

subMIT Overview

Josh Bendavid

Basic Computing Services (subMIT) Workshop

Feb. 2, 2024

Introduction

- subMIT system provides an interactive login pool + scale-out to batch resources
 - Home directories
 - Convenient software environment (CentOS7 native, docker/singularity images, conda)
 - Upgrade in progress from CentOS7 to Alma Linux 9
 - SSH or Jupyterhub access
 - Local batch system with $O(1000)$ cores, >50 GPU's
 - Additional storage for software installation/development, large datasets
 - Convenient access to larger external resources (OSG, CMS Tier-2 and Tier-3, LQCD Cluster, EAPS)
- User support is a key feature of the system
 - Beyond basic troubleshooting
 - Help users make optimal use of the available resources
 - Expert advice on designing/improving workflows
 - Customize and evolve system configuration to accommodate user needs as appropriate

Introduction

- Storage and networking
 - Local storage (1TB/user), 10's of TB for larger group datasets
 - 30TB of ultra-fast NVME storage with room for future expansion
 - Fast networking: 100 Gbps ethernet between all nodes and uplink to IS&T
 - RoCE (RDMA over Converged Ethernet) has been partially tested/commissioned, should be possible for MPI applications
- Additional resources recently or currently being integrated
 - More disk storage (100TB contributed from ABRACADABRA)
 - Integration of existing computing resources from research groups
 - Purchase of several large core count/high memory machines by research groups for additional computing resources and to support specialized workflows and/or R&D where large single node scaling is useful
 - Current “high density” template, Dual AMD EPYC 192 core/384 thread with 0.75-1.5TB of memory

Introduction: subMIT Website



subMIT

Getting physics things done at MIT

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

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

Resources

- >1TB of free storage per user
- 1000s of cores and >50 GPUs available interactively and through Slurm
- Access to OSG, CMS T3 and T2, LQCD Cluster, and EAPS

Software

- Python, anaconda, Julia, Matlab, singularity, and much more!

- Website (with User's Guide/Instructions):

<https://submit.mit.edu/>

- Overview and general information
- Direct JupyterHub access
- User's Guide:

<https://submit.mit.edu/submit-users-guide/>

Introduction: Project Organization

- Formally the project is organized as ***Basic Computing Services*** in the Physics Department
 - **Project Team:** Implementation/Operations/Maintenance of the system
 - **Users Group:** Contact point between the user community and the project team, forum for user feedback, requests, information flow to and from users
 - **Steering Committee:** Faculty oversight, funding, etc
 - See https://submit.mit.edu/?page_id=6

Users Group In Practice

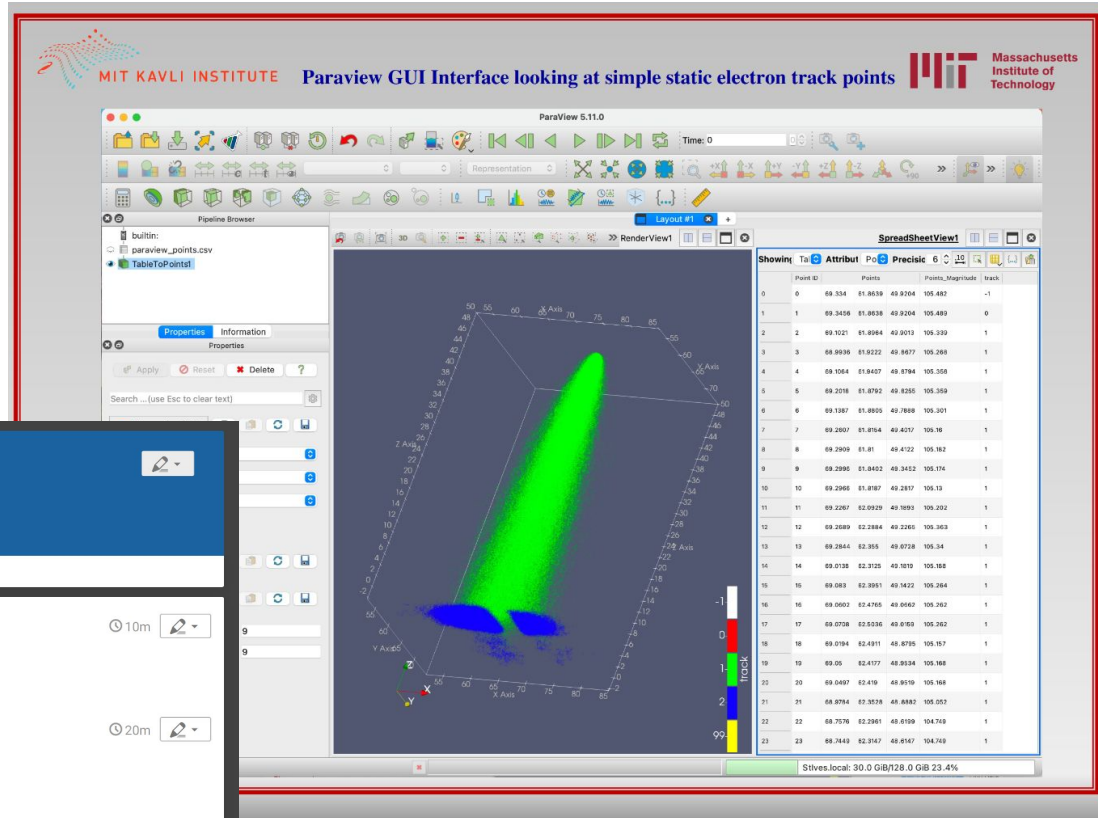
- ~Monthly meetings
 - Advertised and open to the broader community
 - Topical presentations from project team, Users Group representatives, or other users or community members
 - Forum for feedback and information flow between the user community and the project team
 - Regular timeslot: Tuesday 10:00-11:00 EST
 - Next meeting in March TBA
- Users Group representatives
 - Identified representatives from research groups across the department
 - Attend the monthly meetings
 - Provide feedback from your groups/community
 - Distribute information/news from the project team

