

Last updated: 25 May 2026

Introduction to NSCC and High-Performance Computing

Prepared by: Dr. Nicholas CHENG
HPC Application Analyst (Frontier)



01

Introduction



Objectives

The **objectives** of this training module are four-fold.

1. To appreciate **the role** of NSCC and **the support** that NSCC is committed to provide to the user community and beyond.
2. To briefly **introduce HPC** and **why it is needed**.
3. To appreciate the differences between **a local workstation and a supercomputer**.
4. To outline the means of **acquiring access** to NSCC's HPC resources.

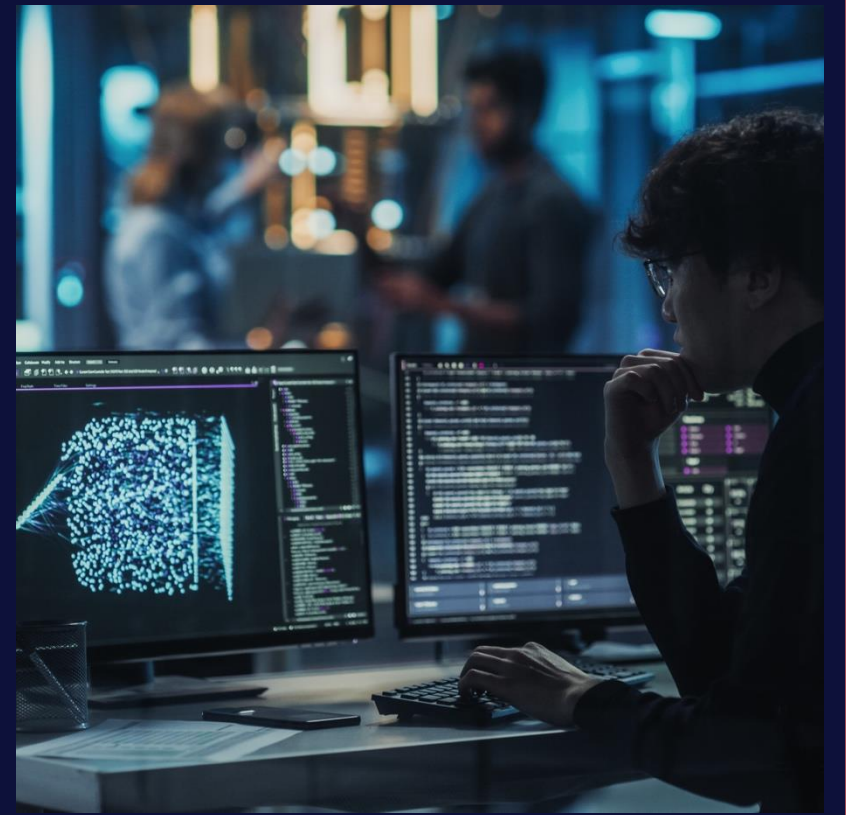
Ultimately, we hope that this training module helps you appreciate HPC. This will inform how you use the resources to achieve quality outcomes for Singapore and beyond.

Agenda

- What is NSCC? What is NSCC's commitment to the user community?
- Why HPC?
- What is the difference between a local workstation and a supercomputer?
- How do I access NSCC's HPC resources?

