


OWENS



JESSE OWENS
OLYMPIC CHAMPION, BEACON FOR EQUALITY, YOUTH ADVOCATE

 Ohio Supercomputer Center
An OH·TECH Consortium Member



Ohio Supercomputer Center

An OH·TECH Consortium Member

An Introduction to OSC Resources and Services


March 15, 2023





Kate Cahill
Education & Training Specialist

“OSC is here to empower your research.”





Outline

- What is OSC?
- High-Performance Computing (HPC) Concepts
- Hardware Overview
- Getting Started at OSC


Break

- User Environment
- Using Software on OSC systems
- Batch Processing
- OnDemand demo





What is the Ohio Supercomputer Center?



About OSC


Founded in 1987, through the Ohio Department of Higher Education

Part of OH-TECH statewide consortium for technology support services

Statewide resource for all higher education institutions in Ohio

high performance computing services

computational science expertise



Ohio Supercomputer Center

An **OH·TECH** Consortium Member




OhioLINK

Connecting Libraries, Learning & Discovery
An **OH·TECH** Consortium Member



OARnet

An **OH·TECH** Consortium Member



Service Catalog



Cluster Computing

A fully scalable center with mid-range machines to match those found at National Science Foundation centers and other national labs.



Research Data Storage

High-performance, large capacity data storage spaces along with others that are perfect for a wide variety of research data.



Education

High performance computing and networking resources come together to create an exciting and innovative teaching and research environment.





Web Software Development

Our expert web development team helps you create custom web interfaces to simplify the use of powerful HPC resources.



Scientific Software Development

Deep expertise in developing and deploying software that runs efficiently and correctly on large scale cluster computing platforms.



Client Services

FY2022



55 Ohio Universities



68 Companies



54 Non-Profits and Government Agencies



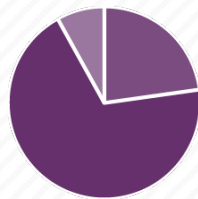
200 Other Educational Institutions



8,494 Clients



1,101 PIs



23% Faculty/Staff
69% Students
8% Other




127 College Courses Used
OSC



29 Training Opportunities




735 Trainees






HPC Concepts



Why Use HPC?



Your simulations or analyses take too long on your personal computer

More (faster) cores
Multithreading
Multi-node parallelization (openmpi & Rmpi)
GPU acceleration (NVIDIA's CUDA)
Distributed computing (Apache Spark with Python and R)




The size of your data is too large to be contained (storage) or accessed (memory) on your computer


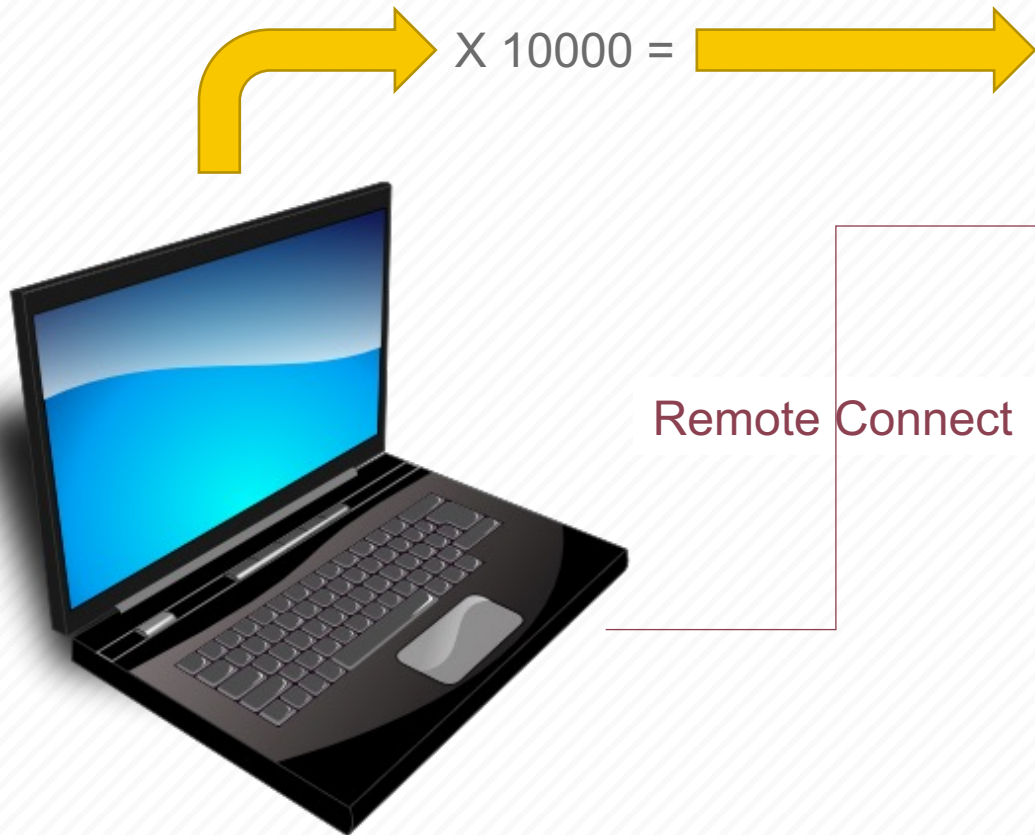
Large memory nodes: 768GB; 1.5TB, 3TB
Distributed computing (Apache Spark with R)
Compute node disk space: 1TB, 4TB, & 24TB
Project storage: TBs range



You need a particular software or package for your work



What is the difference between your laptop and a supercomputer?



HPC Terminology

Compute Node

- Equivalent to a high-end workstation, part of a cluster

Compute Cluster


- A group of computers (nodes) connected by a high-speed network, forming a supercomputer

Core

- A processor (CPU), multiple cores per processor chip

Graphical Processing Unit (GPU)