Users Group Representatives

- Users group has been formed (JB as coordinator)
- Current Users Group representative (associated faculty/group)
 - Yin Lin (Phiala Shanahan)
 - Siddharth Mishra-Sharma (Jesse Thaler)
 - Prajwal Mohan Murthy (Bob Redwine)
 - Kaliroë Pappas (LNS Neutrino/Dark Matter)
 - Sunghan Ro (Julien Tailleur)
 - Yitian Sun (Tracy Slatyer)
 - Molly Taylor (LNS Heavy Ion Group)

Users Group Meetings

- E.g. presentation from April Users Group meeting on Visualization of Geant Simulations on subMIT
- <https://indico.mit.edu/event/752/>



Users Group Meeting

Tuesday Apr 25, 2023, 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 → 10:10 AM News/Introduction

Speakers: Joshua Bendavid (Massachusetts Institute of Technology), Matthew Heine (staff@mit.edu)

UserGroup_04_25_2...

10:10 AM → 10:30 AM Visualization of Geant Simulations

Speaker: Richard Foster (Massachusetts Institute of Technology)

Foster_SubMIT_Par...

Foster_SubMIT_Par...

10:30 AM → 10:45 AM Roundtable

Speakers: Kaliroe Pappas (MIT laboratory for nuclear science), Molly Taylor (Massachusetts Institute of Technology), Prajwal Mohan Murthy (MIT LNS), Siddharth Mishra-Sharma (MIT), Sunghan Ro (MIT), Yin Lin (Massachusetts Institute of Technology), Yitian Sun (Massachusetts Institute of Technology)

10:45 AM → 11:00 AM Discussion

Storage breakdown

- Several different storage areas are available covering different use cases
 - /home/submit/<username>
 - Home directories (nfs server), redundant disk array with backups
 - 5GB quota
 - Use for software development and (small) critical data
 - /work/submit/<username>
 - Work directory (nfs server), no backups (but redundant disk array)
 - 50GB quota
 - Use for software installation (conda or docker/singularity images)
 - /data/submit/<username>
 - Large distributed disk system, no backups, but redundancy against disk failure (“erasure coding”)
 - 1TB user quota, larger quotas available in dedicated group directories
 - Store large datasets here
 - /scratch/submit/<username>
 - Fast NVMe SSD array
 - Commissioned by several groups for high performance data analysis
 - /cvmfs/
 - Read-only distributed storage for distributing software, singularity images, etc
 - Several CERN-related repositories are available
 - Local repository /cvmfs/cvmfs.cmsaf.mit.edu where additional software or data can be added if needed
- Flexible tiered storage system, can accommodate a wide range of user needs
- Larger datasets encouraged to use shared group space, but quotas can be increased when needed

Interactive Use: Terminal or JupyterHub

Select a job profile:

Slurm - Submit - 1 CPU, 500 MB

Start

Quick introduction:

- **Spawn server menu:**
 - Slurm - Submit - 1 CPU, 500 MB: spawns a server on submit slurm partition.
 - Slurm - Submit - 2 CPUs, 1000MB: similar as above, with more resources allocated.
 - Slurm - SubmitGPU - 1 GPU: spawns a server on submit-gpu slurm partition, requesting 1 GPU.
 - Slurm - SubmitGPU1080 - 1 GPU: spawns a server on submit-gpu1080 slurm partition, requesting 1 GPU.
 - Local server - Submit01 - 1 CPU, 500 MB, /home/submit/{username}: spawns on submit01, in your /home/submit/{username}/ directory.
 - Local server - Submit01 - 1 CPU, 500 MB, /work/submit/{username}: spawns on submit01, in your /work/submit/{username}/ directory.
- **GPUs:** you can use GPU resources in your notebooks or Jupyterhub's terminal if you spawn a server on submit-gpu or submit-gpu1080, 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-jupyter.

JupyterLab — Mozilla Firefox

subMIT - Getting physics: x JupyterLab x +

https://submit.mit.edu/jupyter/user/jbendavi/lab?

File Edit View Run Kernel Git Tabs Settings Help

Launcher

Notebook

Python 3 802 802cvmfs distest python3.6

Console

Python 3 802 802cvmfs distest python3.6

Other

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

- Interactive Jupyter session available directly from website with touchstone authentication (subMIT account still required)
- SLURM is used to efficiently share resources between interactive and batch use
- Primary usage is research, but has also been used for several courses in-class and for assignments 10

Communication Channels

- User support mailing list: submit-help@mit.edu
- Large language model application (A2rchi) is used in production to assist with ticket handling
 - Interactive chatbot also available through subMIT web page
- Slack workspace: <https://mit-submit.slack.com>
 - “help-desk” channel
- Monthly Users Group Meetings
 - Open for discussion
 - Open for user contributions: full set of Users Group representatives can be contacted at submit-usersgroup@mit.edu
- Annual subMIT workshop
- In addition to direct interaction with the subMIT project team, users are encouraged to discuss with Users Group representative from their own group or “nearby” group

Linux Distribution Upgrade

- Current CentOS 7 distribution reaches EOL for maintenance updates in June 2024
- Decision by Red Hat to reorganize CentOS project and releases has disrupted logical upgrade path from CentOS 7->8
- Decision taken to upgrade from CentOS 7 to Alma Linux 9, considering:
 - Ease of transition
 - Support lifetime (Alma 9 supported until 2032, though another upgrade probably desirable before that)
 - Functionality
 - Direction being taken at other universities and labs (CERN, Fermilab, etc)
- Discussion has included Users Group and broader community
- Ease transition for users through well-supported and documented use of containers
- Test instances for interactive and batch usage with Alma 9 have been in place and used by power users for several months
- Performance-sensitive services already upgraded (NVME storage)
- New large 384-thread machines using Alma 9 from the start
- Recent change (June 21, 2023) by Red Hat to CentOS source code policy introduces further uncertainty to enterprise linux ecosystem and Alma/Rocky Linux Projects
 - Carefully monitoring developments, but continuing with Alma 9 migration
 - Current plan from Alma Linux project is to maintain ABI compatibility with RHEL but not 1:1 source code or bug compatibility

Linux Distribution Upgrade

- Upgrade is currently in progress
- Compute and GPU nodes in the process of being upgraded
- Test instance of JupyterHub in place
- Test login nodes and slurm partitions in place
- Container support:
 - On Alma9 we support both singularity/apptainer and podman
 - Docker is deprecated in favour of podman (should be a “drop-in” replacement from user perspective)
 - “rootless” operation for all containers, both more flexible and more secure
 - Supported for both CPU and GPU usage
- Default slurm partition, login nodes, etc will be swapped over to Alma9 at the end of the upgrade
- CentOS7 services remain available in the meantime (but with reduced compute resources in the corresponding queues)
- CentOS7 environment will remain accessible via containers on Alma9 machines
- Some further improvements in the pipeline
 - Improved user priority balancing when queue are busy
 - Improved resource request/limit enforcement on batch queues to improve stability
 - Improved stability of CVMFS mounts
- More details on the upgrade in Marianne’s talk

Today's Workshop

- Indico page with timetable and slides:
 - <https://indico.mit.edu/event/956/>
- Overview of subMIT project, resources, software environment
- Discussion of Linux upgrade
- Guest speaker from Fermilab
- Hands-on tutorials
- User talks on research usage
- Kolker Room + Zoom

The screenshot shows the Indico event page for the 'Workshop on Basic Computing Services in the Physics Department - subMIT'. The event is scheduled for Friday, Feb 2, 2024, from 9:00 AM to 5:10 PM in the Kolker Room (26-414). The description states that the subMIT computing facility is a login pool designed to provide access to basic research computing resources. The workshop will provide an overview and updates on the system and plans, along with topical presentations and tutorial sessions. Refreshments will be provided during coffee breaks. The details of individual talks may change as the workshop progresses.

