

OSC Service Catalog

Cluster Computing

- High Performance Computing
- High Throughput Computing
- Data-intensive Computing

Research Data Storage

- Project Storage
- Archival Storage

Client Services

- 24x7 Call Center
- Level 2 Engineering Support

Client Facilitation

- Consultation (in-person and online)
- Training and Education
- Classroom accounts

Scientific Software Development

- Software Development
- Software Parallelization

Web Software Development

- Software Development
- Software Consulting

Partner on Proposals

- Cyberinfrastructure solutions
- Modeling & simulation for industry

Visualization & Virtual Environments

- Visualization Services
- Virtual environments (DSL)




HPC Terminology


- Cluster
 - A group of computers (nodes) connected by a high-speed network, forming a supercomputer
- Node
 - Equivalent to a high-end workstation, part of a cluster
- Core
 - A processor (CPU), multiple cores per processor chip
- FLOPS
 - “FLoating-point Operations (calculations) Per Second”



Structure of a Supercomputer



Structure of a Supercomputer




System Configurations



	Owens (2016)	Ruby (2014)	Oakley (2012)
Theoretical Performance	~750 TF	~144 TF	~154 TF
# Nodes	~820	240	692
# CPU Cores	~23,500	4800	8304
Total Memory	~120 TB	~15.3 TB	~33.4 TB
Memory per Core	>5 GB	3.2 GB	4 GB
Interconnect	EDR IB	FDR/EN IB	QDR IB



Owens Cluster Specifications



Login Nodes – Usage

- Purpose
 - Submit jobs to batch system
 - Edit files
 - Manage your files
 - Interactive work – small scale
- Limits
 - 20 minutes CPU time
 - 1GB memory
- Use the batch system for serious computing!



Filesystem Overview

- Home
 - Store your files here, backed up daily
 - Use \$HOME or *~username* to reference location
- Project
 - Available to Project PIs by request; shared by all users on a project, backed up daily
 - Use */fs/project/project#* to reference location
- Scratch
 - Store large input or output files here
 - Faster I/O than Home or Project
 - Temporary storage, not backed up
- \$TMPDIR
 - Storage on compute nodes, for use during your batch job
 - Be sure to copy any results back to Home at the end of your job, all data purged when job quits



Filesystem	Quota	Backed-Up?	Purged?
Home (\$HOME)	500GB	Yes	No
Project (/fs/project)	By request	Yes	No
Scratch (/fs/scratch)	None	No	Yes – 120 days
Compute (\$TMPDIR)	800GB (Oakley), 1 TB (Ruby & Owens)	No	Yes – when job completes

File Management

- If you are concerned about Home directory quotas:
 - Compress large, rarely used files
 - Use `gzip` or `bzip2` commands
 - Combine large numbers of small files into an archive
 - Use `tar` command
 - Request Project space for your group (PIs only)
 - Large requests are reviewed by allocations committee
 - Contact OSC Help to initiate request



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	4500000000	524288000			631137	950000	1000000

Linux Operating System

- “UNIX-like”
- Widely used in HPC
- Mostly command-line
- Choice of shells (bash is default)
- Freely distributable, open-source software
- Tutorials available
- www.linux.org



Connecting to the clusters

- Connect to OSC machines using **ssh** (secure shell)
 - From a Linux/UNIX machine : At prompt, enter
`ssh userid@oakley.osc.edu`
`ssh userid@ruby.osc.edu`
 - From a Mac: Enter **ssh** command in TERMINAL window
 - From Windows: **ssh** client software needed
 - Both commercial and free versions are available
- New: Connect using OnDemand portal (web-based)



Connecting to an OSC Cluster with Graphics

- Programs on the cluster can have an X-based GUI
 - Display graphics on your computer
- Linux/UNIX and Mac: Use **-X** flag

```
ssh -X userid@oakley.osc.edu
```

- Windows: Need extra software
 - Both commercial and free versions are available
 - Configure your ssh client to tunnel or forward X11
- Primarily used with programs on login node
 - Can also use with interactive batch jobs



OSC OnDemand ondemand.osc.edu

- 1: User Interface
 - Web based
 - Usable from computers, tablets, smartphones
 - Zero installation
 - Single point of entry
 - User needs three things
 - ondemand.osc.edu
 - OSC Username
 - OSC Password
 - Connected to all resources at OSC
 - 2: Interactive Services
 - File Access
 - Job Management
 - Visualization Apps
 - Desktop access
 - Single-click apps (Abaqus, Ansys, Comsol, Paraview)
 - Terminal Access
- Tutorial available at osc.edu/ondemand**



Message of the Day

2016/05/24 - SYSTEM DOWNTIME: JUNE 7TH

A downtime is scheduled for all HPC systems starting June 7th. The downtime will affect all clusters and services. Login shells will not be available during this time.

