, 11, 2024. DOI.
  Division: Astrophysics | **Center:** MKI

Today's agenda

Physics Basic Computing Services (subMIT) Users Meeting

Tuesday Feb 11, 2025, 10:00 AM → 11:00 AM America/New_York

Kolker Room (26-414) (MIT)

Description https://mit.zoom.us/j/96743699673?pwd=b3h2Q3c3cVQwYW12blhMUG5SWXZCZz09

10:00 AM → 10:15 AM **subMIT Overview, Status, & Updates** 15m

Speaker: David Walter

10:15 AM → 10:20 AM **Workshop Summary & Overview** 5m

Speaker: Matthew Heine (Massachusetts Institute of Technology)

10:20 AM → 10:40 AM **Solving Einstein's Equations on the computer** 20m

Speaker: Josu Aurrekoetxea (Massachusetts Institute of Technology)

10:40 AM → 10:50 AM **Roundtable** 10m

Speakers: Amer Al-Hiyasat (MIT), Hans Moritz Guenther (Massachusetts Institute of Technology), Molly Park (Massachusetts Institute of Technology), Prajwal Mohan Murthy (MIT LNS), Yin Lin (Massachusetts Institute of Technology)

10:50 AM → 11:00 AM **Discussion** 10m

Backup

Interactive use

Accessible through website:

- jupyterhub

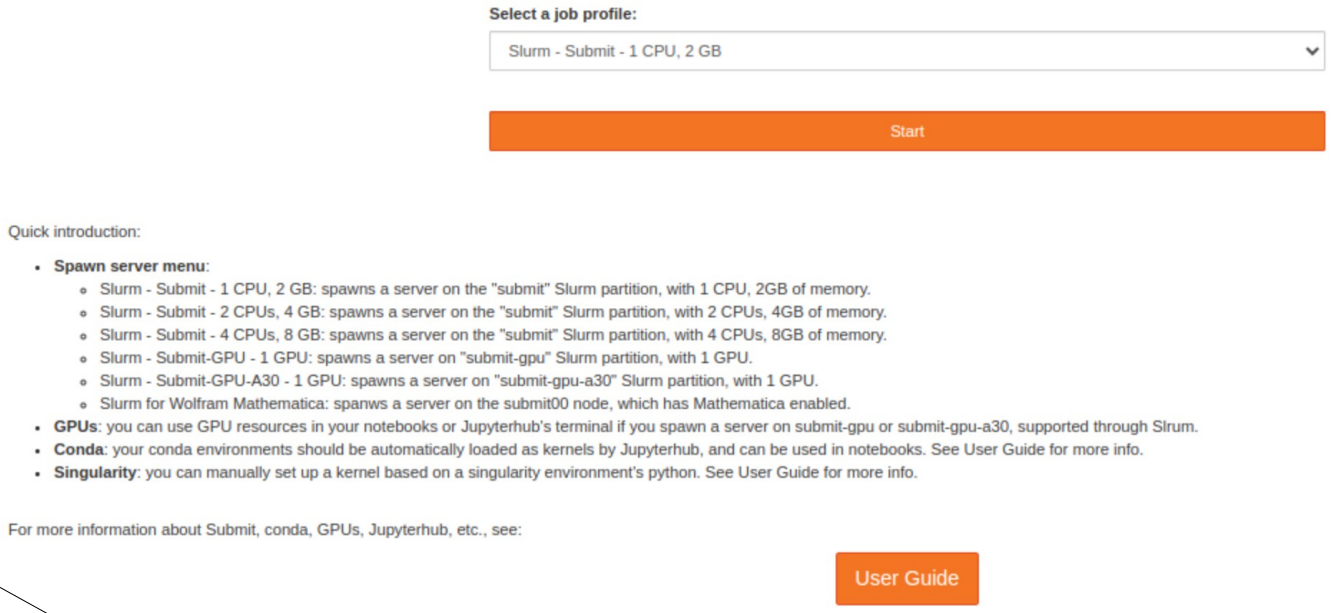

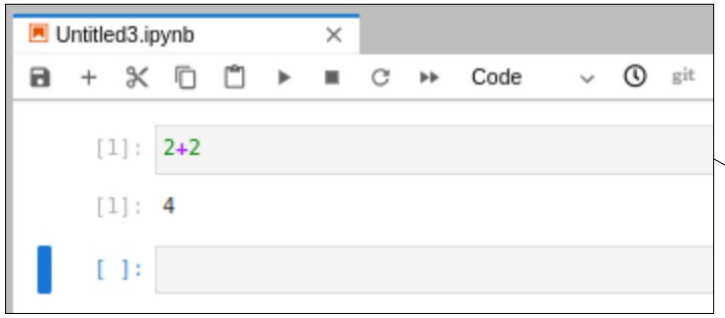

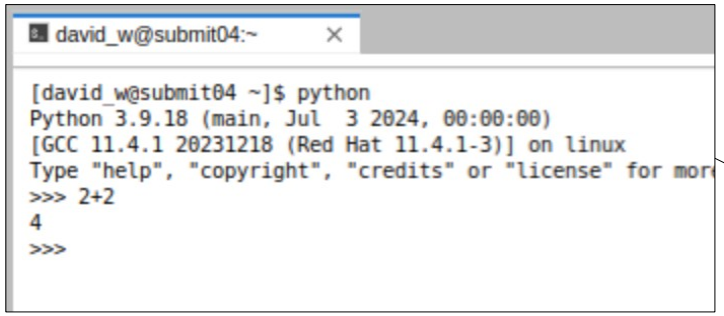

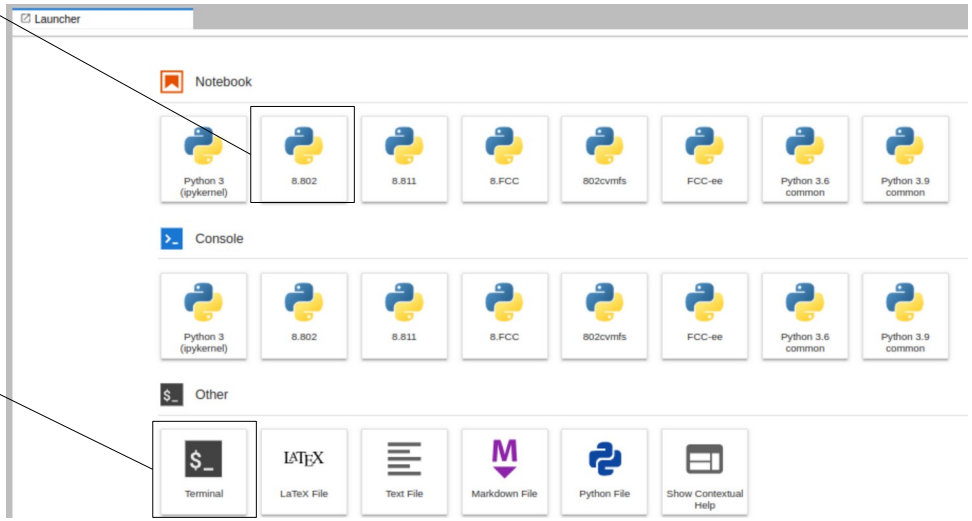

User base

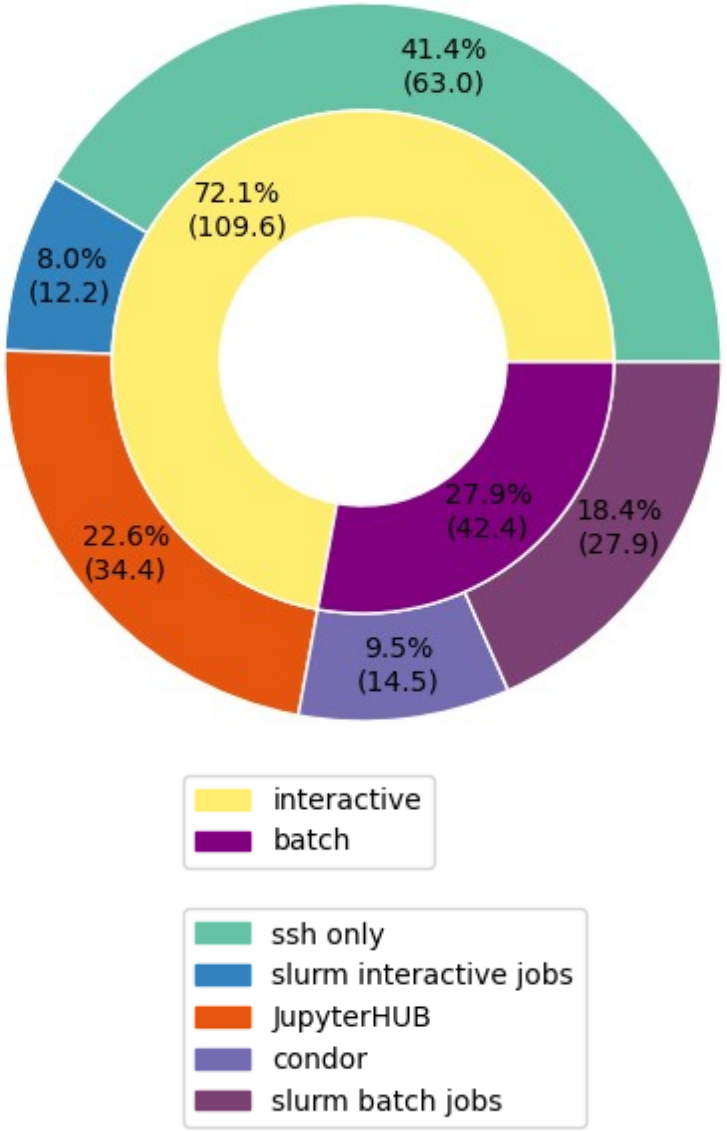