- A separate multi-core processor that can handle many small calculations



Memory



Holds data that is being calculated on, as well as computational instructions



Memory types

Shared memory is local to one node and several process


Distributed memory is on multiple nodes



Each core has an associated amount of memory

Standard nodes: ~4 GB/core

Huge memory nodes: ~75 GB/core



Storage



Storage: different types of “disk” for different needs



Local disk in the compute node



Project storage




Shared scratch

Short-term storage only


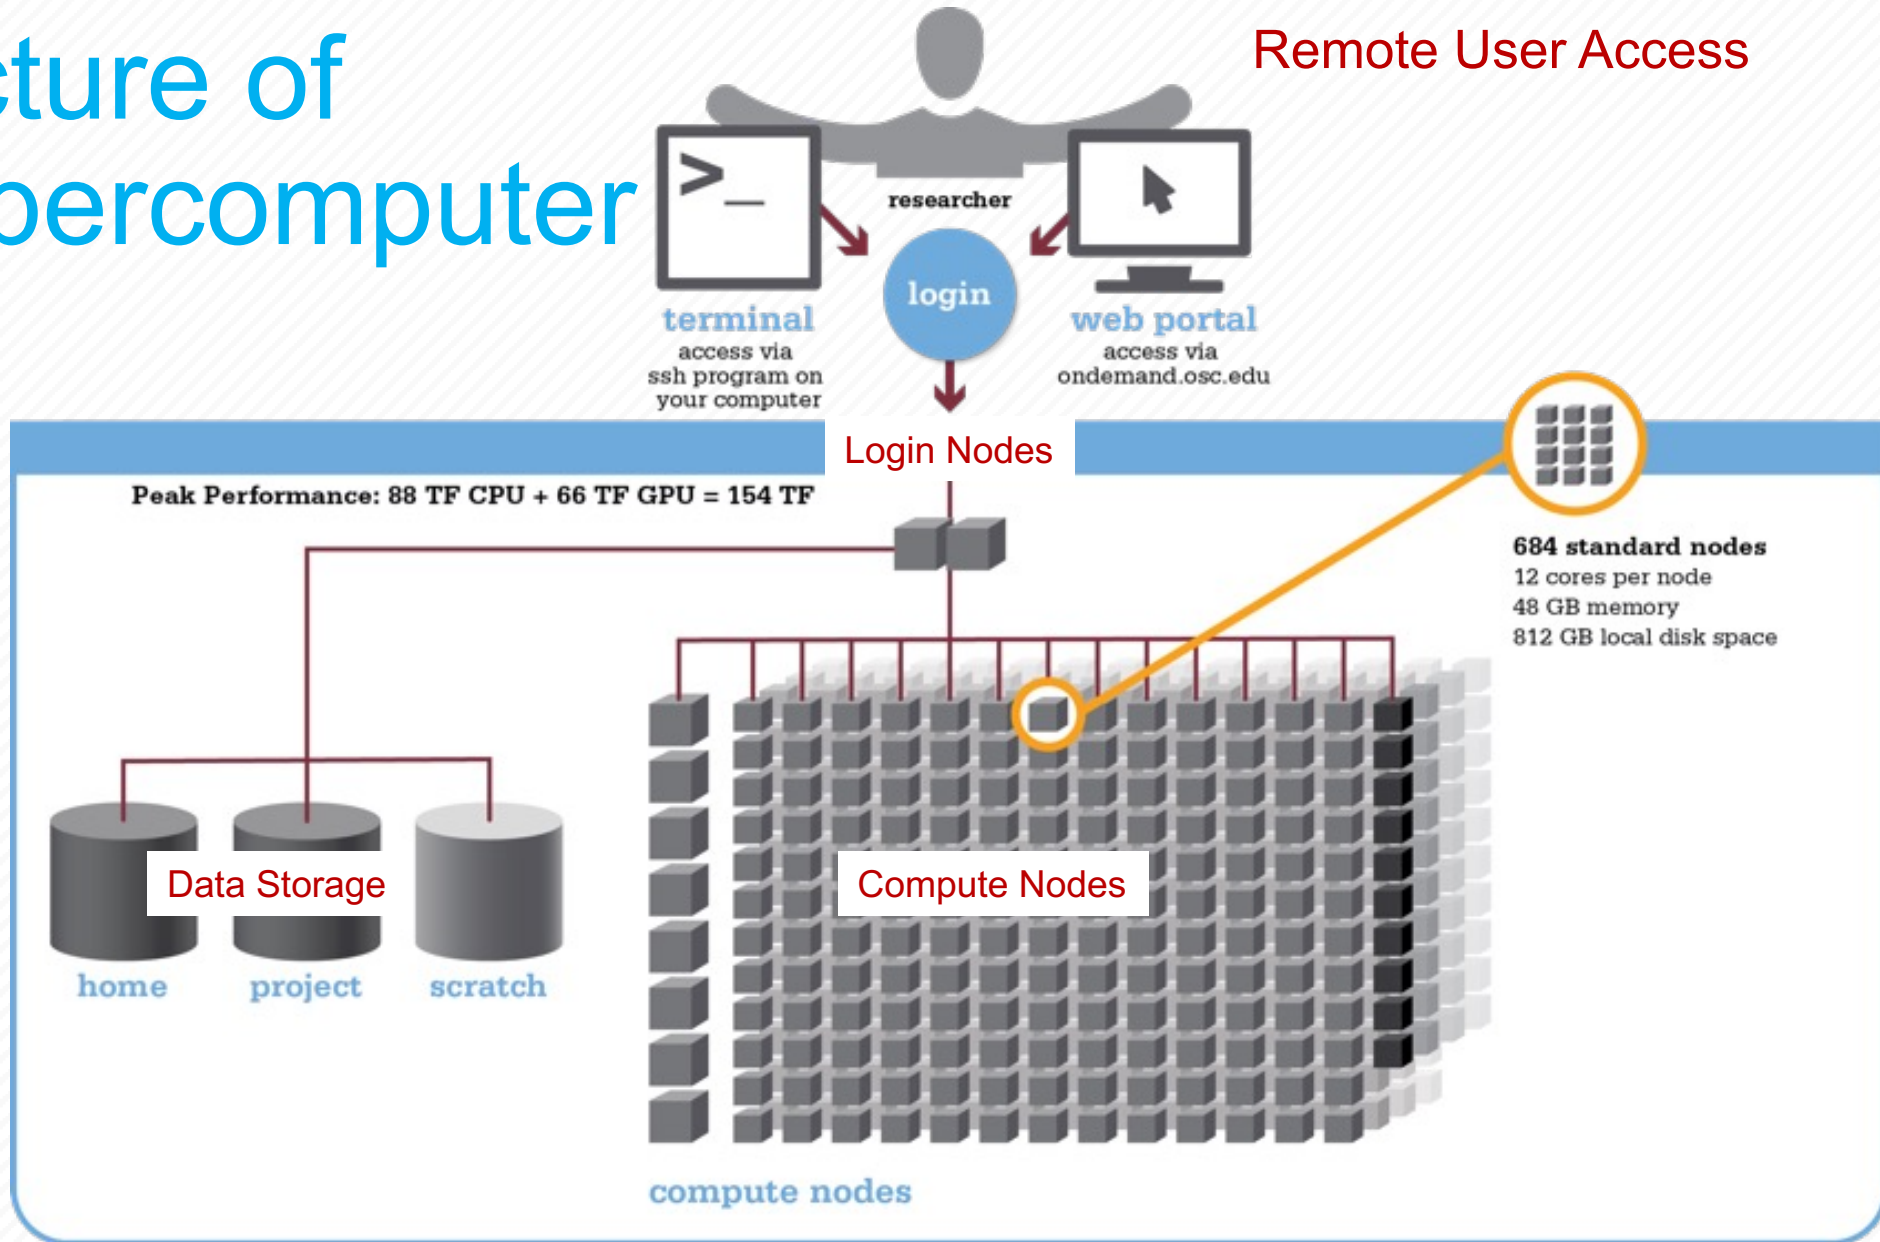


Longer-term or archive




Structure of a Supercomputer

Remote User Access





Hardware Overview




System Status

COMPUTE


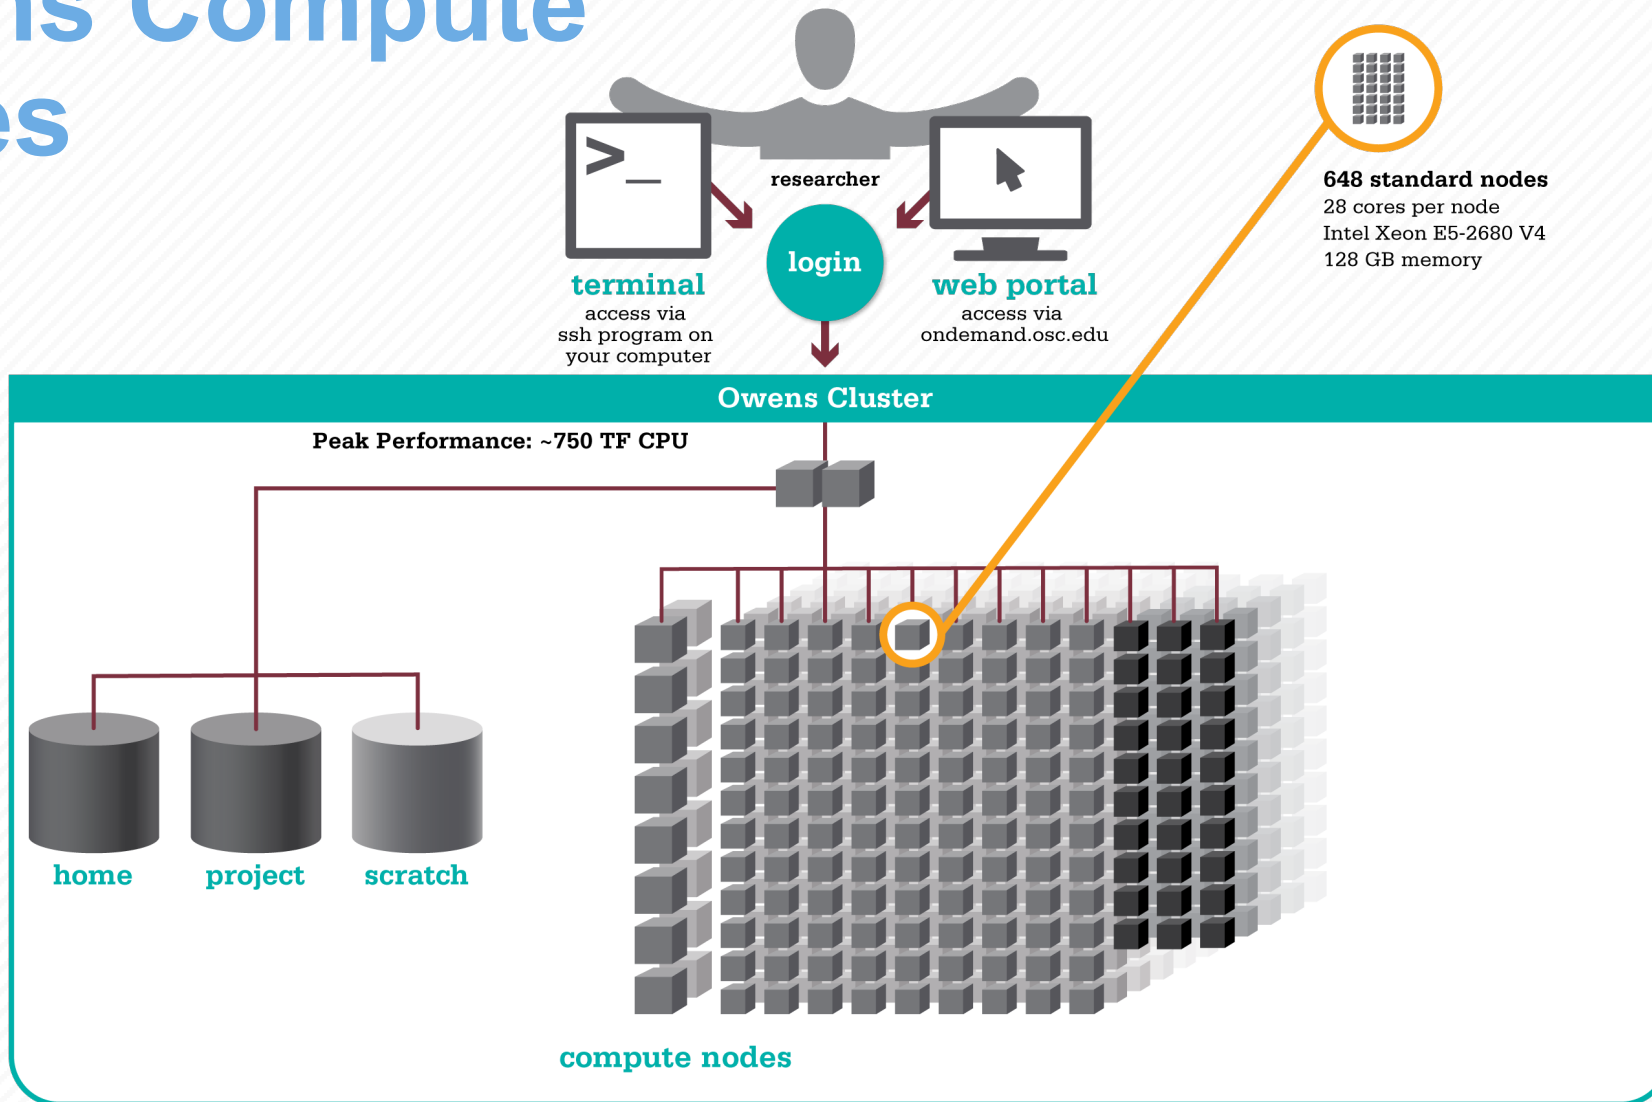
	Owens	Pitzer	Pitzer Expansion	Ascend
Date	2016	2018	2020	2022
Cost	\$7 million	\$3.4 million	\$4.3 million	\$1.8 million
Theoretical Perf.	~1.6 PF	~1.3 PF	~2.6 PF	~2.0 PF
Nodes	824	260	398	24
CPU Cores	23,392 Intel Broadwell	10,560 Intel Skylake	19,104 Intel Cascade Lake	2,304 AMD Milan
RAM	~120 TB	~71 TB	~94 TB	~24 TB
GPUs	160 NVIDIA Pascal P100	64 NVIDIA Volta V100	102 NVIDIA Volta V100	96 NVIDIA Ampere A100
TOTALS	Cost (\$16.5 M) Perf. (~7.5 PF) Nodes (1,506) Cores (55,360) RAM (~309 TB) GPUs (422)			

