

Intro to HPC & Linux

Getting started with high performance computing services in the Department of Statistics

Dayo Ntwari

Senior Computing Specialist
Department of Statistics, Room G.07
University of Oxford
dayo.ntwari@stats.ox.ac.uk

September 2025

Session Overview



- What is HPC?
- Why use HPC in research?
- Intro to Linux basics
- Connecting the the HPC
- Submitting a Slurm job
- Hands-on session
- Wrap-Up and Q&A

What is HPC?



Definition: Many computers working together as a team to solve large problems quickly

Used in: AI/ML model training, genome analysis, crypto mining, high-frequency trading, simulations, 3D animation, etc.

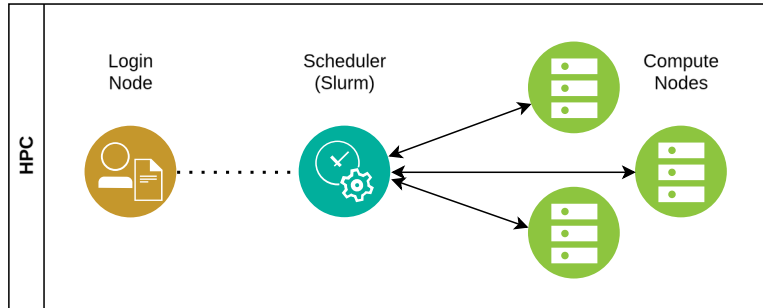
Why Use HPC for Your Research?



- **Fast computation:**
Hours instead of days
- **Parallel processing:**
Multiple tasks at once

What is the HPC System?

You



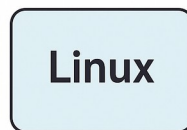
- **Cluster:** Group of servers (nodes)
- **Login Node:** Where you connect to the cluster
- **Scheduler:** Organises who runs what and when
- **Compute Nodes:** Where your jobs run

What is Linux?



- A group of open-source **Operating Systems** (OS) or **Distributions** (distro) based on the Linux **kernel**
- **Kernel?** The core of an OS. It lets your software talk to your computer's hardware.