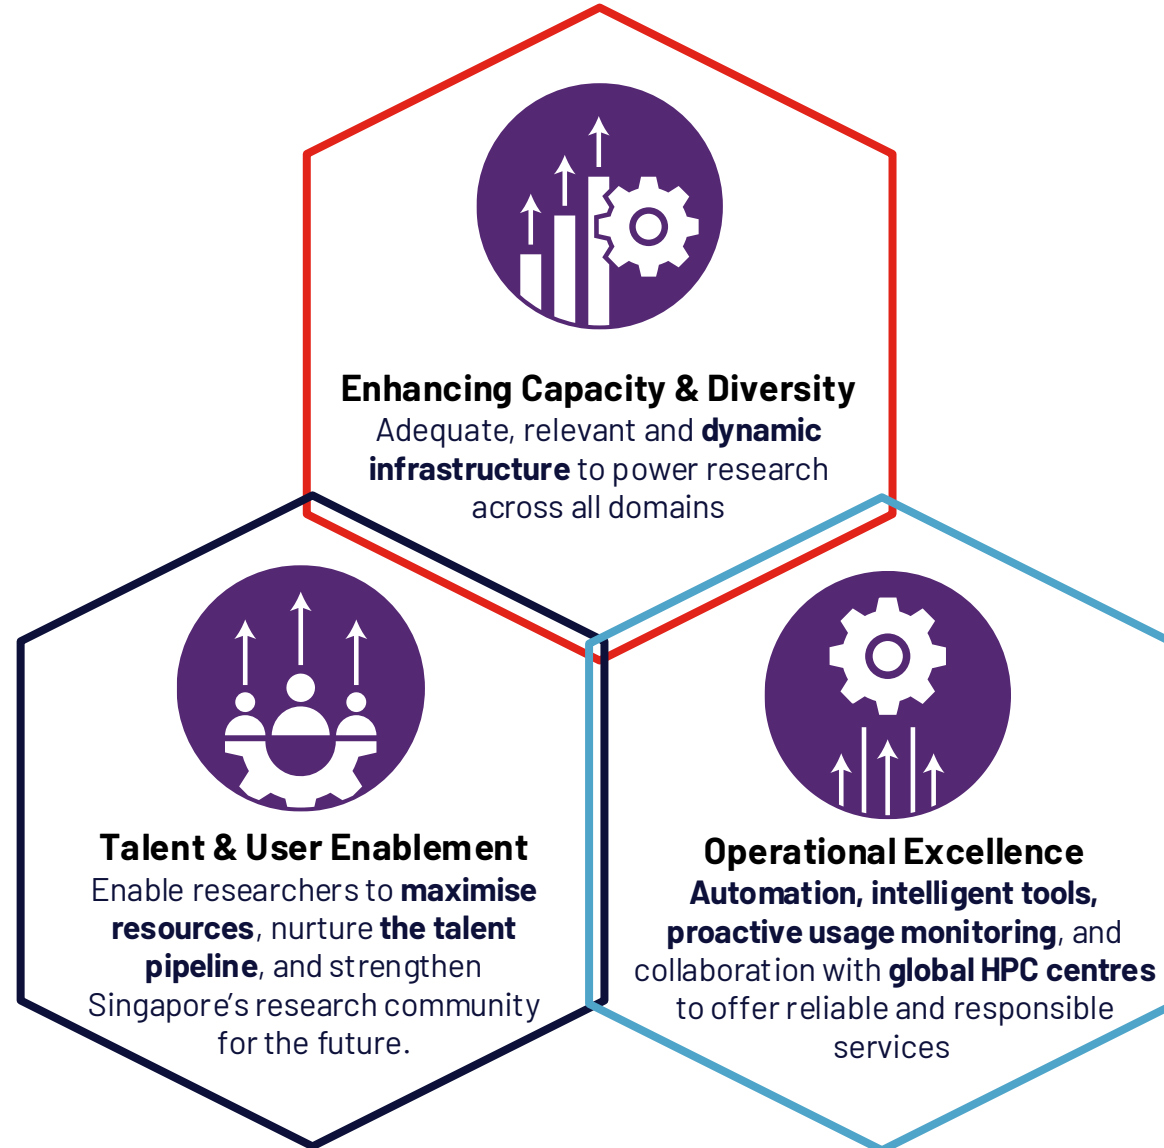
02

What is NSCC?



NSCC as a National Research Infrastructure

- Funded by National Research Foundation and hosted by A*STAR
- **Vision: To be Singapore’s go-to partner for accelerating transformative HPC-driven research.**



Advancing national HPC, AI, and quantum capabilities to support Singapore's research ecosystem

2016



ASPIRE 1
~1 PFLOPS (retired)

Jul 2023



ASPIRE 2A
~10 PFLOPS
100k AMD CPUs
352 Nvidia A100 GPUs

Sep 2024



ASPIRE 2A+
~20 PFLOPS
320 Nvidia H100 GPUs

Q2 2026



ASPIRE 2B
~120 PFLOPS
184K AMD CPU cores
1536 H200 GPUs


Aug/Sep 2023 Two on-premise supercomputers for sensitive healthcare and medical research



Partner HPC Centre Resources



AMD CoE



AWS Quantum Cloud

H2 2026



Quantinuum Helios Quantum Computer

Major workloads on NSCC's systems

ASPIRE 2A (CPU)

Ab initio **VASP**

CFD **OpenFOAM®**

MD **GROMACS** **LAMMPS**

Fusion science **GYSELA**

Weather **LFRic**

Quantum chemistry **CP2K**

Bioinformatics **Nextflow, SAMtools, Snakemake**

ASPIRE 2A (GPU)

AI

Ab initio **VASP**

MD **GROMACS** **LAMMPS**

ASPIRE 2A+ (GPU)

AI

MD **GROMACS**

Most workflows currently rely on CPUs, despite many applications supporting GPU acceleration. To efficiently utilise expanded capabilities of future HPC systems, NSCC is focused on helping users explore transitioning to GPU resources.

Our commitment to the user community: Training & Consultancy

HPC Training

For All Skill Levels: From bringing new users up to speed to helping veterans execute complex workflows.

Persona-driven: Targeted training for HPC software developers.

Expert Consultancy

Performance Engineering: Optimizing and scaling your code for maximum efficiency.

Science Enabling: Translating cutting-edge workflow needs into computational reality.

