



NJIT
New Jersey Institute
of Technology

High Performance Computing



An Introduction to NJIT High Performance Computing (HPC) and Services, Session I

April 10, 2024



Outline

- What is NJIT Advanced Research Computing HPC?
- High-Performance Computing (HPC) Concepts
- Hardware Overview
- Data Storage systems
- Getting Started at Advanced Research Computing
- User Environment
- Using Software on HPC



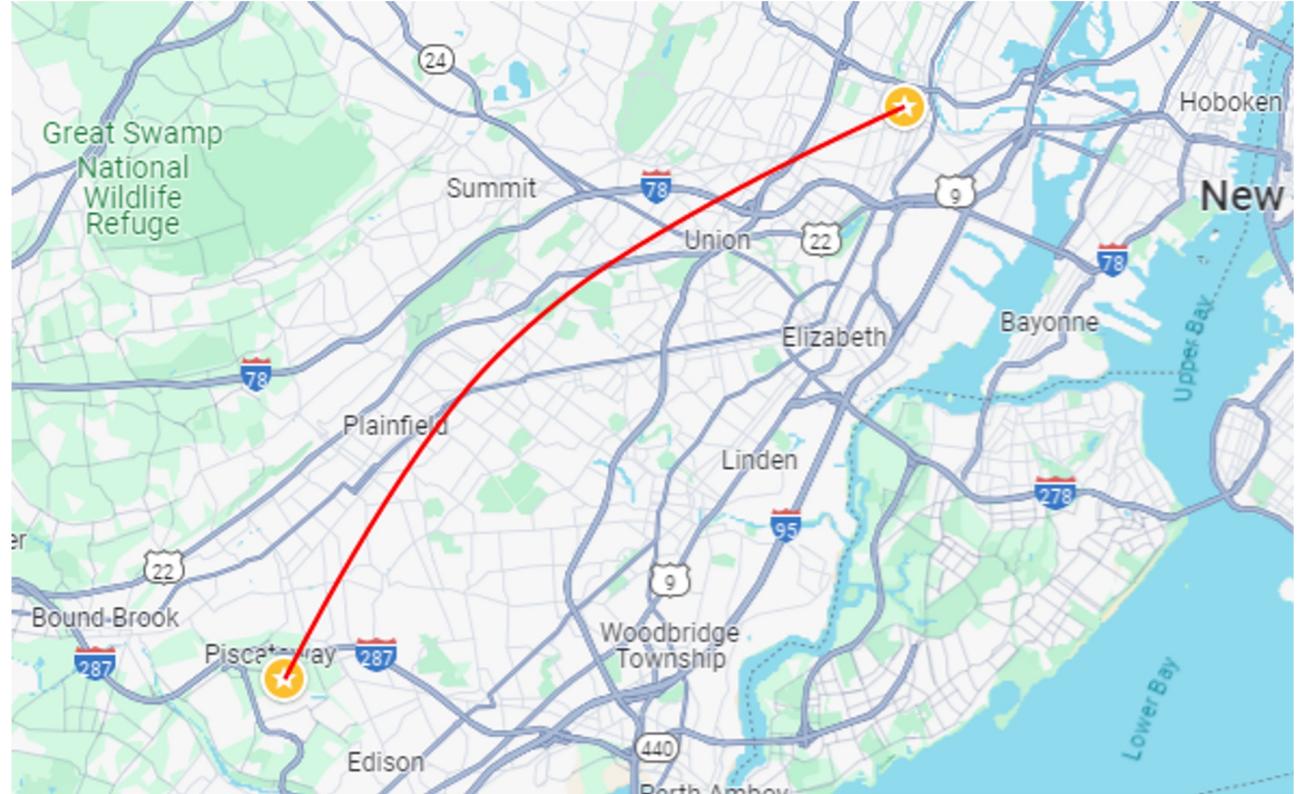
What is the NJIT Advanced Research Computing HPC?

