

ISSUE 1 | OCTOBER 2026

SUPERCOMPUTING ASIA

COMPUTING
THE FUTURE,
FROM ASIA TO
THE WORLD

MERaLiON: BUILDING
EMPATHETIC AI FOR
A MULTILINGUAL REGION

CHARTING THE NEXT HORIZON:
A CONVERSATION WITH
NSCC'S EXECUTIVE DIRECTOR

THE NEXT GENERATION
OF SUPERCOMPUTING
TALENT TAKES SHAPE

PUTTING THE "SUPER"
IN "SUPERCOMPUTING"

COVER STORY

PEERING INTO A DECADE
OF SUPERCOMPUTING IN ASIA

04

COVER STORY

PEERING INTO A DECADE OF
SUPERCOMPUTING IN ASIA

12

CASE STUDY

MERaLiON: Building Empathetic AI
for A Multilingual Region

18

SPECIAL SECTION

Charting the Next Horizon:
A Conversation with
NSCC's Executive Director

26

BEHIND THE BRILLIANCE

The Next Generation of
Supercomputing Talent Takes Shape

31

HPC 101

Putting the "Super"
in "Supercomputing"

EDITOR'S NOTE

A decade ago, Singapore commissioned its first supercomputer. Since then, much has changed — not only in the scale of compute, but in the ambitions now attached to it. This issue of *Supercomputing Asia* arrives at a timely moment. Under the theme *Computing the Future, from Asia to the World*, it considers how far the region has come, and what its growing computational power is ultimately in service of. Performance, scale and speed matter, but so do the human decisions, institutional commitments and cross-border partnerships that determine whether capability becomes contribution, and progress becomes impact.

Our cover story traces a decade of transformation across the Asia-Pacific, from isolated national systems to interconnected scientific powerhouses. Our case study on Multimodal Empathetic Reasoning and Learning in One Network (MERaLiON) shows how the National Supercomputing Centre (NSCC) Singapore's supercomputing infrastructure enabled the training of Southeast Asia's (SEA) first empathetic, multilingual artificial intelligence (AI) model on over a million hours of regional speech. In conversation, NSCC's Executive Director Dr Terence Hung looks at the decade ahead, from quantum-classical hybrid infrastructure and energy efficiency, to the talent pipelines that will determine whether capability translates into lasting impact.

Elsewhere, we meet the student teams who spent 48 sleepless hours crafting and perfecting supercomputing clusters at Supercomputing 2025 (SC25) — a reminder that the future of this field rests on the resilience and ingenuity of those coming into it. For readers newer to the space, our high performance computing 101 (HPC 101) explainer connects the computing foundations to the applications shaping everyday life.

At the end of the magazine, you will find a feedback form. We read every response, and your input helps inform what comes next.

I hope you enjoy this issue of *Supercomputing Asia*.

Dr Julian Tang

Founder and Managing Editor,
Ernest & Kin

PEERING INTO A DECADE OF SUPERCOMPUTING IN ASIA

*From national systems to regional networks,
the pace of progress is redefining science,
industry and the future of AI-enabled innovation.*

From the cooled racks of Singapore's data centres to the sprawling facilities of Japan and China, and the research-driven machines of Australia, Asia-Pacific has been transformed into a supercomputing powerhouse in recent years. Once the preserve of specialist laboratories, HPC has become the engine of growth powering AI, climate forecasts, biomedical research, energy exploration, quantum simulations and more across the continent.

A decade ago, Asia's HPC landscape was uneven. Japan's K computer and China's many petascale systems dominated the scene, while South Korea and Australia gradually expanded research capacity.

In 2016, Singapore had just commissioned ASPIRE 1 — the city-state's first national petascale supercomputer.

Ten years on, transformations abound. Asia is now one of the most dynamic HPC regions globally. Hyperscalers are investing billions in new data centres. Compute has moved from the periphery of research planning to the centre of national digital and scientific strategy. Global HPC spending, estimated at around US\$46.2 billion in 2023¹, continues to rise, with Asia accounting for a big share of the pie as governments and institutions recognise compute as foundational infrastructure.


¹Hyperion Research, HPC Market Update 2023–2024.

