

FOR IMMEDIATE DISSEMINATION

Singapore Launches Next-Generation Supercomputer to Scale AI, Scientific Discovery and Hybrid Quantum Research

- Singapore launches ASPIRE 2B, its next-generation national research supercomputer, significantly boosting the country's AI and advanced computing capabilities as the National Supercomputing Centre (NSCC) Singapore marks its 10th anniversary.
- Delivering up to 115 petaFLOPS of compute power, ASPIRE 2B quadruples the compute capacity currently available under Singapore's national supercomputing infrastructure.

Singapore, 8 June 2026 - Singapore today launched ASPIRE 2B, its next-generation national research supercomputer, significantly advancing the country's high performance computing (HPC) capabilities for AI, scientific computing and emerging hybrid computing approaches.

Launched by Mrs Josephine Teo, Minister for Digital Development and Information and Minister-in-Charge of Cybersecurity and Smart Nation Group, ASPIRE 2B marks a major expansion of Singapore's national compute infrastructure to support increasingly AI and data-intensive workloads. Introduced in NSCC's 10th anniversary year, ASPIRE 2B reflects a decade of national investment in advanced computing infrastructure and Singapore's growing ambitions in AI and scientific innovation.

Delivering up to 115 petaFLOPS (PF)¹ of compute power, capable of more than 100 quadrillion calculations per second, ASPIRE 2B offers close to four times the combined capacity of current national research systems, ASPIRE 2A and ASPIRE 2A+.

Scaling National Compute Capabilities to Meet Growing Research Demands

ASPIRE 2B is equipped with more than 1,500 NVIDIA H200 GPUs², making it the country's largest NVIDIA GPU cluster dedicated to research and public sector use. It will serve as a key testbed for emerging AI-driven research approaches such as agentic AI, where AI systems help automate and orchestrate end-to-end research workflows, from ideation to generative design to lab-based experimental validation. These workloads require AI training, as well as inference – where trained AI models are used to generate predictions and responses from new data – alongside large-scale data analytics and scientific simulations to be performed within a unified compute environment.

ASPIRE 2B is designed to support the new hybrid quantum-classical computing paradigm. It is expected to be linked with Quantinuum's Helios quantum computer, with a targeted installation

¹ A petaFLOP is a measure of computing power. One petaFLOP equals one quadrillion calculations per second, enabling faster processing of complex data and AI workloads. To illustrate the scale of 115 petaFLOPS, if every person on Earth performed one calculation every second, it would take more than 170 days to match what ASPIRE 2B can process in a single second.

² Graphics Processing Units (GPUs) are specialised processors that excel at handling large volumes of parallel calculations. They are critical for AI training, data analytics and accelerated scientific computing.

in Singapore near the end of this year, enabling researchers to experiment with hybrid computing algorithms that could help solve complex problems that are difficult for both systems to handle alone.

Enabling Nationally Significant Research and Applications

NSCC's supercomputers have supported national programmes such as the National Environment Agency's Third National Climate Change Study (V3), led by the Centre for Climate Research Singapore (CCRS), where some of the world's highest-resolution climate projections were developed to support coastal adaptation planning, climate resilience and regional food security. With ASPIRE 2B's enhanced CPU and GPU capabilities, CCRS will be able to support a new generation of AI-powered, higher-resolution climate modelling and weather prediction to strengthen Singapore's resilience against increasingly complex climate risks.

NSCC systems have also powered national AI initiatives such as the Multimodal Empathetic Reasoning and Learning in One Network (MERaLiON) project by A*STAR Institute for Infocomm Research (A*STAR I²R), including Southeast Asia-focused AI models capable of understanding regional accents, dialects, code-switching and emotions for applications across healthcare, finance and digital services. With ASPIRE 2B's expanded AI compute capabilities, researchers will be able to train larger and more advanced multimodal models while accelerating the development of next-generation sovereign AI systems tailored for Singapore and Southeast Asia.

These developments also support Singapore's broader push towards AI for Science, where AI and advanced computing increasingly accelerate scientific discovery. Building on ASPIRE 2A and ASPIRE 2A+, ASPIRE 2B will further enhance the quality and capability of national AI and scientific research, while supporting new initiatives such as Singapore Medical Foundation AI Model (SIMFONI), advanced robotics research and other large-scale AI workloads requiring sustained GPU access.

Prof Robert Morris, Executive Director, Singapore Medical Foundation AI Model (SIMFONI) said:

"Advancing healthcare AI requires the ability to work with large and complex clinical datasets at scale. ASPIRE 2B provides the compute foundation needed to support more advanced AI models for applications such as clinical decision support and patient record analysis, contributing to more data-driven and preventive healthcare."

Supporting Future-Ready Research

Beyond infrastructure, NSCC is evolving its operating model to better support Singapore's long term research and innovation needs. This includes enhancements to its resource provisioning framework to prioritise national Research, Innovation and Enterprise (RIE) programmes while providing greater flexibility for institutional and emerging research demands.