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



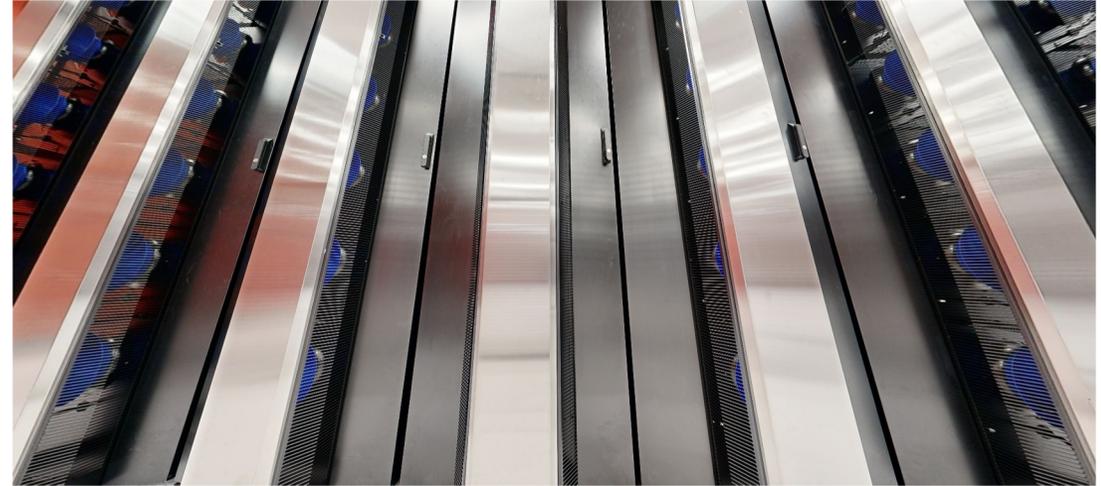
Cluster Computing

Built by Dell, the computing environment “Wulver” provides a total of 127 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