A REGION SUPERCHARGED BY GROWTH

The true hallmark of Asia’s progress is how deeply compute has been embedded into research ecosystems. China now hosts roughly 6% of the world’s TOP500 systems², with the country’s LineShine supercomputer taking the top



Setonix, housed at the Pawsey Supercomputing Research Centre in Perth, Australia, is the country’s most powerful public research supercomputer. Photo: Pawsey Supercomputing Research Centre



The Fugaku supercomputer at the RIKEN Center for Computational Science in Kobe in June 2020. Photo: RIKEN

position in the June 2026 TOP500 ranking.

Japan’s Fugaku remains among the most scientifically productive systems ever built³, supporting breakthroughs in pharmaceuticals, materials engineering

and atmospheric science. Korea’s development of KISTI-6 signals a shift toward national-scale, AI-driven research, while Australia’s National Computational Infrastructure and Pawsey Supercomputing Centre underpin global astronomy and geoscience missions.

In Singapore, ASPIRE 2A and the accelerator-rich ASPIRE 2A+, support

AI-for-science, materials discovery, genomics and high-resolution climate modelling. In the Third National Climate Change Study (V3), researchers ran two-kilometre atmospheric simulations that demanded 10–100 times more compute than earlier models, producing projections that now shape planning for heat stress, flooding and sea-level rise. Similar advances are unfolding across the region, from multi-year typhoon and monsoon studies in Japan and Korea to continent-wide environmental modelling in Australia.

Compared to the United States and Europe, Asia shows a steady growth in supercomputing capabilities over the years. The US leads in exascale systems such as Frontier and Aurora. Europe has built world-class coordination through the Partnership for Advanced Computing in Europe (PRACE) and EuroHPC, supported by strong community software and research networks.

Asia complements these strengths with rapid deployment, strong public investment and a willingness to design systems around hybrid HPC-AI workloads. No longer learning quietly from the sidelines, the region is increasingly shaping global HPC practice.

²TOP500 List, June 2026
³RIKEN R-CCS — Fugaku Highlights & Scientific Output Reports; HPCWire — Fugaku Continues to Lead in HPCG, HPL-AI, and Real-World Workloads.

AI, COLLABORATION
AND THE NEW SHAPE OF
SCIENTIFIC COMPUTE

If one force defines the past decade, it is the rise of AI. Large language models and generative systems have reshaped global compute economics, with some frontier models consuming tens of millions of Graphics Processing Unit (GPU) hours⁴. In SEA alone, AI-ready compute demand is projected to grow five- to eight-fold by 2030⁵. Across the region, HPC centres report that AI workloads are outpacing their refresh cycles, and researchers now expect simulation, machine learning, inference and data assimilation to co-exist within workflows.

Taiwan’s Nano 4, Korea’s KISTI-6 and accelerator-heavy clusters across China reflect a convergence of HPC and AI. AI has not displaced HPC; it has reinforced it. Training, validating and governing AI at scale depends fundamentally on the HPC infrastructure Asia has spent a decade building.

Alongside this technological shift, collaboration has deepened. The Korea-ASEAN Digital Innovation Flagship Project is one of the efforts to establish shared



China’s Sunway TaihuLight was the world’s fastest supercomputer for two years, from June 2016 to June 2018, according to the TOP500 lists. Photo: Xinhua

HPC capability in SEA, deploying a system in Indonesia and training 160 researchers. The Supercomputing Asia conference, co-organised by centres across Japan, Australia, Thailand and Singapore, has become a platform for regional alignment, joint technical development and cross-border scientific exchange.

Talent, however, remains a constraint. While investment in infrastructure has

accelerated, the pipeline of experienced system engineers, research software engineers and application specialists has not kept pace. Long training cycles and still-maturing HPC career pathways continue to hamstring capacity.


In Singapore, universities have begun addressing this gap upstream. Both the National University of Singapore (NUS) and the Nanyang Technological

University (NTU) have launched dedicated scholarship programmes, including the NUS AI Talent Scholarship and NTU’s Turing AI Scholars Programme, to develop home-grown talent equipped to lead in an AI-driven world. Such initiatives help build depth in computational expertise early, before students embark on specialised research or industry roles.