NSCC will also deepen engagement with the research community through enhanced user enablement services, such as application and workflow optimisation, closer technical partnerships, and consultations on algorithm-hardware co-design. Alongside this, NSCC

will strengthen its focus on training and talent development through structured onboarding and curriculum co-development with universities and industry partners.

Dr Terence Hung, Chief Executive, National Supercomputing Centre (NSCC) Singapore, said:

“ASPIRE 2B reflects Singapore’s ambition to strengthen its position in trusted AI and advanced computing. It provides researchers with the compute capabilities needed to tackle larger and more complex challenges in areas critical to Singapore’s future, including healthcare, sustainability and urban resilience. As NSCC marks its 10th anniversary, ASPIRE 2B also strengthens NSCC’s role as a strategic national HPC partner. It will support research innovation, strengthen talent and capability development, and enable the next phase of Singapore’s AI and compute-driven research ambitions.”

ASPIRE 2B forms part of the \$270 million national investment announced by the National Research Foundation (NRF) in 2024 to strengthen Singapore’s supercomputing infrastructure and capabilities.

Enclosed:

Annex A - About ASPIRE 2B

Annex B - National Projects on ASPIRE 2B

Annex C - Institutional Resource Allocation Framework for IHLs

For media queries and clarifications, please contact:

Ms Denise Lee

Marketing, Communications & International Relations

National Supercomputing Centre (NSCC) Singapore

(+65) 9188 9031

Email: deniselee@nscg.sg

Ms Doris Yang

Corporate Communications

Agency for Science, Technology and Research (A*STAR)

(+65) 9367 5336

Email: doris_yang@a-star.edu.sg

About the National Supercomputing Centre (NSCC) Singapore

The National Supercomputing Centre (NSCC) Singapore was established in 2015 to manage Singapore's national petascale facilities and high performance computing (HPC) resources. A National Research Infrastructure funded by the National Research Foundation (NRF), and hosted by A*STAR, the HPC resources that we provide help support the research needs of the public and private sectors, including research institutes, institutes of higher learning, government agencies and companies.

As a national strategic technological platform, NSCC has the mission to enhance competence, capacity and competitive advantage in the use of HPC in all relevant fields such as computational science, analytics, engineering, advanced manufacturing, genomics, biomedicine, healthcare, AI and quantum computing, among many others. With the support of our research partners, NSCC catalyses national research and development initiatives, develops HPC skillsets and applications, and enhances Singapore's research capabilities. For more information, please visit: <https://www.nscg.sg/>

About the National Research Foundation (NRF)

The National Research Foundation, Singapore (NRF), set up on 1 January 2006, is a department within the Prime Minister's Office. The NRF sets the national direction for research and development (R&D) by developing policies, plans and strategies for research, innovation and enterprise. It also funds strategic initiatives and builds up R&D capabilities by nurturing research talent. For more information, please visit: <https://www.nrf.gov.sg/>

Annex A - About ASPIRE 2B



ASPIRE 2B is Singapore's next-generation national supercomputer, designed to operate alongside ASPIRE 2A and ASPIRE 2A+ as part of the country's high performance computing (HPC) ecosystem. Together, these systems support a broad range of research needs, with each optimised for different workloads.

ASPIRE 2B is a heterogeneous system that combines advanced CPU and GPU capabilities on a single platform to support large-scale workloads spanning simulation, data analytics and artificial intelligence. Delivering up to 115 petaFLOPS (PF) of compute power, ASPIRE 2B can perform more than 100 quadrillion calculations per second. To put this into perspective, if every person on Earth performed one calculation every second, it would take more than 170 days to match what ASPIRE 2B can process in a single second.

While ASPIRE 2A supports CPU-intensive workloads such as climate modelling, materials science and engineering simulations, and ASPIRE 2A+ strengthens capabilities for AI and accelerated computing, ASPIRE 2B brings these capabilities together at greater scale. This enables more complex and integrated research across climate and weather science, healthcare, materials discovery, advanced manufacturing, next-generation AI and emerging hybrid quantum-classical research approaches.

ASPIRE 2B delivers:

- Four times the combined computing capacity of ASPIRE 2A and ASPIRE 2A+
- More than double Singapore's national GPU capacity
- Enhanced support for integrated workflows combining simulation, data and AI within a single system
- Enhanced readiness for emerging hybrid quantum-classical computing approaches

Specifications of ASPIRE 2B:

- CPU: 184,320 cores (AMD EPYC™ 9655)
- CPU Performance: ~12 PFLOPS (theoretical peak performance)
- GPU: 1,536 NVIDIA H200 GPUs (141 GB variant)
- GPU Performance: ~103 PFLOPS (theoretical peak performance)
- System Memory: 1,072 TB
- Storage: 63.5 PB
- Interconnect: 400 Gbps Slingshot

Annex B - National Projects on ASPIRE 2B

a) SIMFONI (Singapore Medical Foundation AI Model)

SIMFONI was launched in 2025 to drive the development and deployment of foundational models for healthcare. SIMFONI is a programme of the Consortium for Clinical Research and Innovation, Singapore (CRIS). SIMFONI is supported by the Singapore Ministry of Health through the National Medical Research Council (NMRC) Office, MOH Holdings Pte Ltd. under the NMRC SIMFONI Funding Initiative (MOH-002026).