Active users of last 90 days (153)

- from different physics departments
- using subMIT in different ways

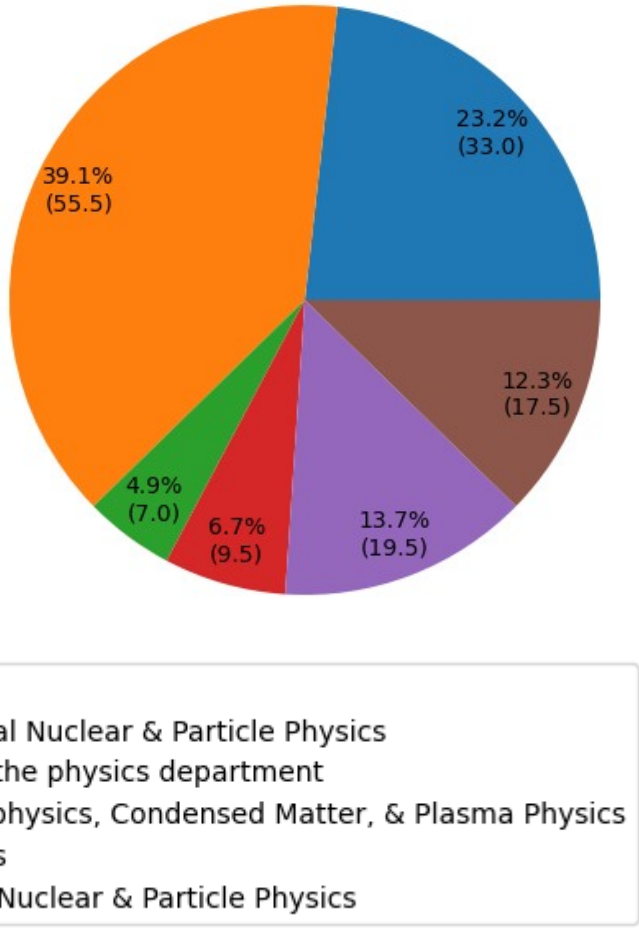

Classroom usage

Introductory undergraduate courses: 8.01, 8.02

- Technology-Enabled Active Learning (TEAL)