STORAGE


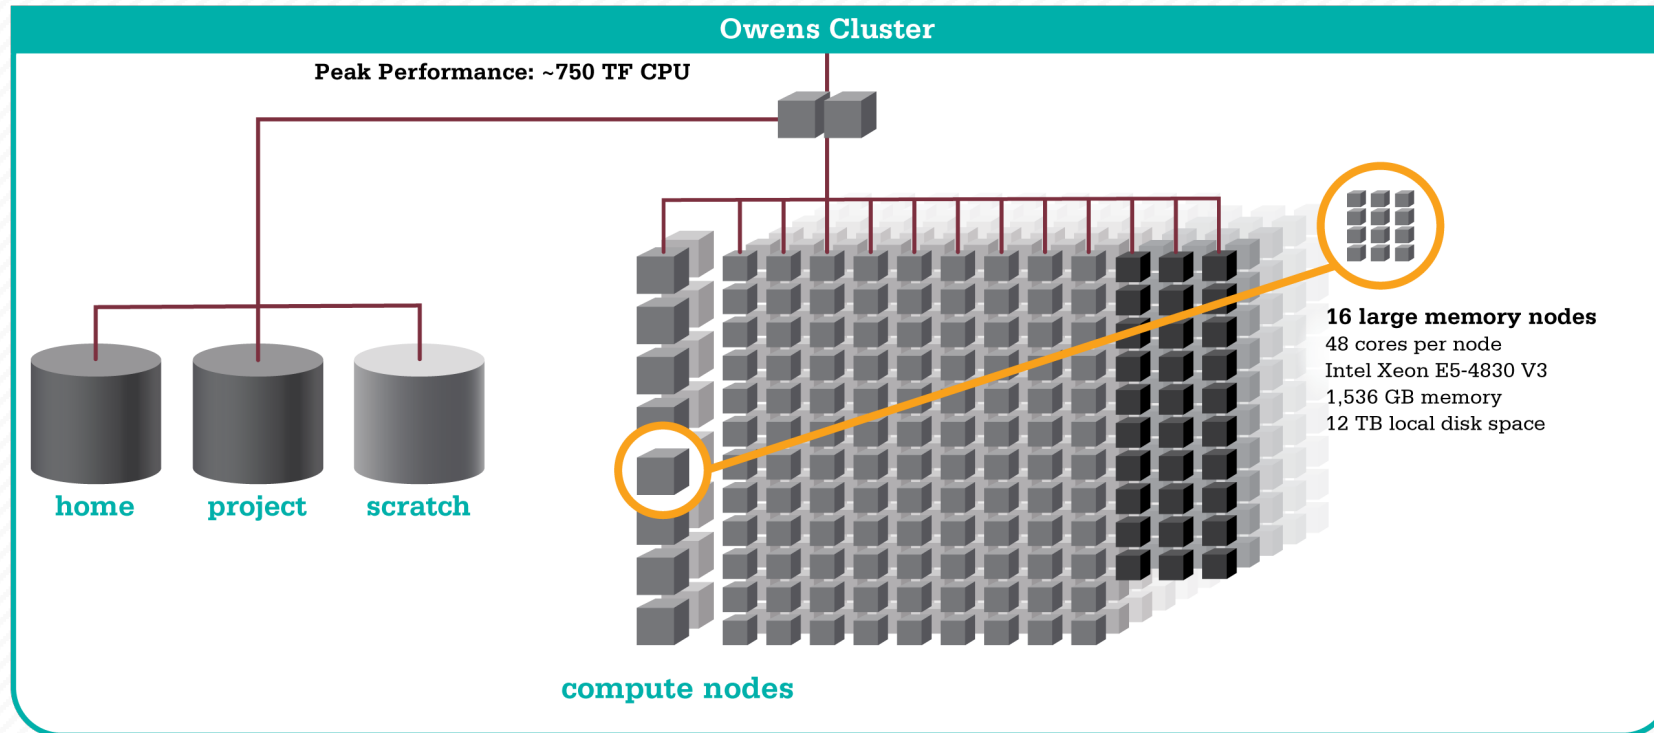
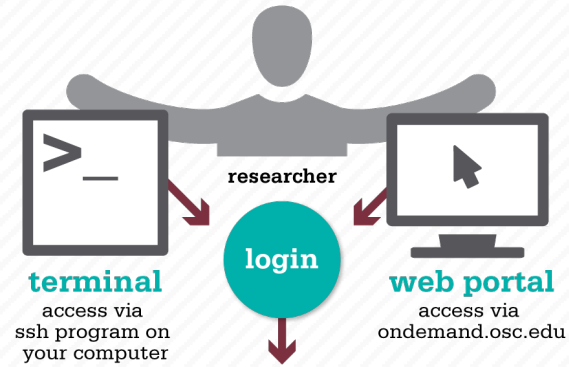
	IBM TAPE	IBM ESS	NetApp
Year Acquired	2016 – 2021	2020, 2022	2021
Cost	\$0.8 million	\$3.9 million	\$0.7 million
Capacity	23.5 PB (x2)	~16 PiB	1.8 PiB
TOTALS	Cost (\$5.4 M) Disk Capacity (~17.8 PiB or ~20 PB)		