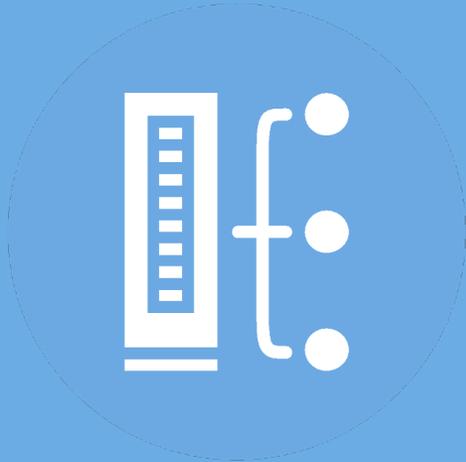
HPC Facilitation Service

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



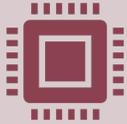
Scientific Software Development

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



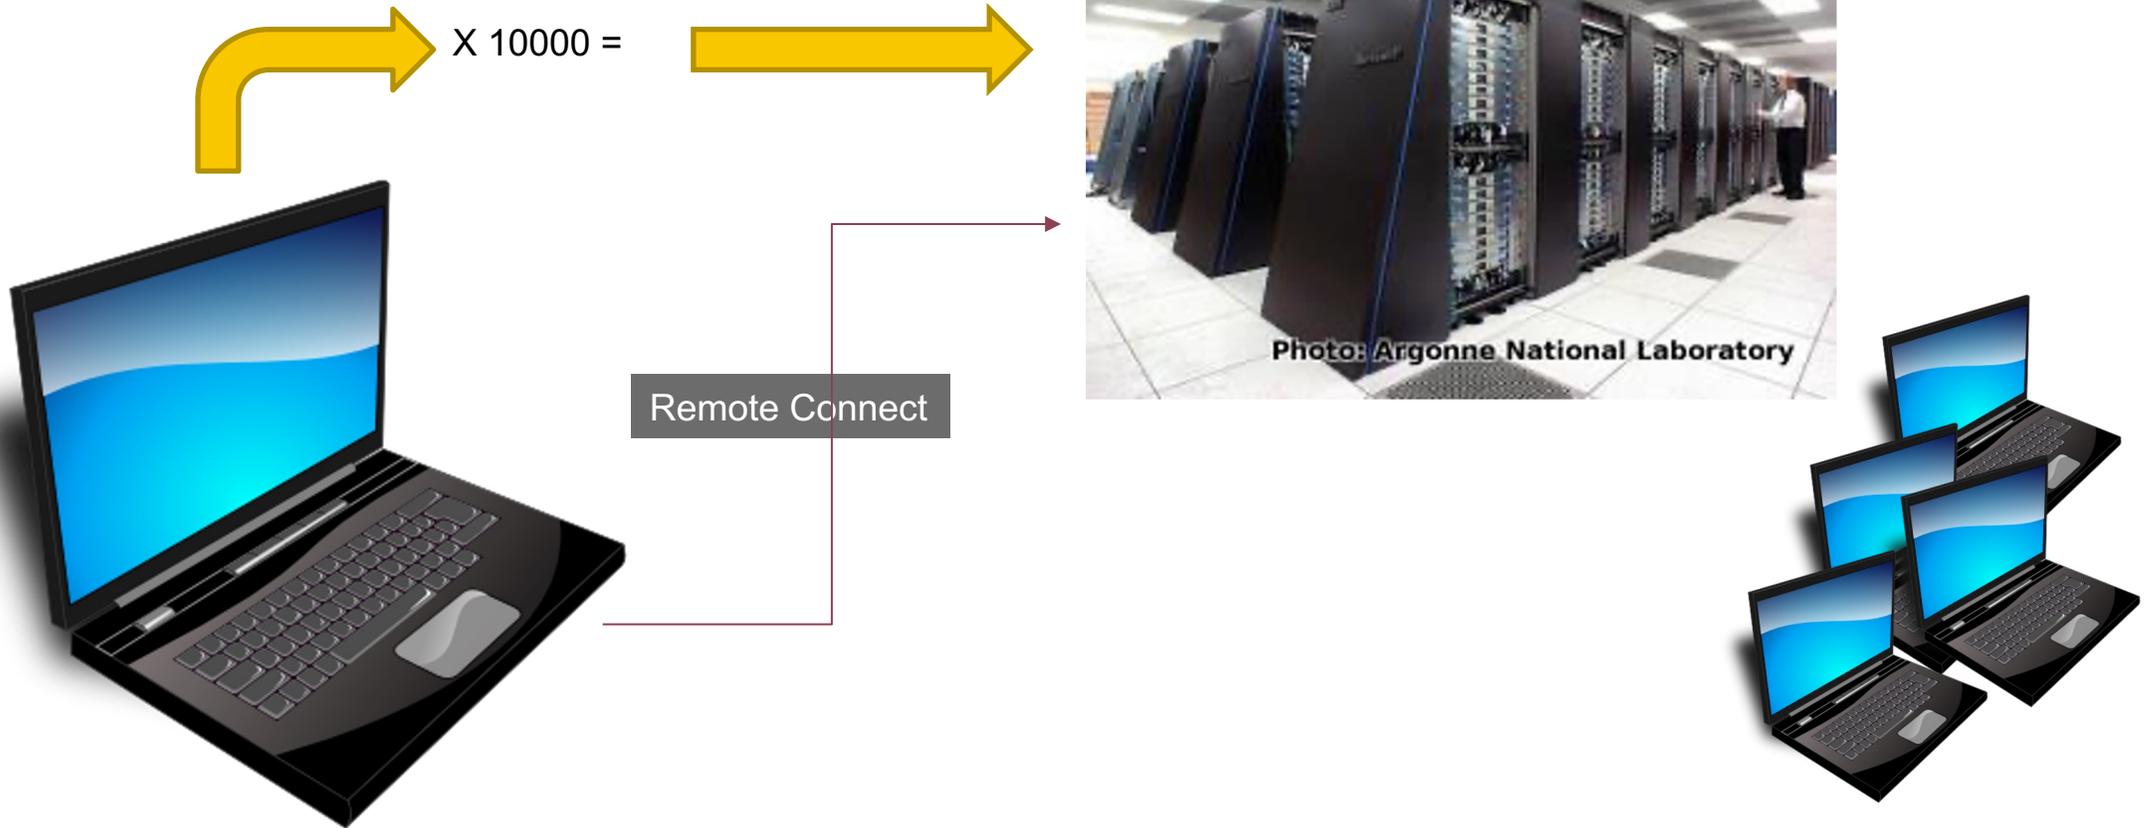
The size of your data is too large to be contained (storage) or accessed (memory) on your computer

