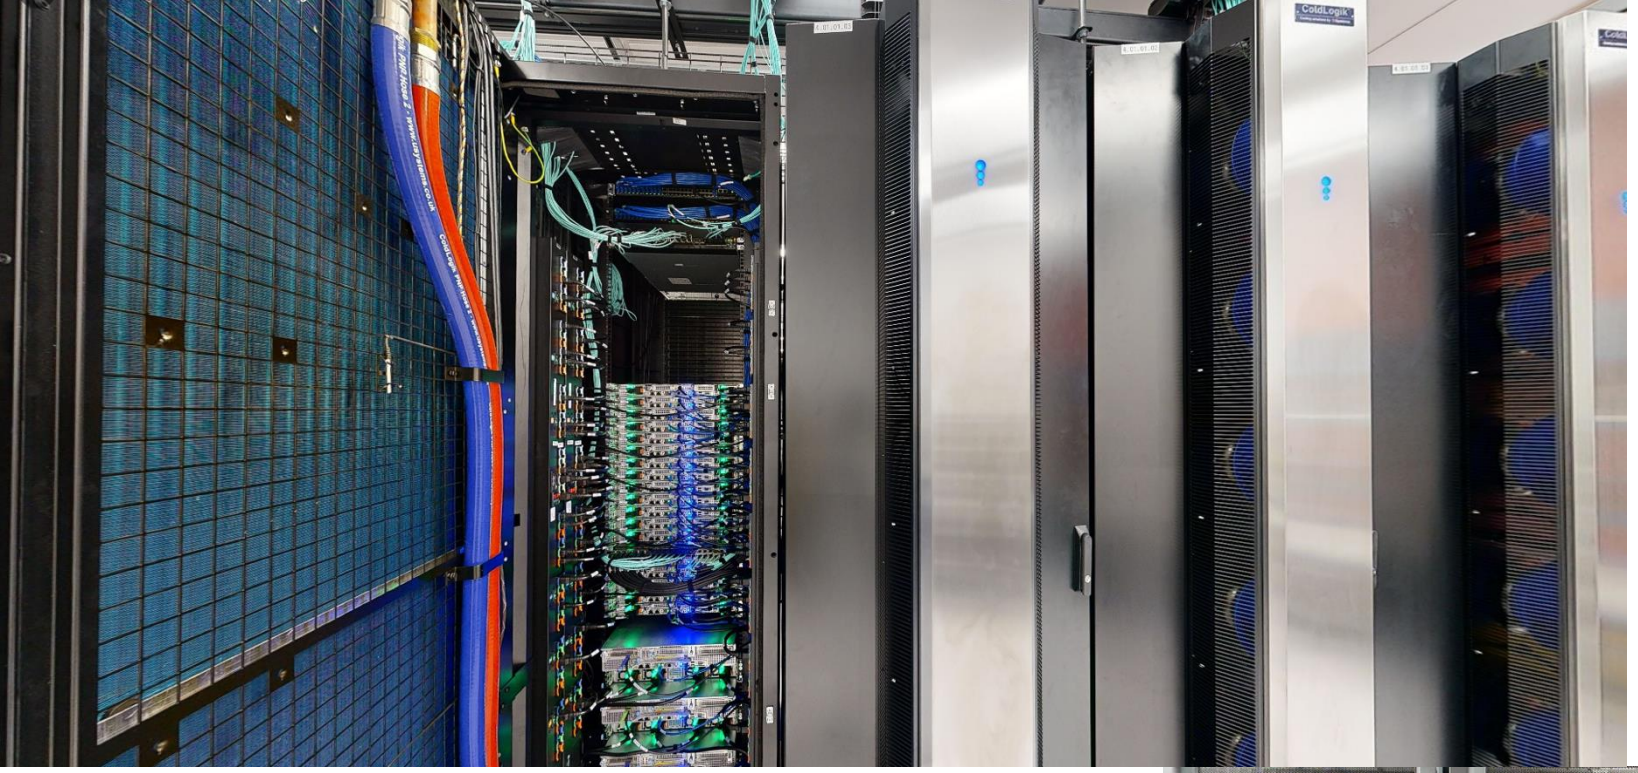

NJIT
New Jersey Institute of Technology

High Performance Computing

Introduction to Wulver: Resources & HPC

Jan 28, 2026

Outline

- What is NJIT Advanced Research Computing HPC?
- High-Performance Computing (HPC) Concepts
- Hardware Overview
- Getting a Login
- Allocations
- Data Storage systems
- Computing Resources
- User Environment
- Using Software on HPC
- Contact Us

What is the NJIT Advanced Research Computing HPC?

About NJIT HPC