That foundation feeds into more targeted HPC training. Regional initiatives such as the Association for Computing Machinery (ACM) Asian HPC School, alongside national HPC training programmes, Student Cluster Competitions and cross-centre optimisation workshops, provide hands-on exposure to production systems. However, these efforts remain modest relative to the scale of demand unleashed by AI, underscoring the need for sustained, long-term investment in HPC manpower development.

Asia must also confront the environmental realities of large-scale compute. Modern supercomputers can draw 20–30 megawatts of power, and in tropical climates, cooling alone can account for up to 40% of data-centre energy use⁶. With Asia’s data-centre electricity

⁴Stanford HAI – AI Index Report 2024; Epoch AI — Compute Trends Across Large Models.
⁵IDC & ASEAN Digital Masterplan 2025 (ADM2025) analytics; Google-Temasek-Bain e-Economy SEA projections on AI infrastructure growth.
⁶Uptime Institute — Global Data Center Survey



demand expected to roughly double by 2030, driven significantly by AI and HPC workloads, sustainability has become a design imperative.

Responses are emerging. Singapore’s Sustainable Tropical Data Centre Testbed is pioneering tropical-optimised cooling approaches; Japan is prioritising energy efficiency in its next-generation architectures; and China is siting new compute campuses in hydropower-rich western provinces.

THE DECADE AHEAD

Asia’s HPC scene is no longer playing catch up. A region that once built systems in parallel is now converging to exchange expertise, talent and, increasingly, compute itself. It is converting capability into global influence.

The next decade could see the advent of shared climate and ocean models, distributed AI training tuned to Asia’s languages and environments, hybrid quantum-HPC workflows, tropical-optimised data-centre designs, and many more. Crucially, it will test whether the region can turn momentum into lasting scientific, economic and societal advantage.



MERaLiON: BUILDING EMPATHETIC AI FOR A MULTILINGUAL REGION

A look at how NSCC's supercomputing power enabled Singapore to train a multimodal model on millions of hours of regional speech, laying the groundwork for AI built around SEA's linguistic complexity.

Singapore's linguistic landscape is more layered than it appears. Conversations toggle easily between English, Mandarin, Malay, Tamil, Singlish and various dialects, often in the same breath. Across SEA, this complexity deepens. Speech carries local rhythms and accents; meaning is shaped

by tone and context; emotion could be implicit rather than explicit.

Most language learning models (LLMs) are not built for this convolution. Trained predominantly on Western datasets, they interpret speech as text alone,

overlooking the tonal and cultural cues that influence how people across the region communicate. A new class of AI was needed — one that could understand people as they sound and speak in daily life.

That became the impetus for Multimodal Empathetic Reasoning and Learning in One Network (MERaLiON), developed by the A*STAR Institute for Infocomm Research (A*STAR I²R) under the National Multimodal LLM Programme (NMLP). It is an AI model that can listen, reason and respond with sensitivity to language, accent, tone and emotion across SEA. Achieving this required access to computing power far beyond conventional research infrastructure.

TRAINING EMPATHETIC AI AT A SCALE NEVER ATTEMPTED BEFORE

Developing MERaLiON meant training models on an immense breadth of speech and linguistic patterns. The team worked with high volumes of audio spanning English, Chinese, Malay, Tamil, Singlish and a growing list of regional languages, including Bahasa Indonesia, Thai and Vietnamese. Beyond vocabulary, the model needed to recognise intonation, pitch, emotional registers and code-switching.

Capturing these nuances required building unified multimodal architectures and merging components from Whisper, Gemma-SEA-LION and Gemma. It also



Mrs Josephine Teo, Minister for Digital Development and Information, unveiled MERaLiON Version 2.0 and launched the MERaLiON Consortium at the ATxSummit 2025.

meant pre-training a foundational speech encoder on more than one million hours of multilingual audio.

Such a model required serious computing muscle. Datasets alone reached 300–600 terabytes. Training had to run in distributed environments, with iteration cycles measured in days rather than months. Without HPC, MERaLiON would have remained as just a vision.

HOW NSCC ACCELERATED MERaLiON's DEVELOPMENT

NSCC's ASPIRE 2A+ supercomputer provided the computational backbone that enabled MERaLiON to take shape at scale.

High-speed distributed training reduced processes that normally stretch across months to as little as 14 hours to six days, depending on the model. This shortened experimentation cycles, enabling the team to test configurations and fine-tune architectures at a pace beyond the reach of traditional systems.

