Summary

Josh Bendavid

Basic Computing Services (subMIT) Review

July 18, 2023

Introduction

- subMIT system provides an interactive login pool + scale-out to batch resources
  - Home directories
  - Convenient software environment (CentOS7 native, docker/singularity images, conda)
  - 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

Highlights: Recent/Ongoing Development and Upgrades

- 100 Gbps uplink installed in B24
- Two new high density compute nodes (192 core / 384 thread) installed
- Ultra-high performance data analysis commissioned with NVMe storage: 75 Gbps over nfs
- System is fully ipv6 enabled
- Alma Linux 9 upgrade in progress with test instances available for users imminently

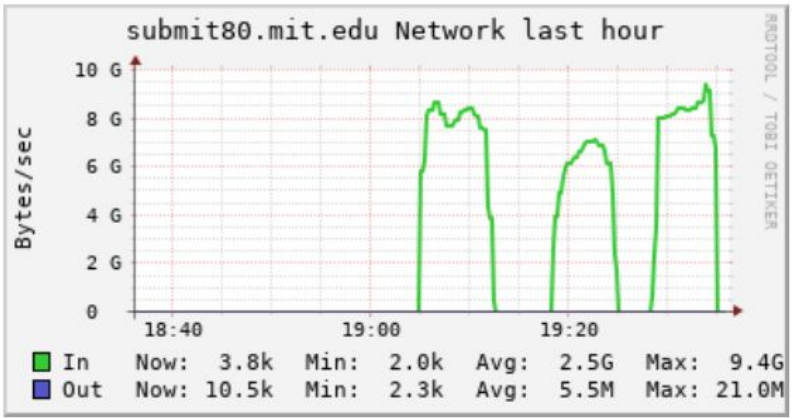

Usage Statistics

- >300 total users
- 40-60 active users in any given week
- Detailed usage statistics are available
- Broad representation of users across the department