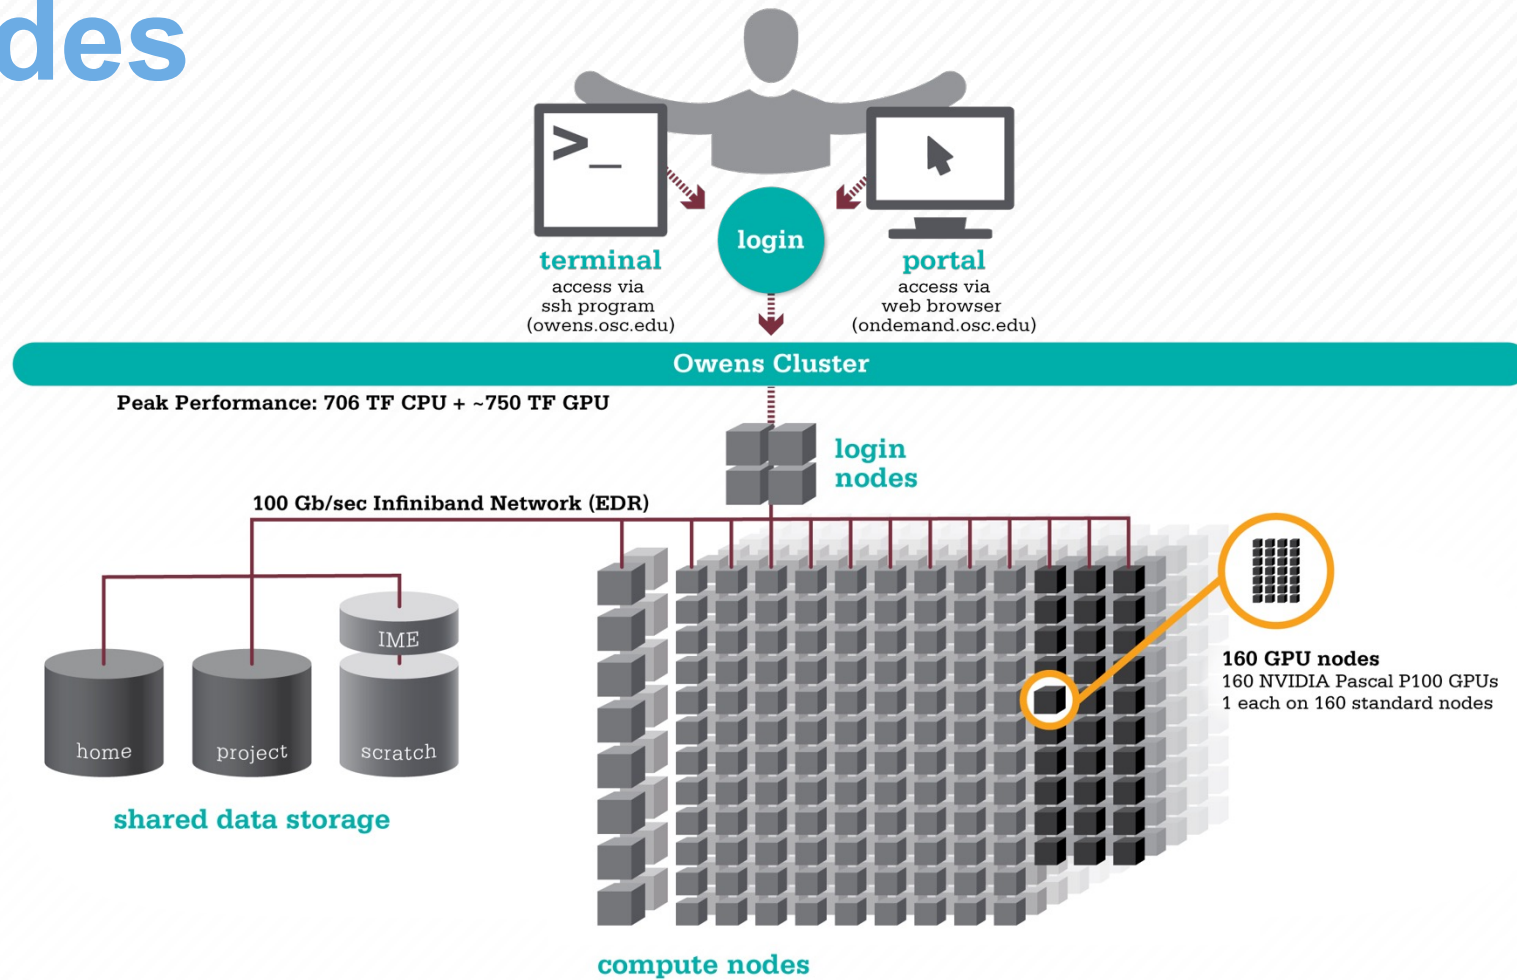
Owens Compute Nodes



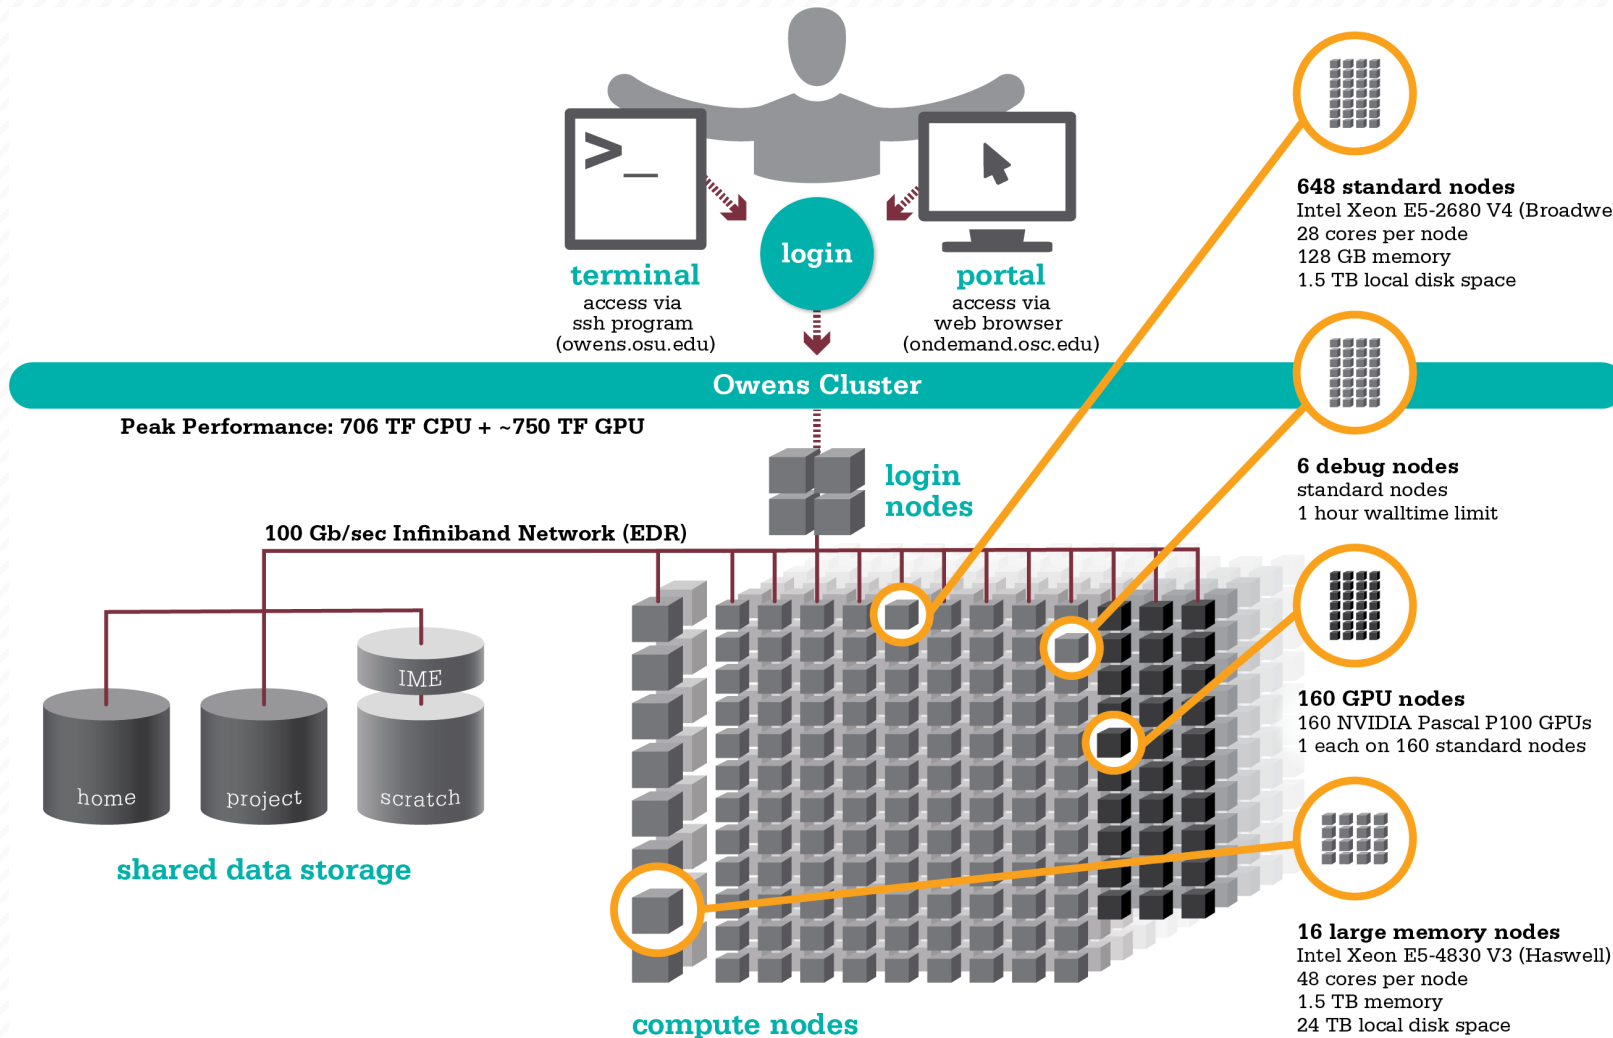
Owens Data Analytics Nodes




Owens GPU Nodes



Owens Cluster Specifications


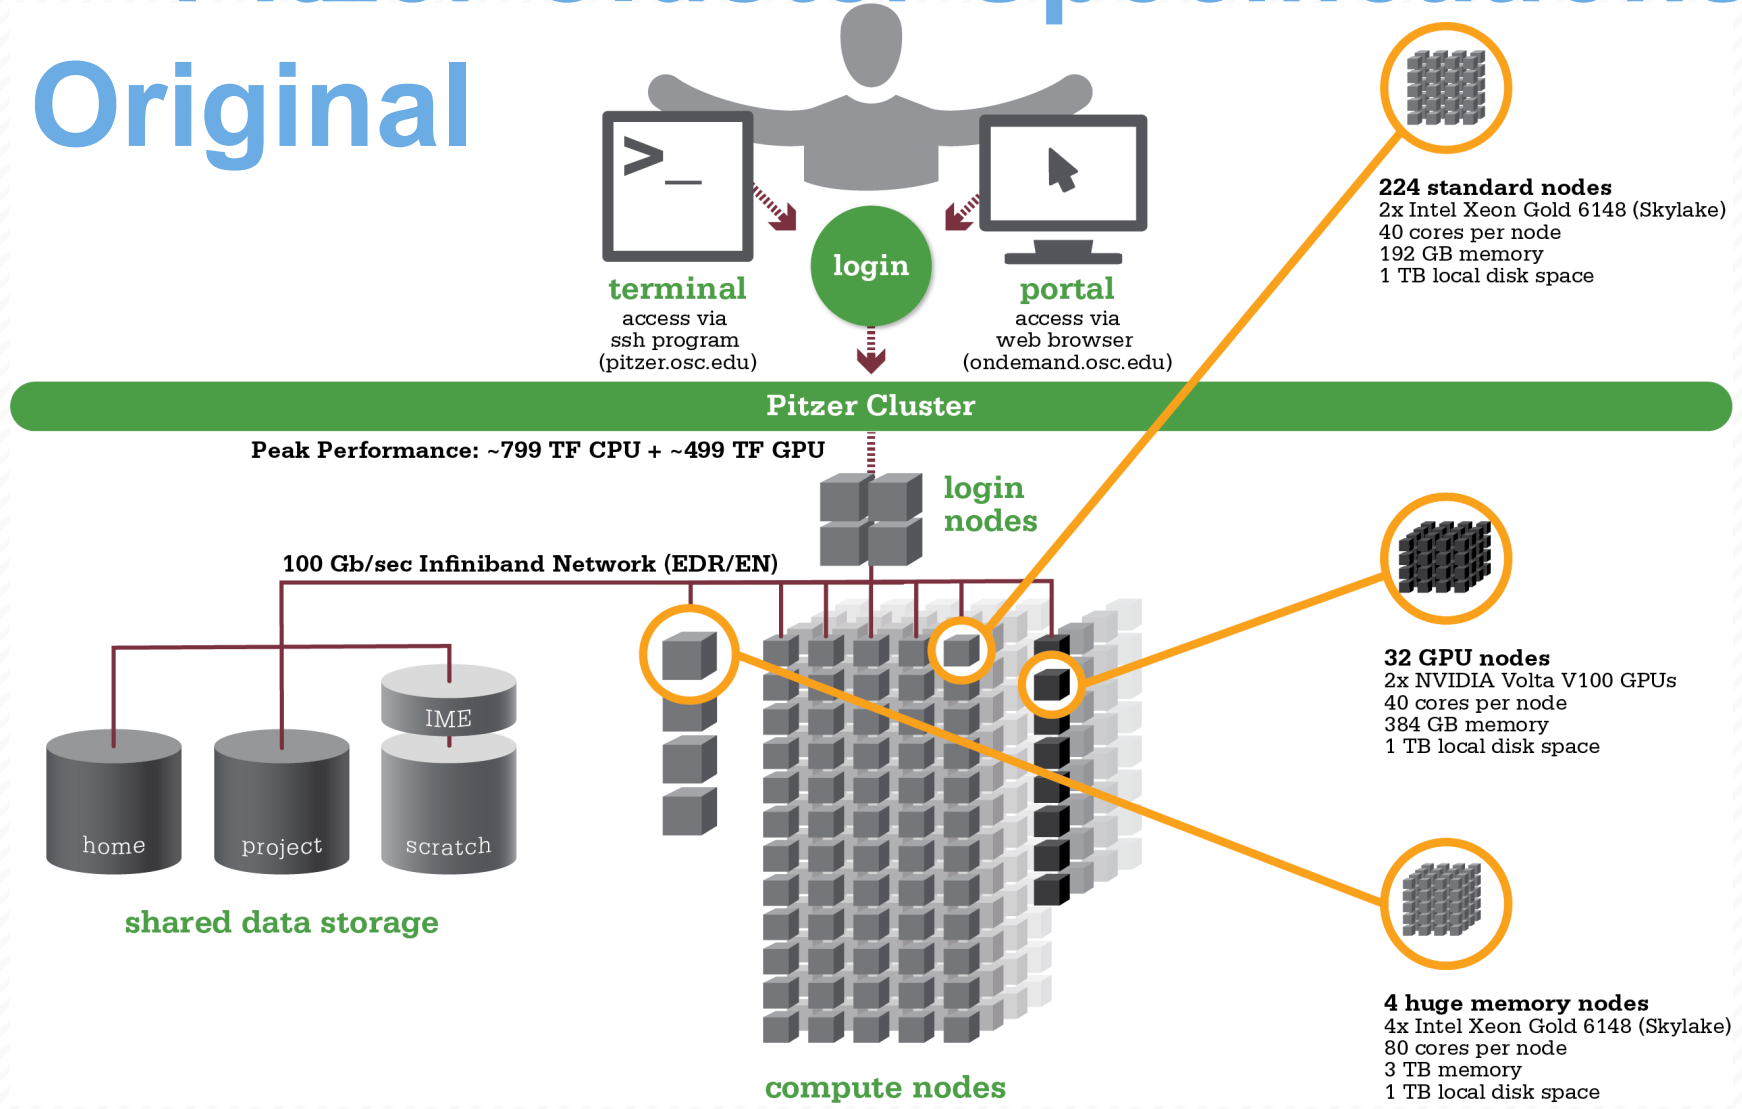


https://www.osc.edu/resources/technical_support/supercomputers/owens

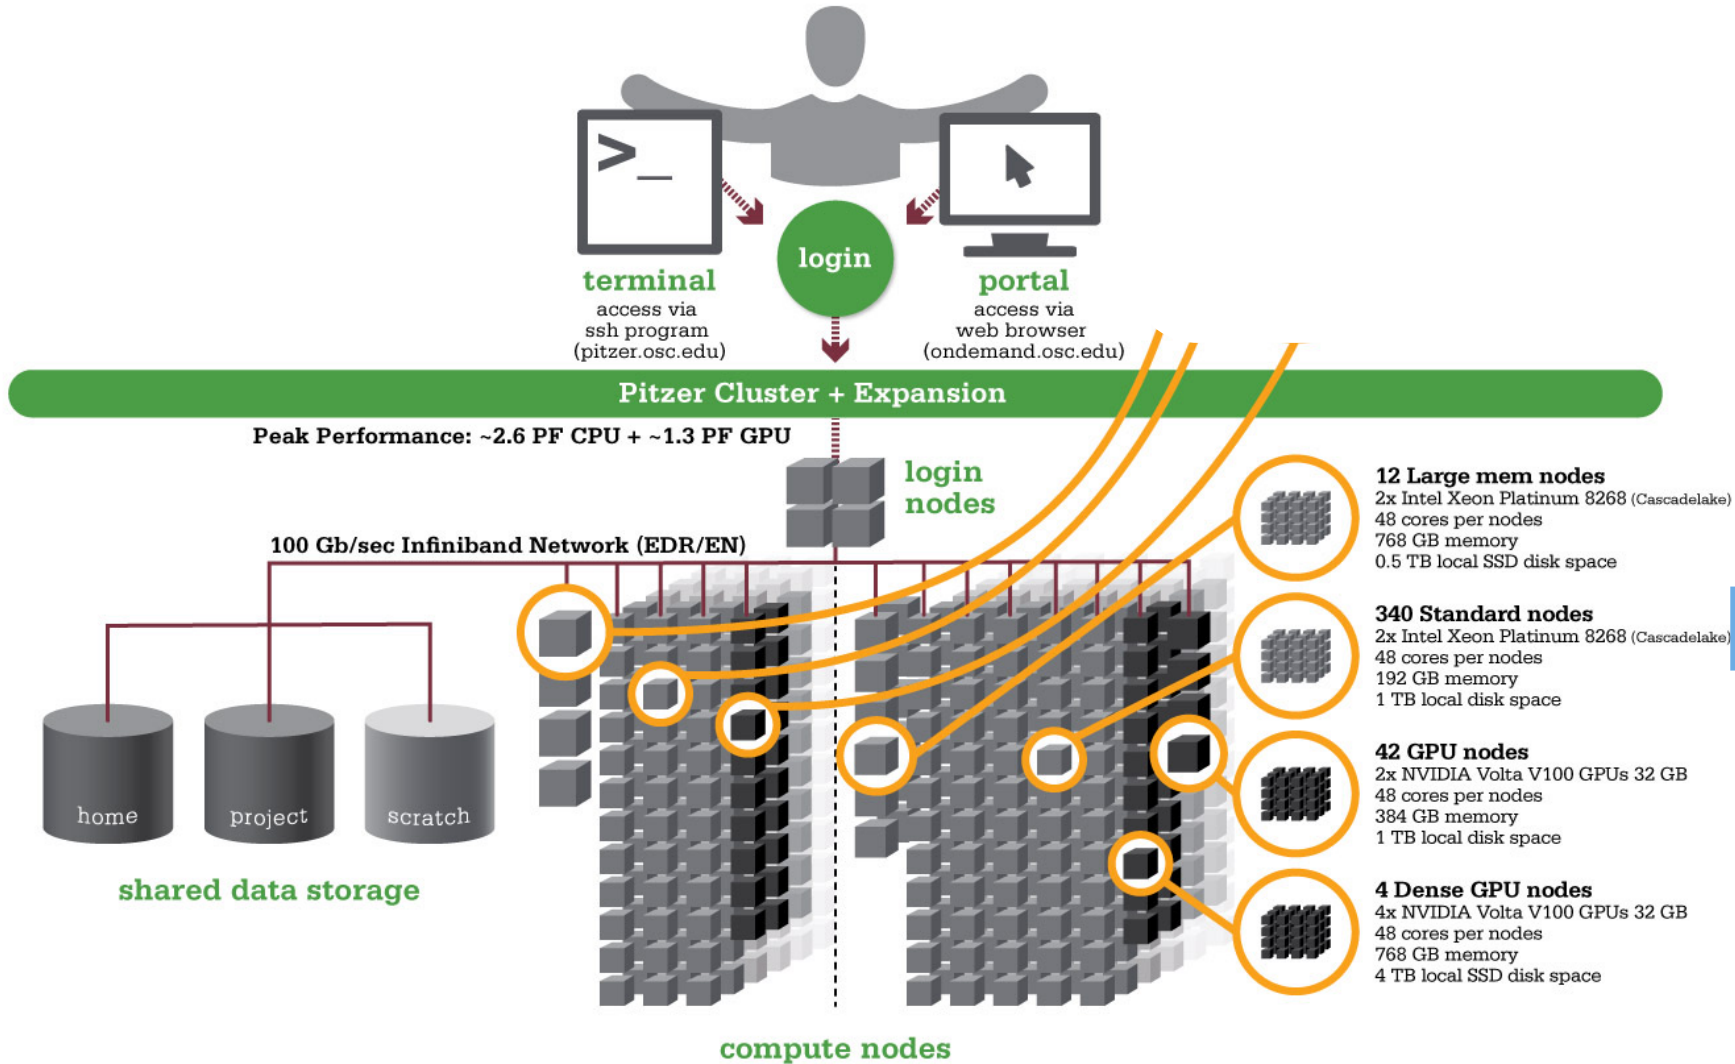


Pitzer Cluster Specifications

Original




Pitzer Cluster Specifications


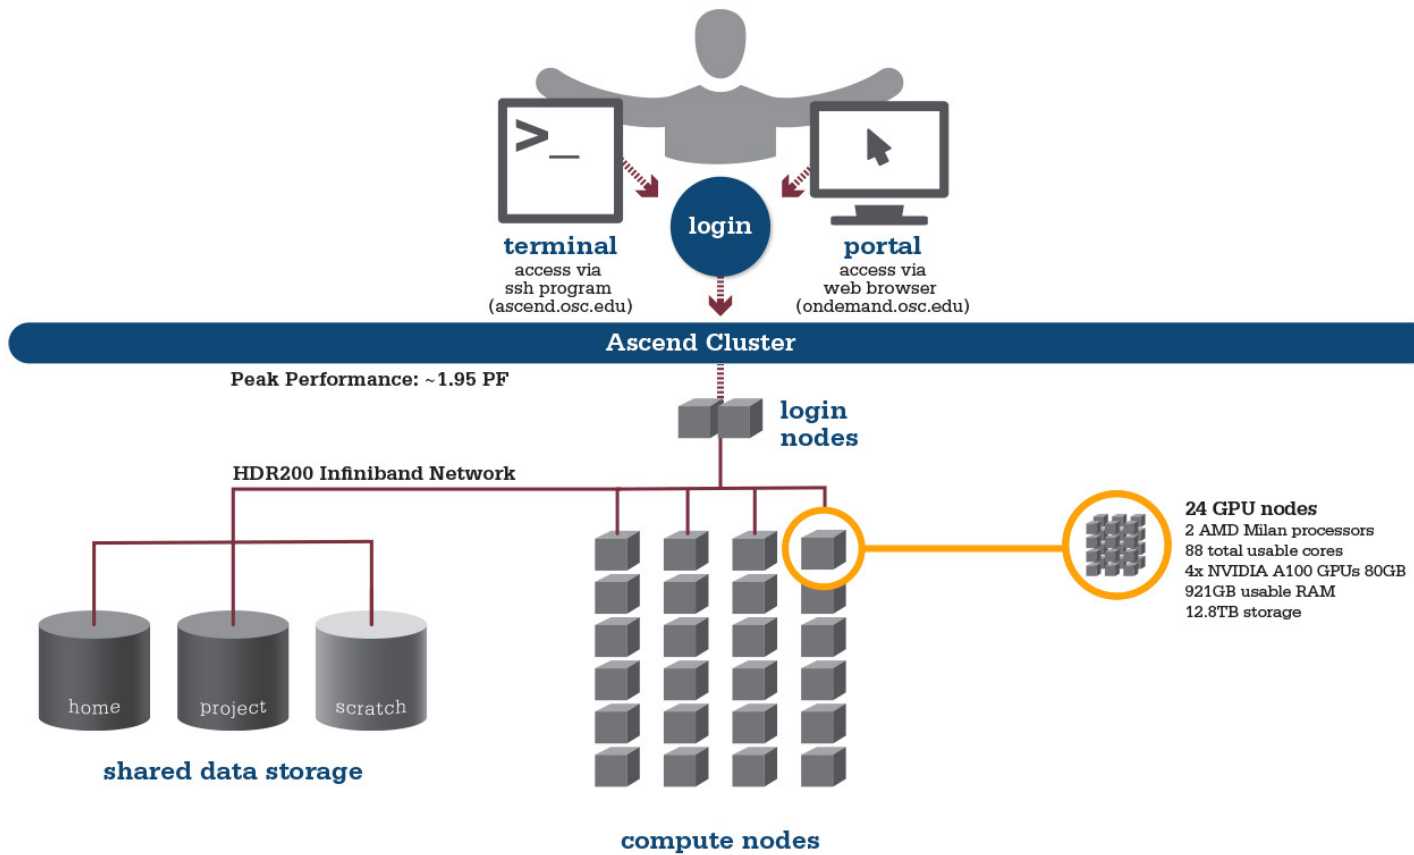


Expansion

https://www.osc.edu/resources/technical_support/supercomputers/pitzer




