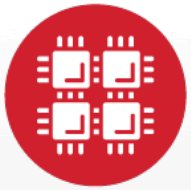

Ohio Supercomputer Center

An OH·TECH Consortium Member

Computing Services to Accelerate Research and Innovation

Kate Cahill
Education & Training Specialist

Outline

- Overview
  - What is OSC?
  - HPC Concepts
  - Hardware Overview
- How to use our systems
  - User Environment
  - Storage
  - Batch Processing
  - Policies
- OSC News

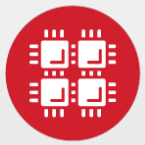

Ohio Supercomputer Center

An OH·TECH Consortium Member

What is the Ohio Supercomputer Center?

The OH-TECH Consortium

Ohio Supercomputer Center provides high performance computing, software, storage and support services for Ohio's scientists, faculty, students, businesses and their research partners.

OARnet connects Ohio's universities, colleges, K-12, health care and state and local governments to its high-speed fiber optic network backbone. OARnet services include co-location, support desk, federated identity and virtualization.

OhioLINK serves nearly 600,000 higher education students and faculty by providing a statewide system for sharing 50 million books and library materials, while aggregating costs among its 90 member institutions.

eStudent Services provides students increased access to higher education through e-learning and technology-enhanced educational opportunities, including virtual tutoring.

Research & Innovation Center will operate, when opened, as the proving grounds for next-generation technology infrastructure innovations and a catalyst for cutting-edge research and collaboration.

Ohio Innovates with the World

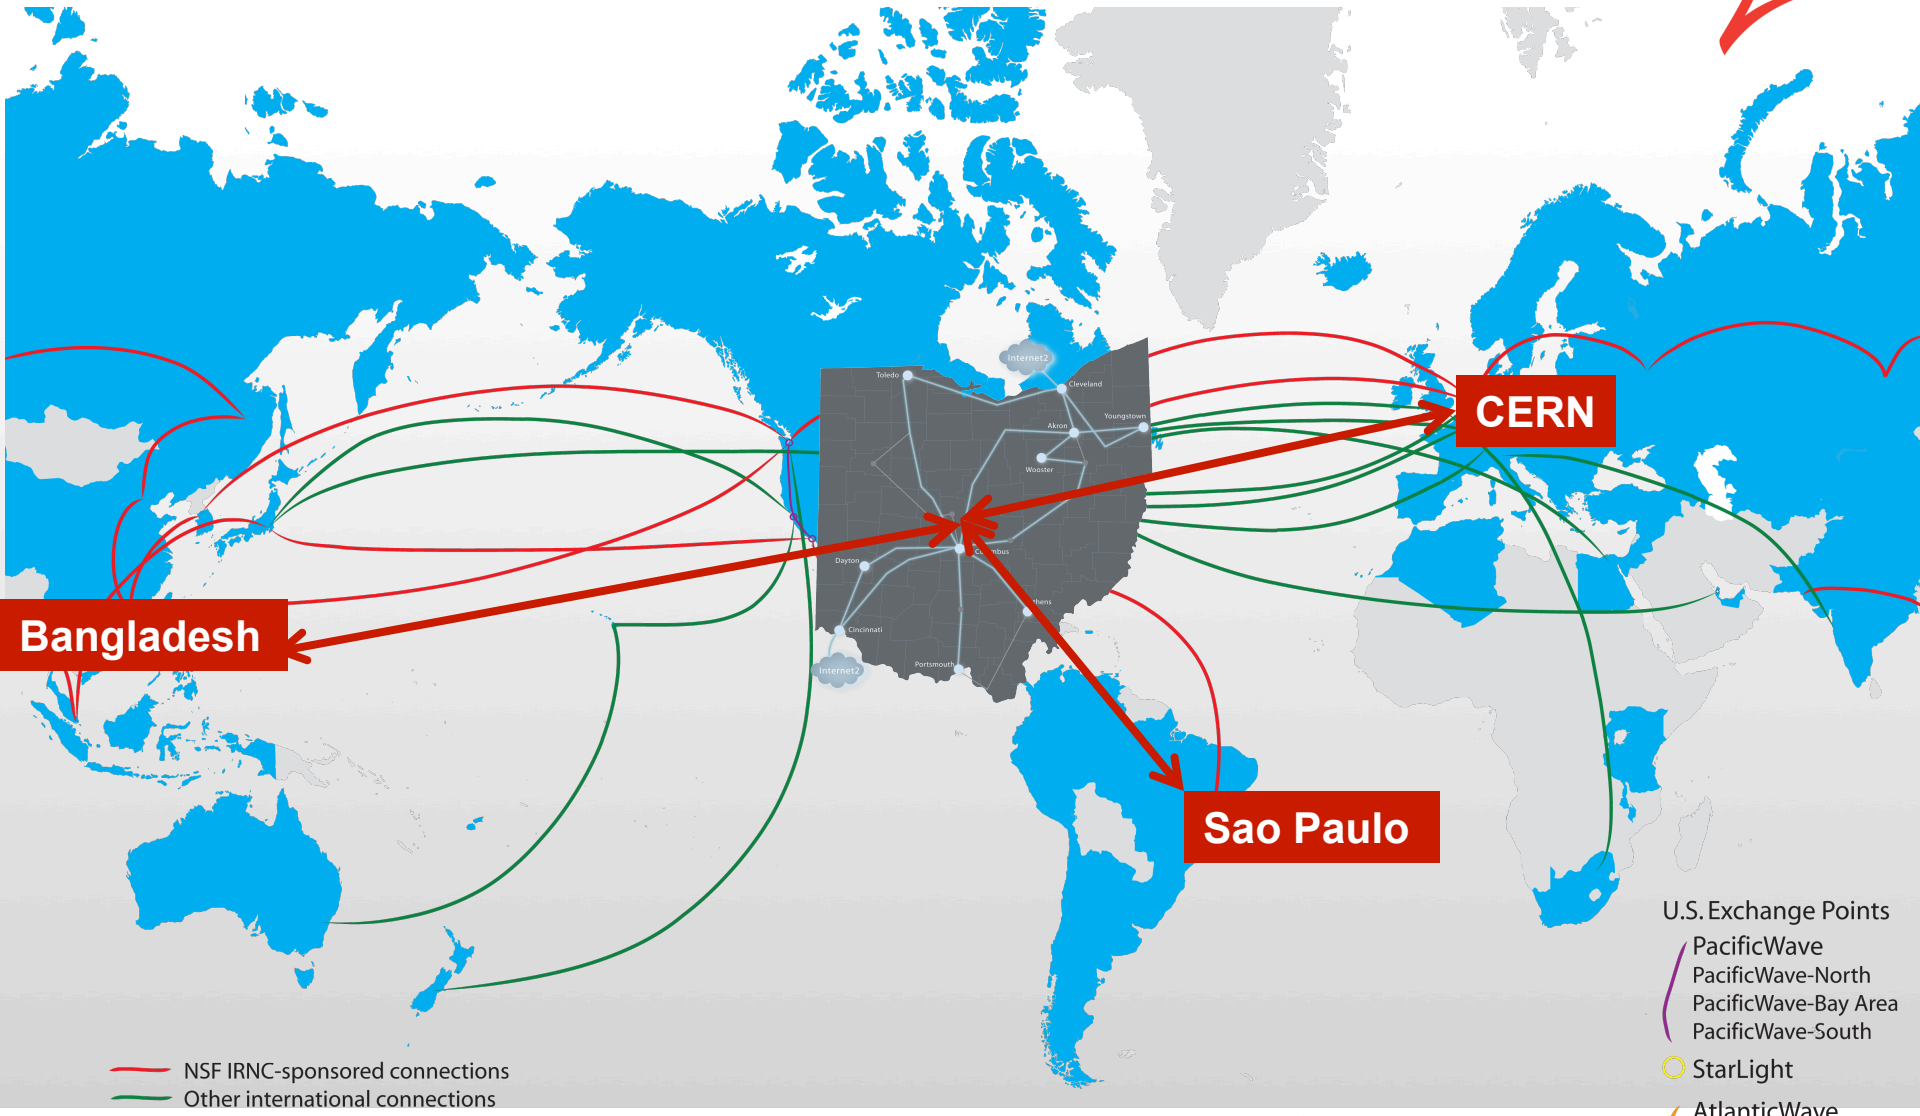

About OSC

- Founded in 1987
- Statewide resource for all universities in Ohio
  - high performance computing services
  - computational science expertise
  - " … propel Ohio's research universities and private industry to the forefront of computational based research."
- Funded through the Ohio Department of Higher Education
- Reports to the Chancellor
- Located on OSU's west campus
- Fiscal agent is OSU

Empowering Clients: Organizational Impact CY2015

Client Services


25 Ohio-based universities


32 companies


1,267 clients


124 trainees


10 educational opportunities


451 awards made


459 projects served


19 academic courses used OSC's supercomputers

Production Capacity


161,000,000+ core-hours consumed


4,000,000+ computational jobs


842 TB data stored


99.7% up-time (target: 96%)