NJIT new high performance computing environment, built through a partnership with DataBank, a leading provider of enterprise-class colocation, connectivity and managed services, is live in DataBank's Piscataway, N.J. data center (EWR2) and will support NJIT's research efforts.

The services NJIT HPC provides

- High performance computing services
- Computational science expertise

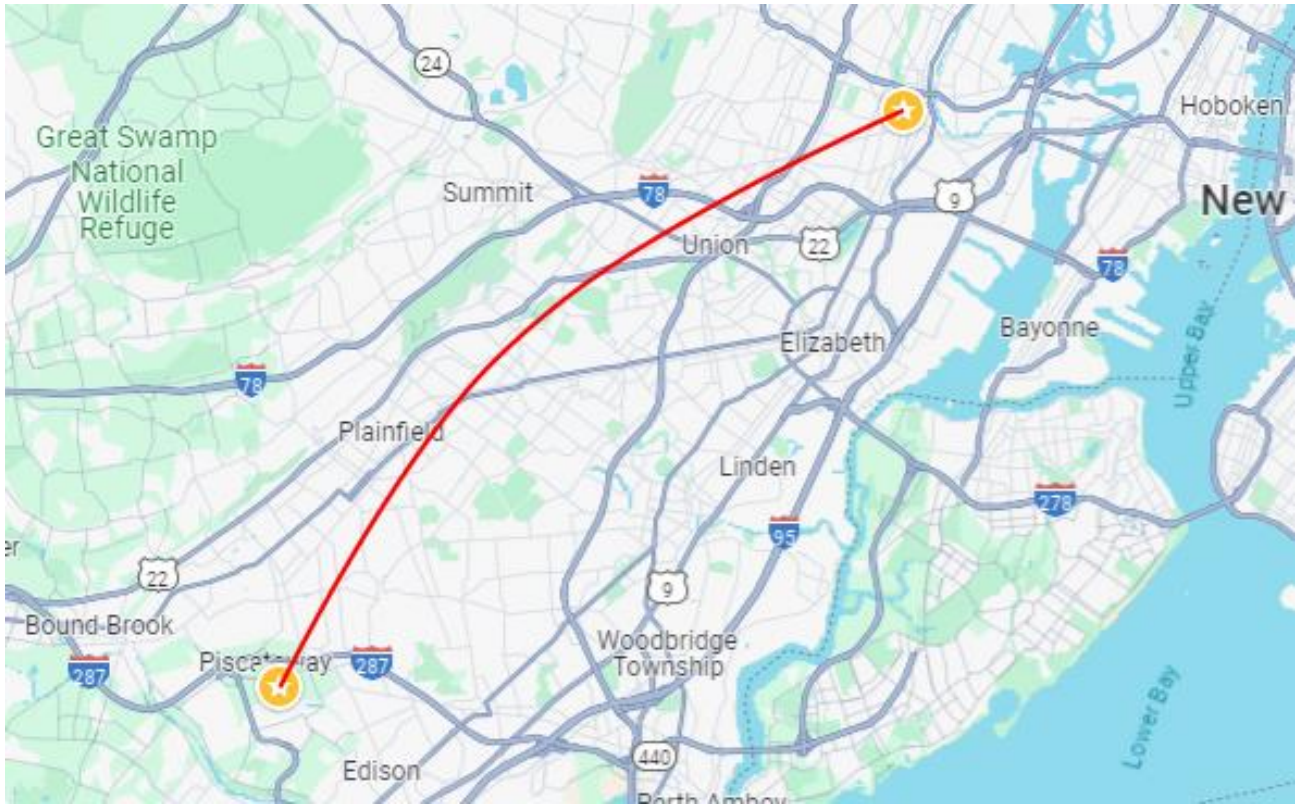

Service Catalog

Cluster Computing

Built by Dell, the computing environment “Wulver” provides a total of 197 compute nodes or servers