NSCC's storage capacity supported continuous ingestion of multilingual datasets, while dedicated data pipelines enabled the processing more than 600,000 hours of audio without bottlenecks.



Minister Josephine Teo (pictured in photo above) visited the MERaLiON booth at ATxSummit. Initially equipped with English, Mandarin and Singlish, MERaLiON Version 2.0 includes Malay, Vietnamese, Thai, Tamil, Bahasa Indonesia, as well as Chinese dialects for future releases.



GPU infrastructure was another critical factor. To date, the project has accumulated more than 2 million GPU-card hours of compute time — measured similarly to how electricity usage is tracked in kilowatt-hours. This included MERaLiON-AudioLLM, trained on 120,000 to 260,000 hours of speech, and MERaLiON-SpeechEncoder-v2, trained on 1.4 million hours of multilingual audio. Capable of interpreting speech as an expression — with underlying nuances from words, tone, rhythm and affect expertly captured — a multilingual multimodal model was born.

In parallel, NSCC's Frontier team worked closely with the MERaLiON researchers to optimise workloads for scale. Since late 2024, the team has provided hands-on guidance on job configuration, sizing and performance tuning, including profiling to identify bottlenecks and improve resource utilisation. Early challenges in launching large, distributed jobs were addressed through system-level refinements, enabling larger workloads and nearly doubling compute efficiency. This collaboration continues as MERaLiON prepares for even more demanding training workloads in the future.



*At the HPC Users Symposium held on 25 September 2025, Dr Jeremy Wong, Senior Scientist at the A*STAR Institute for Infocomm Research (A*STAR I²R), shared the development progress of MERaLiON-SpeechEncoder-2 — a general-purpose speech encoder tailored to Singaporean and Southeast Asian speech. The Symposium, organised by NSCC, provides a platform for local HPC users to come together to present their research findings, share new ideas and discuss the future of HPC, while spotlighting the impact of HPC and the benefit that it brings to Singapore.*

“MERaLiON is a key step in advancing Singapore’s AI research, especially in understanding language nuances,” says Aw Ai Ti, Principal Investigator and Head of the Aural & Language Intelligence department at A*STAR I²R. “The support from NSCC’s computational resources has been essential in accelerating our research and bringing us closer to creating AI solutions tailored to the region’s needs.”

UNLOCKING NEW POSSIBILITIES FOR SINGAPORE AND THE REGION

With NSCC’s computing resources spurring its development, MERaLiON has become one of Singapore’s most significant AI contributions to SEA.

The savvy AI linguist enables new forms of customer interaction, especially for small and medium enterprises (SMEs)

that serve multilingual communities and older adults. In healthcare and social services, MERaLiON can detect signs of distress or fatigue to support earlier intervention. Public service and education benefit from real-time transcription, summarisation and analysis across diverse linguistic settings. The model also supports strategic decision-making by enabling sentiment analysis, anomaly detection, trend identification and technology foresight for public agencies and enterprises.

Encouragingly, adoption has been swift. MERaLiON models and datasets have recorded more than half a million downloads. The MERaLiON Consortium — comprising banks, media organisations, transport operators, healthcare agencies, as well as technology and engineering

firms, with NSCC as a founding member — is now developing practical, regionally tuned AI applications.

Work is also underway with the Monetary Authority of Singapore and financial institutions to create a variant of MERaLiON customised for automatic speech recognition across Singlish and commonly spoken local languages and dialects.

As its second version extends support to eight Southeast Asian languages and deepens its emotion-aware capabilities, MERaLiON shows how AI can be shaped by regional context rather than global defaults. At the core of this progress is NSCC’s supercomputing infrastructure, which enabled SEA’s first multimodal empathetic AI model to be built at scale — and to grow with the region it serves.



CHARTING THE NEXT HORIZON: A CONVERSATION WITH NSCC'S EXECUTIVE DIRECTOR

Dr Terence Hung reflects on how the changing nature of scientific research is reshaping the role of national advanced computing — and NSCC's ambitions for the decade ahead.

As NSCC marks its 10th anniversary, its Executive Director, Dr Terence Hung, is looking ahead to what the next decade of advanced computing will demand. The question, he says, is not simply how much computing power Singapore will need, but how advanced computing must evolve to support a rapidly changing research landscape.