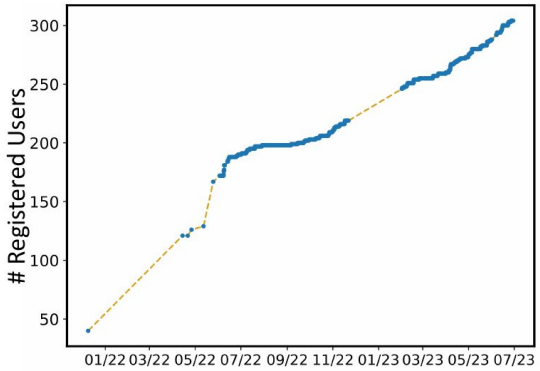

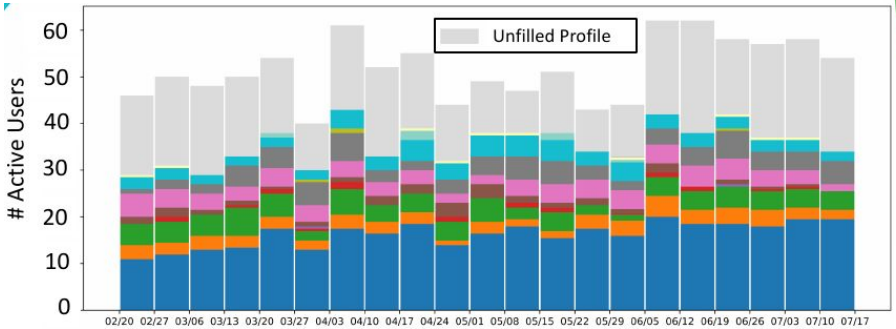

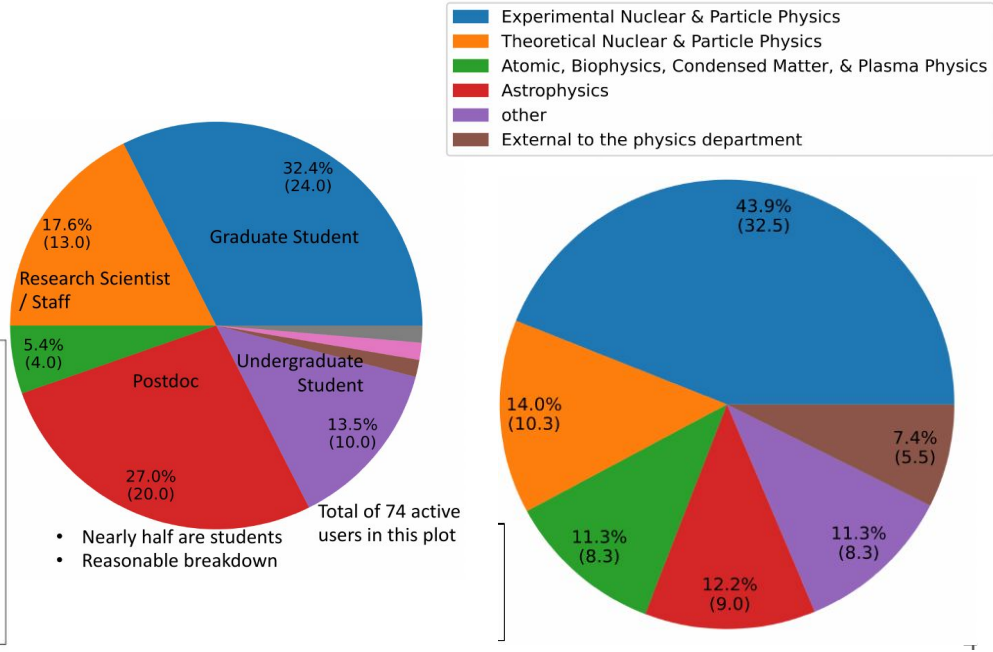

Introduction: subMIT Website

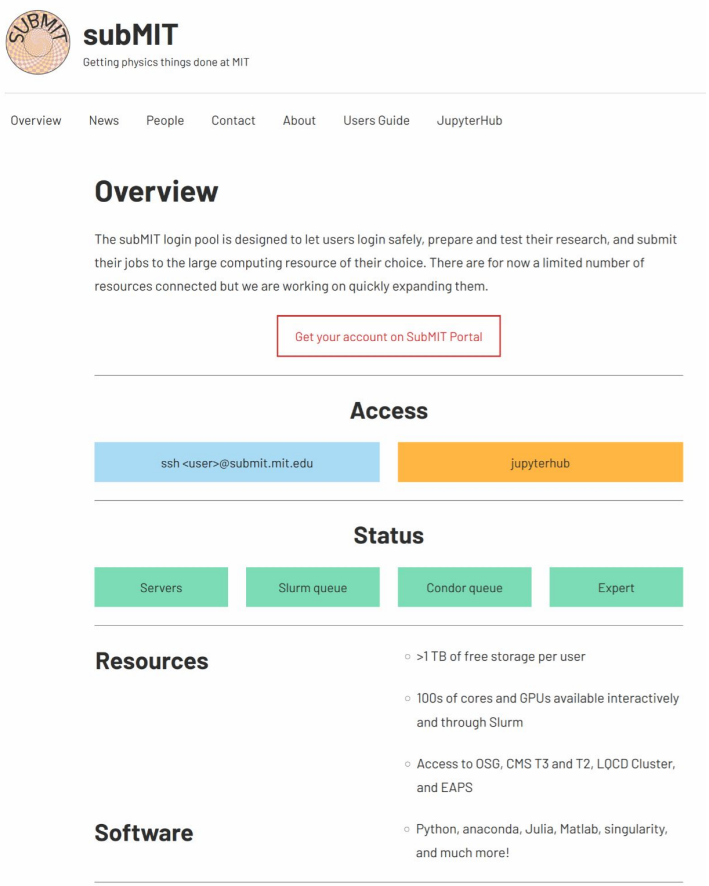

- Website (with User's Guide/Instructions): https://submit.mit.edu/
  - Overview and general information
  - Direct JupyterHub access
  - User's Guide: https://submit.mit.edu/submit-users-guide/
- Experimental large language model application under development
  - For interactive use from website + augmented support ticket handling
  - Joint project with College of Computing, with dedicated funding
  - Dedicated talk at LLM workshop on Friday https://indico.mit.edu/event/759/

