

SubMIT: getting physics things done at MIT

Project overview

David Walter

30th January 2025

Workshop on Basic Computing Services in the Physics
Department - subMIT





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 experienced 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

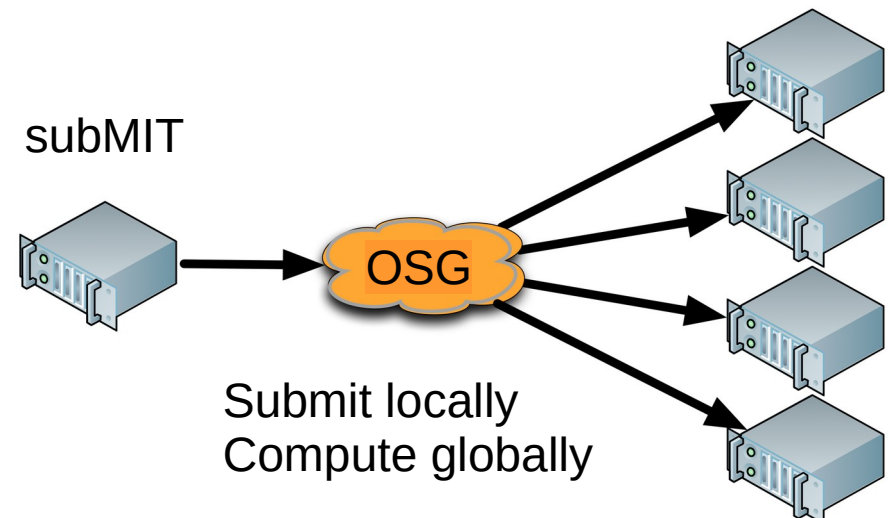
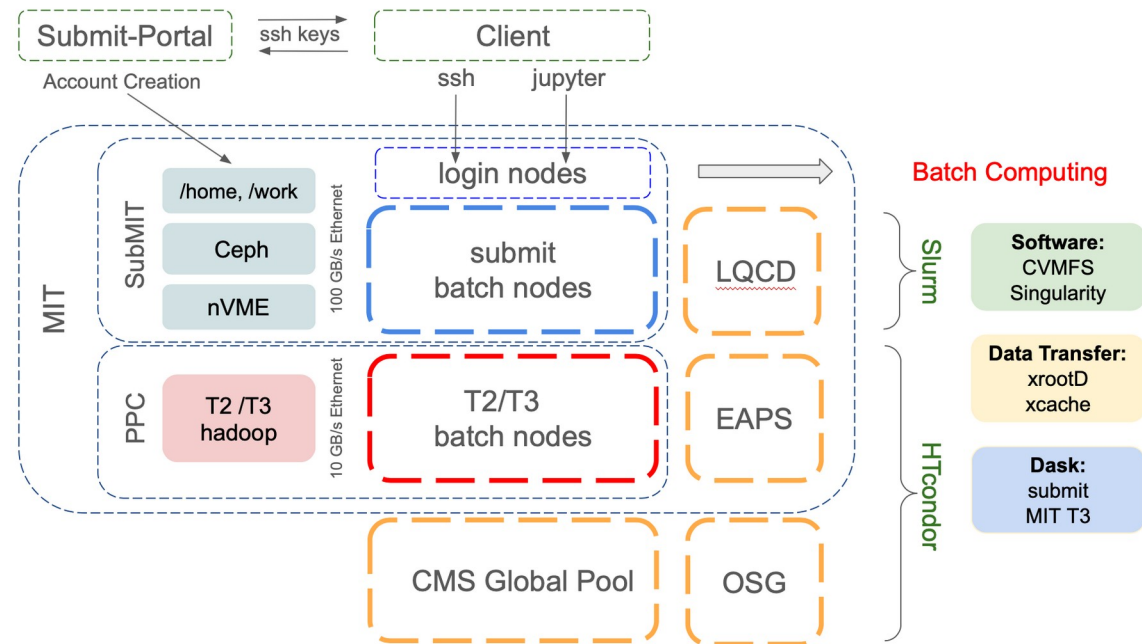
- O(1000 cores); ~30GPUs
- Interactive SSH login pool
- Jupyterhub access

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



subMIT

Getting physics things done at MIT

[Overview](#) [News](#) [People](#) [Contact](#) [About](#) [Users Guide](#) [A2rchi](#) [Jupyter](#)

Overview

The subMIT login pool is designed to let users login safely, prepare and test their research, and submit their jobs to the large computing resource of their choice. There are for now a limited number of resources connected but we are working on quickly expanding them.