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



You need a particular software or package for your work

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: ~15 GB/core

Storage



Storage: different types of “disk” for different needs



\$HOME directory

Limited quota: not intended for primary storage



Project storage

Long term storage



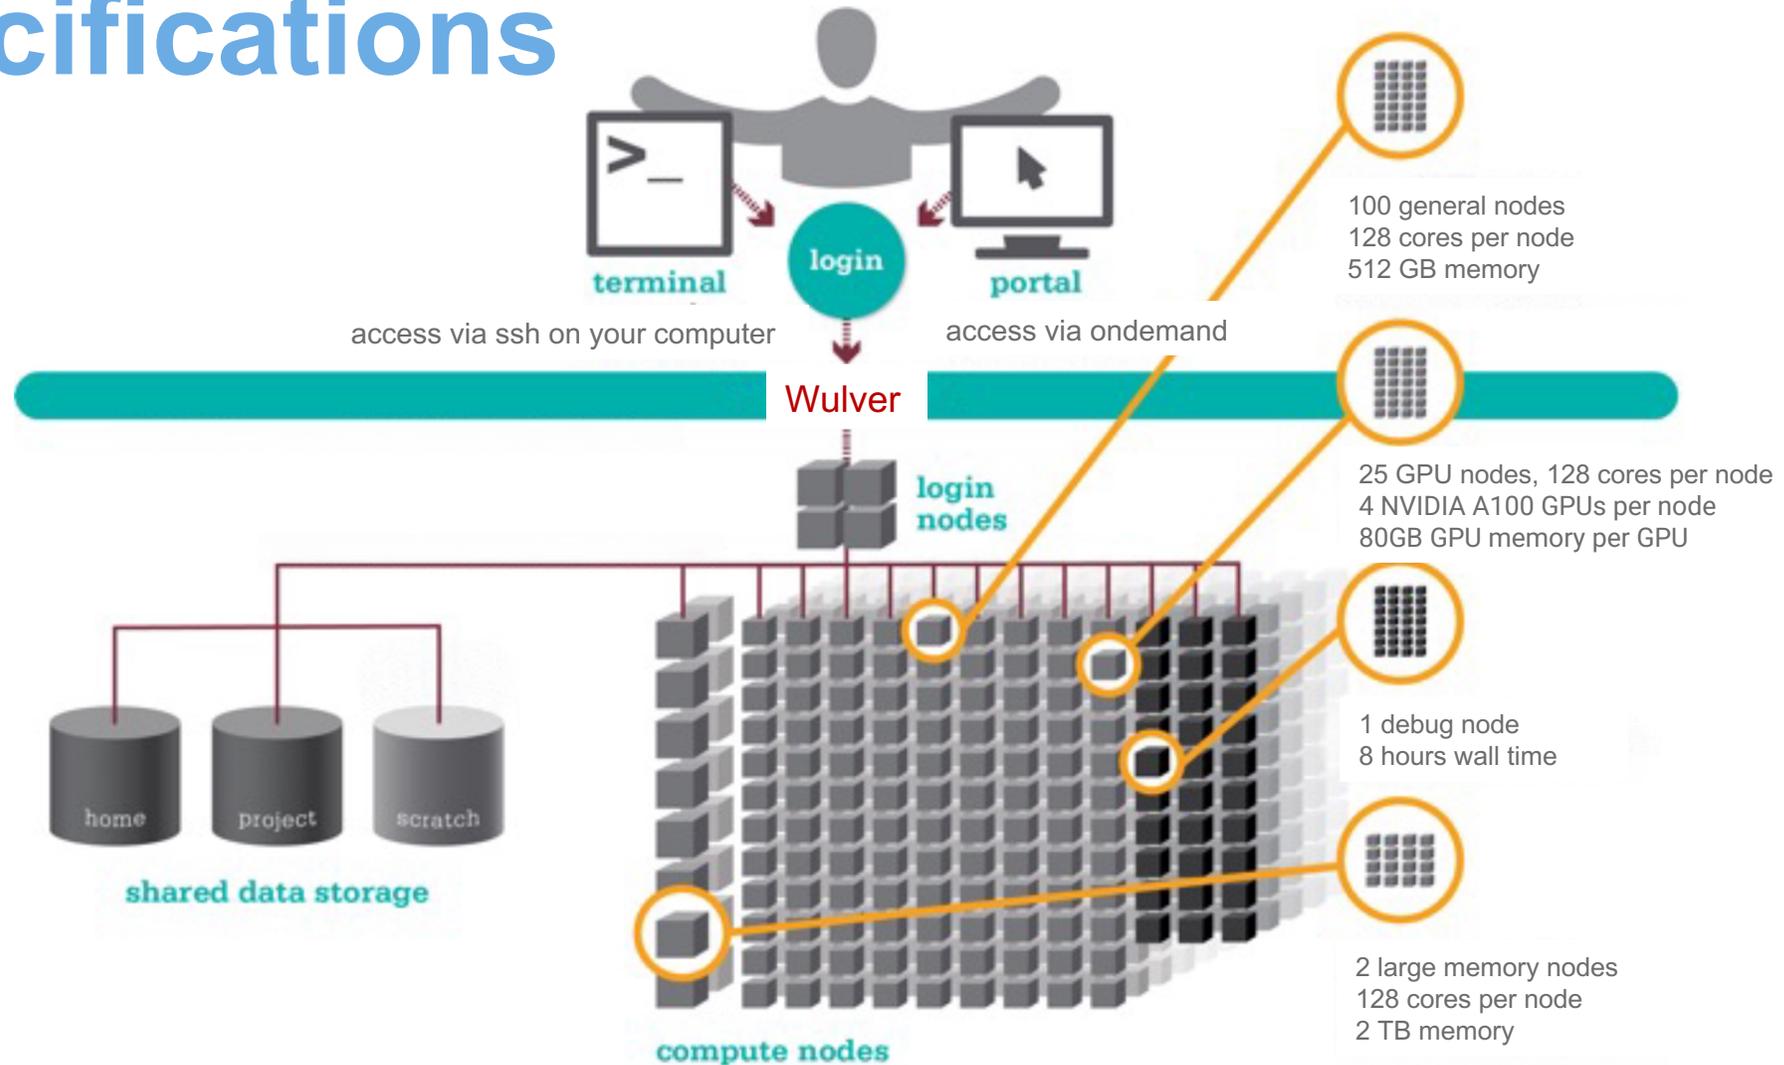
Shared scratch

Short-term storage only



Hardware Overview

Wulver Cluster Specifications



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
- **Use the batch system for serious computing**



Data Storage Systems

Research Data Storage

- **Home**
 - Store your small sized files here.
 - 50GB per user
 - Backed up
- **Project**
 - Store the large files here.
 - 2TB/PI Group,
 - Backed up
 - Additional storage can be purchased if needed.
- **Archive**
 - Long term storage, by request
- **Scratch**
 - Temporary storage, not backed up
 - Files will be deleted once they 30 days old.
 - Users need to move the important files to project before deletion.
 - Quota per PI group: 10 TB
- **Research**
 - Storage on Lochness, 500GB/PI
 - Are remounted on Wulver

hpc.njit.edu/clusters/get_started_on_Wulver/#wulver-file 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.	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



Getting Started at Wulver

Getting Access to Wulver

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

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 (\$500)
 - Wall time maximum - 72 hours
- 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
- High Priority (`--qos=high_$PI`)
 - Not charged against SU allocation
 - Wall time maximum - 72 hours – can be increased based on PI's request
 - Only available to contributors
 - `listqos` – To see the list of high priority queues

Usage Charges

- Charges are in terms of SU (Service Units) core hours