Research Data Storage

High-performance, large capacity data storage spaces 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

HPC Facilitation Service

Empowering users to perform essential research computing projects through training and effective user support

Scientific Software Development

Deep expertise in developing and deploying software

HPC Concepts

Why Use HPC?

Your simulations or analyses take too long on your personal computer

- More (faster) cores
- Multithreading
- Multi-node parallelization (openmpi & intelmpi)
- GPU acceleration (NVIDIA's CUDA)
- Distributed computing

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

- Large memory nodes: 512GB; 2 TB
- Distributed computing
- Project storage: TBs range

You need a particular software or package for your work

When HPC is not useful

- **Small-scale tasks:** Suitable for desktops or laptops (e.g., basic data analysis, simple simulations).
- **Real-time processing:** Tasks requiring immediate interaction or response (e.g., interactive graphics, live streaming).
- **Lack of parallelism:** You are running a serial code.
- **Running databases**

What is the difference between your laptop and a supercomputer?

X 10000 =

Remote Connect

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: ~15 GB/core

Hardware Overview

Wulver Cluster Specifications

ssh $UCID@wulver.njit.edu

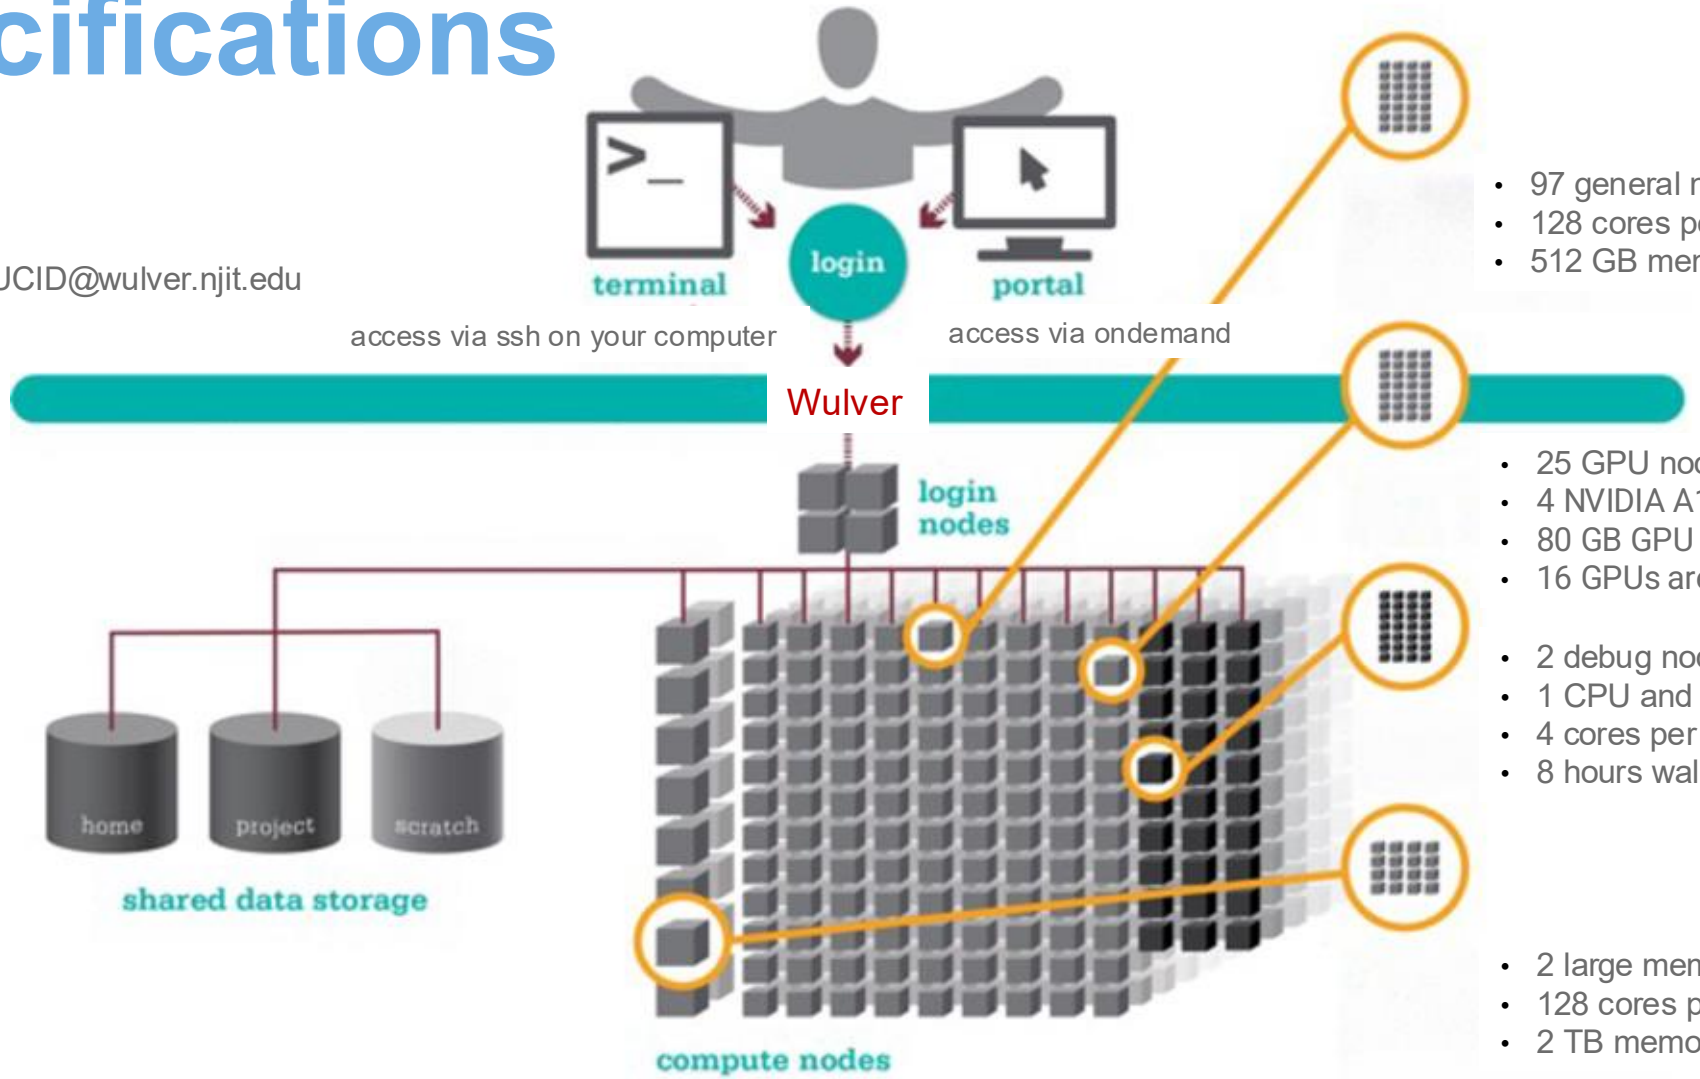

- 97 general nodes
- 128 cores per node
- 512 GB memory/Node

- 25 GPU nodes, 128 cores per node
- 4 NVIDIA A100 GPUs per node
- 80 GB GPU memory per GPU
- 16 GPUs are sliced to 64 MIGs