Interactive Use: Terminal or JupyterHub

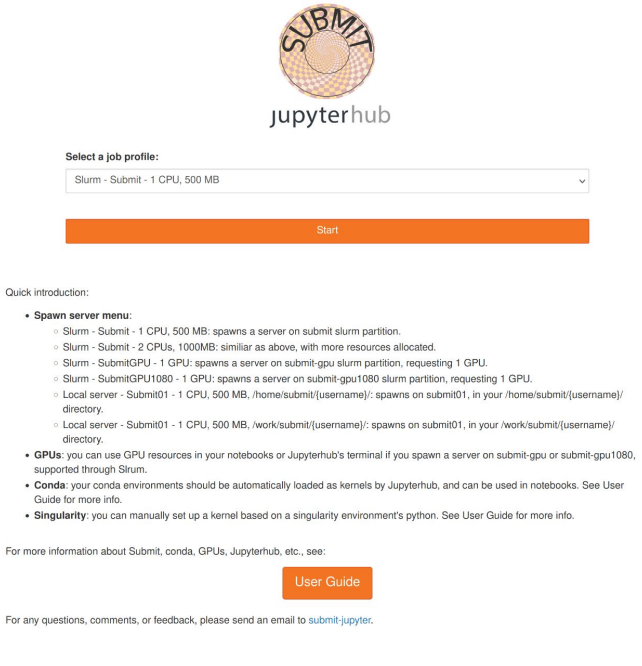

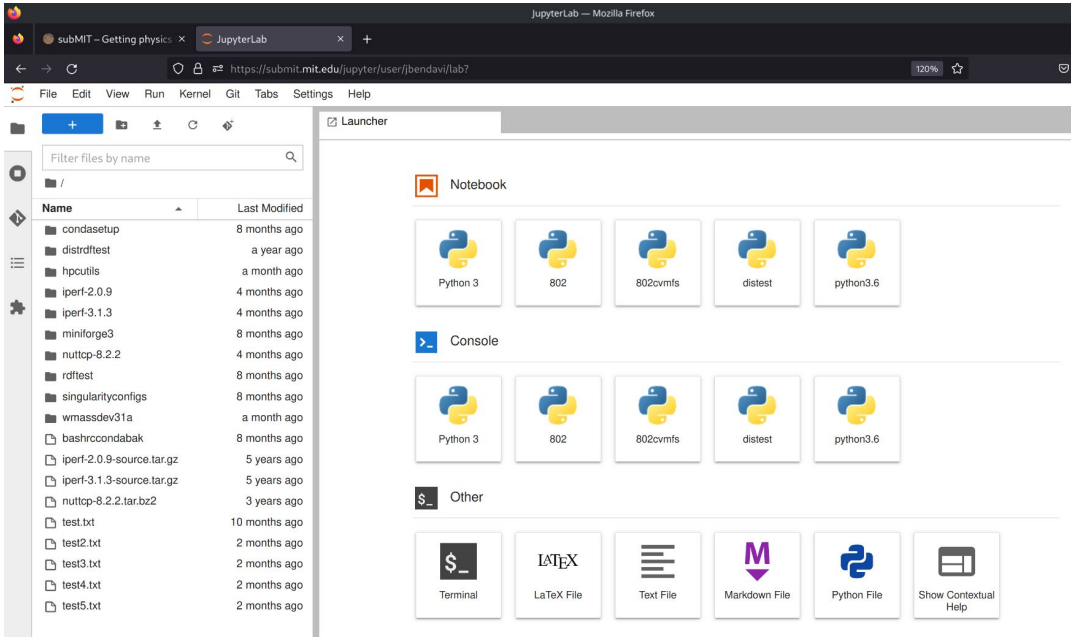

- Interactive Jupyter session available directly from website with touchstone authentication (subMIT account still required)
- SLURM is used to efficiently share resources between interactive and batch use
- Primary usage is research, but education applications also possible
  - Limited trial has been conducted for 802 exercises