- Services, e.g. compute and storage
- PI will be charged if requesting more space in Project

Partition	Nodes	Cores /Node	CPU	GPU	Memory	SU charge
<code>--partition=general</code>	100	128	2.5G GHz AMD EPYC 7763 (2)	NA	512 GB	1 SU per hour per cpu
<code>--partition=debug</code>	1	4	2.5G GHz AMD EPYC 7763 (2)	NA	512 GB	No charges, must be used with <code>--qos=debug</code>
<code>--partition=gpu</code>	25	128	2.0 GHz AMD EPYC 7713 (2)	NVIDIA A100 GPUs (4)	512 GB	3 SU per hour per cpu
<code>--partition=bigmem</code>	2	128	2.5G GHz AMD EPYC 7763 (2)	NA	2 TB	1.5 SU per hour per cpu

Check SU and Quota usage

- `quota_info $LOGNAME`

```
Usage for account: xy1234
```

```
SLURM Service Units (CPU Hours): 277557 (300000 Quota)
```

```
PROJECT Storage: 867 GB (of 2048 GB quota)
```

```
User ab1234 Usage: 11 GB (No quota)
```

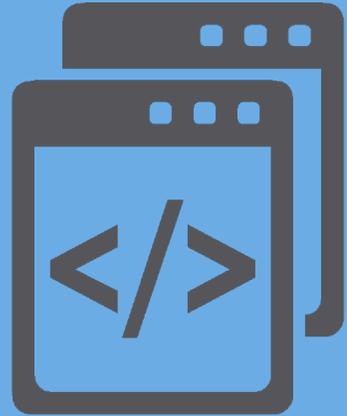
```
SCRATCH Storage: 791 GB (of 10240 GB quota)
```

```
User ab1234 Usage: 50 GB (No quota)
```

```
HOME Storage ab1234 Usage: 0 GB (of 50 GB quota)
```

Wulver Maintenance

- 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



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



Basic Linux Shell Commands

List of basic shell commands	
cd	rm
mkdir	echo
pwd	export
ln	alias
ls	grep
cp	scp
mv	rsync
cat	man

Available software on Wulver

- **General programming software** (⌘ licensed)

- gnu compilers and debugger
- ⌘ Intel compilers
- ⌘ ANSYS
- ⌘ COMSOL
- ⌘ MATLAB
- ⌘ VASP
- Python
- 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

Resources to get your questions answered

Getting Started: [Access to Wulver](#)

List of Software: [Wulver Software](#)

HOW TOs: [Conda Documentation](#)

Installing Python packages via Conda

Request Software: [HPC Software Installation](#)

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

 twitter.com/njit

 facebook.com/NewJerseyInstituteofTechnology

 linkedin.com/school/njit/