Active Awards
Total: 459

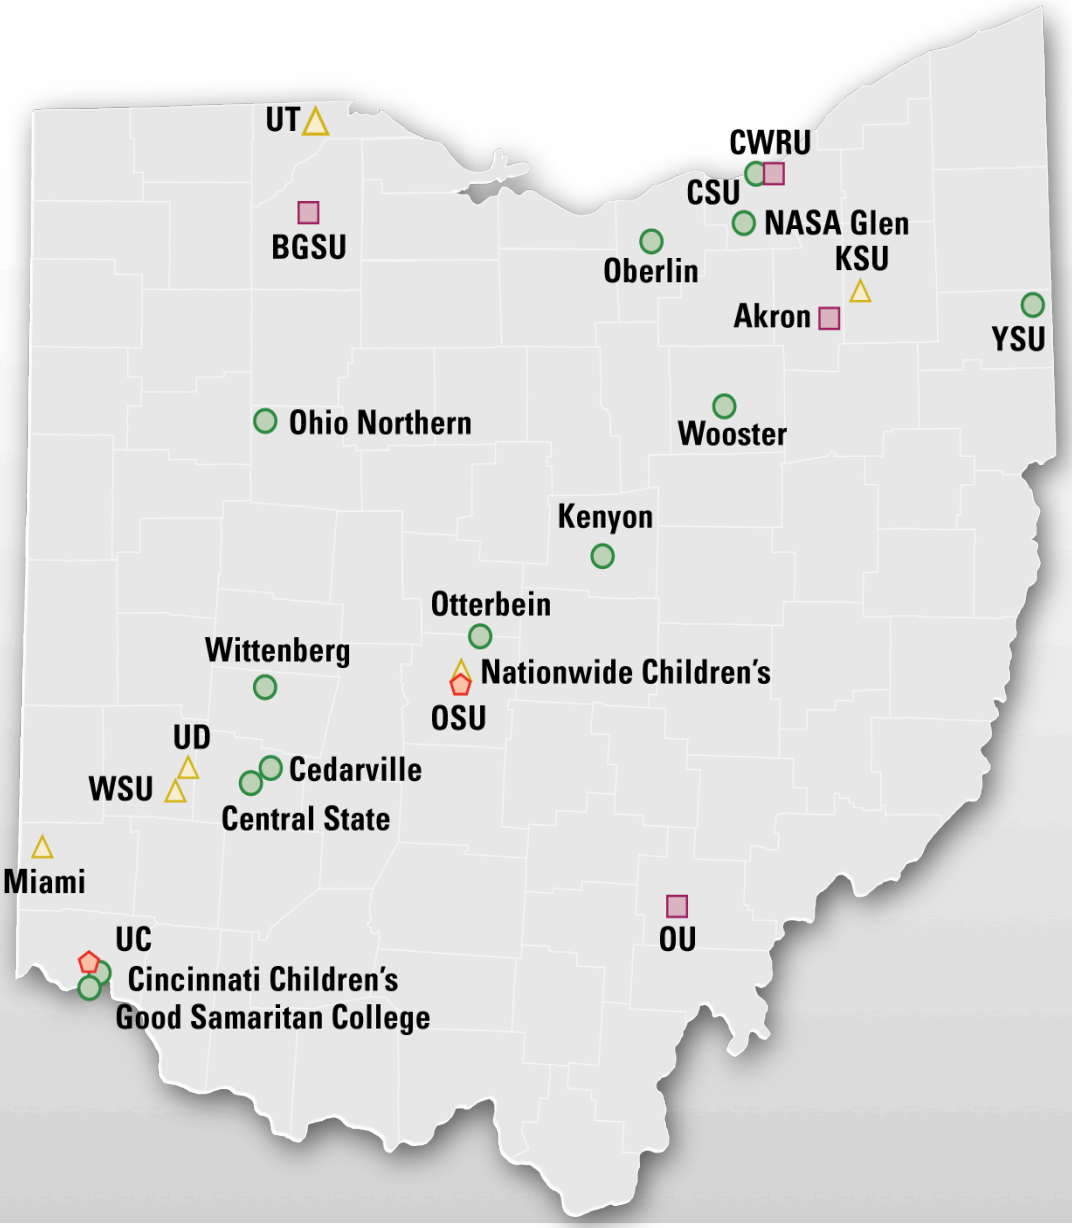

1 – 5

6 – 10

11 – 20

21+

Computing Resource Usage by Field of Science (FoS)

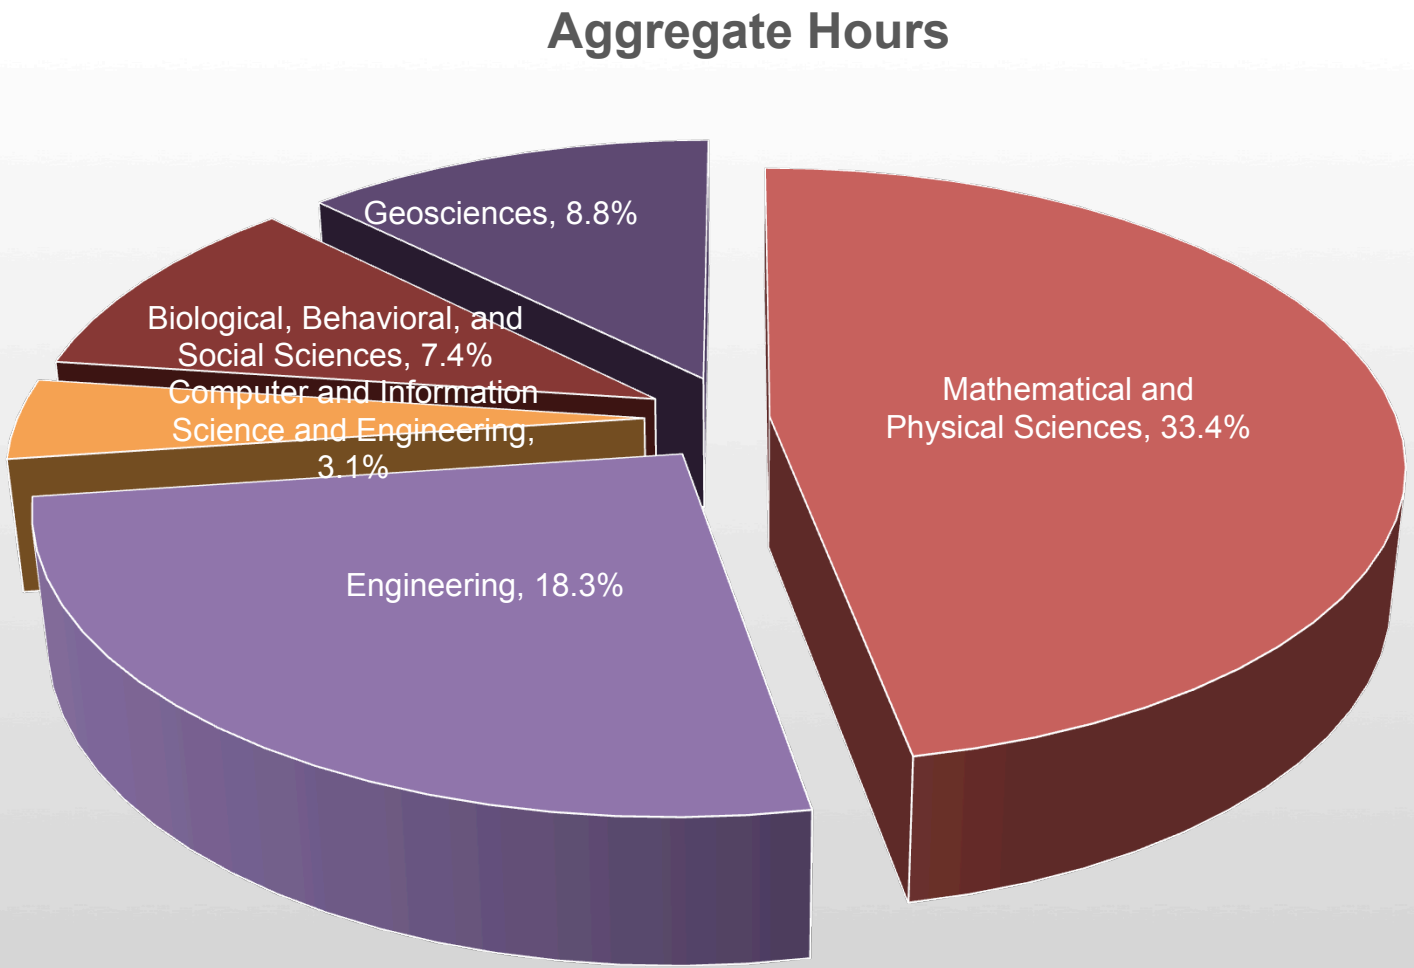

OSC Service Catalog

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 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?

- You can complete your research for less cost.
- You can do more science for the same cost.
- You can get to solution faster.

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.

Statewide Licensed Software

- Use the software in your lab or office
  - Connect to license server at OSC
- Software available
  - Altair Hyperworks
  - Totalview Debugger
  - Intel Compilers, Tools, Libraries
  - Portland Group Compilers
- Contact OSC Help
  - Provide your IP address

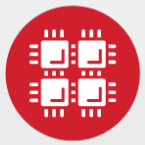

Ohio Supercomputer Center

An OH·TECH Consortium Member

HPC Concepts

Supercomputers become history quickly!

Smartphone - 2015

$740	Cost	**$20,000,000**
4 GB	Memory	**128** MB
64 GB	Storage	**30** GB
302 GFLOPS	Speed	**2** GFLOPS

Supercomputer - 1989

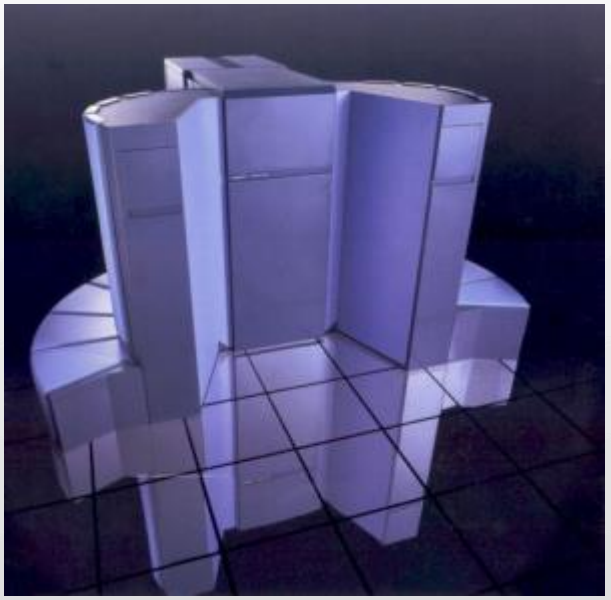

Big Numbers

Prefix	Example: bytes
• K – kilo, 10^3, thousand • M – mega, 10^6, million • G – giga, 10^9, million • T – tera, 10^{12}, trillion • P – peta, 10^{15}, quadrillion • E – exa, 10^{18}, quintillion	• 1KB – very small • 12MB L2 cache per core • 48GB memory per node • .5 TB disk space per user • 4 PB aggregate storage • Exascale systems – current research area

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”