In preparation for the downtime the batch scheduler will be restarted on June 7th. Jobs that are held will be scheduled after the system

Virtual Desktop Interface

🖥️ Oakley VDI

📺 Paraview

Compute Node Session

🖥️ Oakley Desktop

📊 Abaqus/CAE

🔧 ANSYS Workbench

📐 COMSOL

scheduled to finish by 5PM. All systems will not be available during this time.

Jobs will be complete before 6AM June 7th.

Oakley Cluster Status

600 of 602 Nodes Active

99.7%

7765 of 8072 Processors Active

96.1%

2617 Total Jobs

772 Active

497 Eligible

1348 Stalled

Ruby Cluster Status

272 of 246 Nodes Active

110%

4032 of 4032 Processors Active

100%

291 Total Jobs

24 Active

104 Eligible

43 Stalled

Transferring Files to and from the Cluster

- Most file transfers to and from OSC machines use `sftp` or `scp`
 - Linux and Mac have them built in
 - Windows needs extra software - FileZilla
- For small files, connect to login node `oakley.osc.edu`
- For large files, transfer may fail due to shell limits
 - Connect to `gridftp01.osc.edu` (file transfer only)



File Permissions


- By default all files are readable by all users
- Check permissions using `ls -l`

```
-rw-r--r-- 1 osu7824 PAS0925 10839 Jan 13 2015 triarm_VVVacid.sdf
-rw-r--r-- 1 osu7824 PAS0925 11667 Jan 13 2015 triarm_VVvester.sdf
drwxr-xr-x 8 osu7824 PAS0925 4096 Jan 16 2014 tutorial
-rw-r-x--- 1 osu7824 PAS0925 9917889 Jan 15 2015 ValBaskEst32_gopt.log
-rw-r--r-- 1 osu7824 PAS0925 12818 Jan 15 2015 ValBaskEst32_gopt.mol2
-rwxr-xr-x 1 osu7824 PAS0925 453376 Feb 26 2015 ValBaskEst_c0_ValBaskEst0-CyHexPO-2OMe-cl1_md1.mdcrd
```

- `-rwxr-xr-x` User, Group, Others
- Change file permissions using `chmod`
`chmod u=rw,g=r file`

`chmod -R u=rw,g=r directory`






Ohio Supercomputer Center

An OH·TECH Consortium Member

Batch Processing

Why do supercomputers use queueing?



Batch System at OSC

- Compute nodes are allocated through the batch system
 - PBS – Portable Batch System
 - Torque – resource manager
 - Moab – scheduler
- Documentation at
www.osc.edu/supercomputing/batch-processing-at-osc



Idea Behind Batch Processing

- Whatever you would normally type at the command prompt goes into your batch script
- Output that would normally go to the screen goes into a log file (or files)
- The system runs your job when resources become available
- Very efficient in terms of resource utilization



Steps for Running a Job on the Compute Nodes

1. Create a batch script for a job
2. Submit the job
3. Job gets queued
4. Job runs when resources become available
5. Get your results when the job finishes



Specifying Resources in a Job Script

- Nodes and cores (processors) per node
- Memory
- GPUs
 - See “Batch Processing at OSC” on OSC website
- Walltime
 - Overestimate slightly – job will be deleted if it hits limit
 - Shorter job may start sooner due to backfill
- Software licenses
 - See specific software page on OSC website



Sample Batch Script

```
#PBS -N serial_fluent
#PBS -l walltime=1:00:00
#PBS -l nodes=1:ppn=28
#PBS -j oe
#PBS -l software=fluent+1
# Set up the FLUENT environment
module load fluent
# Move to directory job was submitted from
cd $PBS_O_WORKDIR
# Copy input files to compute node
cp run.input $TMPDIR
cd $TMPDIR
# Run fluent and copy results back to home
fluent 3d -g < run.input
cp 'results*' $PBS_O_WORKDIR
```

Job setup information
for PBS

This is a comment

Commands
to be run

Put all this into a text file!