With ASPIRE 2B, SIMFONI will be able to train and fine-tune models on larger, more diverse and multimodal clinical datasets, including both structured and unstructured health records. This will support more advanced and compute-intensive workflows, such as training larger models, handling longer clinical context and integrating multiple data modalities, which are increasingly important for real-world healthcare AI applications.

These capabilities are expected to support more robust and scalable model development, helping to strengthen the use of AI in areas such as diagnostics, clinical workflow support and treatment planning. Over time, this could also contribute to more personalised and preventive approaches to healthcare by enabling better analysis of longitudinal patient data and earlier identification of potential risk signals.

b) Multimodal Empathetic Reasoning and Learning in One Network (MERaLiON)

MERaLiON is a Singapore multilingual and multimodal large language model (LLM) initiative, developed by A*STAR Institute for Infocomm Research (A*STAR I²R) to advance sovereign AI capabilities tailored to Singapore and Southeast Asia. As part of Singapore's national AI efforts, it supports the development of locally relevant foundation models for future applications across public services and industry.

Powered by NSCC's ASPIRE 2A+, the MERaLiON team has released MERaLiON AudioLLM v3 – a Southeast Asia-focused voice-first multilingual model capable of transcription, translation, summarisation and question answering, with the ability to interpret regional accents, dialects, colloquialisms and code-switching. The v3 series also includes the Text to Speech capability for local Hokkien, contributing to use cases in healthcare and social services. MERaLiON supports 21 products across sectors ranging from finance, healthcare and public safety to enterprise applications. An example is with Axiom and Lion Befrienders, a social service agency supporting seniors in the community, where it helps automate routine check-in calls and enhance outreach operations. To date, the MERaLiON team has released twelve models and two datasets on Hugging Face, collectively achieving close to 1 million downloads, and established an 18-member consortium to drive real-world applications across healthcare, banking, mobility, media and homeland security.

With ASPIRE 2B, the team will be able to train larger and more advanced multimodal models, covering more underserved languages, accelerate development of next-generation end-to-end model and complex AI systems and pursue richer capabilities including deeper audio-visual understanding and emotion-aware contextual reasoning, advancing Singapore's ability to translate sovereign AI research into impactful national applications.

c) Singapore's Future Climate Modelling Study (V4)

CCRS will continue to advance the next phase of its climate research on ASPIRE 2B, focusing on AI-driven, physics-based and hybrid modelling approaches for applications such as rainfall nowcasting, local weather prediction and regional climate projections. Rainfall nowcasting leverages data-driven AI methods to deliver rapid, high-frequency forecasts, while for local weather prediction, CCRS is exploring emerging AI approaches alongside established physics-based models, including evaluating them against the current 'SINGV' system, a next-generation regional weather and climate modelling system designed for the tropics, while developing its next-generation 'SINGV_NG' capabilities.

These efforts also reflect the growing integration of AI into climate and weather research, where AI is increasingly used alongside traditional simulations to accelerate modelling, improve forecasting accuracy and support more advanced climate analysis.

Building on the Third National Climate Change Study (V3), CCRS is also preparing for the fourth iteration (V4), which will integrate both AI and physical approaches to produce higher-resolution projections at regional and urban scales. These efforts aim to deliver improved long-term regional climate projections to support Singapore's climate resilience planning, benefiting sectors such as coastal protection, heat resilience, biodiversity, flood management and maritime operations.

ASPIRE 2B will be a key enabler for these efforts. Its enhanced CPU capabilities will support large-scale simulations required for regional climate projections, while its GPU-based architecture will enable the training of advanced AI models for weather prediction and climate downscaling. Together with efficient access to CCRS's extensive weather and climate datasets, ASPIRE 2B will support increasingly compute-intensive and AI-driven climate research workflows.

Annex C - Institutional Resource Allocation Framework for IHLs

The Institutional Resource Allocation Framework (Tier 2) is designed to support growing demand from Singapore's Institutes of Higher Learning (IHLs) by providing a structured pathway to access national supercomputing resources beyond traditional RIE-funded programmes. Through this framework, NSCC can better aggregate institutional demand and provide compute resources in a more cost-efficient and coordinated manner, helping institutions access advanced capabilities without duplicating infrastructure.

This tier will support a broader base of researchers, including emerging teams and new projects that may not yet fall within major national funding categories, while enabling institutions to explore new computational approaches, pilot innovative workflows, and build capabilities in advanced computing. By complementing existing allocation models, Tier 2 strengthens Singapore's overall HPC ecosystem by broadening access, improving flexibility, and ensuring national compute resources remain responsive to evolving research, talent development, and innovation needs.

NSCC is also engaging Institutes of Higher Learning (IHLs) and research institutions on broader participation in this shared national HPC resource model, with more institutions expected to come on board over time.