

subMIT Users Group: Kickoff Meeting

Josh Bendavid
Nov. 22, 2022

Introduction

- subMIT system provides an interactive login pool
 - 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
 - Local storage (1TB/user), 10's of TB for larger group datasets
 - Convenient access to larger external resources (OSG, CMS Tier-2 and Tier-3, LQCD Cluster, EAPS)
- Website (with User's Guide/Instructions): <https://submit.mit.edu/>
- 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
- 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

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, 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
 - Currently experimental/under development
 - /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
- Please ask if you have any questions about which storage to use or if more space is needed for a particular purpose and we will find a good solution for you

Communication Channels

- User support mailing list: submit-help@mit.edu
- Slack workspace: <https://mit-submit.slack.com>
 - “help-desk” channel
- This meeting (monthly)
 - Open for discussion
 - Open for user contributions: contact the full set of Users Group representatives at submit-usersgroup@mit.edu if you have a presentation or discussion topic for an upcoming meeting
- Discuss with Users Group representative from your group or “nearby” group
- Discuss with your group and identify a Users Group representative

Users Group Kickoff

- Users group has been formed with initial set of representatives (JB as coordinator)
 - Users Group representative (associated faculty)
 - Yin Lin (Phiala Shanahan)
 - Siddharth Mishra-Sharma (Jesse Thaler)
 - Prajwal Mohan Murthy (Bob Redwine)
 - Sunghan Ro (Julien Tailleur)
 - Yitian Sun (Tracy Slatyer)
 - Molly Taylor (LNS Heavy Ion Group)
 - More representatives to come in the near future
- Today: Short introduction from some of the Users group representatives

Important Topic Requiring User Feedback

- 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 created uncertainty for upgrade paths and disrupted logical transition from CentOS 7->8
 - Decision on next distribution to be carefully considered
 - Ease of transition
 - Support lifetime
 - Functionality
 - Consider direction being taken at other universities and labs
 - Discussion to be held in coordination with Users Group
 - Ease transition for users through well-supported and documented use of containers

Workshop in January

- Full day workshop Jan. 6, 2023
 - <https://indico.mit.edu/event/647/>
 - Developing agenda, announcement email soon
 - Mix of contributions from users, Users Group, Project Team
 - Proposed contributions welcome
- See also agenda of previous workshop from Jan. 2022
 - <https://indico.cern.ch/event/1108229/>