- 2 debug nodes
- 1 CPU and 1 GPU node
- 4 cores per node
- 8 hours wall time

- 2 large memory nodes
- 128 cores per node
- 2 TB memory/Node

hpc.njit.edu/clusters/get_started_on_Wulver

Getting Started at Wulver

Getting Access to Wulver

- Getting a Login
  - Faculty (PI) can obtain a login to NJIT's HPC by sending an email to hpc@njit.edu
  - Students can obtain a login by asking their research adviser to contact on their behalf.
  - For course requiring HPC resources students need to contact their course instructor.

Login Nodes – Usage

- Purpose
  - Submit jobs to batch system
  - Edit files
  - Manage your files
  - Interactive work – small scale
- Limits
  - 20GB memory
  - CPU usage is limited to 25% per user
  - Processes are limited to 100 per user
- **Use the batch system for serious computing**

Allocations

HPC Allocations

Storage

- Home (~50GB/user) - Limited quota: not intended for primary storage
- Project (2TB/PI Group) – Long term storage
- Scratch - Short-term storage only

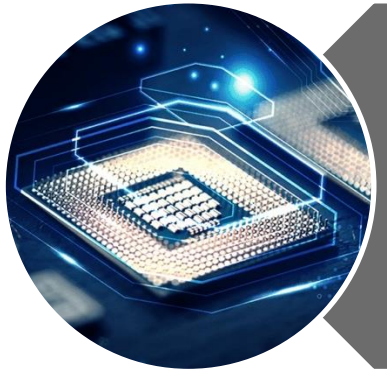

Computing Time

- 1 SU = Number of CPUs x Walltime in hours x usage factor
- No limit for low priority
- Standard annual allocation – 300,000 SU's per year

Check Allocations usage

- quota_info

```
Usage for account: xy1234
SLURM Service Units (CPU Hours): 277557 (300000 Quota)
   User ab1234 Usage: 1703 CPU Hours (of 277557 CPU Hours)
PROJECT Storage: 867 GB (of 2048 GB quota)
   User ab1234 Usage: 11 GB (No quota)
SCRATCH Storage: 791 GB (of 10240 GB quota)
   User ab1234 Usage: 50 GB (No quota)
HOME Storage ab1234 Usage: 12 GB (of 50 GB quota)
```

Data Storage Systems

Filesystem	Purpose	Quota	Backed-Up?	Purged?
Home ($HOME)	Non-research such as profile, history	50GB	Yes, daily	No
Project (`/project/$PI_UCID/$LOGIN/`)	Active research by groups. PI will be charged if requesting more space in `/project`	2 TB/ PI Group	Yes, daily	No
Scratch (`/scratch/$PI_UCID/$LOGIN`)	Temporary space for intermediate results, downloads, checkpoints, and such. MOVE YOUR RESULTS & IMPORTANT FILES TO /project or /research	10 TB/ PI Group	No	Yes – 30 days
Compute (/tmp)	Very high-speed temporary storage	Varies (~1 TB)	No	Yes – after job ends
Research (`/research/$PI_UCID`)	Long term archive. Users can buy as much as they need. Existing purchases/quotas will be kept over from Lochness.		Yes, daily	TBD

https://hpc.njit.edu/clusters/get_started_on_Wulver/#wulver-filesystems

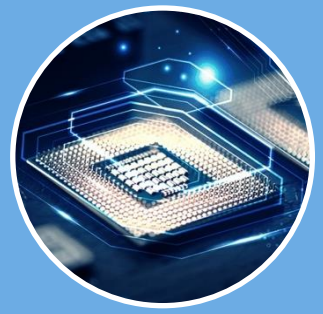

Computing Resources

Partitions

Partition	Nodes	Cores/ Node	CPU	GPU	Memory
`--partition=general`	97	128	2.5G GHz AMD EPYC 7763 (2)	NA	512 GB
`--partition=debug`	1	4	2.5G GHz AMD EPYC 7763 (2)	NA	512 GB
`--partition=debug_gpu`	1	4	2.0 GHz AMD EPYC 7713 (2)	• Two 10G MIG • One 20G MIG • One 40G MIG	512 GB
`--partition=gpu`	25	128	2.0 GHz AMD EPYC 7713 (2)	• NVIDIA A100 GPUs (4) • 64 MIGs	512 GB
`--partition=bigmem`	2	128	2.5G GHz AMD EPYC 7763 (2)	NA	2 TB