Linux
Distributions



Windows

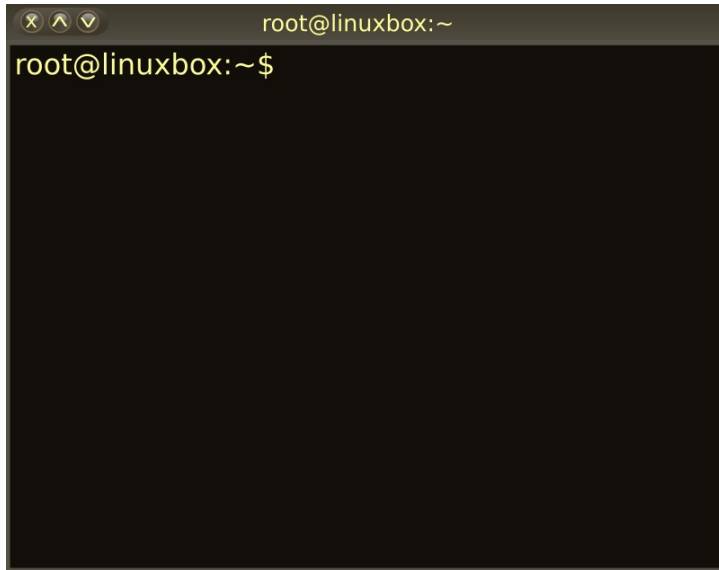


macOS



Basic Linux Commands

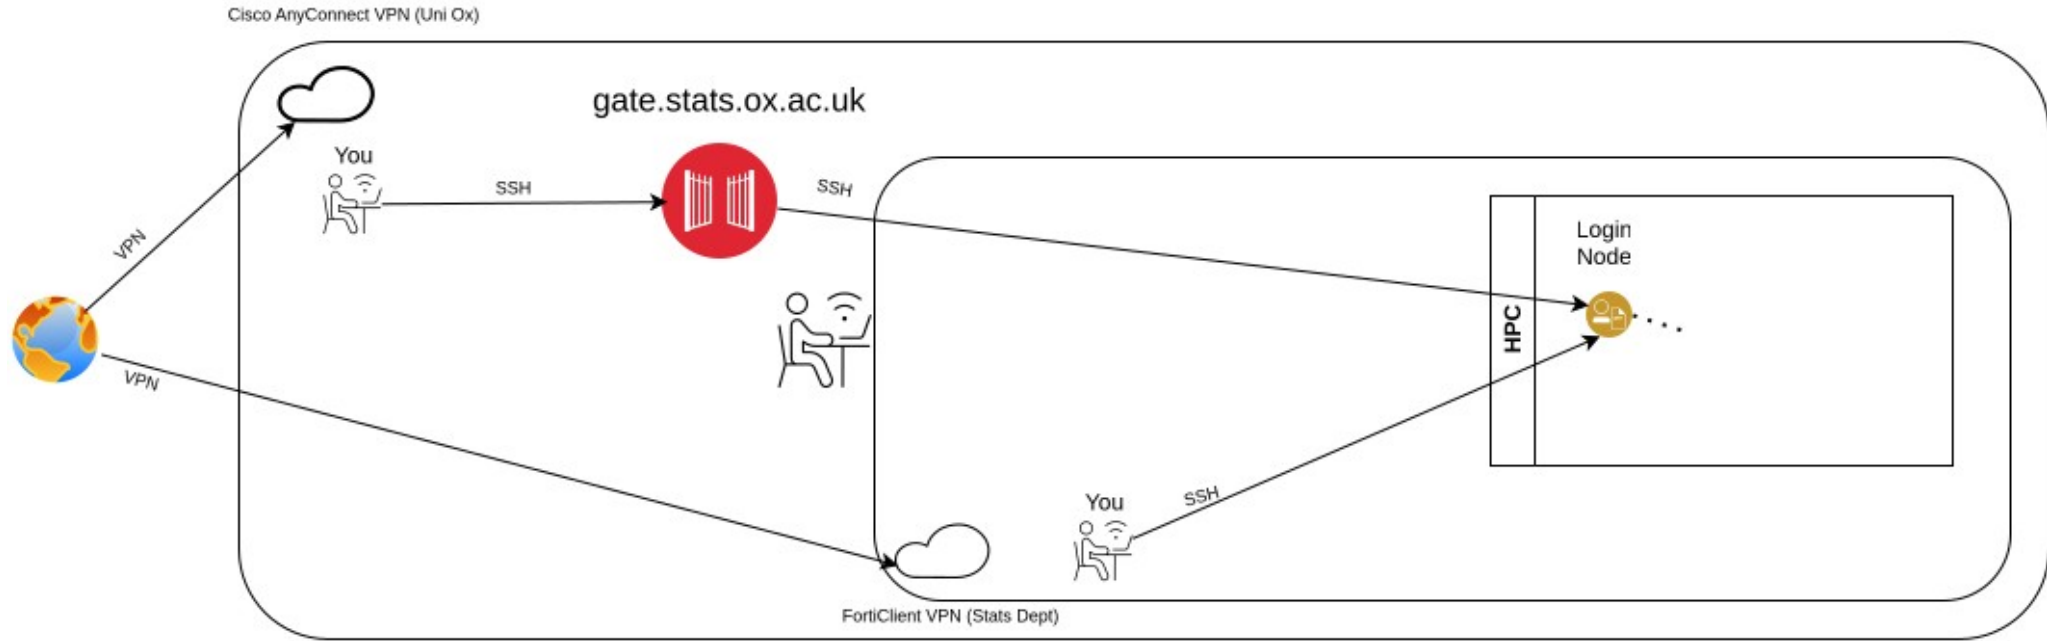
Linux Terminal



```
root@linuxbox:~$
```

- **man**: system reference manual
- **who**: show who is logged on
- **pwd**: print name of current/working directory
- **ls**: list directory contents
- **cd**: change directory ← *builtin, no man page!*
- **cp**: copy files and directories
- **mv**: move (rename) files and directories
- **rm**: remove files or directories ← **CAUTION!**
- **rsync**: copy remote files and directories
- **nano/vim**: Editors

Connecting to the Cluster



Connecting to the Cluster

- 1) Connect to the Cisco VPN
- 2) SSH to `gate.stats.ox.ac.uk` with your Stats username

- **SSH:**

```
$ ssh alice@gate.stats.ox.ac.uk
```

NOTE: Windows users, please use the **cmd** or **MobaXterm** for SSH (see the 'Moving files' page of this presentation)

Moving your files into the HPC

- **Linux/MacOS** (rsync or scp)

From your laptop's Terminal:

Make sure you're already connected to the Cisco VPN.

```
$ rsync my/Slurm/files  
<StatsUsername>@gate.stats.ox.ac.uk:
```

Example:

```
$ rsync myProjects alice@gate.stats.ox.ac.uk:
```

Moving your files into the HPC

- Windows:

MobaXterm
(Portable or
Installer)

How to use:

- 1)Download from:
<https://mobaxterm.mobatek.net/>
- 2) Choose either the portable version (no install) or
installer
- 3)Open a terminal tab and type
the same commands as in
the Linux example, eg rsync

Let's login!

1) From gate ssh to slurm-hn03

```
$ ssh <yourStatsUsername>@slurm-hn03
```

What is Slurm?



- **Job scheduler** that manages HPC resources
- You write a script, Slurm runs it when **resources** are free
- **Benefits:** Fair sharing, efficient usage

Submitting your first Slurm job

1) `ssh to slurm-hn03`

```
$ ssh <yourStatsUsername>@slurm-hn03
```

2) `cp -r /opt/hpcintro .`

Getting Help



- Internal documentation
- Ask your supervisor
- Ask Stats IT:
ithelp@stats.ox.ac.uk
- Mailing list?

Recommended Linux/HPC beginner guides

- **Linux:** The Unix Shell
<https://swcarpentry.github.io/shell-novice/>

(Software Carpentry <https://software-carpentry.org/about-us/>)
- **Slurm:** Quick Start User Guide
<https://slurm.schedmd.com/quickstart.html>
- **Checkpointing:** An example: PyTorch Lightning and SLURM
<https://mjenrungrot.com/blog/5e2c76d700cb445bb54e09c9b1ca43b9>