Proven Impact on NSCC Systems:


 **20x Speedup:** Achieved via job configuration advisory for a material science project, saving **24M CPU core hours**.

 **3x Faster Training:** Enabled large-scale, multi-node LLM training to drastically reduce iteration time.

Contact us if you're interested in advancing your projects further on NSCC's resources.


03

Why HPC?

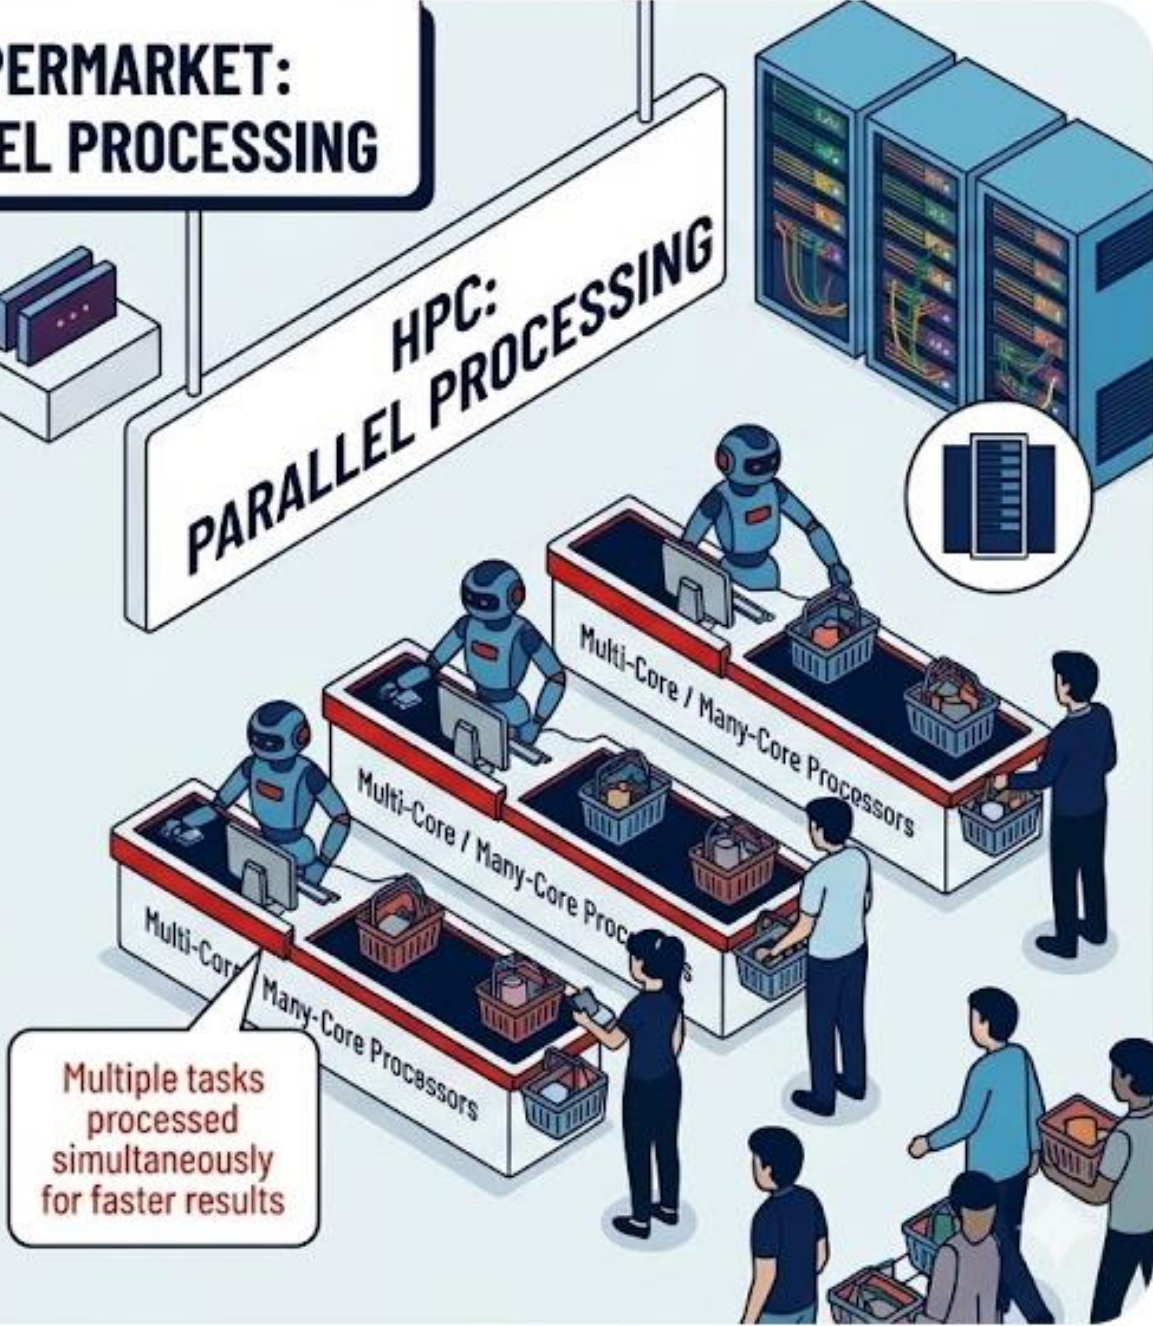


COMPUTING SUPERMARKET: SERIAL VS. PARALLEL PROCESSING

TRADITIONAL COMPUTING: SERIAL



HPC: PARALLEL PROCESSING



HPC: The promise and the reality

THE PROMISE: Unmatched Speed

Calculations that take **decades** on a standard laptop can be solved in **hours** on HPC infrastructure.

THE REALITY: It is NOT a Giant Laptop

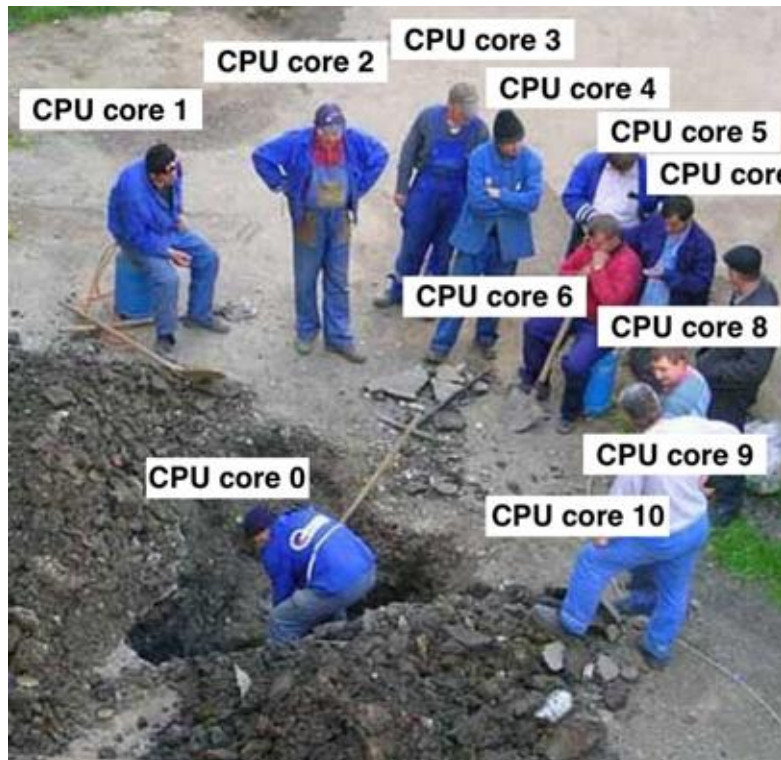
The Architecture Shift: We must understand how supercomputers work differently from laptops.

The Risk of Waste: If you try to use it like a personal laptop, you will bottleneck the system and burn through expensive resources.


(See next page for the cost of inefficiency)

Efficiency


When a user requests many CPU resources, but only one CPU core is doing work...



Interventions needed to ensure efficient use of HPC resources



Single GPU card running efficiently



Multiple GPU cards running inefficiently

Hundreds of thousands of dollars

Wastage due to inefficient 2-node GPU runs

04

What is the difference between a local workstation and a supercomputer?