Submitting a Job and Checking Status

- Command to submit a job
 - `qsub script_file`
- Response from PBS (example)
 - `123456.oak-batch.osc.edu`
- Show status of batch jobs
 - `qstat -a jobid`
 - `qstat -u username`
 - `qstat -f jobid`

[List of Batch commands](#) on osc.edu



Scheduling Policies and Limits

- Walltime limit
 - 168 hours for serial jobs (single node)
 - 96 hours for parallel jobs (multiple nodes)
- Per-user limits
 - 128 concurrently running jobs
 - 2040 processor cores in use
 - 1000 jobs in the batch system, running or queued
- Per-group limits
 - 192 concurrently running jobs
 - 2040 processor cores in use



Waiting for Your Job To Run

- Queue wait time depends on many factors
 - System load
 - Resources requested
 - nodes, cores, large memory, gpus, software licenses
 - Fair share limits (if load is high)
 - reduced priority for users or groups using a lot of resources



Job Output

- Screen output ends up in file `job_name.ojobid`
 - Copied to your working directory when job ends
 - Example: `testjob.o1234567`
- To see screen output while job is running
 - `qpeek jobid`
 - Example: `qpeek 1234567`



Interactive Batch Jobs

- Interactive, but handled through batch system
 - Resource limits same as standard batch limits
- Useful for tasks forbidden on login nodes
 - Debug parallel programs
 - Run a GUI program that's too large for login node
- May not be practical when system load is high
 - Long wait, same as standard batch job
- To submit an interactive batch job (example)
 - `qsub -I -X -l nodes=2:ppn=28 -l walltime=1:00:00 -m abe`



Batch Queues

- The three clusters have separate batch systems
 - Submit job and check status on the same cluster
- Debug reservation
 - A few nodes on each system are reserved for short jobs (≤ 1 hour)
 - Special flag required on Ruby and Owens: `-q debug`



Parallel Computing

- Each processor is fast, but real speed comes from using multiple processors
- Multithreading
 - Use multiple cores on a single node
 - Shared memory
- Message passing (MPI)
 - Use one or multiple nodes
 - Distributed memory



Modules for Software access

- How modules work
 - Modify environment variables like `$PATH` and `$MANPATH` within your shell
- Default set of modules loaded at login
 - module system, batch system (do not unload)
 - default compiler and MPI modules
- Do NOT completely replace `$PATH` in your `.bash_profile` or `.bashrc`
- DO prepend directories to the existing `$PATH`
 - Type: `export PATH=$HOME/bin:$PATH`



Module Commands

- What modules do you have loaded?
 - `module list`
- What modules are available?
 - `module spider` or `module avail`
- Multiple versions of the same software
 - `module avail intel`
- Add a software module to your environment
 - `module load cuda`
- Remove a software package from your environment
 - `module unload intel`
- Load a different software version
 - `module swap intel intel/13.1.3.192`



Adding or Removing Software from Your Environment

- Load the module for the software you need, e.g.,
 - `module load comsol`
- Allows multiple versions of software to coexist on our system
- Allow us to make changes without affecting you
 - PLEASE DON'T HARDCODE PATHS!
- Can load modules at command prompt or in your `.bash_profile` or `.bashrc` file
- Also load modules in your job (batch) scripts



Third party applications

- **General programming software** (statewide licensed)
 - gnu compilers and debugger
 - Intel compilers
 - Totalview debugger
 - Allinea profiler
 - MPI library
 - HDF5
 - NetCDF
 - Java, Java Virtual Machine
 - Python



Access to Licensed Software

- Most software licenses for academic use only
- Some software requires signed license agreement
 - Check website
 - Contact OSC Help
- List of applications can be found at Software page:
<http://www.osc.edu/supercomputing/software/>



OSC doesn't have the software you need?

- Commercial software
 - Fill out a request form (see our FAQ)
 - SUG will consider it
- Open-source software
 - You can install it yourself in your home directory
 - If there's enough demand, we can install it for shared use
- Have your own license?
 - Contact OSC Help



Resources to get your questions answered

- FAQs & HOWTOs on osc.edu
 - https://www.osc.edu/resources/getting_started/supercomputing_faq
 - https://www.osc.edu/resources/getting_started/howto
- System updates
 - Read Message of the Day on login
 - Follow @HPCNotices on Twitter