Much has changed since NSCC was established. Today, modelling and simulation remain important to scientific research, but the emergence of machine learning, generative AI and data-intensive approaches has broadened the range of computational methods available to researchers.

"The future of research is no longer just theory, experiment and computation," says Dr Hung. "AI is rapidly becoming a native element of research. At the same

time, scientific workflows are becoming more interconnected, data-intensive and increasingly dependent on how effectively researchers can combine simulation, AI, data analytics and experimentation across multiple computational environments."


As AI becomes more deeply embedded across the scientific workflow, its role is also evolving. Researchers are exploring agentic AI — intelligent software capable of planning a sequence of tasks and carry out parts of the discovery process autonomously. Researchers are exploring how such systems could coordinate cycles of proposing hypotheses, running simulations, analysing results and organising experimental validation, all of which opens new avenues for scientific discovery.



FROM CYCLES TO CAPABILITY

In NSCC’s earlier years, most workloads centred on traditional HPC use cases such as climate modelling, engineering and materials science. The priority was ensuring researchers had access to the computing resources needed to run large-scale models and simulation.

Today, researchers increasingly draw on different architectures to suit research requirements — from central processing units (CPUs) and GPUs to AI accelerators and quantum systems. Initiatives such as Hybrid Quantum-Classical Computing (HQCC 1.0) are already preparing Singapore’s research community to work across these technologies.



ASPIRE 2B combines advanced CPU and GPU capabilities, delivering up to 115 petaFLOPS of compute power to tackle larger and more complex research challenges. Its expanded capabilities support integrated simulation, data analytics and AI workflows, accelerating research in areas such as climate science, healthcare, materials discovery and next-generation AI. ASPIRE 2B is also expected to be integrated with Singapore’s incoming quantum computer, enabling researchers to explore emerging hybrid quantum-classical computing approaches.



Guest-of-Honour and Minister for Digital Development and Information Mrs Josephine Teo announced the new Hybrid Quantum Classical Computing (HQCC 1.0) initiative in her speech at the opening event of the Supercomputing Asia 2025 conference on 11 March 2025.

“Providing compute cycles is still essential, but it is no longer enough on its own,” he says.

“HPC systems are becoming more powerful, but the challenge is how we ensure they are used effectively. We need to enable researchers to harness these growingly sophisticated technologies.”

The deployment of ASPIRE 2B, Singapore’s latest national supercomputer, reflects this shift. Building on ASPIRE 2A and ASPIRE 2A+, it significantly strengthens

Singapore’s national advanced computing capabilities. Over the past decade, NSCC’s compute capacity has grown 100-fold to support complex research.

Alongside this growth in infrastructure, NSCC has strengthened the support and make effective use of it. Through its Frontier team, the Centre’s HPC specialists work closely with researchers to optimise workloads and applications. In addition, introductory workshops and tutorials help users build their HPC capabilities and navigate NSCC’s systems

more effectively, while Communities of Practice bring domain scientists together to deepen their expertise and share best practices.

That focus extends to building Singapore's future HPC talent pipeline. Through outreach, internships and support for student teams competing in international cluster competitions, NSCC provides opportunities for hands-on experience. These efforts help to build the interest and skills Singapore will need for its future advanced computing workforce.

DEMOCRATISING ACCESS

As HPC and AI become progressively important to researchers from traditionally non-computing domains, the technologies they need to navigate are also becoming more complex. Making computing resources more accessible and easier to use is central to what Dr Hung describes as “democratising HPC.”

“Our vision is for HPC to become a utility, much like electricity and water — something researchers can access when they need it without having to understand all the complexity underneath,” he says.

For researchers, that would mean intuitive interfaces, reusable workflows, models and tools that provide access to the capabilities



1 NSCC conducts regular workshops to familiarise researchers and industry users with HPC tools, best practices and emerging technologies, helping them translate advanced computing capabilities into impactful outcomes.

2 NSCC opens its doors to students from the National University of Singapore, offering them a firsthand look at the supercomputers housed in its data centre and introducing them to the scale and possibilities of national HPC infrastructure.