ARCHITECTURAL DIFFERENCES: LAPTOP VS. SUPERCOMPUTER


(BEYOND COMPUTE POWER)

LAPTOP

CPU & GPU
(Limited Cores/Performance)



Internal, small-scale processors for single user tasks.



Slow Internal Motherboard Connections/Bus

Limited internal data paths for components within one device, not optimized for external clustering.
(Compare with High-Speed Interconnect)


Local HDD/SSD Storage
(Not Parallel Shared Storage)

Accessible only locally by the laptop's processors, relatively slower access.
(Compare with Parallel Shared Storage)

VS


SUPERCOMPUTER

Compute Nodes




Thousands of powerful CPUs & GPUs across distributed nodes for massive parallelism.

VS



High-Speed, Low-Latency Interconnect
(e.g., InfiniBand/HPE Slingshot)

Enables large-scale parallel processing by allowing rapid communication, synchronization, and data sharing between all compute nodes in the cluster with minimal delay.



Parallel Shared Storage
(e.g., Lustre/GPFS Parallel Storage)

Provides high-capacity and high-throughput concurrent data access across the high-speed network for all compute nodes simultaneously, supporting massive datasets for parallel applications.

Interconnects



Home broadband

~1 - 10 Gbps



HPC interconnects

~100 - 400 Gbps

Compared with home broadband (~10 Gbps), HPC interconnects have massive bandwidth (100 - 400 Gbps) but come with a hefty price tag.


Therefore, it is important to have **large workloads** running on HPC systems to effectively utilise interconnects!