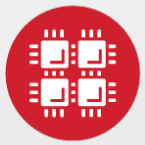

Ohio Supercomputer Center

An OH·TECH Consortium Member

Hardware Overview

Organization of a Supercomputer

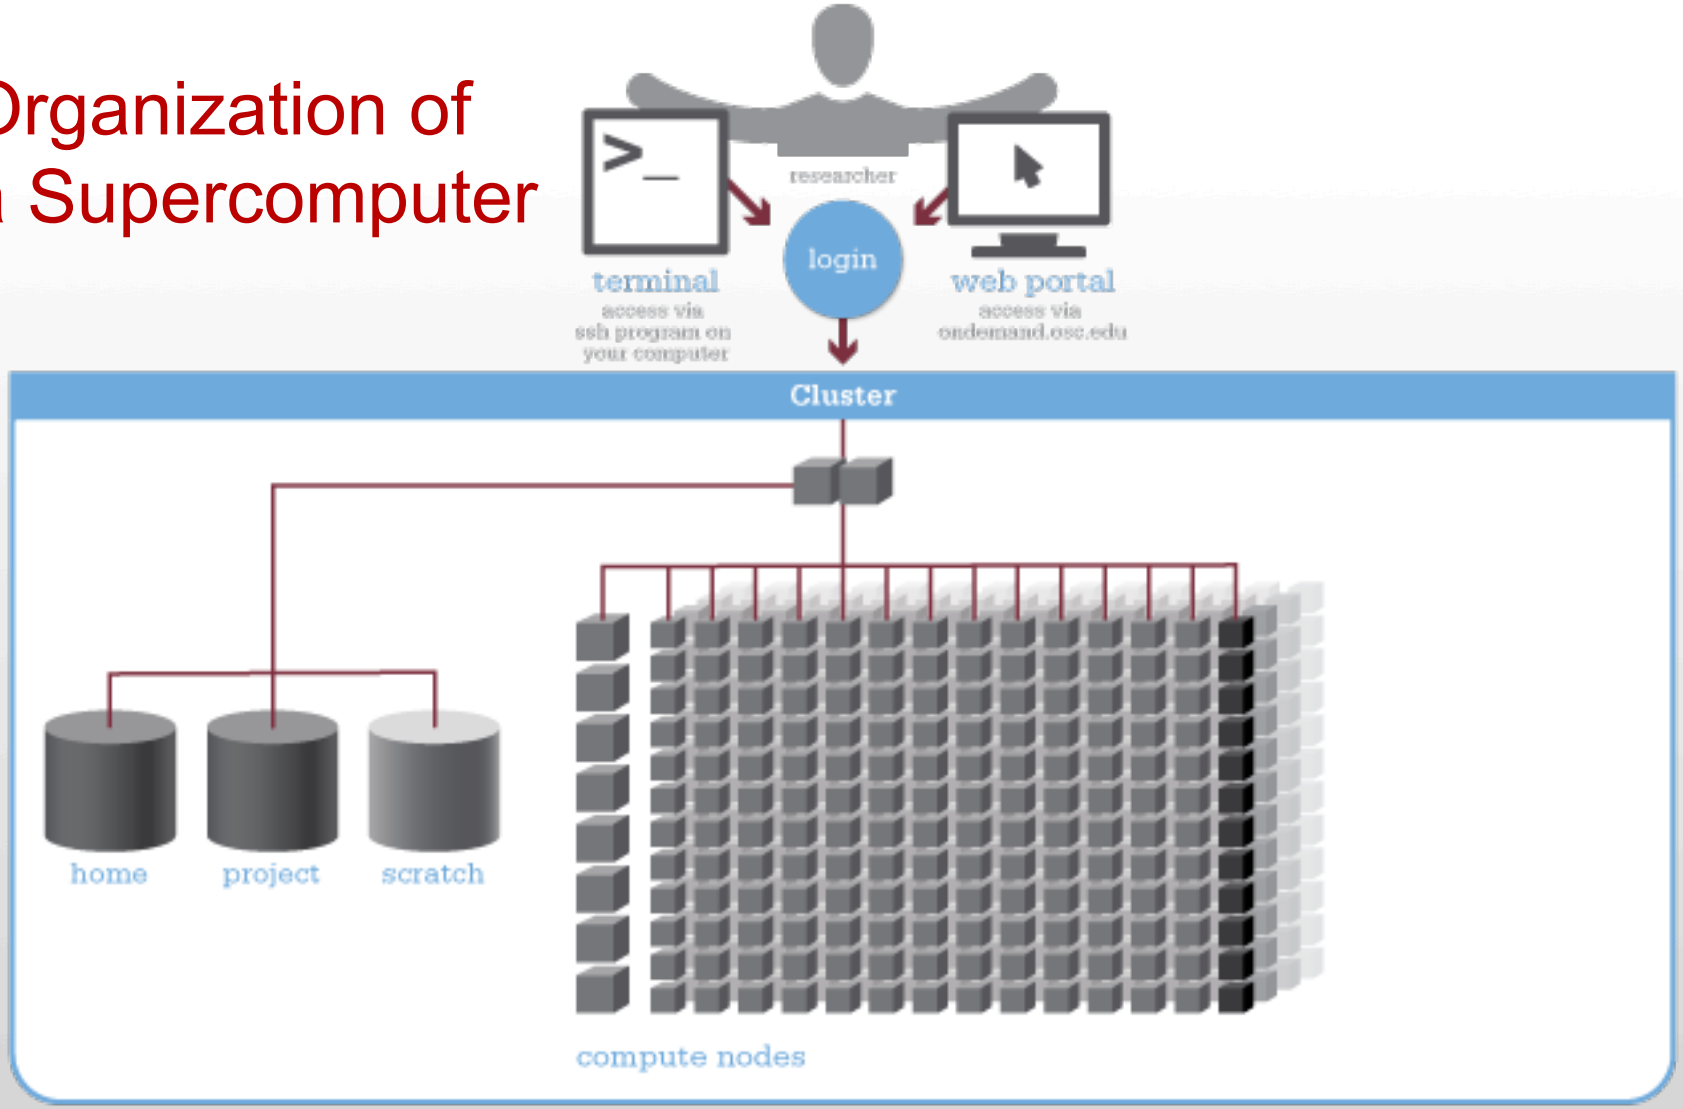

Login Nodes – Usage

- 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!

Supercomputers at OSC

- Ruby cluster (small cluster, limited access)
  - Online March 2015
  - Named for Ruby Dee, actress, poet, playwright, screenwriter, journalist and activist. She was born in Cleveland.
  - HP system, Intel Xeon processors, 4800 cores
- Oakley cluster
  - Online March 2012
  - Named for Annie Oakley, famous Ohio sharpshooter
  - HP system, Intel Xeon processors, 8280 cores
- Glenn cluster
  - "Glenn phase II" online July 2009 – **retired March 2016**
  - Named for John Glenn, Ohio astronaut and senator
  - IBM 1350, AMD Opteron processors, 3500 cores

Oakley Cluster

Login Nodes – Configuration

- Oakley
  - 2 general-purpose login nodes
  - 12 cores, 124 GB memory each
  - Connect to oakley.osc.edu
- Ruby
  - 2 general-purpose login nodes
  - 16 cores, 132 GB memory each
  - Connect to ruby.osc.edu
- Glenn
  - 2 general-purpose login nodes
  - 16 cores, 64 GB memory each
  - Connect to glenn.osc.edu

Compute Nodes – Oakley

- 684 standard nodes
  - 12 cores per node
  - 48 GB memory (4GB/core)
  - 812 GB local disk space
- 8 large memory nodes
  - 12 cores per node
  - 192 GB memory (16GB/core)
  - 812 GB local disk space
- Network
  - Nodes connected by 40Gbit/sec Infiniband network (QDR)