[Get your account on SubMIT Portal](#)

Access

ssh <user>@submit.mit.edu

jupyterhub


Status

Servers

Slurm queue

Condor queue

Expert



Interactive use



Accessible through website:

- [jupyterhub](#)

Select a job profile:

Slurm - Submit - 1 CPU, 2 GB

Start

Quick introduction:

- **Spawn server menu:**
 - Slurm - Submit - 1 CPU, 2 GB: spawns a server on the "submit" Slurm partition, with 1 CPU, 2GB of memory.
 - Slurm - Submit - 2 CPUs, 4 GB: spawns a server on the "submit" Slurm partition, with 2 CPUs, 4GB of memory.
 - Slurm - Submit - 4 CPUs, 8 GB: spawns a server on the "submit" Slurm partition, with 4 CPUs, 8GB of memory.
 - Slurm - Submit-GPU - 1 GPU: spawns a server on "submit-gpu" Slurm partition, with 1 GPU.
 - Slurm - Submit-GPU-A30 - 1 GPU: spawns a server on "submit-gpu-a30" Slurm partition, with 1 GPU.
 - Slurm for Wolfram Mathematica: spans a server on the submit00 node, which has Mathematica enabled.
- **GPUs:** you can use GPU resources in your notebooks or Jupyterhub's terminal if you spawn a server on submit-gpu or submit-gpu-a30, supported through Slurm.
- **Conda:** your conda environments should be automatically loaded as kernels by Jupyterhub, and can be used in notebooks. See User Guide for more info.
- **Singularity:** you can manually set up a kernel based on a singularity environment's python. See User Guide for more info.

For more information about Submit, conda, GPUs, Jupyterhub, etc., see:

User Guide

For any questions, comments, or feedback, please send an email to [submit-help](#).

```
Untitled3.ipynb x
[1]: 2+2
[1]: 4
[ ]:
```

```
 david_w@submit04:~ x
[david_w@submit04 ~]$ python
Python 3.9.18 (main, Jul 3 2024, 00:00:00)
[GCC 11.4.1 20231218 (Red Hat 11.4.1-3)] on linux
Type "help", "copyright", "credits" or "license" for more
>>> 2+2
4
>>>
```

Launcher

Notebook

- Python 3 (pykernel)
- 8.802
- 8.811
- 8.FCC
- 802cmfs
- FCC-ee
- Python 3.6 common
- Python 3.9 common

Console

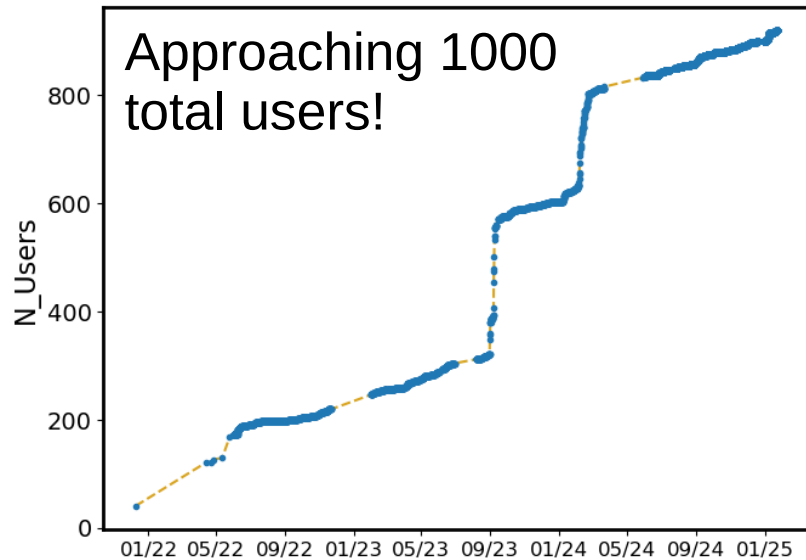
- Python 3 (pykernel)
- 8.802
- 8.811
- 8.FCC
- 802cmfs
- FCC-ee
- Python 3.6 common
- Python 3.9 common