What do you use this bandwidth for?

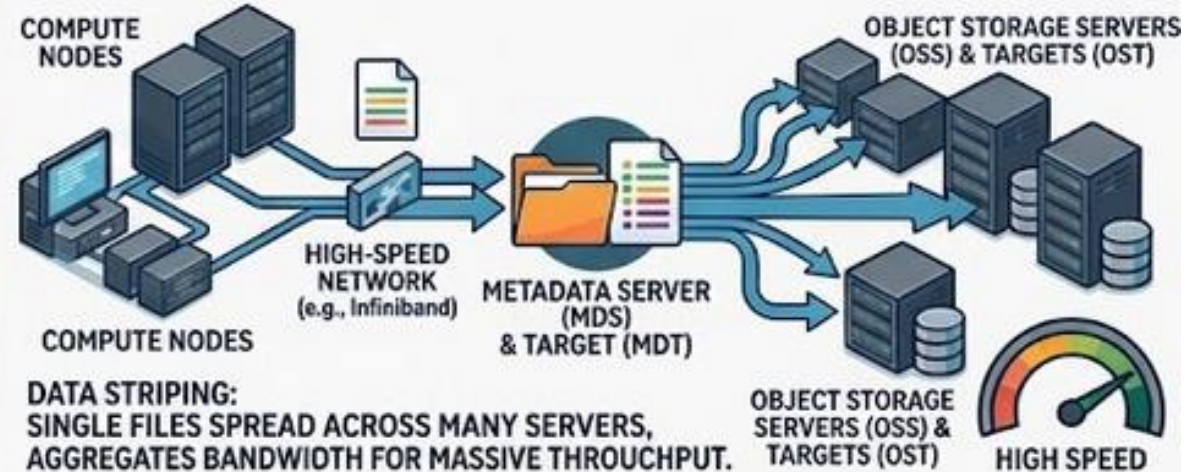
- Streaming massive amounts of data **between compute nodes**
- **Data transfer** from storage to memory
- **Hiding communication latency** to ensure continuous computation

UNDERSTANDING HPC STORAGE: PARALLEL FILE SYSTEMS VS. LOCAL DRIVES

LAPTOP HDD/SSD (SINGLE DRIVE)

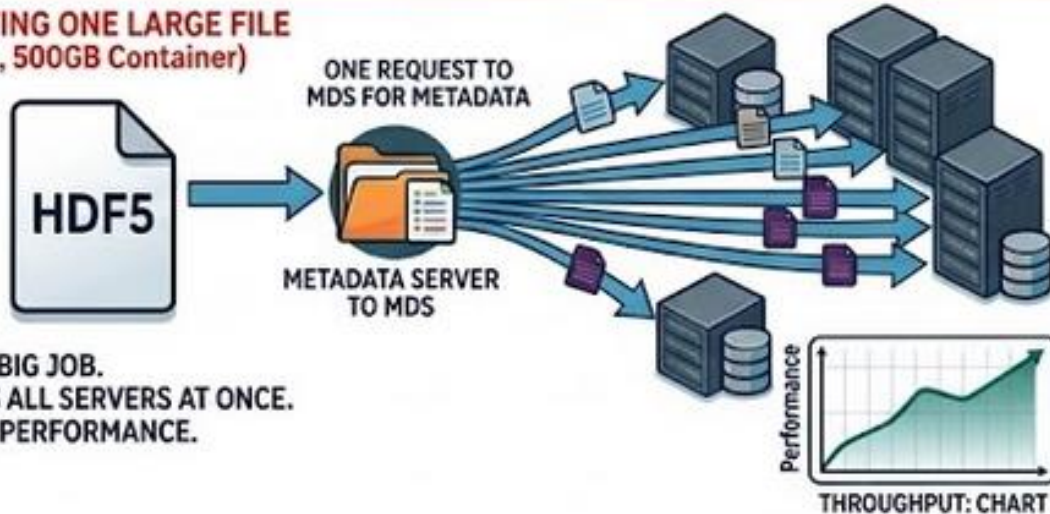


LUSTRE PARALLEL FILE SYSTEM (CLUSTER)