Oakley Cluster Specifications

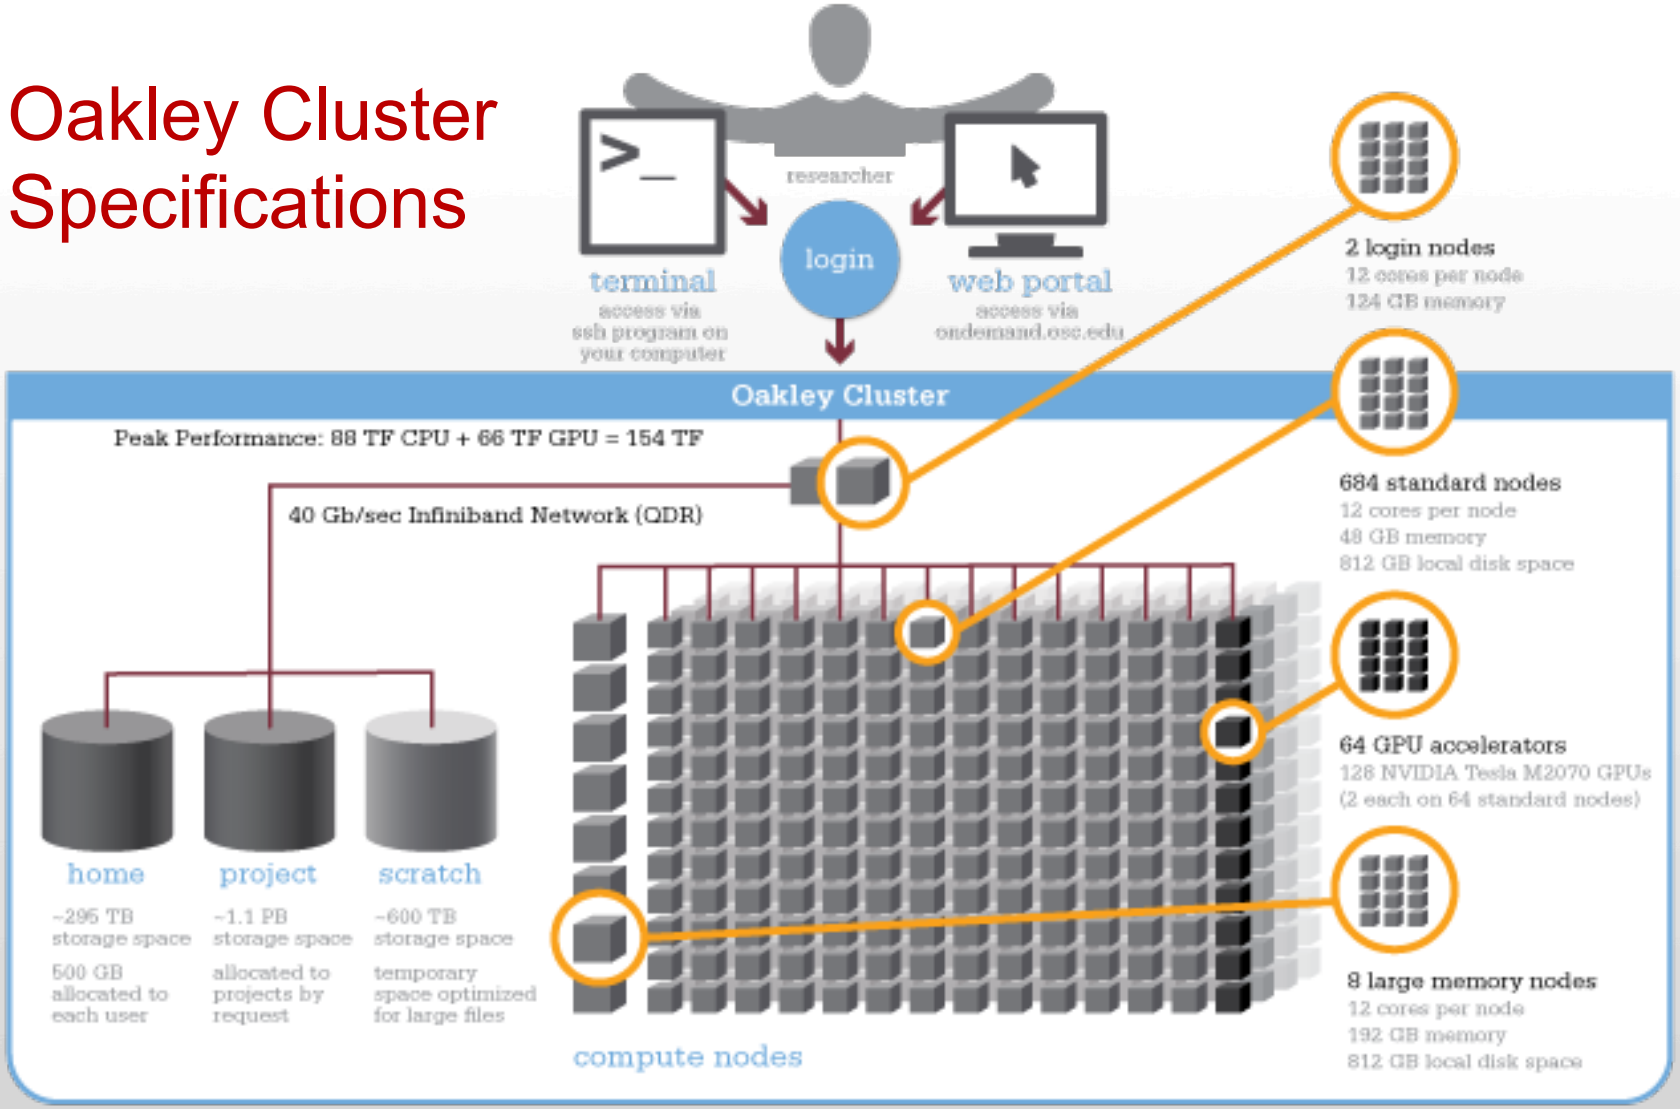

Specs: Oakley Cluster vs. Top 500 Systems in the World

Metric	June 2012	June 2012	November 2012	November 2012	June 2013
	Performance Ranking	Efficiency Ranking	Performance Ranking	Efficiency Ranking	Ranking
Overall Ranking in the World	180th	37th	460th	30th	Not Listed
Overall Ranking in US	89th	8th	235th	8th	Not Listed
Overall Academic Ranking in the World	40th	9th	91st	13th	Not Listed
Overall Academic Ranking in US	11th	2nd	23rd	2nd	Not Listed