3 Through outreach sessions at universities, NSCC engages schools and universities interested in HPC, raising awareness of HPC as a discipline and encouraging the next generation to explore careers in advanced computing and AI-enabled research.



they need without navigating individual systems and architectures. However, simplifying the user experience is only part of the equation. Meeting diverse research needs will also require the ability to draw on multiple sources of compute, Dr Hung believes.

NSCC is therefore looking towards a multi-tier, federated model spanning its own infrastructure, institutional resources, qualified cloud providers and overseas partners. Researchers would have greater flexibility to match computing resources to their workloads.

Access models would also evolve, combining traditional fair-share scheduling — which balances competing demands for shared computing resources — with more flexible, cloud-like access, dedicated compute and responsive provisioning for different workload requirements.

BUILDING THE CAPABILITIES BEHIND THE TECHNOLOGY

Delivering this more seamless model of advanced computing will require more than changes to infrastructure and access. Dr Hung sees a growing role for specialists who can translate increasingly complex technologies into practical solutions for researchers: system



Boosting HPC capabilities, one clinic at a time. NSCC's HPC Clinics provide hands-on guidance and direct engagement with users, forming a key part of its strategy to strengthen and grow Singapore's national HPC ecosystem.

architects who design computing environments, software engineers who build applications and HPC performance engineers who optimise them for efficiency and performance.

NSCC's longer-term ambition is to provide more integrated support and embed this expertise within research projects where needed. By working alongside research teams, its specialists could help design computational approaches and determine how best to use different resources.

Research environments will also need to evolve to support different requirements. Trusted Research Environments (TREs), for example, can provide secure, governed spaces with controls over data access and analysis, particularly for research involving sensitive data in field such as healthcare.

Building this broader capability will require collaboration beyond NSCC, including partnerships with universities and leading international computing centres to develop talent, exchange expertise and explore emerging technologies.

A STRATEGIC RESEARCH PARTNER

Advanced computing is increasingly a strategic national capability. But the value of that capability will ultimately be measured not by the scale of the infrastructure, but by what researchers are able to achieve with it. Its impact can extend well beyond research itself, from advancing healthcare and developing new materials to accelerating technological innovation and strengthening industry competitiveness.

"Ultimately, our ambition is not simply to operate one of the world's leading supercomputing centres," says Dr Hung. "We aim to become a trusted strategic partner for Singapore's research community."

That partnership would mean engaging researchers earlier and understanding the scientific problems they are trying to solve before assembling the computing resources and technical expertise needed to address them.

"We laid a solid foundation in our first decade," Dr Hung says.

"The next decade is about turning that foundation into greater impact. Our role is to ensure Singapore's researchers can harness emerging technologies to pursue ambitious science."

"That is how we can strengthen Singapore's research capabilities and remain at the forefront of scientific discovery."

THE NEXT GENERATION OF SUPERCOMPUTING TALENT TAKES SHAPE

Meet the APAC teams that took on the 2025 Student Cluster Competition (SCC), an intensive, 48-hour challenge where students design, build and optimise supercomputing clusters within strict power and time constraints.

NATIONAL UNIVERSITY OF SINGAPORE

Team Kent Ridge¹ from NUS came together through a sharp curiosity for HPC. Many members were drawn in from a parallel computing course or word of mouth, looking for a chance to move beyond coursework and into systems-level practice.

As preparation progressed, students gravitated toward areas where they could contribute most — machine learning, physics-based applications, systems engineering, or cluster administration. Weekly technical clinics, peer-led sessions and practice runs became the backbone of the team's preparation.

¹Team Kent Ridge won third place in the virtual category of the ISC High Performance 2025 Student Cluster Competition.


The steepest challenge was building and optimising a physical cluster entirely from scratch, without vendor support and under tight timelines. Cabling, BIOS settings, networking, filesystems and GPU-CPU scheduling all had to be mastered before meaningful benchmarking could begin. Challenging applications such as MLPerf spurred the team to devise custom parallelism

strategies and contribute their solution back to the wider community. However, not every application was seamless. One resisted repeated attempts to run reliably, even after extensive patching. The episode revealed a core truth about HPC: getting software to work can be harder than optimising it. "What really matters is how workloads, data movement, memory and I/O are managed — not just how fast the hardware is," observed Dr Cristina Carburanu, the team's advisor from the Department of Computer Science, School of Computing, NUS. "Apart from systems-level thinking, SCC taught our students that HPC rewards patience, curiosity and problem-solving more than perfection on the first attempt."

