Internships with NSCC

PROJECT GALILEO

Proactive Performance Observability for Supercomputing Systems

How To Apply

Send your CV, academic transcript and a brief introduction to **recruit@nscc.sg** with the subject line: **Internship Application - [Your Name] [Project Title]**

Please include your intended internship period within the email. Our team will be in touch with shortlisted candidates.

The Challenge

The National Supercomputing Centre (NSCC) Singapore manages Singapore's national high performance computing (HPC) resources, a critical asset that drives the advancement of cutting-edge research, technology, and innovation. Maintaining peak performance in such an elite supercomputing environment is paramount.

Subtle "performance drift"—caused by hardware degradation, software updates, or routine maintenance—can secretly degrade throughput, elongating user runtimes and causing costly job failures.

We are building a Proof of Concept (POC) for an automated observability pipeline to run regular benchmarks and detect this drift early, separating true system degradation from normal cluster noise.

What You Will Do

You will take ownership of designing and validating a modular component of this POC based on your skill set:

- **Systems & Data Engineering:** Evaluate benchmarking frameworks (e.g. ReFrame, JUBE), configure proxy applications that mirror our most demanding user workloads, and design the pipeline to securely ingest and manage this output data.
- **Data Analytics:** Build statistical models to filter out cluster noise, establish reliable performance baselines, and trigger automated early warnings.
- **Dashboarding:** Develop a prototype frontend and backend integration to visualise these critical system insights.

What You Will Deliver

A functional POC for your specific module and a technical evaluation of its scalability and feasibility. Your prototype and findings will provide the foundational data to shape our long-term supercomputing monitoring strategy.

Required Qualifications

- Currently pursuing a degree in Computer Science, Computer Engineering, Data Science, or a related field
- Proficiency in scripting and programming languages, particularly **Python** and **Bash**
- Comfortable navigating and working within a **Linux/Unix command-line** environment
- Fundamental knowledge of data manipulation and building basic statistical models

Bonus / Preferred Skills

- Familiarity with systems monitoring or data visualisation tools (e.g. Grafana, Prometheus)
- Prior exposure to benchmarking frameworks (like ReFrame or JUBE) or general software testing
- A foundational understanding of HPC architecture