Access to external resources

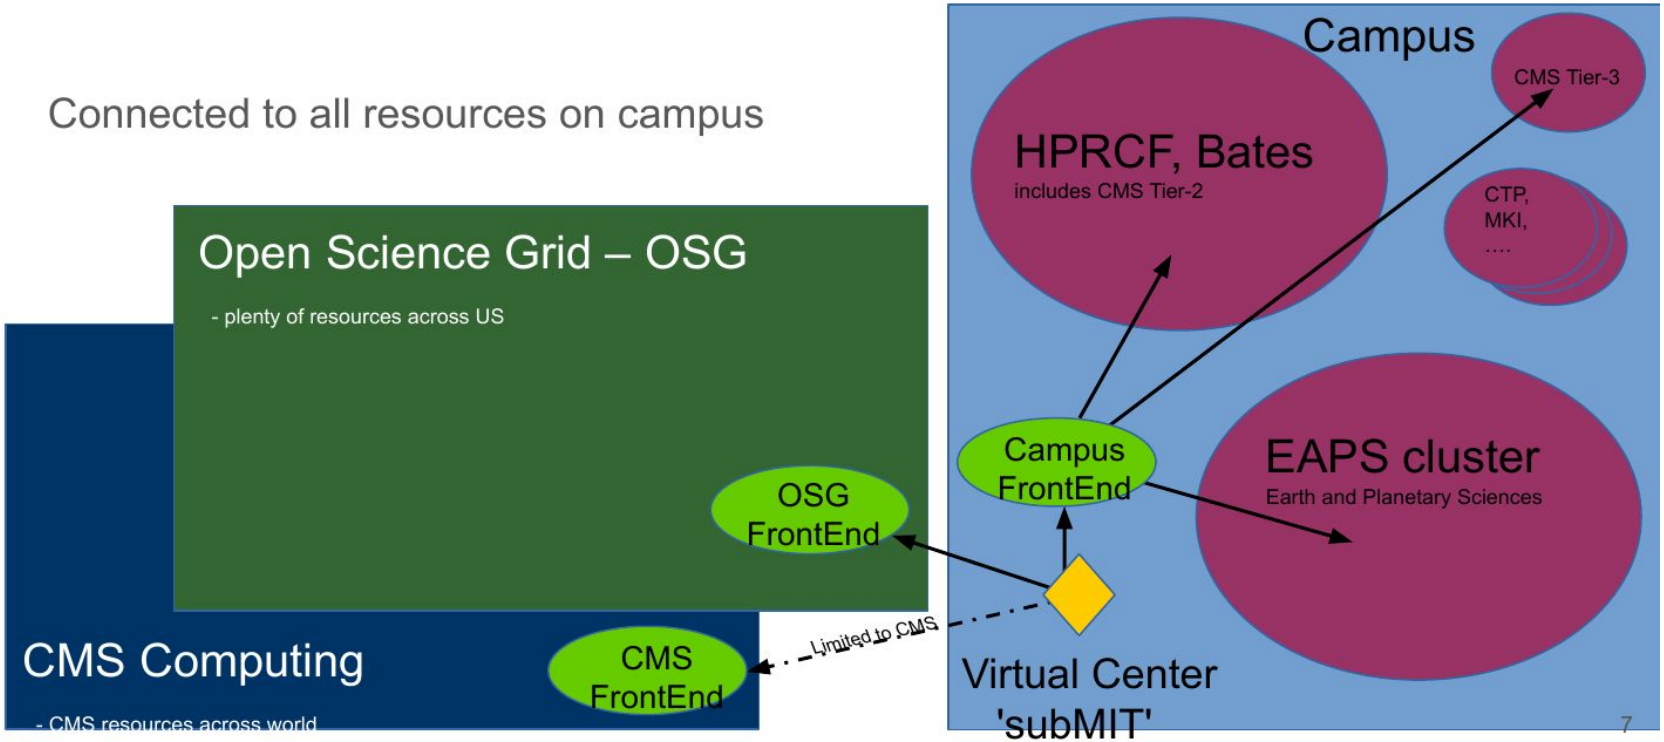

Resource Availability and Monitoring

- Both CPU and GPU resources are available, with access shared between interactive (terminal or Jupyter) and batch usage
- Robust monitoring of batch system usage, network, machine load, etc