NANYANG TECHNOLOGICAL UNIVERSITY

For NTU's Team Supernova², SCC was an exercise in disciplined problem-solving. Members came from different academic backgrounds, some moved towards system administration and cluster deployment, while others worked on benchmarking, performance tuning and application optimisation. But a common thread weaves through them: a common curiosity for HPC and an openness to wrestle with difficult questions.

The team found that achieving stable, reproducible performance under strict



Team Kent Ridge from the National University of Singapore comprised students Zhi Shen Yeo, Amy Ling, Gabriella Gloria, Way Yan Win, Xinyang Yu and Taanish Bhardwaj, and was led by advisors Dr Cristina Carburanu and Dr Sriram Sami.

constraints was the main challenge. Doing so requires a granular understanding the hardware-software stack, from operating systems and compilers to MPI libraries and memory bandwidth. Adding to the test was the fact that minor configuration tweaks could lead to outsized swings in results, making systematic testing and careful documentation essential.

For a challenge of this scale, missteps were inevitable. During preparation, configuration changes occasionally led to regressions or instability. The nimble-fingered team expertly reverted to known baselines, isolated variables and rebuilt

confidence through controlled experiments, reinforcing trust and keeping problem-solving firmly collective during the process.

Ultimately, SCC reshaped the lens through which the team viewed HPC. As one of the team’s advisors, Dr Loke Yuan Ren from the College of Computing & Data Science at NTU, added, “For us, HPC has shifted from theory into a practical, interdisciplinary engineering discipline.” The experience sharpened both technical judgement and collaboration, leaving the team better prepared for the complexity and ambiguity of HPC in the real world.

Whether hard power limits or stubborn software or challenging configurations, SCC offers something classrooms alone cannot: hands-on experience with real systems under pressure. In that process, the next generation of HPC talent gains the confidence, exposure and dexterity needed to grapple with the constraints, complexity and capriciousness of production computing.

²At the time of publication, Team Supernova had also won joint first place onsite and second place online at the ISC 2026 Student Cluster Competition. These results build on NTU’s longstanding record of strong performance in international HPC competitions.

NATIONAL TAIWAN UNIVERSITY

National Taiwan University entered SCC at SC25 for the very first time in 2025, and was honoured to receive the championship title. With no prior SCC experience to build upon, the team's journey was as much a learning experience as it was a competition.

The team was assembled from the ground up in early 2025, bringing together students from different year levels and technical backgrounds. With no institutional precedent to draw on, roles emerged naturally through hands-on preparation: some members gravitated towards cluster architecture, power management and cooling, while others focused on benchmarking and application optimisation. One student stepped in to coordinate timelines and communication. Rather than assigning titles, the team found its shape through practice and collaboration.

Managing the competition's strict 10,000-watt power limit turned out to be the team's steepest learning curve. Performance depended on carefully balancing CPU and GPU settings, memory configurations and thermal behaviour across different workloads. Through extensive trial and error,



Team Supernova (Bryan Guanrong Shan, Ruey Yee Hong, Yoong Ken Tan, Yichi Zhang, Boyuan Gu and Zhanyan Zhu), led by Dr Loke Yuan Ren and Dr Wang Chen, posing with the NSCC team at SC25.



National Taiwan University students Tso-Fei Yen, Jui-Chien Tsou, Hsuan-Chi Liu, Wei-Chin Wang, Chia-Yi Chin and Kuan-Hsun Tu, led by Professor Chun-Yi Lee, claimed the championship title at the SC25 Student Cluster Competition.

PUTTING THE "SUPER" IN "SUPERCOMPUTING"

From lightning-fast processors to seamless interconnects and precision-tuned software, supercomputers crunch oceans of data into insight — powering discovery at scales the world has never seen.



the team developed a monitoring system that could reconfigure the cluster between applications while staying within budget.

The 48-hour competition itself was an intense test of coordination and endurance. With multiple benchmarks running back-to-back under tight time constraints, every decision had to be made quickly and collectively. Team members rotated responsibilities, communicated constantly and supported one another through unexpected challenges. It was this spirit of close collaboration and shared

commitment that carried the team through to the end.