Ascend 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**




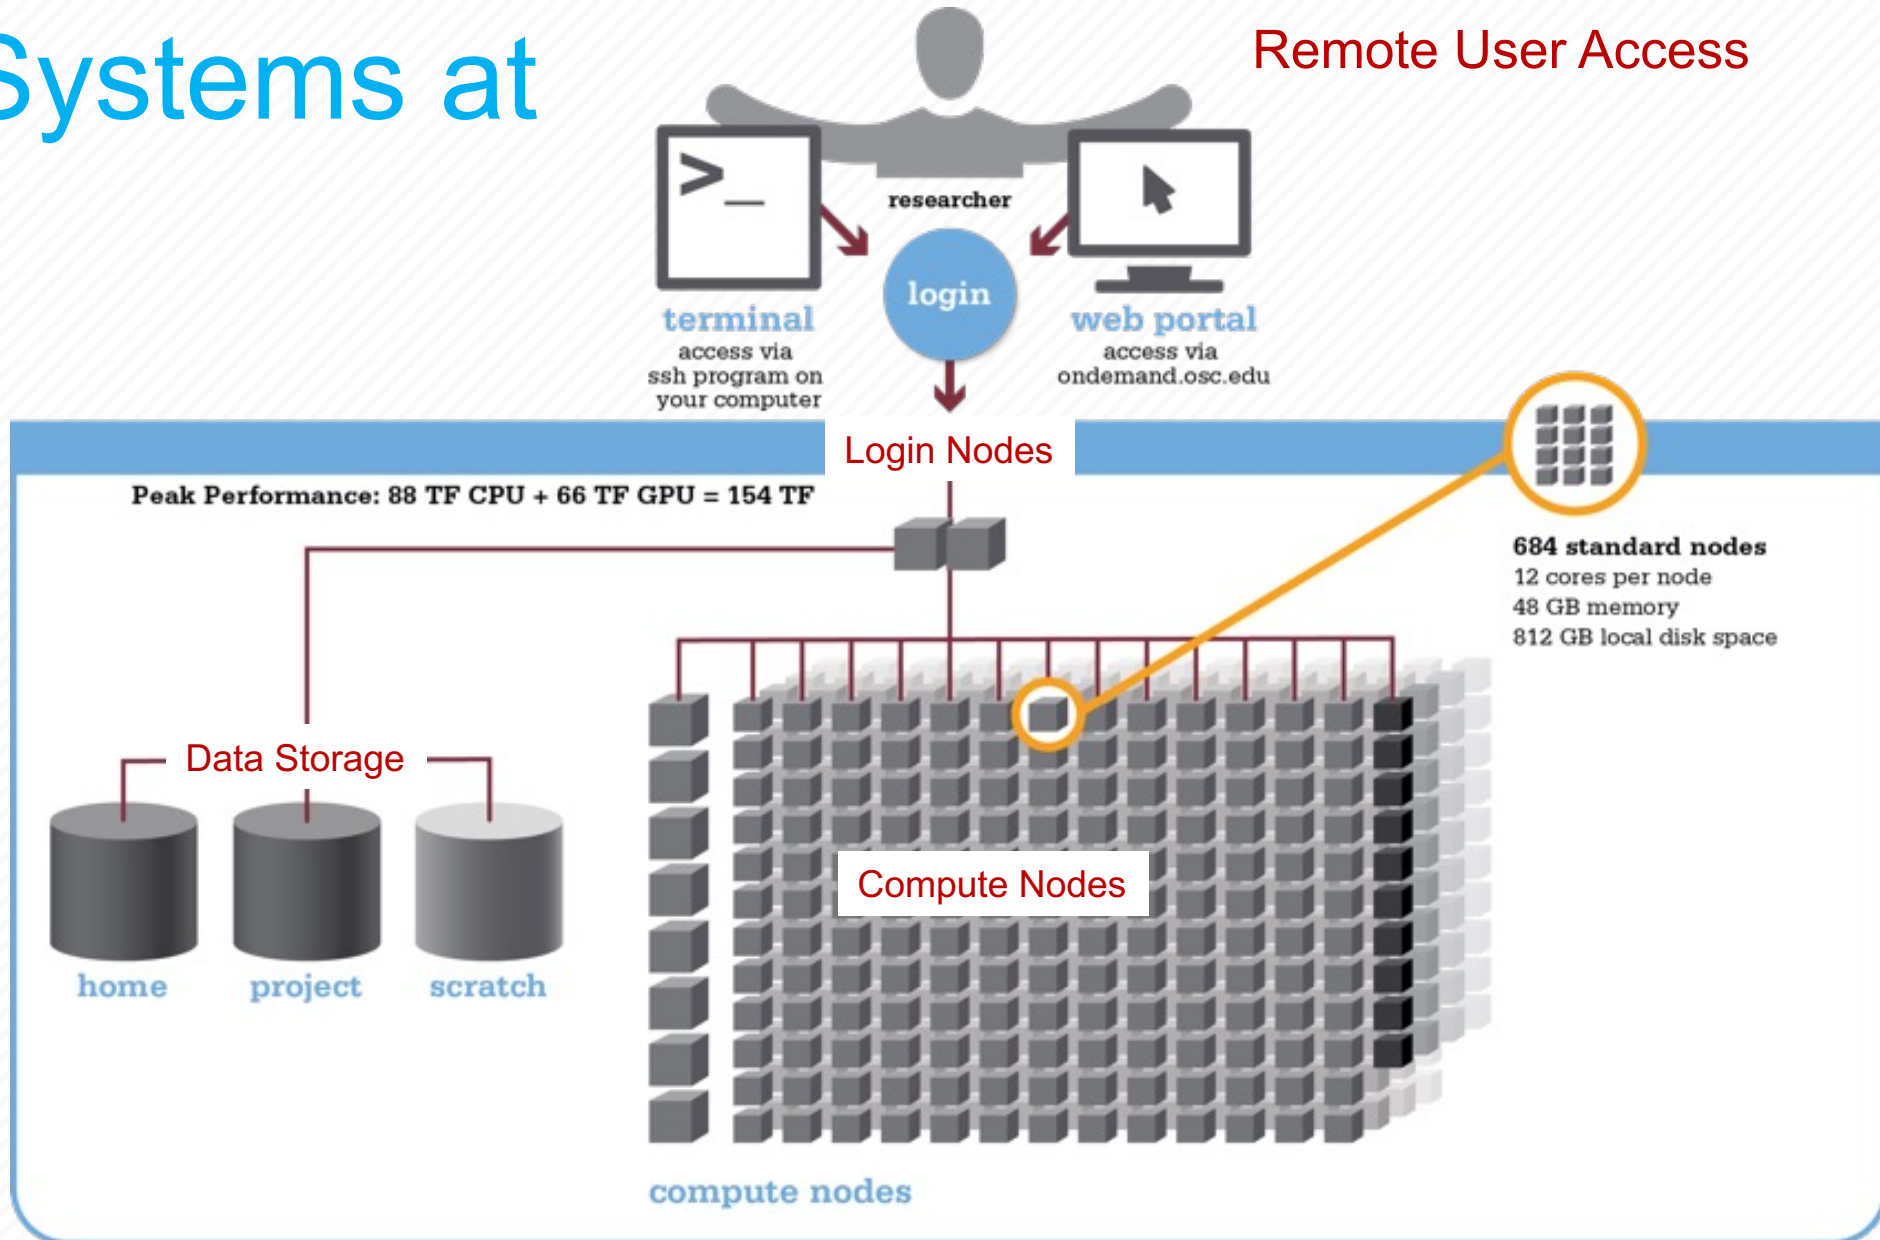


Data Storage Systems



File Systems at OSC


Remote User Access



Research Data Storage


- Home
 - Store your files here, backed up daily
 - 500GB per user
- Project/ESS
 - Available to Project PIs by request; shared by all users on a project, backed up daily
 - 1-5TB standard request
- Scratch
 - Temporary storage, not backed up
- Compute Nodes
 - Storage on compute nodes, for use during your batch job
 - All data purged when job quits
- Archive
 - Long term storage, by request


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




Filesystem <small>(Reference)</small>	Quota	Network	Backed-Up?	Purged?
Home (\$HOME)	500GB	40 GB/s	Yes	No
Project (/fs/ess/ <Project>)	By request (Typically 1-5 TB)	50 GB/s	Yes	No
Scratch (/fs/scratch/ <Project>)	100TB	170 GB/s	No	Yes – 90 days
Compute (\$TMPDIR)	Varies (~1 TB)	Varies	No	Yes – when job completes

https://www.osc.edu/supercomputing/storage-environment-at-osc/storage-hardware/overview_of_file_systems






Getting Started at OSC




Who can get an OSC project?