OSC's Newest HPC System: Ruby Cluster

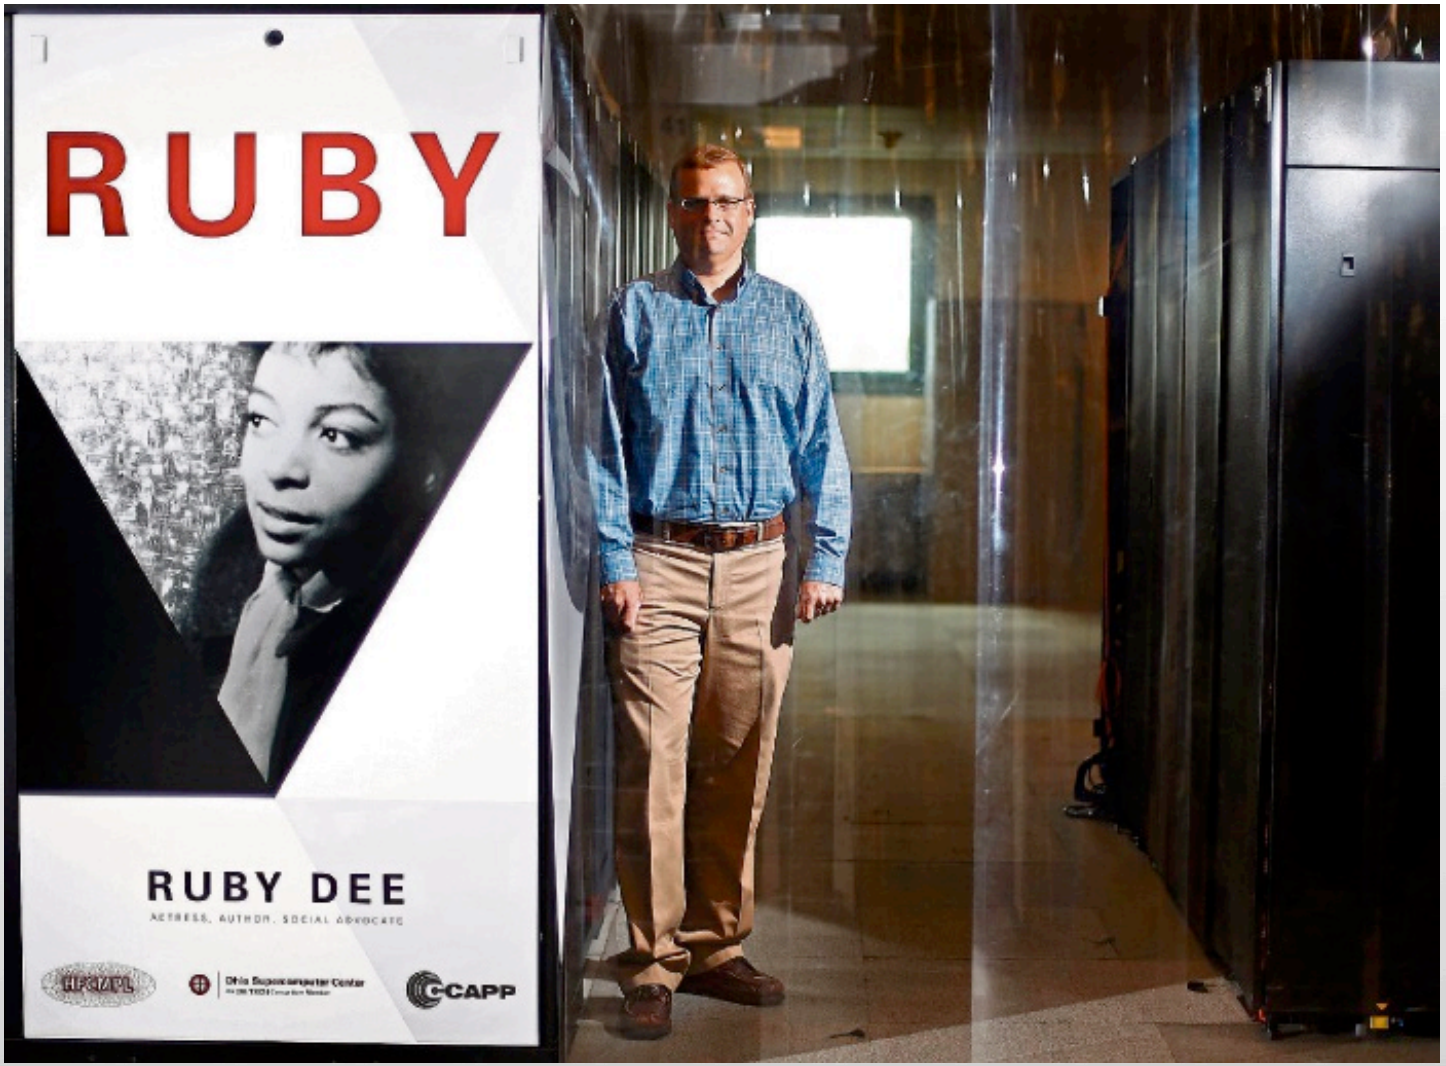

Ruby Cluster Specifications

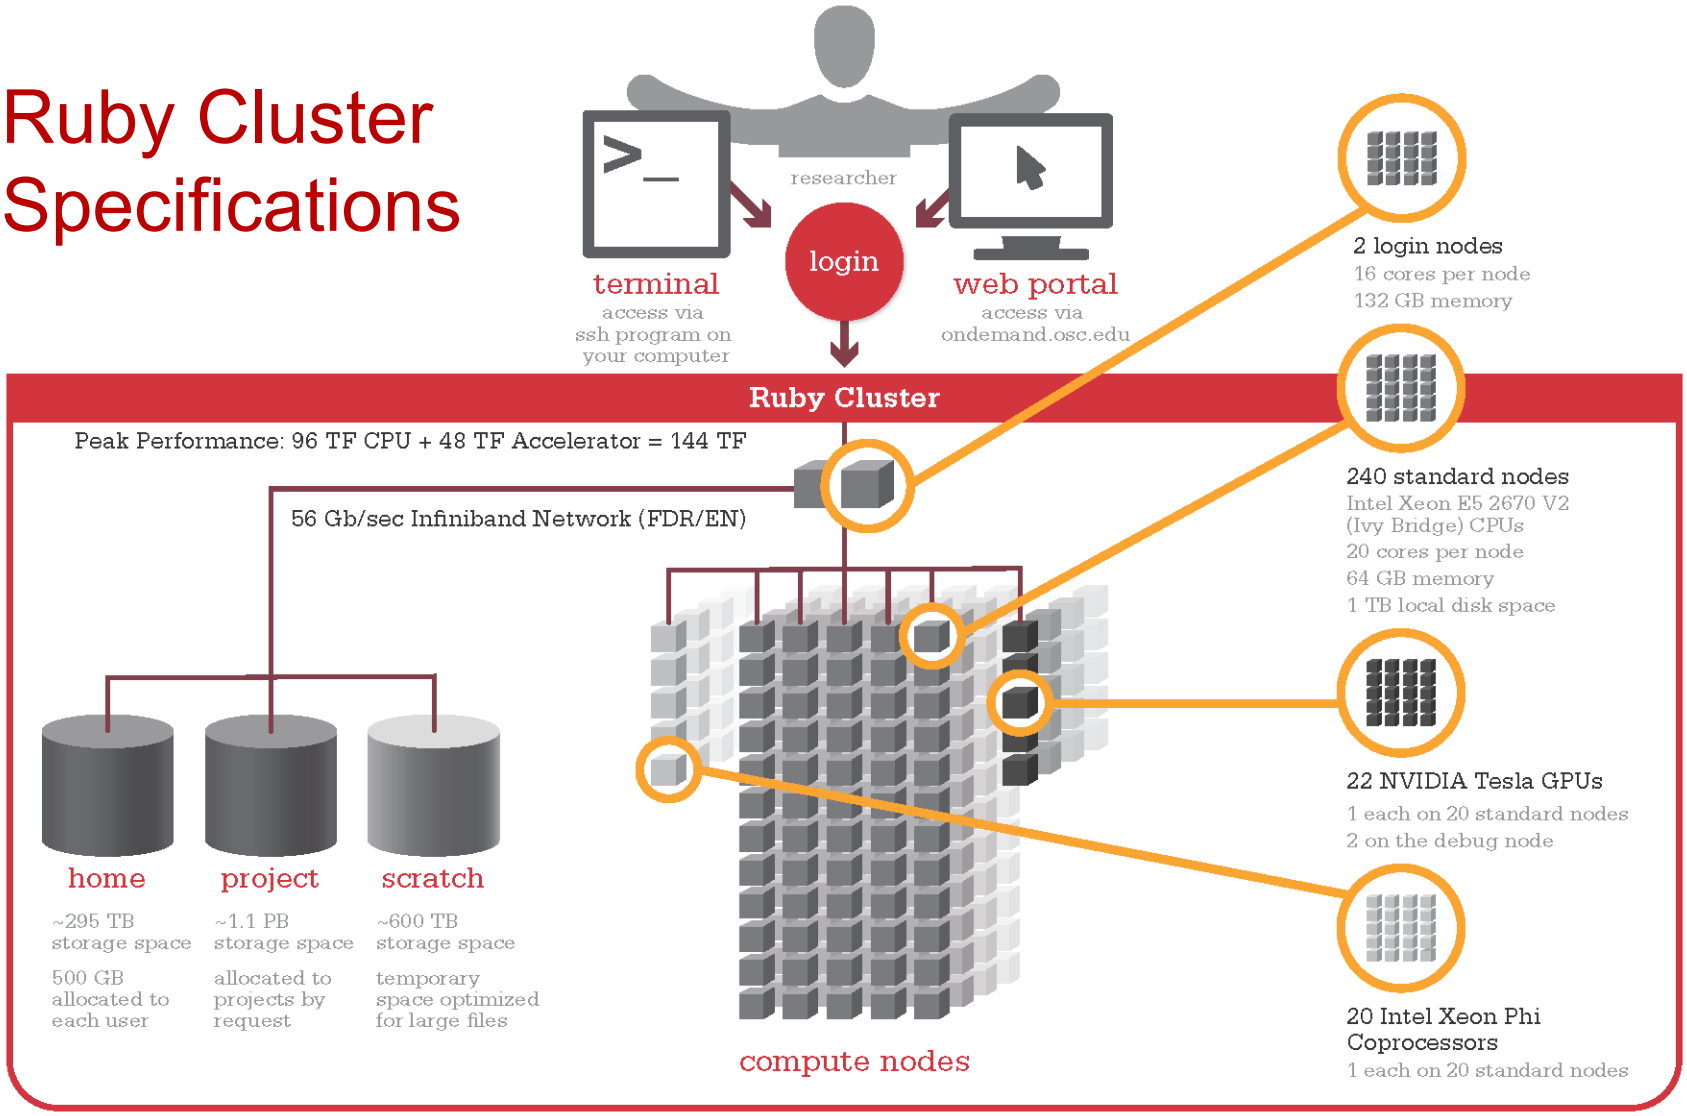

OSC File Space Information

- Scratch - Lustre – Parallel File System ~570 TBs (all disk)
- Project - 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 Lustre storage
- 1.8 PBs tape

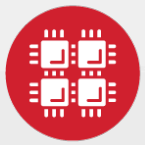

Ohio Supercomputer Center

An OH·TECH Consortium Member

Resource Grants and Accounts at OSC

Who can get an account?

- 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
- Commercial accounts
  - Commercial organizations may purchase time on OSC systems

Accounts and Projects at OSC

- Project
  - Headed by a PI
  - May include other users
  - Basis for accounting at OSC
  - Submit proposal for computing resources for a project
- Account
  - Username and password to access HPC systems
  - Each account associated with one project
  - Each account used by one person (please!)
  - If you work on multiple projects, you will have multiple accounts.

Allocations and Charges

- Charges are in terms of resource units
- Resource units
  - 1 resource unit (RU) = 10 CPU hours
  - CPU hour = walltime x (total # of cores requested)
- Project receives an allocation of RUs
- Jobs are charged to a project

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)
- Condo model for HPC

Citing OSC

- Please cite OSC in your publications:
  - Details at www.osc.edu/citation
- These publications should be reported to OSC

ARMSTRONG Researcher Portal

- https://armstrong.osc.edu
- Manage your project and accounts
  - Monitor resource utilization on all your projects
  - Add authorized users (request accounts) – PIs only
- View current information
  - OSC system notices
  - Research opportunities
- Post publications

ARMSTRONG Researcher Portal

https://armstrong.osc.edu

MyOSC

- Site for managing your identity at OSC
- Update your email
- Change your password
- Recover access to your account
- Change your shell
- And a lot more in the future
  - Project reporting
  - Authorized user management
  - Requesting services (e.g. software access)

Your Contact Info

- Keep your contact information current
  - Use my.osc.edu to manage your account details.
- If your student continues to use OSC after graduation, make sure email address is updated
  - Acceptable if still collaborating with you
- May need to contact you about problems
- Will need to contact you about regular password changes
- You can opt out of routine notifications

Connecting to the Oakley or Ruby Cluster

- 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)

OSC OnDemand

- 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

File Transfer & Editing

Visualization

Common Catalog

Job Submission & Monitoring

Command Line

System Status

- Check system status on:
  - https://www.osc.edu/supercomputing (bottom of page)
  - Message of the day (/etc/motd) – displayed at login
  - Twitter: @HPCnotices
  - Email for major outages or problems
- Scheduled downtimes
  - Quarterly maintenance – one day outage
  - Jobs held for up to two weeks prior

Statewide Users Group (SUG)

- The Statewide Users Group (SUG) is made up of OSC users
  - Provides program and policy advice to OSC
  - Meets twice a year
  - Headed by a chairperson elected yearly
- Standing committees
  - Allocations
  - Software and Activities
  - Hardware and Operations
- Get involved!
  - Next meeting is October in Columbus

Demo

- Website tour: www.osc.edu
- ARMSTRONG: https://armstrong.osc.edu
- MyOSC: https://my.osc.edu/

Demo

- Website tour: www.osc.edu
- MyOSC: https://my.osc.edu/
- OnDemand ondemand.osc.edu

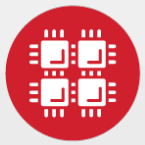

Ohio Supercomputer Center

An OH·TECH Consortium Member

What's New at OSC

C16 – available Oct. 1st (partial Aug. 1st)
Compute Nodes

- 648 standard nodes
  - 28 cores per node (2 14-core Intel Xeon processors)
  - 128 GB memory
  - 1 TB local disk space
- 160 GPU Ready nodes
  - Waiting for NVIDIA's next generation Pascal GPUs
- Network
  - Nodes connected by 100Gbit/sec Infiniband network (EDR)

C16 – available Oct. 1st (partial Aug. 1st) Special Resources

- 8 large memory nodes
  - 48 cores per node
  - 1536 GB memory (32GB/core)
  - 1 TB local disk space
- 8 large memory/large disk nodes
  - 48 cores per node
  - 1536 GB memory (32GB/core)
  - 48 TB local disk space