Advanced courses, junior lab: 8.13, 8.14

Workshops/ Hackathons

- FCC month, Gaia Hackathon, ...

Resource reservation via slurm

Software distribution and robust usage for O(100) students

- Kernel with customized python environment accessible through JupyterHub
- CVMFS for specialized programs or environments

Previous work

Moved to server room in B24

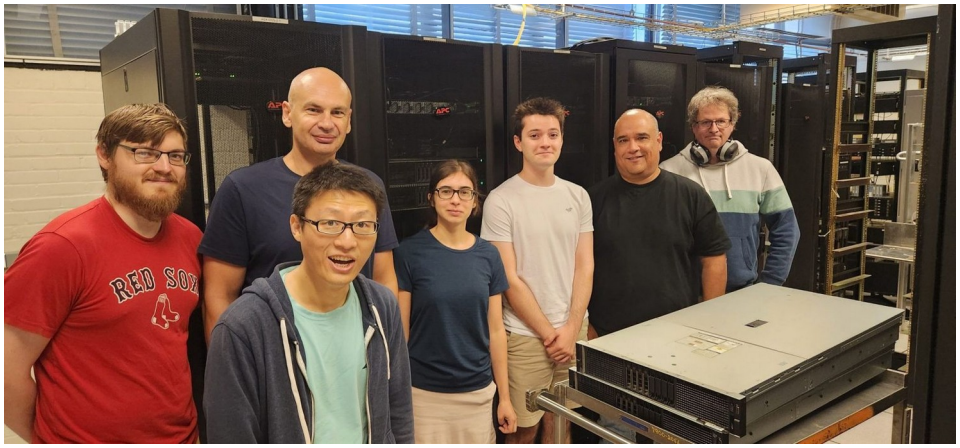

Migrated operating system of all machines

- CentOS7 (end of life) → Alma Linux 9

Migrated distributed file system

- GlusterFS (end of life) → CephFS
- Upgraded ceph to most recent stable release: 19.2.0 squid

Migrated from Docker to Podman

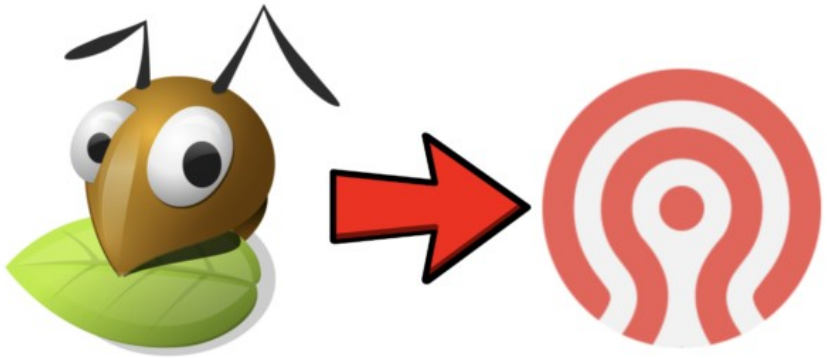

New hardware: ~500TB hdd for file storage

Whats next

Ensure stable long term operations

Streamline account creation

Add O(500) CPU cores to slurm

Add software support for

- Dask gateway, OpenMPI, Globus, ...

Analyze user experience

- Understand frequent causes why jobs/ jupyterhub sessions/ ... fail
- Find cases of inefficient use of resources

→ Dedicated actions: Give recommendations, adapt system configuration, ...

Please let us know what you need