Other

- Terminal
- LaTeX File
- Text File
- Markdown File
- Python File
- Show Contextual Help

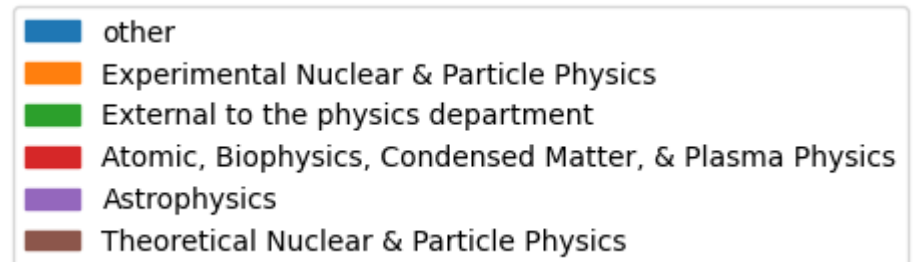

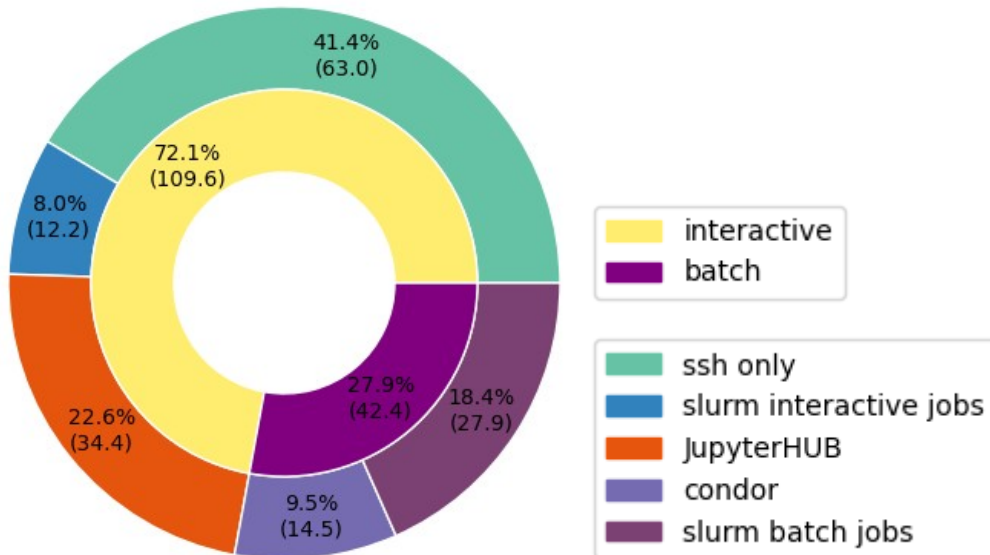


User base



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



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



Monthly user group meetings

Information flow between user community and project team

- Advertised and open to the broader community
- Regular time slot: Tuesdays 10:00-11:00 EST
- Hybrid: Kolker room + Zoom

Topical presentations from

- Users and community members
- User group representatives
- Project team

Room for discussions and feedback

Latest meeting:

- **November 19, 2024**

Physics Basic Computing Services (subMIT) Users Meeting
Tuesday Nov 19, 2024, 10:00 AM - 11:00 AM America/New_York
Kolker Room (26-414) (MIT)
Description: <https://mit.zoom.us/j/96743699673?pwd=b3h2Q3c3cVQwYW12blhMUG5SWXZCZz09>

Time	Topic	Speaker(s)	Duration
10:00 AM - 10:20 AM	subMIT Overview, Status, & Updates	Mariarosaria D'Alfonso (Massachusetts Institute of Technology)	20m
10:20 AM - 10:35 AM	What I'd like to use submit for	Hans Günther (MIT)	15m
10:35 AM - 10:45 AM	Roundtable	Amer Al-Hiyasat (MIT), Kaliroe Pappas (MIT laboratory for nuclear science), Molly Park (Massachusetts Institute of Technology), Prajwal Mohan Murthy (MIT LNS), Siddharth Mishra-Sharma (MIT), Yin Lin (Massachusetts Institute of Technology), Yitian Sun (Massachusetts Institute of Technology)	10m
10:45 AM - 11:00 AM	Discussion		15m