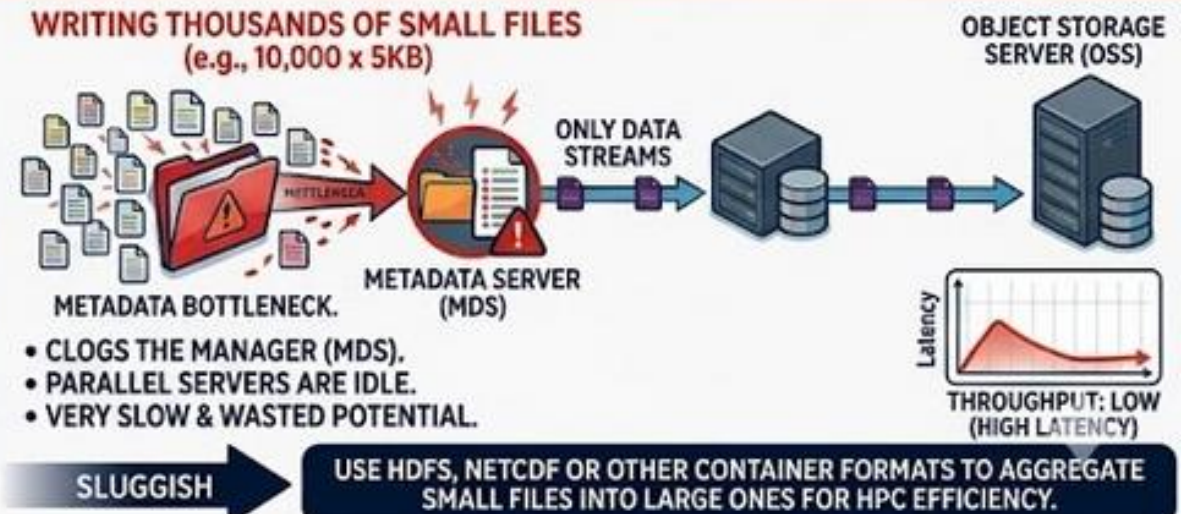


PERFORMANCE PITFALL: SMALL FILES VS. LARGE FILES ON LUSTRE

WRITING ONE LARGE FILE (e.g., 500GB Container)



WRITING THOUSANDS OF SMALL FILES (e.g., 10,000 x 5KB)





Storage: Parallel file systems

| Characteristic | Local Workstation (Laptop / PC) | Supercomputer (HPC Cluster) |
|---------------------------------|--|--|
| Scale & Architecture | <p>Gigabytes to Terabytes (1 User).</p> <p>Physically inside or physically attached to the machine.</p> | <p>Petabytes to Exabytes (100s of Users).</p> <p>Housed in data centre racks connected via high-speed networking.</p> <p>Hundreds of users and thousands of compute nodes access the exact same filesystems simultaneously.</p> |
| I/O Speeds | <p>Optimized for Low Latency.</p> <p>Designed to open applications and small files instantly. Peak throughput limited to physical limits of the drive (e.g. 3 - 7 GB/s for fast SSD).</p> | <p>Optimized for massive data transfer bandwidth.</p> <p>Designed to move large datasets in parallel across multiple servers with potential TB/s throughput. Individual operations (like opening a tiny file) are limited by network bandwidth.</p> |
| Impact of Poor Usage | <p>Isolated to you.</p> <p>Your machine may stall or crash, and only you are inconvenienced.</p> | <p>Everyone is affected.</p> <p>Poor I/O practices overload metadata servers or the network, stalling data access and causes lags for every user.</p> |



What is the difference between a local workstation and a supercomputer?

Effectively using HPC is not a trivial task! NSCC is dedicated to help you use HPC effectively.