- Academic project
 - 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 projects are also available
- Commercial projects
 - Commercial organizations may purchase time on OSC systems



Accounts and Projects at OSC


- Project (OSC defined: PZZXXXX)
 - Headed by a PI
 - May include other users
 - Oversees computing resources for a project
- Account (User defined)
 - Username and password to access HPC systems
 - Each account used by one person
 - If you work on multiple projects, you will have one account that can access all of them



Usage Charges

- Charges are in terms of core hours, GPU hours, TB months
- Project has a dollar balance
- Services, e.g. compute and storage, are charged to a project
- General Compute, GPU, Huge Memory, Storage costs are still partially subsidized and highly competitive


https://www.osc.edu/content/academic_fee_model_faq



Ohio Academic Projects

- Standard Projects
 - Each PI can receive \$1,000 grant annually to cover OSC services
 - PI can set a budget so no unexpected charges
 - No more proposal submissions
 - 500GB storage, 1,000,000 files (can request more storage)
- Classroom projects are fully subsidized
- Request at my.osc.edu

<https://www.osc.edu/supercomputing/support/account>




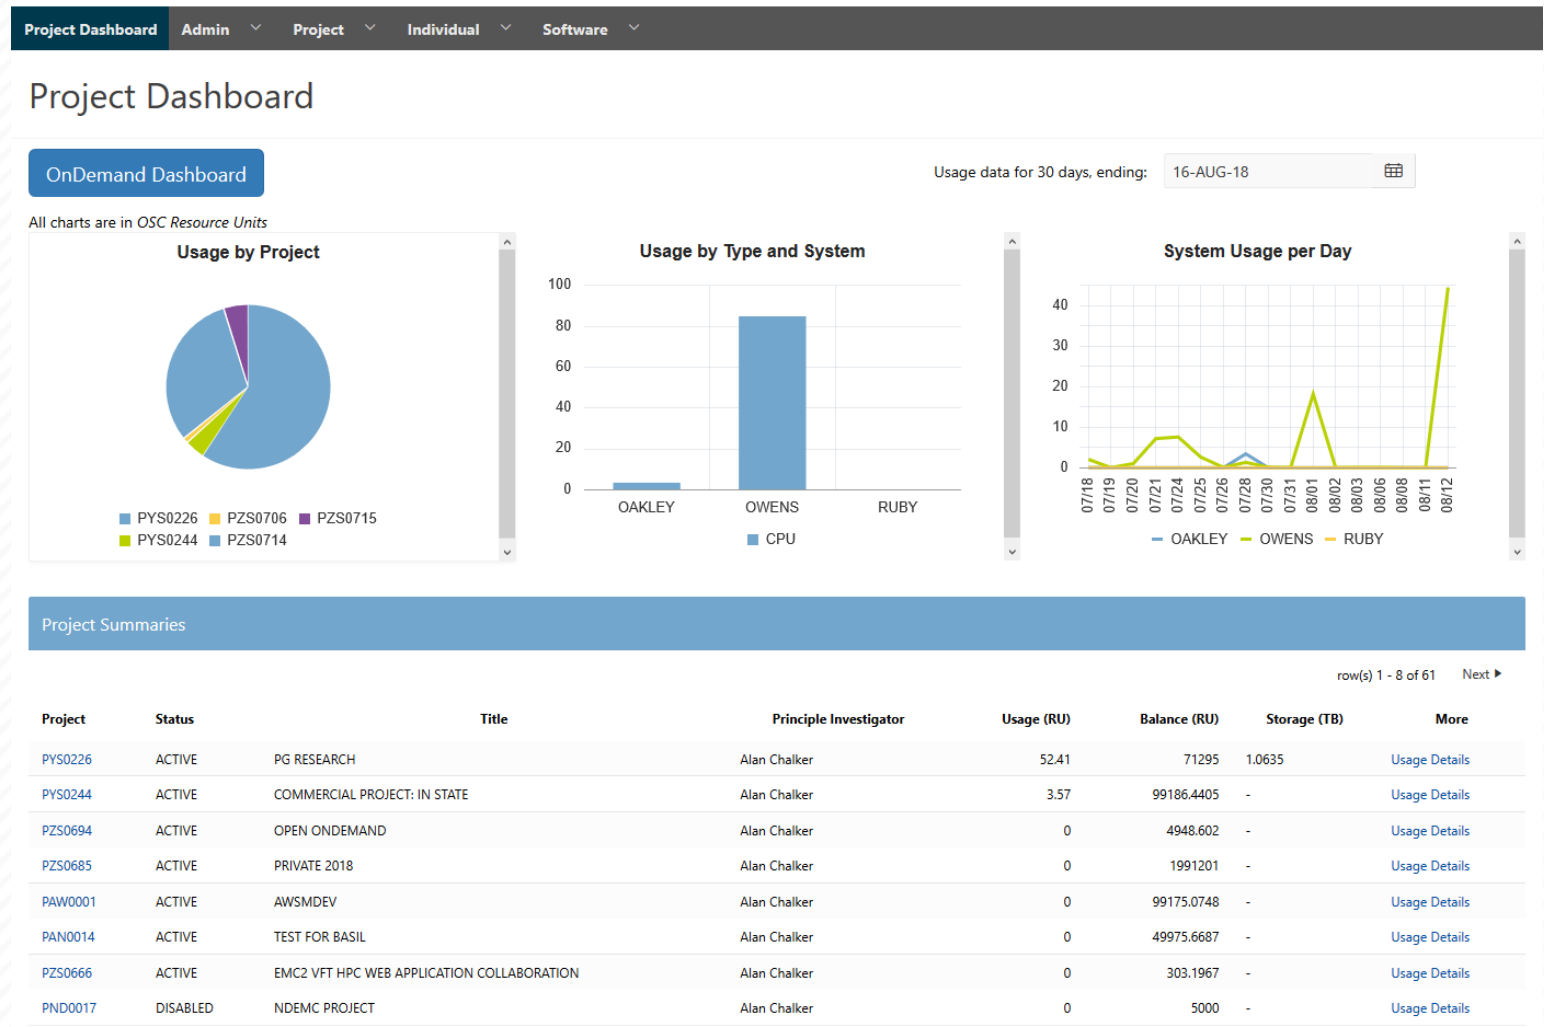
Client Portal– my.osc.edu

Features

- Create your account
- Update your email
- Change your password
- Recover access to your account
- Change your shell


PI resources

- Project reporting
- Authorized user management
- Requesting services (e.g. software access)




Statewide Users Group (SUG)

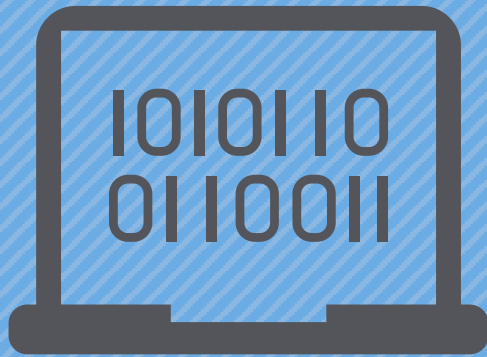
- The Statewide Users Group (SUG) is made up of OSC users
 - Provides program and policy advice to OSC
 - Meets twice a year
 - Headed by a chairperson elected yearly
- Standing committees
 - Software and Activities
 - Hardware and Operations
- Get involved!
 - Twice a year- April and October: [OSC calendar](#)




Communications & Citing OSC


- Keep your email updated
 - We send regular user emails
 - notifications of regular and unplanned maintenance
 - quarterly downtimes
- Please cite OSC in your publications:
 - Details at www.osc.edu/citation
- These publications should be reported to OSC






Short Break






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: <https://www.osc.edu/content/linux-command-line-fundamentals>
- www.linux.org



Connecting to an OSC Cluster


- Connect to OSC machines using **ssh** (secure shell)
 - From a Linux/UNIX (and Mac) terminal: At prompt, enter
`ssh userid@owens.osc.edu`
 - From Windows: **ssh** client software needed
 - Both commercial and free versions are available
- Connect using OSC OnDemand portal
 - 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



Transferring Files to and from the Cluster

- File transfers to and from OSC machines can use `sftp` or `scp` from a terminal window
- For small files, connect to a login node
`owens.osc.edu`
- For large files, transfer may fail due to shell limits
 - Connect to `sftp.osc.edu` (file transfer only)
- OnDemand drag and drop file transfer up to 10GB files
- GLOBUS-- a large file transfer system

https://www.osc.edu/resources/getting_started/howto/howto_use_globus_overview



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

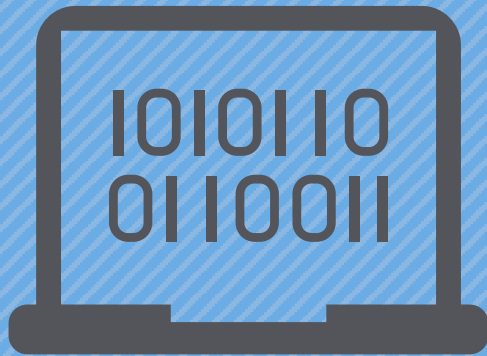


- 2: Interactive Services


- File Access
- Job Management & Monitoring
- Visualization Apps
 - Desktop access
 - Single-click apps (Abaqus, Ansys, Comsol, Paraview)
- Terminal Access

Tutorial available at
osc.edu/ondemand






Using and Running Software at OSC




Software Maintained by OSC

- 235+ software packages maintained for users
- **Always** first check software page on https://www.osc.edu/resources/available_software/browse_software
 - Version information for all clusters
 - License information – some software you must request access
 - Usage examples




Third party applications

- **General programming software** (⌘ statewide licensed)
 - gnu compilers and debugger
 - ⌘ Intel compilers
 - ⌘ Arm DDT debugger
 - ⌘ Arm MAP profiler
 - ⌘ ANSYS
 - MPI library
 - HDF5
 - NetCDF
 - Java, Java Virtual Machine
 - Python
 - R Statistical & Programing environment




Third party applications

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



Access to Licensed Software


- Most software licenses for academic use only
- Some software requires signed license agreement
 - Check software page at OSC.edu
 - 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

*https://www.osc.edu/resources/getting_started/howto/howto_locally_installing_software