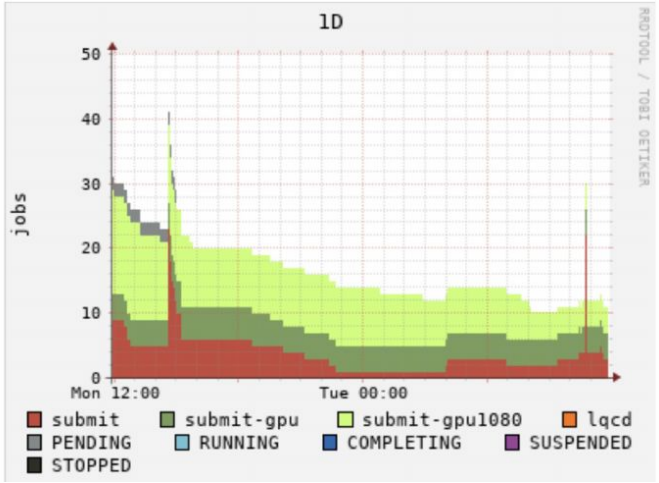

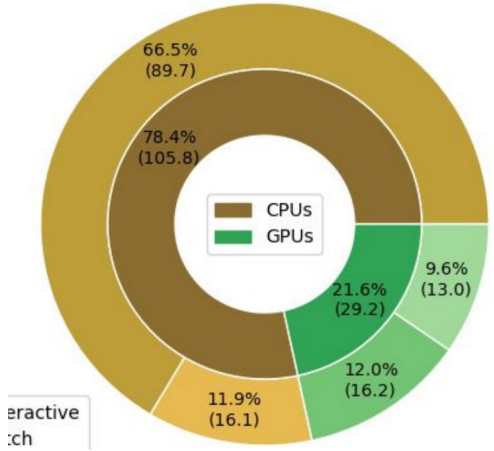

User Outreach and Education

- IAP Seminar in February on usage of computing resources
- Tutorials as part of the User's Guide, including hands-on session in January Users Workshop
- Next tutorial session planned for Users Group meeting in September (synchronized with new arrivals in the department)
- Dedicated user survey of MKI community

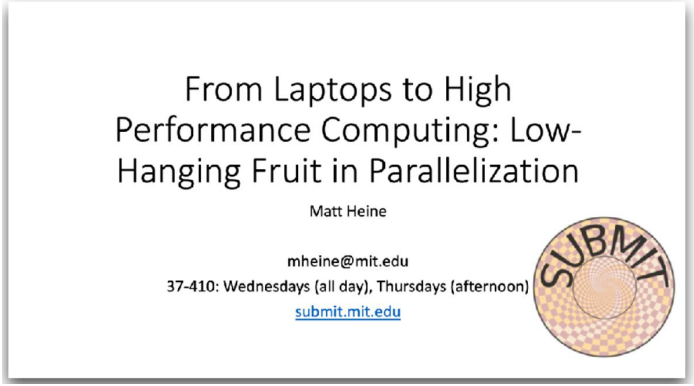

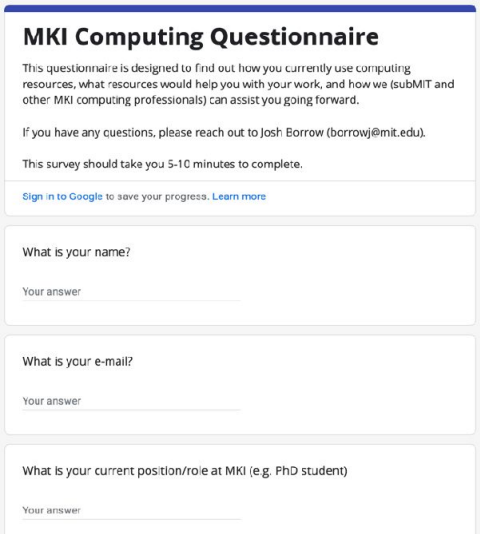

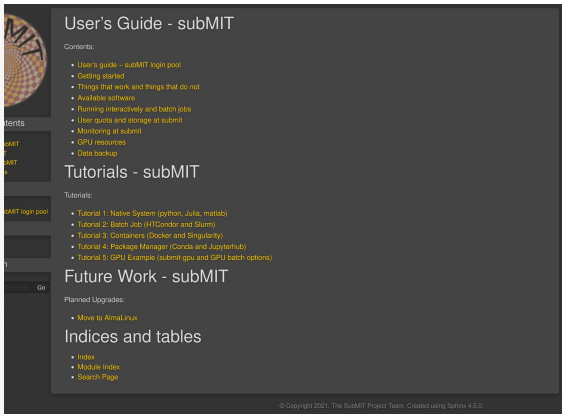

Conclusions

- subMIT system is already up and running and successfully supporting research across the department
- Strong focus on user support by project team with direct connection and expertise in the physics department
- Ongoing upgrades/consolidation/evolution of infrastructure and services in consultation with users
- Longer term relationship of subMIT with ORCD and balance between institute-wide and department-level services to be defined in coming year(s)