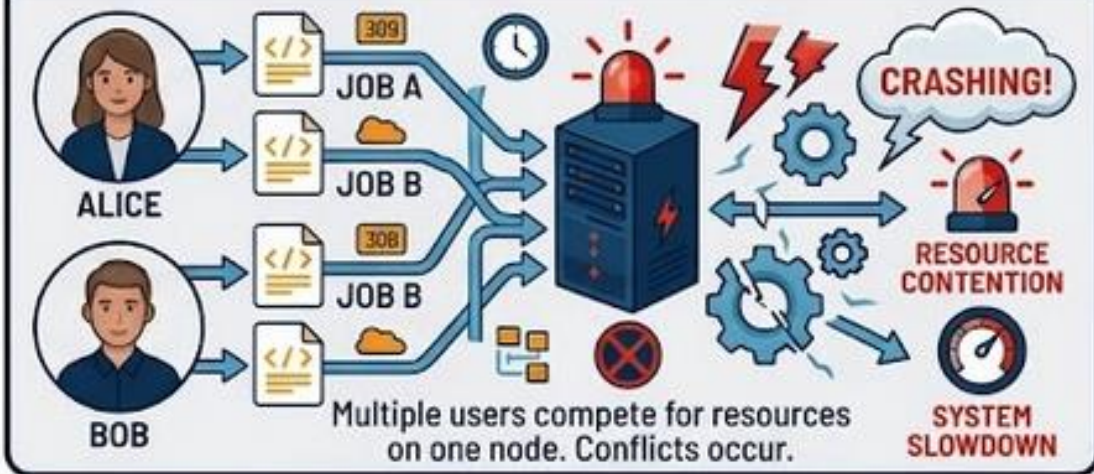


Compute: Appreciating the differences

| Component & Feature | Your Laptop | The Supercomputer | What it means for us |
|---|---|---|--|
| CPU | | | |
| Core Count | 4 to 16 cores (single chip). | 64 to 128+ cores per socket, spanning hundreds of nodes. | Supercomputers are not just more powerful laptops. Parallelism is key to unlocking the power of supercomputers. Familiarity with parallel programming frameworks helps build confidence. |
| Clock Speed | Higher base clocks (~3.0- 4.0 GHz) which can boost up to ~5.0 GHz for high performance. | Lower base clocks (2.0GHz - 3.0GHz) running continuously to manage heat and power consumption. | |
| GPU | | | |
| Scale & Cooling | 1 consumer GPU | 4 to 8 enterprise GPUs per node | Efficiently use the GPUs for compute. First step is ensuring the code is actually running on GPUs. Data movement is the enemy. High-speed links alleviate it and enable large-scale calculations! |
| Architecture | Optimised for graphics rendering. | Optimised for heavy compute. | |
| GPU Memory & inter-GPU links | GDDR memory; PCIe connection to the CPU. | High-Bandwidth Memory (HBM); NVLink/AMD Infinity Fabric for high-speed inter-GPU communication. | |

HPC JOB SCHEDULER: ORGANIZING THE COMPUTE CLUSTER


1. THE PROBLEM: CLUSTER WITHOUT A SCHEDULER



2. THE SOLUTION: INTRODUCING THE JOB SCHEDULER



3. THE ANALOGY: RESTAURANT QUEUING SYSTEM

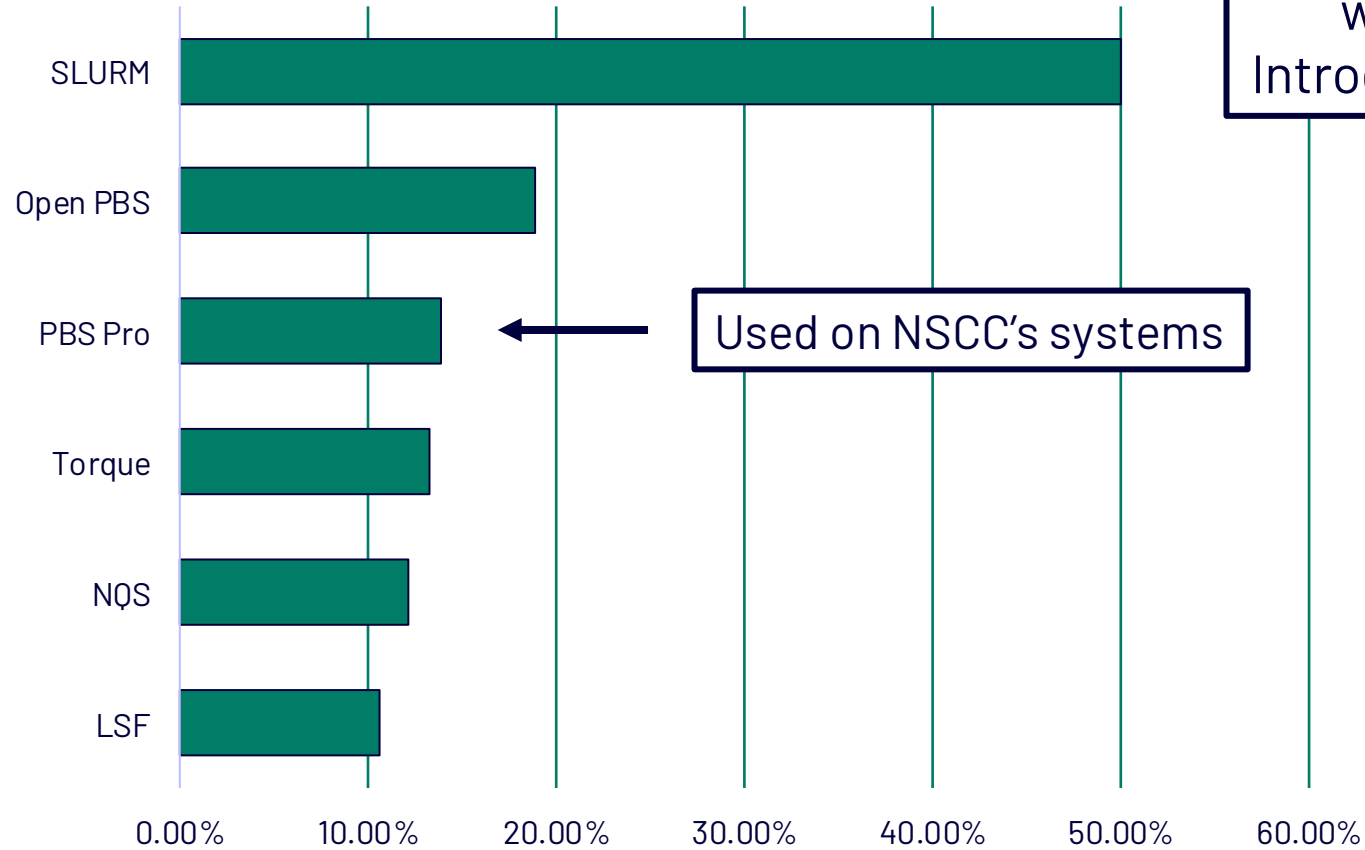


ORDERLY WORKFLOW → SUCCESSFUL COMPUTATION ←

A shared environment: Job scheduling on HPC



Percentage share of popular job schedulers



Usage of job schedulers will be covered in Introductory Workshops.