Loading Software Environment



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

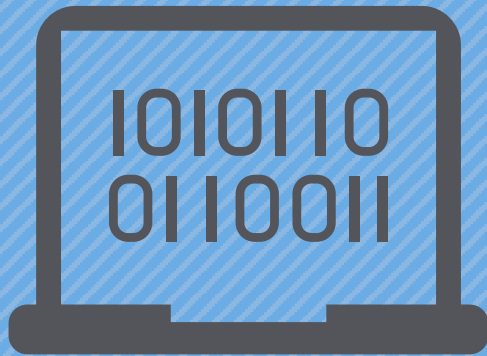
`module unload intel`




Load a different software version

`module swap intel intel/13.1.3.192`


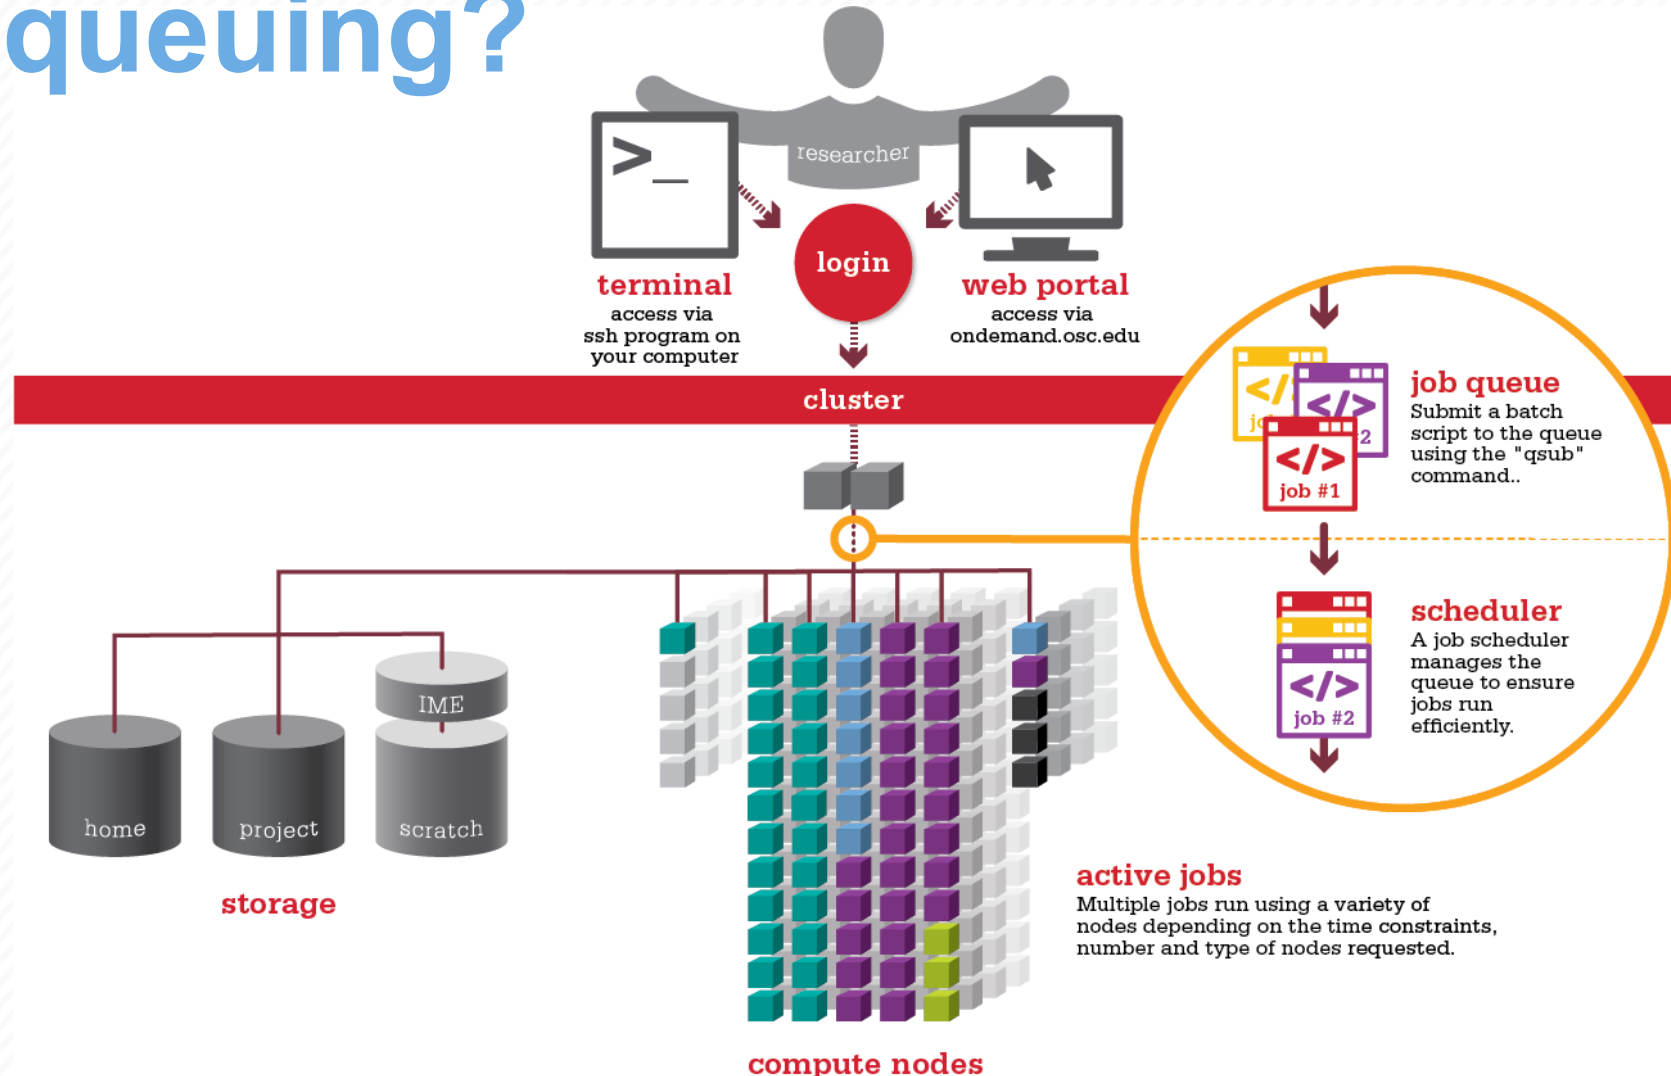




Batch Processing




Why do supercomputers use queuing?




Steps for Running a Job on the Compute Nodes

1. Create a batch script for a job
2. Prepare and gather input files in your home directory
3. Submit the job
4. Job gets queued
5. Job runs when resources become available
6. Get your results in your home directory when the job finishes



Specifying Resources in a Job Script

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



Sample Slurm Batch Script

```
#!/bin/bash
#SBATCH --time=1:00:00
#SBATCH --nodes=2 --ntasks-per-node=40
#SBATCH --job-name=hello
#SBATCH --account=PZSXXXX
#SLURM already starts job in working directory
cd $SLURM_SUBMIT_DIR


#Set up software environment
module load intel
#Move input files to compute node
cp hello.c $TMPDIR

mpicc -O2 hello.c -o hello
srun ./hello > hello_results
#Copy results back to working directory
cp hello_results $SLURM_SUBMIT_DIR
```

Job setup information
for SLURM

Commands
to be run

Put all this into a text file!



Submit & Manage Batch Jobs


Slurm Directive	Description
<code>sbatch <jobscript></code>	Submit job script
<code>scancel <jobid></code>	Cancel a job
<code>scontrol hold <jobid></code>	Put job on hold
<code>scontrol release <jobid></code>	Release job from hold
<code>squeue -u <user></code>	View information about job(s) of a user

Submitted job Slurm response:

Submitted batch job 35484

Learn more at Batch System training – check OSC calendar

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




Scheduling Policies and Limits

- **Walltime limit**
 - 168 hours for serial jobs (single node)
 - 96 hours for parallel jobs (multiple nodes)
- **Per-user limits**
 - 384 concurrently running jobs
 - 3080 processor cores in use
 - 1000 jobs in the batch system, running or queued
- **Per-group limits**
 - 576 concurrently running jobs
 - 3080 processor cores in use


**Limits for
Owens**

Find limits for other systems: https://www.osc.edu/services/cluster_computing



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



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)

```
sinteractive -N 1 -n 4 -t 00:10:00 -J test
```



Batch Queues

- Our 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)



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 interface (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



Use *mpiexec* for multiple nodes



Can't just request more nodes or cores and expect your job to run faster



New - Online Training Available!

FREE OWENS

Ohio Supercomputer Center
An OH-TECH Consortium Member

JESSE OWENS

Batch System at OSC

An introduction to submitting batch jobs at Ohio Supercomputer Center clusters

THE OHIO STATE UNIVERSITY

Self-paced
FREE

Introduction to OSC High...


An introduction to services and resources for new researchers using OSC clusters

THE OHIO STATE UNIVERSITY

Nov 1, 2022 - Jul 1, 2023
FREE

- ScarletCanvas hosted courses and modules
- Free and available to anyone (need to create a ScarletCanvas account)
- Covers material from OSC Intro and Intro to Batch System webinars
- Includes Hands-on activities & quizzes
- More training materials will be added (let us know what you would like to see)

<https://www.osc.edu/supercomputing/training>



Resources to get your questions answered

Getting Started: https://www.osc.edu/resources/getting_started/new_user_resource_guide

FAQs: https://www.osc.edu/resources/getting_started/supercomputing_faq

HOW TOs: https://www.osc.edu/resources/getting_started/howto

Installing Software


Installing R packages

Tutorial materials: https://khill42.github.io/OSC_IntroHPC/

Office Hours: Virtual, every other Tuesday, 1:00pm – 4:00pm [Sign up](#) in advance


System updates

- Read Message of the Day on login
- Follow @HPCNotices on Twitter



Key OSC Websites

- **OSC.EDU**
 - Main informational site – event calendar, HOWTO's, help pages
- **MY.OSC.EDU**
 - Client portal – manage accounts & projects, request services
- **ONDEMAND.OSC.EDU**
 - Cluster portal – access clusters, transfer files, submit and review jobs, run interactive apps





OH·TECH

Ohio Technology Consortium
A Division of the Ohio Department of Higher Education

 info@osc.edu

 twitter.com/osc

 facebook.com/ohiosupercomputercenter

 osc.edu

 oh-tech.org/blog

 linkedin.com/company/ohio-supercomputer-center