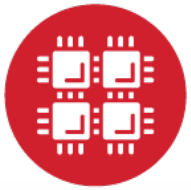

Ohio Supercomputer Center

An OH·TECH Consortium Member

SUG Meeting: Face-to-Face with OSCHelp

Summer Wang
xwang@osc.edu
December 3 2015

Storage: OSC File Space Information

- Scratch – Parallel File System (Temporary Space)
  - ~570 TBs (all disk)
- GPFS
  - ~1.1PB total usable (Disk)
  - Hierarchical storage capable to tape subsystem
  - Allocated to projects in TBs, for limited time periods
- Home Directory Space / NFS
  - ~295 TBs usable (Disk)
  - Allocated to each user, 500 GB quota limit

Mass Storage Overview

- 2 Petabytes (PBs) of usable disk
- 1100 TBs GPFS storage
- 570 TBs Scratch storage
- 1.8 PBs tape

Home Directories

- Each user has a home directory
- Visible from all OSC systems
- Backed up daily – “permanent storage”
- Quotas
  - 500GB of storage per user account
  - 1,000,000 files maximum
  - Cannot create new files if over quota
  - Quota and usage info displayed at login

Project Directories

- PI may request project directory if more space needed
  - Send request to OSC Help
  - Large requests are reviewed by SUG Allocations Committee (>5TB)
  - Shared by all users in the project
- Backed up daily
- Visible from all OSC systems
- Project quota is separate from the home directory quota

Sample Quota Display

Quota display at login (information collected nightly):

```
As of 2010 Jul 15 04:02 userid usr1234 on /nfs/06 used 28GB of
quota 500GB and 41374 files of quota 1000000 files
As of 2010 Jul 16  04:02 project/group PRJ0321 on /nfs/proj01
used 27GB of quota 5000GB and 573105 files of quota 1000000
files
```

Output from **`quota`** command (run manually):

```
Disk quotas for user usr1234 (uid 11059):
Filesystem
     blocks   quota   limit   grace   files   quota   limit   grace
fs06-oak.ten.osc.edu:/nfs/06/osc
    201698292  450000000 524288000          631137  950000 1000000
```

Parallel File System – Scratch

- Designed to handle heavy parallel I/O load
- Faster access than home and project directories
- NOT good for small files
- Visible from all cluster nodes (shared)
- Suitable for short-term storage (up to 180 days) of large amounts of data
- Also useful as batch-managed temporary storage
- **NOT backed up**

https://www.osc.edu/supercomputing/storage-environment-at-osc/available-file-systems

- Scratch file system: purge every Wednesday
- GPFS (project space storage): OSC will move to a sliding scale pricing model for project space allocations in 2016. This change is necessary to address the increasing demand and associated cost with providing project space allocations to our users. As an added side effect, we hope to curb uncontrolled usage of the service and incentivize responsible data management.
- Tape: used for archive/inactive data (starting from 1TB)
- HOWTO: Reduce Disk Space Usage (https://www.osc.edu/content/howto_reduce_disk_space_usage)

GPU Computing

- Oakley GPU computing
  - 128 NVIDIA Tesla M2070 GPUs
  - 64 of the standard nodes have 2 GPUs each
- Ruby GPU computing
  - 20 NVIDIA Tesla K40 GPUs
- CUDA Support
  - https://www.osc.edu/supercomputing/software/CUDA
  - NCIS online training: http://www.nics.tennessee.edu/hpc-seminar-series

Help optimize code at OSC

- We do support it
  - What language is your software written in?
  - Is it parallelized?
  - What operating system(s) has it been run on?

HPC Client Services

- Technical Assistance
  - Help desk and basic consulting
  - Contact by phone or email (oschelp@osc.edu)
- Facilitation
  - Meet with OSC staff to discuss your research needs
  - Get recommendations on services, connections to subject matter experts, and specialized projects initiated
- Project Administration
  - Manage allocations
  - Add/Remove authorized users
  - Utilization reports
- Training
  - Usually three workshops per semester on a variety of topics
- Advanced consulting
  - Code parallelization & optimization
  - Software development, algorithm research
- Website
  - www.osc.edu/supercomputing

What can OSC provide you?

- "Capability computing" (High Performance Computing)
  - Computation too large to run on laptop/desktop
- "Capacity computing" (High Throughput Computing)
  - Takes too long on laptop, need to make many runs
- Data Analytics
  - Massive memory requirements
- Access to licensed software
  - Have academic licenses for many commercial packages
- Expertise, collaboration
  - Parallel computing, algorithms, web portals, etc.

Software Support at OSC

- https://www.osc.edu/supercomputing/software-list

Apply Accounts

- Academic accounts
  - Principal investigator (PI) must be a full-time faculty member or research scientist at an Ohio academic institution
  - PI may authorize accounts for students, post-docs, collaborators, etc.
  - Classroom accounts are also available
  - No cost to Ohio academic users
- https://www.osc.edu/supercomputing/getting-started/allocations-and-accounts

Getting an Account

- Startup grant
  - One per PI per lifetime
  - Provide contact info, institution, department
  - 5000 RUs
- Additional allocations for a project
  - Submit a proposal for more RUs
    - Standard: 10,000
    - Major: 30,000
    - Discovery: >30,000
  - Peer-reviewed
  - Grants awarded by Statewide Users Group (SUG)
- Classroom account:
  - 5000 RUs

Scheduling Issue

- Our scheduling policy of job priorities was changed early October. One outcoming is that serial job (requesting n=1) with large walltime will have much longer queue time than before. To shorten the queue wait time of your jobs:
  - Try to split the jobs into several shorter ones if possible
  - Make the code/jobs parallel and shorten the walltime if possible. Jobs requiring more nodes with smaller walltime usually have shorter queue time than jobs requiring 1 node with larger walltime
  - Make the walltime as accurate as possible. If you could not make the job parallel, try to provide the walltime as accurate as possible. A job requiring 1 node with a couple of hours will get started significantly faster
  - Switch to Glenn cluster