System Configurations

	C16 (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

C16 – available Oct. 1st (partial Aug. 1st) Performance Increase

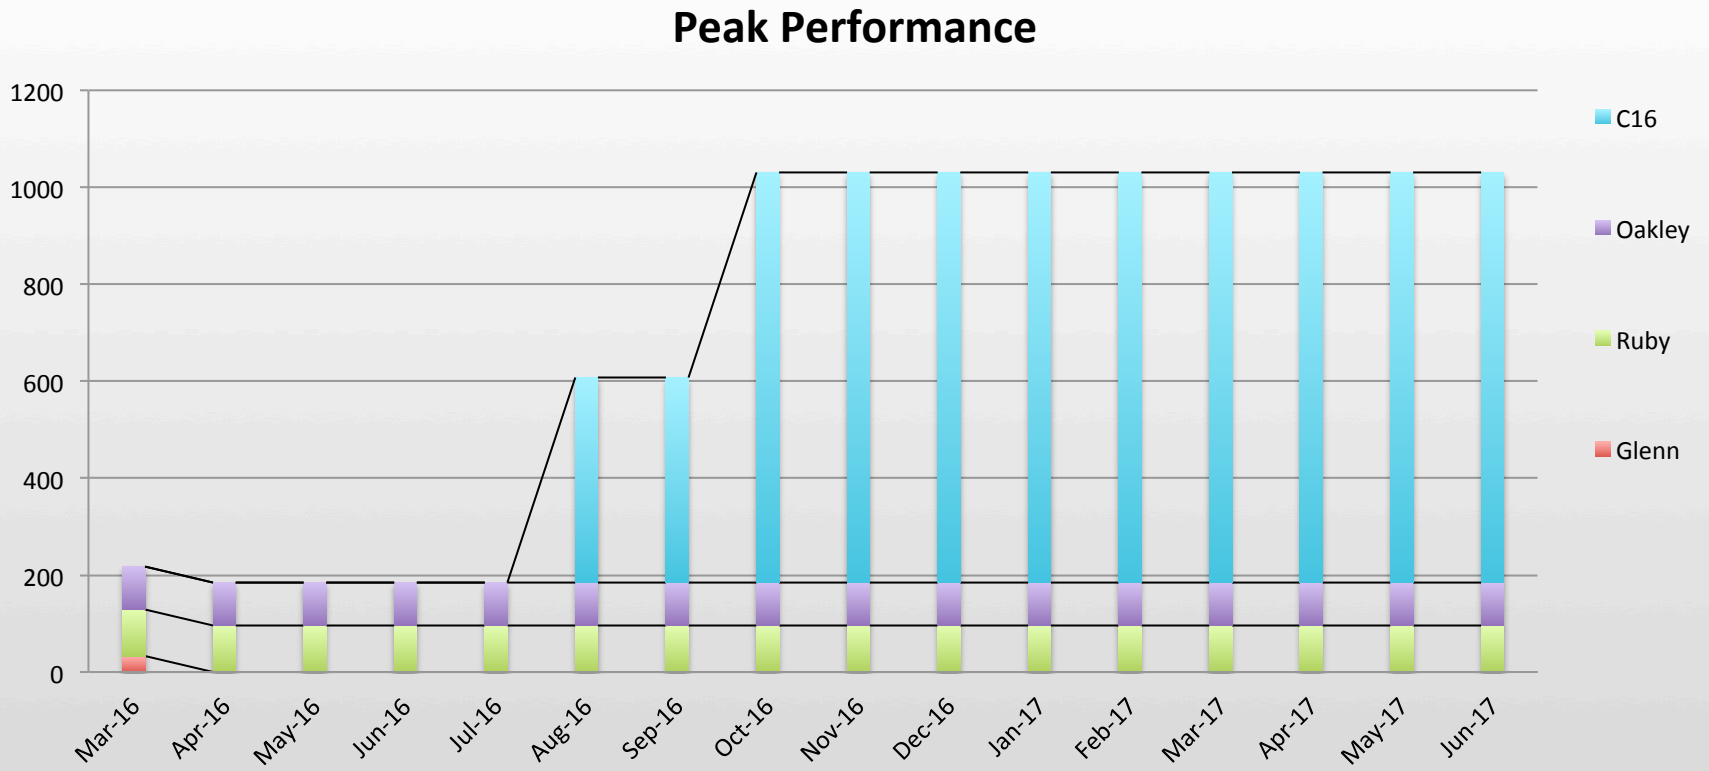

OSC File Space Information Update

- Scratch – DDN GPFS
  - 1 PB (~570 TB today)
  - 40-50 GB/s peak performance (~10 GB/s today)
- Project – DDN GPFS
  - 3.4 PB usable space (~1.1 PB today)
  - 40-50 GB/s peak performance (8-9 GB/s today)
- **Expected June/July**
- Home Directory Space / (Net App) NFS
  - **New system online in May**
  - ~900 TB usable (Disk) (~300 TB today)
  - Allocated to each user, 500 GB quota limit

Mass Storage Update

- >5 Petabytes (PBs) of usable disk
- 3.4 PB Project storage
- 1 PB Scratch space available
- 1.8 PBs tape – will be expanded

Questions

Kate Cahill
Education & Training Specialist
Ohio Supercomputer Center
kcahill@osc.edu

1224 Kinnear Road
Columbus, OH 43212
Phone: (614) 292-9623

ohiosupercomputercenter

ohiosupercomputerctr

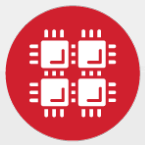

Ohio Supercomputer Center

An OH·TECH Consortium Member

User Environment

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 Oakley or Ruby Cluster

- 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)

OSC OnDemand

- 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

File Transfer & Editing

Visualization

Common Catalog

Job Submission & Monitoring

Command Line

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

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
- 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)

Text editing

- Traditional Unix editors
  - **`vi`**
  - **`emacs`**
  - Many books and web pages about **`vi`** and **`emacs`**
- GUI editor
  - **`gedit`**
- Simple editor
  - **`nano`**
- Can also edit on your computer and transfer files back and forth
  - **`dos2unix, unix2dos, mac2unix`**