NVIDIA Multi-Instance GPU (MIG)

- Multi-Instance GPU (MIG) is a hardware-based virtualization technology introduced with NVIDIA's Ampere architecture
- MIG enables the division of a single GPU into multiple independent instances.
- MIG partitions a physical GPU into up to seven smaller GPU instances, each with dedicated compute cores, memory, cache, and bandwidth.
- Each MIG instance ensures complete resource isolation, preventing interference and allowing safe and predictable performance for multi-tenant workloads or diverse applications on the same hardware.
- On Wulver each GPU is sliced into
  - One 40G MIG
  - One 20G MIG
  - Two 10G MIGs

Job Priorities

- Standard Priority (`--qos=standard`)
  - Faculty PIs are allocated 300,000 Service Units (SU) per year on request at no cost
  - Additional SUs may be purchased at a cost of $0.005/SU.
  - The minimum purchase is 50,000 SU ($250)
  - Wall time maximum - 72 hours
  - SUs will reset every fiscal year with no carryover.
- Low Priority (`--qos=low`)
  - Not charged against SU allocation
  - Wall time maximum - 72 hours
  - Jobs can be preempted by those with higher and standard priority jobs when they are in the queue
- Debug Priority (`--qos=debug`)
  - Not charged against SU allocation
  - Wall time maximum - 8 hours
  - Must be used with `--partition=debug` or `--partition=debug_gpu`
  - Only one job per user is allowed at a time
- High Priority (`--qos=high_$PI_UCID`)
  - Not charged against SU allocation
  - Wall time maximum - 72 hours – can be increased based on PI's request
  - Only available to contributors
  - Use `listqos` command

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.hostinger.com/tutorials/linux-commands**
- www.linux.org

Available software on Wulver

- **General programming software (₵ licensed)**
  - gnu compilers and debugger
  - ₵ Intel compilers
  - ₵ ANSYS
  - ₵ COMSOL
  - ₵ MATLAB
  - Python
  - ₵ VASP
  - MD software - GROMACS, LAMMPS
  - CFD – OpenFOAM
  - Visualization Software – ParaView, ₵ Tecplot

Do you use a specific software package?

- Open-source software packages can be installed
- If you have a license, we can help you use it on Wulver

https://hpc.njit.edu/Software

Reminder

- Wulver will be temporarily out of service for maintenance once a month, specifically on the 2nd Tuesday, to perform updates, repairs, and upgrades.
- During the maintenance period, the logins will be disabled
- Jobs that do not end before the maintenance window begins will be held until the maintenance is complete
  - Reduce the walltime in the job script to run the job

- Date: Every Monday and Wednesday
- Time: 2:00–4:00 p.m.
- Location: **GITC 5302N**
- Meet with our student consultants and ask any questions you have about using HPC resources.
- There's no need to create a ticket in advance.

Resources to get your questions answered

- HPC Documentation: https://hpc.njit.edu/
- Getting Started: Access to Wulver
- List of Software: Wulver Software
- HOW TOs: Conda Documentation
  - Installing Python packages via Conda
- MIG: MIG
- Request Software: HPC Software Installation
- Access to OnDemand Open OnDemand
- Contact: Please visit HPC Contact
- Open a ticket: email to hpc@njit.edu
- Consult with Research Computing Facilitator: Facilitator Calendar Appointment
- System updates
  - Read Message of the Day on login
  - Visit NJIT HPC News

NJIT

hpc@njit.edu

hpc.njit.edu