The agenda is as follows:

- 9:00 AM → 9:30 AM: Introduction/subMIT Project Overview**
Speaker: Joshua Bendavid (Massachusetts Institute of Technology)
- 9:30 AM → 9:40 AM: Getting started on subMIT: Available Resources**
 - Computational Resource
 - Documentation: Users Guide, GitHub Examples
 - Support: Help Desk & Chat BotSpeaker: Zhangkai Wang
- 9:40 AM → 9:55 AM: Getting started on subMIT: How to interact with subMIT**
 - ssh Terminal
 - JupyterHub
 - X2Go
 - Visual Studio Code (Remote Development)
 - 1 minute summary of batch jobs (see also later tutorial)Speaker: Matthew Heine (Massachusetts Institute of Technology)
- 9:55 AM → 10:05 AM: Getting started on subMIT: Installing / Managing Software**
 - native environment
 - installing software with conda
 - software environment management tools
 - containers overviewSpeaker: Chad Freer (Massachusetts Institute of Technology)
- 10:05 AM → 10:20 AM: Navigating the recent Linux Upgrade**
In this session we will outline key changes in the migration to the new version of Linux as well as how to get things back to a portion of time to answer user questions about the upgrade, so please bring any questions.
Speaker: Marianne Moore (MIT)
- 10:20 AM → 10:40 AM: Coffee Break**
- 10:40 AM → 11:00 AM: Experience thus far running a multi-experiment OpenShift Cluster at FNAL**
Speaker: Lindsey Gray (Fermi National Accelerator Laboratory)
- 11:00 AM → 11:30 AM: Tutorial: Batch Job / Workflow Management: SLURM & HTCondor**
This session will show how to use the batch scheduler/resource managers SLURM & HTCondor to manage across the shared resources. Depending on your workflow, this may provide an easy way to run your calculations, eliminating some manual tasks.
Speaker: Chad Freer (Massachusetts Institute of Technology)
- 11:30 AM → 12:30 PM: Tutorial: GPUs: Introduction to Likelihood-Free Inference with PyTorch Lightning**
This tutorial will give an overview of using pytorch lightning for building and training neural networks. A simple problem in inference: measuring the parameters of a line will be presented. The tutorial will also introduce distributed training with pytorch lightning.
Speaker: Deep Chatterjee (Massachusetts Institute of Technology)
- 12:30 PM → 2:00 PM: Lunch**
- 2:00 PM → 2:15 PM: Parameter Estimation of Un-modeled Gravitational-wave Signals using Likelihood-free Inference**
Speaker: Deep Chatterjee (Massachusetts Institute of Technology)
- 2:15 PM → 2:35 PM: DM21cm: a GPU-accelerated simulation of dark matter energy injection in 21cm cosmology**
Speaker: Yitai Suft (Massachusetts Institute of Technology)
- 2:35 PM → 2:45 PM: Embarrassingly parallel ray-tracing for the Arcus X-ray spectrometer**
Speaker: Hans Guenther (Massachusetts Institute of Technology)
- 2:45 PM → 2:55 PM: Employing snakemake to perform measurements with the LHCb experiment at CERN**
Speaker: Blaise Delany (Massachusetts Institute of Technology)
- 2:55 PM → 3:05 PM: Tutorial: Introduction to the Snakemake workflow manager**
Speaker: Blaise Delany (Massachusetts Institute of Technology)
- 3:05 PM → 3:30 PM: Coffee Break**
- 3:30 PM → 3:45 PM: CMS Analysis on subMIT with crab and dask jobqueue**
Speaker: Simon Rothman (Massachusetts Institute of Technology)
- 3:45 PM → 4:00 PM: Adaptive Brillouin-zone integration for optical conductivity**
Speaker: Lorenzo Van Munoz (Massachusetts Institute of Technology)
- 4:00 PM → 4:10 PM: Axion and high-frequency gravitational wave searches with ABRACADABRA and DMRadio**
Speaker: Kalrloe Pappas (MIT Laboratory for Nuclear Science)
- 4:10 PM → 4:20 PM: Training spiking neural networks via adjoint sensitivity analysis**
Speaker: Enrique Tolosa (Massachusetts Institute of Technology)
- 4:20 PM → 4:30 PM: ML for Science: Best Practices**
Speaker: Denis Boyda (Massachusetts Institute of Technology)
- 4:30 PM → 5:00 PM: Open Discussion & Closing Remarks**