Demo

- OSC OnDemand
- ssh
- sftp
- Linux
- Home directory tree
- Text editor: nano

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

Modules and your shell environment

- 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`

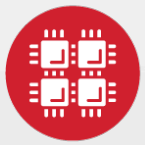

Ohio Supercomputer Center

An OH·TECH Consortium Member

Storage

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
  - 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
```

File Management

- Compress large, rarely used files
  - Use **`gzip`** or **`bzip2`** commands
- Combine large numbers of small files into an archive
  - Use **`tar`** command

Parallel File System – Lustre

- 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 6 months) of large amounts of data
- Also useful as batch-managed temporary storage
- **Scratch storage – NOT backed up!**

Local Disk – $TMPDIR

- Local file system on each compute node
  - 812 GB on each Oakley node
  - 1 TB on each Ruby node
- Fast – use for intermediate or scratch files
- Not shared between nodes
- Not backed up
- Managed by the batch system
- Data removed when job exits

Overloading the File Servers

- *"A supercomputer is a device for turning compute-bound problems into I/O-bound problems."* --Ken Batcher (parallel computing pioneer)
- One user's heavy I/O load can affect responsiveness for all users on that file system
- **Never** do heavy I/O in your home or project directory!
- Use $TMPDIR, copying files in and out as necessary
- Don't let large numbers of jobs run in lockstep.

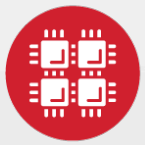

Ohio Supercomputer Center

An OH·TECH Consortium Member

Batch Processing

Organization of an OSC cluster

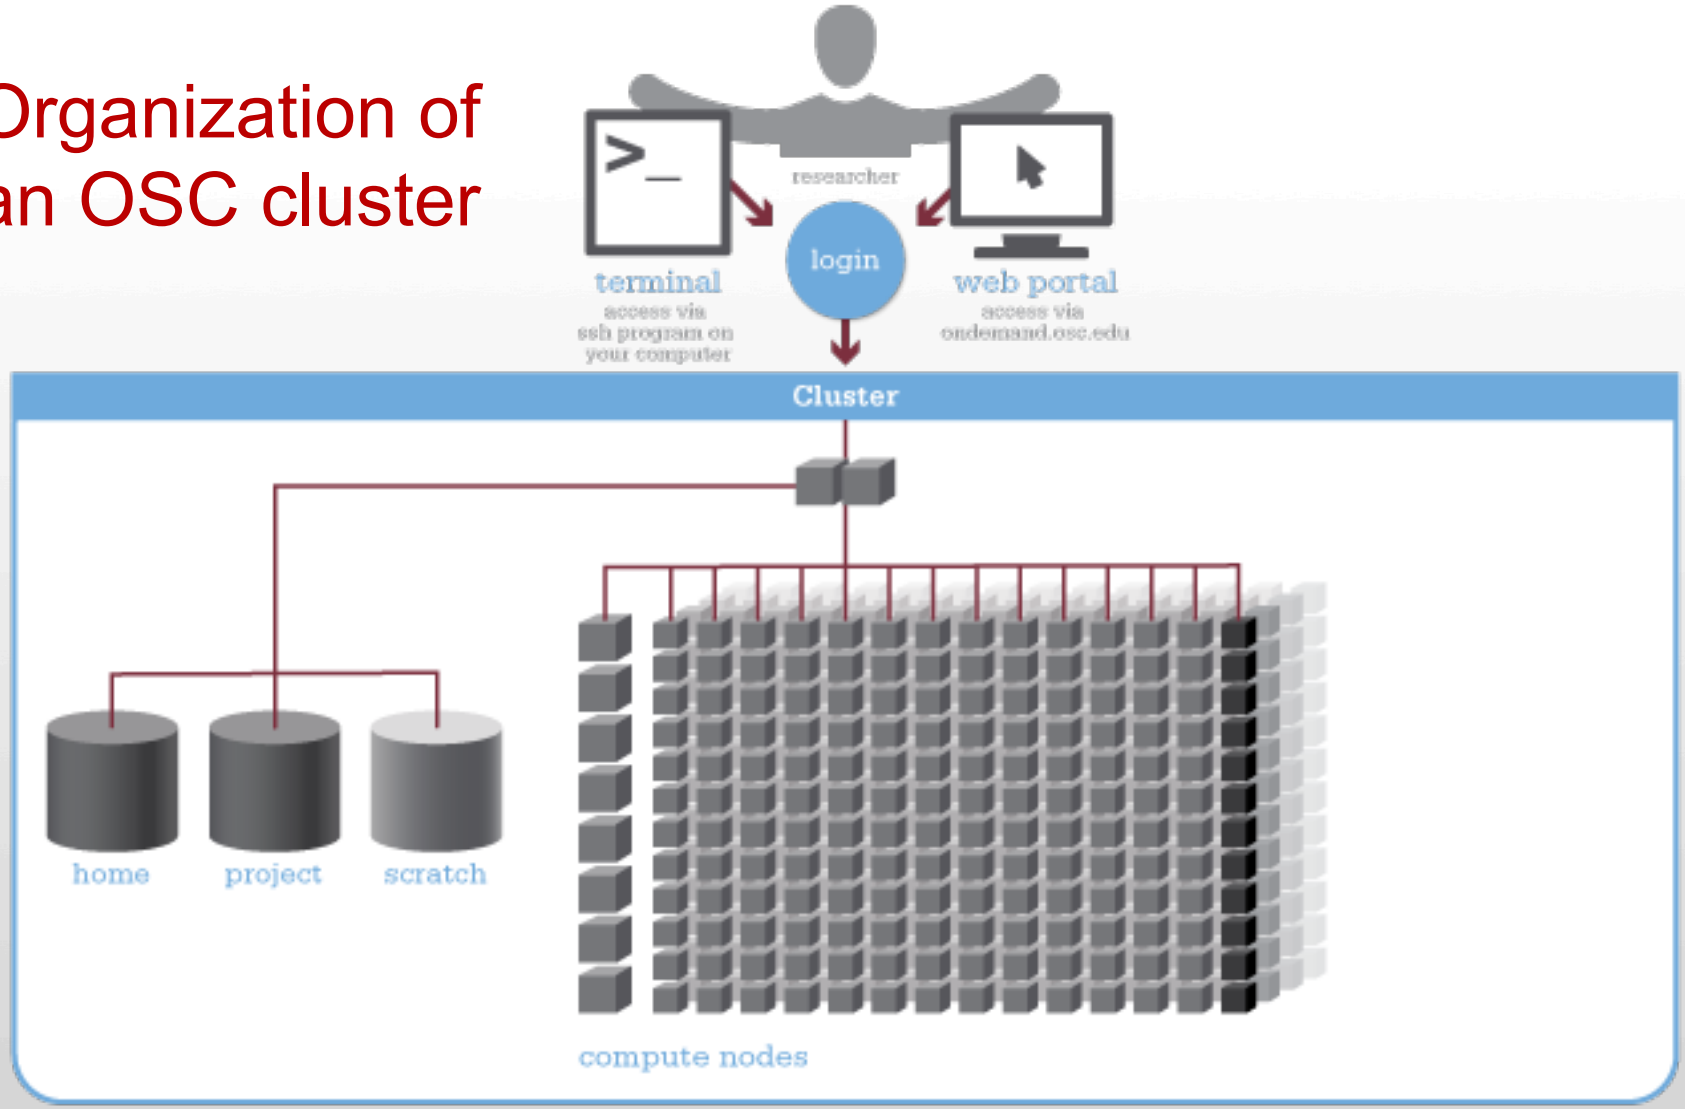

Batch System at OSC

- 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

Why do supercomputers use queueing?

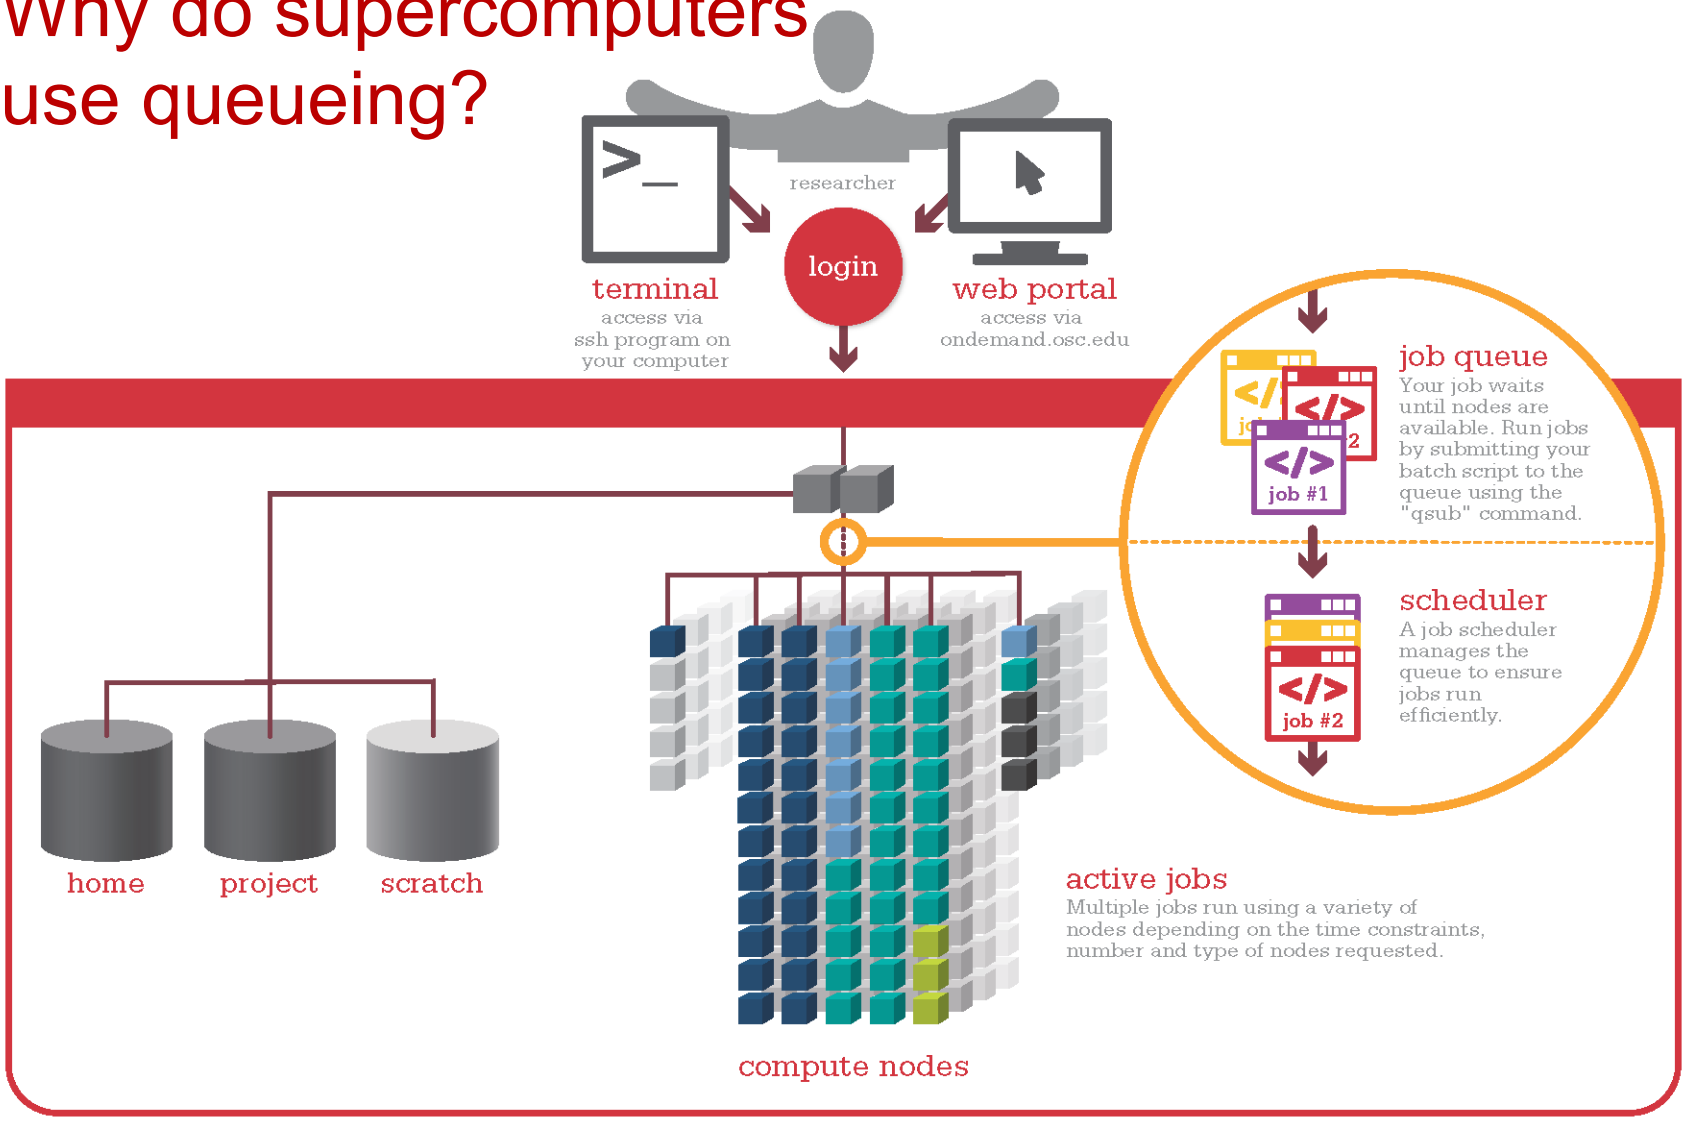

Queuing Jobs

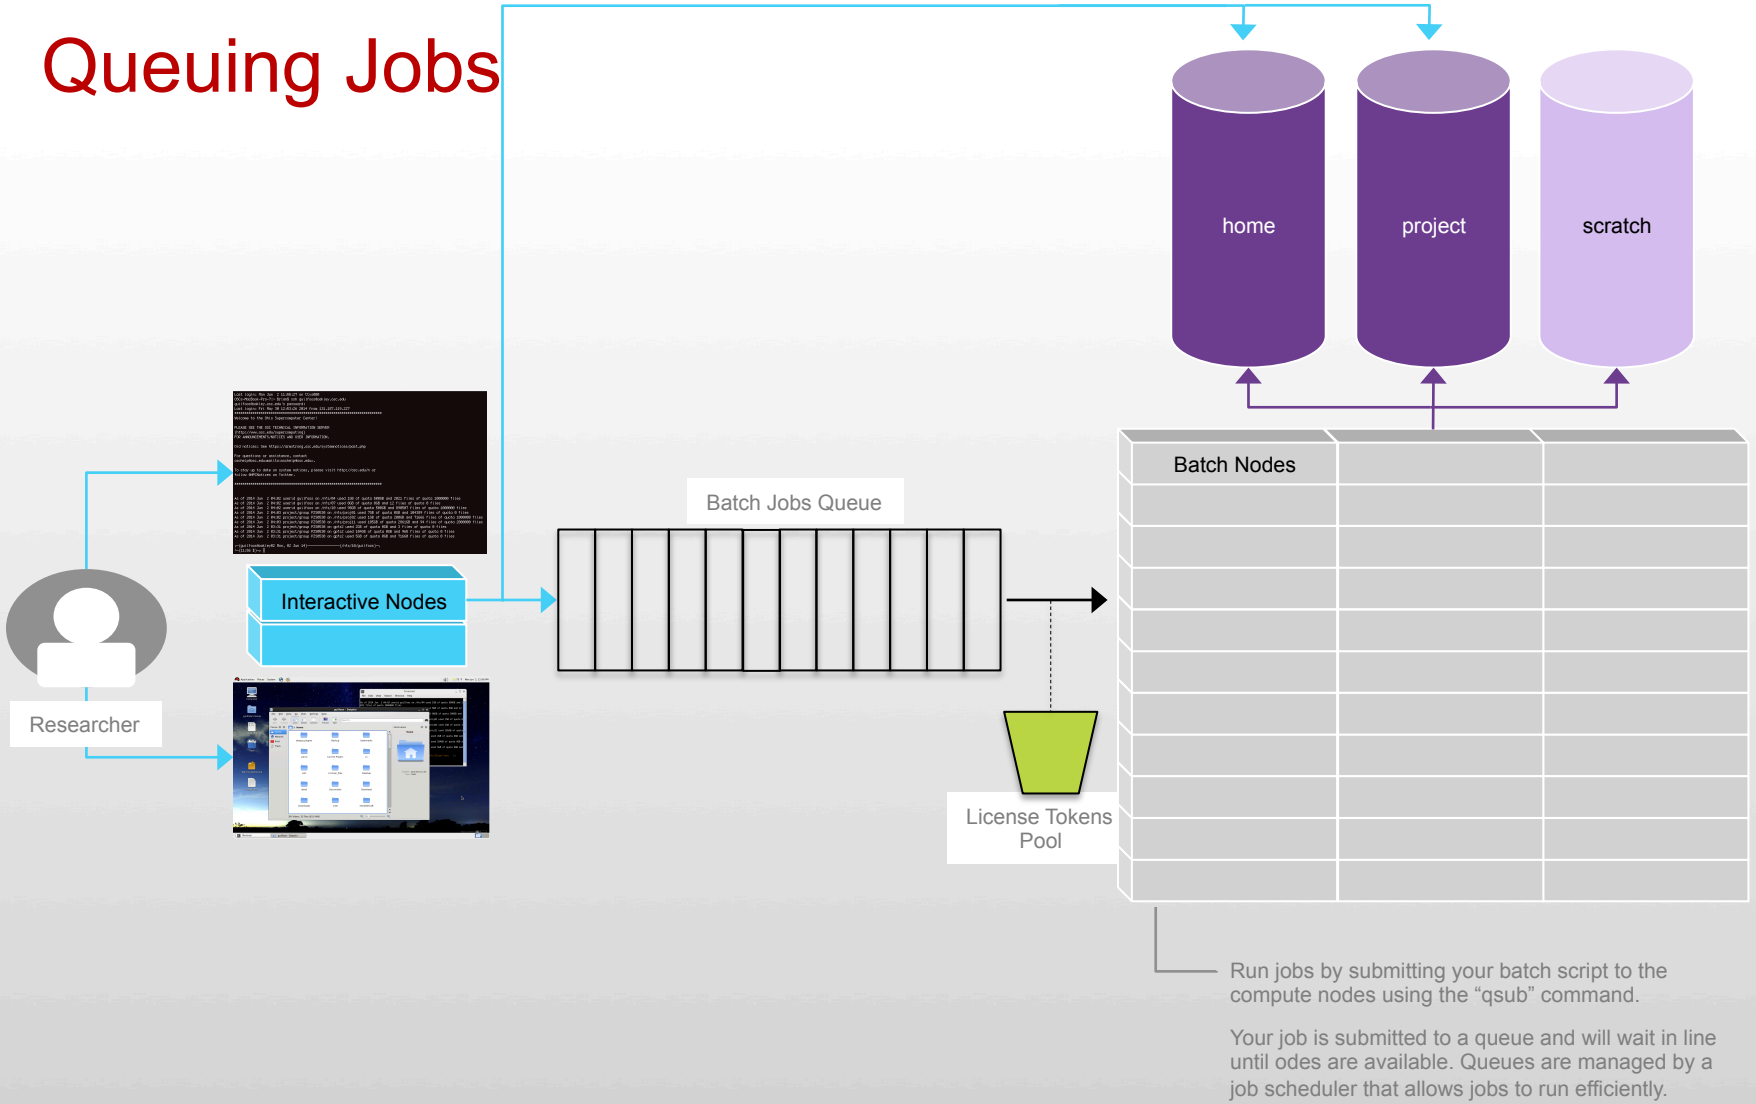

Run jobs by submitting your batch script to the compute nodes using the “qsub” command.

Your job is submitted to a queue and will wait in line until odes are available. Queues are managed by a job scheduler that allows jobs to run efficiently.

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

Running a Job on the Compute Nodes

- Create a batch script for a job
- Submit the job
- Job gets queued
- Job runs when resources become available
- 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=1
#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***

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=12 -l walltime=1:00:00`

Batch Queues

- Oakley and Ruby 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: -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

To Take Advantage of Parallel Computing

- Program must be written to take advantage of multiple cores and/or multiple nodes
- Many commercial applications have multithreaded or parallel versions
- Must use `mpiexec` for multiple nodes
- **Can't just request more nodes or cores and expect your job to run faster**

Specifying Resources in a Job Script for GPUs

- Nodes and cores (processors) per node
- Memory
- GPUs
  - See “Batch Processing at OSC” on OSC website