Used on NSCC's systems

A shared environment: Login nodes and compute nodes

Login nodes

Access points for users to interact with the HPC cluster, intended for file management and job preparation

Not intended for heavy computations!



User Computer



Internet



Login Nodes



Login Nodes



ASPIRE 2A



ASPIRE 2A+

Compute nodes

Dedicated servers within the HPC cluster that perform the actual computations





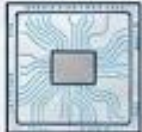




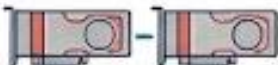

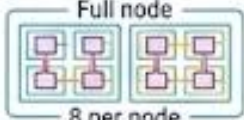

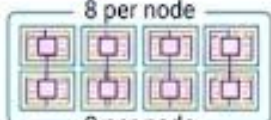




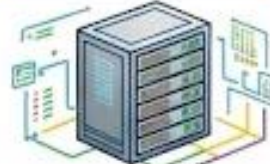
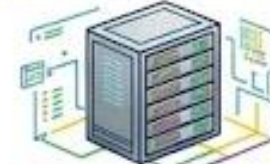
Accessed by submitting jobs to the job scheduler

05

How to access NSCC's resources?



NSCC SINGAPORE SUPERCOMPUTING SYSTEMS: TECH SPECS AT A GLANCE

| | ASPIRE 2A | ASPIRE 2A+ | ASPIRE 2B |
|----------------|--|--|--|
| VENDOR |  HPE Cray |  NVIDIA DGX |  HPE Cray |
| CPU |  AMD EPYC 3rd gen "Milan"  ~100k CPU Cores |  Intel Xeon Platinum 8480C 112 CPU Cores per Node |  AMD EPYC 5th gen "Turin"  ~150k CPU Cores |
| GPU |   4 or 8 per node 352 x NVIDIA A100 (40GB GPU Memory) |   Full node 8 per node 320 x NVIDIA H100 (80GB GPU Memory) |   8 per node 8 per node 1,536 x NVIDIA H200 (141GB GPU Memory) |
| INTER-CONNECT |  HPE Slingshot 100 Gb/s |  InfiniBand 400 Gb/s |  HPE Slingshot 400 Gb/s |
| SHARED STORAGE |  Lustre, GPFS |  Lustre |  Lustre |

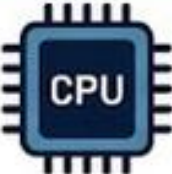





Resource allocation



Resource allocation follows the [Strategic Resource Allocation Policy](#) to manage HPC as a national strategic R&D resource.

| Categories | Overseeing Body | Key Objectives |
|--|-----------------------------------|--|
| RIE-Funded Projects | SRAC | Supports RIE-funded projects of national priorities and complexity |
| Tier 2 (Institutional) Include YISP | Respective IHLs | Support institutional R&D needs. |
| HPC Research | NSCC | Promote and invest in HPC R&D to enhance HPC capabilities - tools, models, GPU/CPU optimisation. |
| Training and Promotion | NSCC | Increase the pool of HPC talents through competency development projects. |
| SME/Startups | Committee to be led by IMDA & ESG | Improve the competitiveness of SMEs/Startups through R&D. |
| Non-RIE Funded Projects | NSCC | Support research projects that are not funded by RIE. |

Service Units (SU)

 
1 CPU core-hour = **1 SU**

 
1 A100 GPU card-hour = **64 SU**

 
1 H100 GPU card-hour = **128 SU**



Allocation



Project Allocation
(Total SUs Granted)

Maximum resources (CPU, GPU, Storage) a project can consume.



Scheduler



Limit Enforced

Once exhausted, no new job submissions will run.

Indicative values, subject to adjustment.

Next steps

- **Applying for a project?** [Project Calls](#)
 - Applications for RIE-funded projects are accepted during the bi-annual **Call for Projects** windows.
- **Need an account?** Eligible users may register via the [NSCC User Portal](#).
- **Getting started with HPC?**
 - Attend the **Introductory Workshops** to learn how to navigate NSCC's systems effectively
 - **QuickStart guides** for [ASPIRE 2A](#) and [ASPIRE 2A+](#)
- **Further help and training resources** can be found at [NSCC User Guide](#)

Let us have your feedback on the training offerings!

Thank you.

Email: contact@nsc.sg
Website: www.nsc.sg