"We are grateful for the opportunity to have participated in SCC and deeply honoured by the outcome. Competing alongside internationally recognised teams was a humbling and invaluable experience for our students. It gave them a much clearer understanding of what production HPC environments demand, and reinforced the value of persistence and teamwork," said the team's advisor, Professor Chun-Yi Lee from the Department of Computer Science and Information Engineering at National Taiwan University.



Beneath the soft, captivating electric glow of a data hall, rows of sleek black racks whirl and buzz like the pulse of a digital metropolis. Each rack holds thousands of processors, wired together to think as one. This is where supercomputing takes place — where machines can perform quadrillions of calculations per second, outpacing personal computers by factors of millions.

To picture that scale, imagine your laptop attempting a weather simulation. It might struggle for days or even weeks, tied up until the calculation ends. A supercomputer can complete it in a fraction of the time, then move on to billions more. Its strength lies not only in raw speed, but in how it is architected to think in parallel. Instead of a single processor tackling one task at a time, supercomputers slice colossal problems into smaller ones solved simultaneously across tens or hundreds of thousands of cores. Each processor contributes a piece of the puzzle, and together they assemble the full picture with extraordinary speed and accuracy.

HPC is the discipline that makes this possible — an ecosystem of hardware, algorithms, software, data pipelines and human expertise that transforms computation into discovery. It is not simply the machine itself, but the ingenuity behind it.

Achieving this performance requires a mind-boggling level of precision and engineering. Data moves across high-speed networks faster than the blink of an eye, threading through layers of memory built to store and retrieve information without pause. Advanced chips — from CPUs and GPUs to dedicated AI accelerators — handle specialised workloads for a multitude of sectors, from aviation and AI to biotechnology and banking. Meanwhile, sophisticated software coordinates millions of operations, balancing loads and ensuring every ounce of computational power is put to work.

Yet what truly makes supercomputing super extends beyond its machinery. In an increasingly digitised and digitalised world where over 400 million terabytes of data are generated every day⁷ — the equivalent of every person on Earth taking 6,000 photos daily — HPC has become the indispensable workhorse of modern science. It extracts clear insights from convoluted datasets.

Climate scientists use it to model monsoons, track atmospheric shifts and project sea-level change with resolution once unachievable. In healthcare, supercomputers help scientists design new drugs, unravel genomes and simulate precision treatments, shaving off years of painstaking, laborious lab work. During the COVID-19 pandemic, HPC simulations modelled how droplets dispersed in the air, helping shape public safety measures. Engineers across aerospace, automotive and manufacturing sectors now design safer vehicles, optimise electric batteries and test new materials using digital twins instead of costly physical prototypes.

What is perhaps less visible is how deeply HPC underpins everyday life. Banks use advanced modelling to detect fraud and manage risk. Urban planners simulate entire cities to reduce congestion and optimise energy usage. Semiconductor

⁷<https://explodingtopics.com/blog/data-generated-per-day>

manufacturers use HPC-driven predictive modelling to shorten development cycles. And the AI systems shaping daily experiences, from language translation to personalised healthcare, are all trained in supercomputing infrastructure, the only kind capable of managing their immense computational load.

Asia has emerged as one of the world's most dynamic regions for HPC, home to some of the fastest supercomputers ever built. Governments and institutions across the region are investing heavily in national compute programmes, AI strategies and talent development, pushing the frontiers of science and spurring economic growth.

Above all, the “super” in supercomputing is accentuated by the real-world impact it sets in motion. It makes the previously impossible possible, accelerates breakthroughs that once took decades and lays the groundwork for what comes next, from quantum technologies to the pursuit of clean fusion energy.

Supercomputers may not wear capes or masks. But in their quiet, relentless work, they turn complexity into clarity, and possibility into progress. In every meaningful way, they are unmistakably super.

WE WANT TO HEAR FROM YOU!

What resonated with you in this issue?
What would you like to see more of in future editions?

Your input helps shape the stories, perspectives and topics we bring to *Supercomputing Asia*.

If you have a few minutes,
we would appreciate hearing from you.

[Click here to share your feedback.](#)

Thank you for being
part of the conversation!