Monthly user group meetings

Information flow between user community and project team

- Advertised and open to the broader community
- Regular time slot: Tuesdays 10:00-11:00 EST
- Hybrid: Kolker room + Zoom

Topical presentations from

- Users and community members
- User group representatives
- Project team

Room for discussions and feedback

Next meetings on

- February 11
- March 11

Contact us if you like to present!

- submit-help@mit.edu

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	→	10:35 AM	User talk	20m
Speaker: Josu Aurrekoetxea (Massachusetts Institute of Technology)				
10:35 AM	→	10:45 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)				
10:45 AM	→	11:00 AM	Discussion	15m

Physics Basic Computing Services (subMIT) Users Meeting

Tuesday Mar 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	→	10:35 AM	User talk	20m
Speaker: Richard Nally (Massachusetts Institute of Technology)				
10:35 AM	→	10:45 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)				
10:45 AM	→	11:00 AM	Discussion	15m



Previous work

Moved to server room in B24

Migrated operating system of all machines

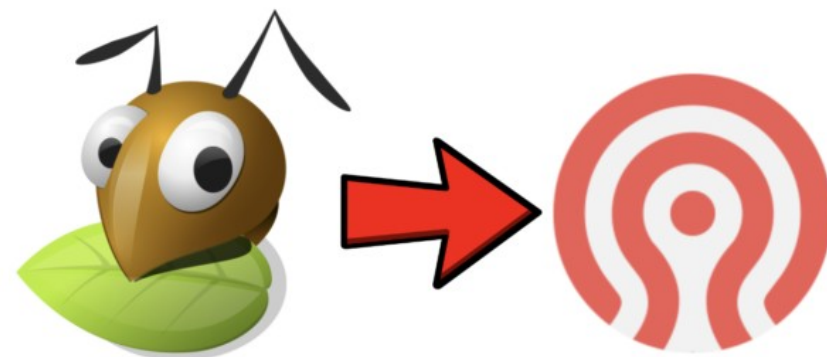

- CentOS7 (end of life) → Alma Linux 9

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





Today's Workshop

Indico agenda

- Overview of subMIT project, resources, software environment
- Hands-on tutorials
- User talks on research usage

9:00 AM	9:05 AM	Introduction to the workshop Speaker: Christoph Paus (MIT)	5m	26-414 (Kolker room)
9:05 AM	9:30 AM	subMIT Project Overview Speaker: David Walter	25m	26-414 (Kolker room)
9:30 AM	9:45 AM	Getting started on subMIT: Available Resources <ul style="list-style-type: none">Computational ResourcesDocumentation: Users Guide, GitHub ExamplesSupport: Help Desk & Chat Bot Speaker: Mariarosaria D'AiFonso (Massachusetts Institute of Technology)	15m	26-414 (Kolker room)
9:45 AM	10:00 AM	Getting started on subMIT: How to Interact with subMIT <ul style="list-style-type: none">ssh / terminalJupyterHubX2GoVisual Studio Code (Remote Development)1-minute summary of batch jobs (see also later tutorial) Speaker: Matthew Heine (Massachusetts Institute of Technology)	15m	26-414 (Kolker room)
10:00 AM	10:15 AM	Getting started on subMIT: Installing, Managing, & Sharing Software (Overview) Speaker: Xuejian Shen (Massachusetts Institute of Technology)	15m	26-414 (Kolker room)
10:15 AM	10:30 AM	Coffee Break	15m	26-414 (Kolker room)
10:30 AM	11:00 AM	Tutorial: Managing & Installing Software with Conda Speaker: Marianne Moore (MIT)	30m	26-414 (Kolker room)
11:00 AM	11:30 AM	Tutorial: Containers for Portable Software Environments Speaker: Luca Lavezzo (MIT)	30m	26-414 (Kolker room)
11:30 AM	12:00 PM	Tutorial: Batch Job / Workflow Management: SLURM & HTCondor This session will show how to use the batch schedulers/resource-managers SLURM & HTCondor to manage your computational tasks, distributing them across the shared resources. Depending on your workflow, this may provide an easy way to run your calculations in parallel, shortening time-to-result and eliminating some manual tasks. Speaker: Zhangqier Wang (Massachusetts Institute of Technology)	30m	26-414 (Kolker room)