```
#PBS -l walltime=01:00:00
#PBS -l nodes=1:ppn=1:gpus=1
#PBS -N compute
#PBS -j oe
module load cuda
cd $HOME/cuda
cp mycudaApp $TMPDIR
cd $TMPDIR
./mycudaApp
```

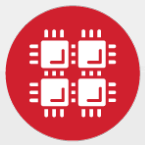

Ohio Supercomputer Center

An OH·TECH Consortium Member

Third-Party Software Applications

Access to Licensed Software

- Most software licenses for academic use only
- Some software requires signed license agreement
  - Check website
  - Contact OSC Help

Third party applications

- **Chemistry** (*license agreement required)
  - *AMBER
  - ChemTools
  - COLUMBUS
  - *CSD (Cambridge Structural Database)
  - ESPRESSO
  - GAMESS
  - *Gaussian
  - GROMACS
  - LAMMPS
  - MacroModel®
  - MEAD
  - NAMD
  - NWChem
  - Open Babel
  - *Turbomole

Third party applications

- **Bioinformatics**
  - Bioperl
  - BLAST
  - BLAT
  - Bowtie
  - Clustal W
  - EMBOSS
  - Fitmodel
  - HMMER
  - MrBayes
  - NAMD
  - PAML
  - PAUP
  - RAxML
  - RepeatMasker
  - TreeBeST

Third party applications

- **Structural Mechanics** (*license agreement required; ₵statewide licensed)
  - *ABAQUS
  - ₵Altair HyperWorks
  - *ANSYS
  - COMSOL Multiphysics
  - *LSDYNA
  - LS-PREPOST

Third party applications

- **Fluid Dynamics** (*license agreement required)
  - *Fluent
  - OpenFOAM

Third party applications

- Mathematics/Statistics (₵statewide licensed)
  - MATLAB (special licensing restrictions)
  - Octave
  - R
  - Stata
  - FFTW
  - ScaLAPACK
  - MINPACK
  - sprng2
  - ₵ Intel MKL
  - ACML (Glenn only)

Third party applications

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

Third party applications

- **Parallel programming software (**$statewide licensed)
  - MPI library (mvapich, mvapich2)
  - OpenMP
  - CUDA
  - OpenCL
  - OpenACC

Third party applications

- **Visualization software**
  - GNUplot
  - Jmol
  - VTK

- More 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

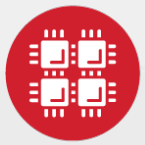

Ohio Supercomputer Center

An OH·TECH Consortium Member

OSC Policies

Firefox

Supercomputing | Ohio Supercomputer ...

https://www.osc.edu/supercomputing

Google

Home | About OSC | Supercomputing | Cyberinfrastructure | Research | Education | News | Contact Us

www.osc.edu/supercomputing

Home

Supercomputing

As a leader in high performance computing and networking, the Ohio Supercomputer Center (OSC) is a vital resource for Ohio's scientists and engineers. OSC is a fully-scalable center with mid-range machines to match those found at National Science Foundation centers and other national labs. The Center's flagship supercomputing system is the Oakley Cluster, an 8,300+ core HP Intel Xeon machine.

OSC also is home to the Ohio Interface Labratory, which provides state-of-the-art graphics equipment and output services for computer graphics and animation. Members of the Science and Technology Support Group provide much of the information for this server.

Oakley HP Intel Xeon Cluster (8,328 Cores)

Getting Started

Welcome to OSC! If you are new to supercomputing, new to OSC, or simply interested in getting an account (if you don't already have one), we have some resources to help you.

- Read about the basics of High Performance Computing (highly recommended for new users).
- Learn more about our account allocation process, and how to get started.
- Apply for expanded resources beyond those for a start-up project.
- Get help connecting to our resources.
- Read our policies, FAQ, code of ethics for users, and glossary.

OSC Policies

Supercomputing

- Getting Started
- HPC Environments
- Available Software
- Portals
- Tutorials & Training
- Support Services
- Search Documentation

Upcoming Events

August 8, 2013 - Statewide User Group (SUG) General Meeting (Columbus, OH)

August 23, 2013 - Your Plan For Health Biometric Screenings (Columbus, OH)

September 3, 2013 - Intro to OSC (Columbus, OH)

More

Recent News

Columbus, Ohio (July 23, 2013) - OSC OnDemand gives computational

OSC Policies

- OSC-1, OSC Data Lifecycle Management Policy
  - Use of home directory, project directory and **`$TMPDIR`**
  - Storage and file quotas
  - Backup and recovery

OSC Policies

- OSC-11, OSC User Management Policy
  - Who can get an account
  - Charges for accounts
  - Types of accounts
  - Account restrictions
  - Account resource units
  - Inappropriate system use

For More Information

- Key webpages

  www.osc.edu/supercomputing - general documentation

  https://www.osc.edu/supercomputing/batch-processing-at-osc

- Contact the help desk (OSC Help) 24/7

  oschelp@osc.edu

  614-292-1800

  1-800-686-6472

Questions

Kate Cahill
Education & Training Specialist
Ohio Supercomputer Center
kcahill@osc.edu

1224 Kinnear Road
Columbus, OH 43212
Phone: (614) 292-9623

ohiosupercomputercenter

ohiosupercomputerctr