1:30 PM	3:00 PM	User Talks: Session 1 Session of User-Contributed talks	1h 30m	26-414 (Kolker room)
		Emulating the atomic nucleus Speaker: Antoine Belle (Massachusetts Institute of Technology)	20m	
		Symbolic Learning Nuclear Relations related to https://arxiv.org/abs/2404.11477 Speaker: Jose Miguel Munoz Arias (Massachusetts Institute of Technology)	20m	
		The needle in a haystack problem of homology-directed DNA repair related to https://www.nature.com/articles/s41594-023-01065-w Speaker: Henrik Pinholt (Massachusetts Institute of Technology)	20m	
		Lattice-QCD 🏆 Software 🏆 Hardware Speaker: Joshua Lin (student@mit.edu)	15m	
		Soft unclustered energy patterns in CMS using dask Speaker: Pietro Lugato (student@mit.edu)	15m	
3:00 PM	3:15 PM	Coffee Break	15m	26-414 (Kolker room)
3:15 PM	4:30 PM	User Talks: Session 2 Session of User-Contributed talks	1h 15m	26-414 (Kolker room)
		Topological Chiral Superconductors beyond Fermi liquid pairing related to https://arxiv.org/abs/2409.18067 Speaker: Luke Kim (Massachusetts Institute of Technology)	20m	
		A qubit that corrects itself related to https://arxiv.org/abs/2412.03650 Speaker: Max Geier (Massachusetts Institute of Technology)	20m	
		subMIT for cosmological uses Speaker: Mikhail Ivanov (Massachusetts Institute of Technology)	20m	
		Assessing a new analysis with batch jobs Speaker: Sumita Ghosh (Massachusetts Institute of Technology)	15m	
4:30 PM	5:00 PM	Open Discussion & Closing Remarks	30m	26-414 (Kolker room)



Backup



Documentation

User's guide

- How to interact with the system
- Recommendations
- Tutorials & examples

A2rchi (chatbot)

- Ragging of users guide
- Interactive user support
- Support ticket handling



subMIT v1 documentation » User's Guide - subMIT

User's Guide - subMIT

Contents:

- [User's guide – subMIT login pool](#)
- [Getting started](#)
- [Access to subMIT](#)
- [Best practices](#)
- [Available software](#)
- [Batch computing](#)
- [User quota and storage at submit](#)
- [Monitoring at submit](#)
- [GPU resources](#)
- [Data backup](#)
- [Conda and its benefits beyond python](#)
- [Acknowledging subMIT](#)

[Back to the main subMIT webpage.](#)

Tutorials and Examples

Tutorials:

- [Tutorial 0: Introduction to the UNIX terminal](#)
- [Tutorial 1: Common software packages \(python, Julia, MATLAB\)](#)
- [Tutorial 2: Batch Job \(HTCondor and Slurm\)](#)
- [Tutorial 3: Containers \(Podman and Singularity\)](#)
- [Tutorial 4: Source Control \(Git/Github\) with Visual Studio Code \(VSCode\)](#)
- [Tutorial 5: Debugging Fortran code with Visual Studio Code \(VSCode\)](#)
- [Tutorial 6: Introduction to Pytorch Lightning](#)
- [Tutorial 7: Introduction to Snakemake](#)

Examples of scripts can be found on our [submit-examples GitHub repository](#).

Table of Contents

- [User's Guide - subMIT](#)
- [Tutorials and Examples](#)
- [Search tool and tags](#)
- [Contributing](#)
- [Archive - subMIT](#)

Next topic

- [User's guide – subMIT login pool](#)

This Page

- [Show Source](#)

Quick search

A2rchi

Start a conversation and explore the power of A2rchi, specially trained on the SubMIT Cluster.
Your chat history will be displayed here.

By using this website, you agree to the [terms and conditions](#).

Enter your question for A2rchi