




Sustainable Metal Cloud



Build Sustainable AI.



NVIDIA PREFERRED NPN CSP PARTNER

The world's most energy efficient GPU cloud. Built with sustainable technology.

Cloud computing accounts for over 2.5% of the world's carbon emissions, with high-powered GPUs pushing power demands further. As AI workloads keep growing, so does our harm to the environment.

Powered by advanced cooling technology, Sustainable Metal Cloud delivers large-scale GPU clusters for AI workloads with nearly half the carbon footprint.

48%

GENUINE REDUCTION IN CO₂ EMISSIONS
THROUGH ADVANCED TECHNOLOGY

>70%

BETTER TCO THAN LEGACY CLOUDS



Engineered unlike any other cloud. Ever.

Our performance, scale and cost benefits are entirely driven by the advanced GPU hosting infrastructure we've spent the past decade perfecting.

The result? NVIDIA-designed GenAI training clusters delivered via the world's first hyperscale immersion-cooled platform — providing full-rail InfiniBand, H100 SXM, BlueField-3. Zero compromises on performance, with a significant reduction of the energy required to run AI in the cloud.

HOW?

By pushing recirculated fluid through each cluster, instead of inefficient refrigerated air, our engineers redesigned every element of infrastructure to increase cluster density by 4× and lower energy use by half. Cost savings and environmental benefits that we pass directly to our customers.

<1.03 PUE

A STANDALONE PUE OF LESS THAN 1.03, BUILT FOR MULTI-MEGAWATT ENVIRONMENTS AND INDEPENDENTLY CERTIFIED TO RUN NVIDIA HGX SYSTEMS; IT'S NOT JUST SUSTAINABLE; IT'S A GENERATION AHEAD OF EVERYTHING ELSE.



Powered by


HyperCube

SINGAPORE 2 H100 AVAILABILITY ZONE [↑]
>2,700 GPUS PER CLUSTER



Breakthrough CO₂ savings, from any region on earth

Our hyperscale technology is deployed via a fleet of sustainable AI factories. No matter what the energy source, SMC availability zones use up to half the energy as legacy GPU clouds.



The world's most powerful AI infrastructure

Access dedicated scalable NVIDIA-designed GPU clusters built to train and inference AI models. Purpose-built to handle complex AI workloads, sustainably.

L40S

PERFORMANCE AND SCALE

Delivers next-generation graphics, compute, and AI capabilities to address increasingly complex workloads.

H100

LARGE-SCALE HPC & AI WORKLOADS

Designed to maximise AI throughput, it provides organisations with a highly refined, systemised, and scalable platform to achieve breakthroughs in natural language processing, recommender systems, data analytics, and much more.

A100

UNIVERSAL SYSTEM FOR ALL AI WORKLOADS

Tackle the largest and most complex AI workloads — from analytics to training to inference.

PFS

HI-IOPS PARALLEL FILE STORAGE

Connected to every node at up to 200GBPS non-blocking, SMC deploys WEKA ultra-high-performance parallel file storage to handle data saturation across any type of AI workload. Aligned with SMC's sustainable mission, WEKA's efficiency yields a CO₂ footprint up to 80% smaller than legacy storage infrastructure.





H100 80GB SXM

- | Full-rail ConnectX-7 InfiniBand
- | BlueField-3 DPU, Sapphire Rapids 8462Y+ 128vCPU
- | 30.72TB NVMe
- | 2,048GB DDR5 Memory
- | 200GBPS storage fabric

The sustainable advantage

Genuine sustainability from the ground up changes everything. Years of developing greener technology leads to significant wins for our customers.

32,000

UP TO 32,000 GPUS PER CLUSTER

24/7

CRITICAL SUPPORT, YEAR-ROUND

HIGHER DENSITY = QUICKER SCALING

We fit 4x more GPU compute in our tanks than traditional data centre infrastructure. Go-to market quicker, scaling from a single GPU to clusters of up to 32,000 GPUs.

LESS ENERGY = COST SAVINGS

With a near 50% energy saving for any workload, our infrastructure delivers customers industry-best TCO that scales with projects.

MODULAR TECHNOLOGY = GLOBAL AVAILABILITY

We're backed by one of the world's largest data centre groups. As a member of the ST Telemedia Global Data Centres group, we are scaling rapidly in over 85 availability zones across the globe. GPU Metal Instances are hosted in secure data centres, ensuring the confidentiality and integrity of data.



A global network of AI Factories

As a member of the ST Telemedia Global Data Centres group, we are scaling rapidly across a platform of over 85 availability zones in Asia and Europe, democratising worldwide access to sustainable AI compute.

This expansion introduces local Availability Zones in strategic regions: Singapore, India, Thailand, Australia, and Europe. Our Tier III data centers, upgraded to accommodate Sustainable AI Factories, are key to a worldwide fabric that ensure seamless connections to global CSPs, private data caches, and IP transit.

It's about making AI that's universally accessible and sustainably delivered, aligning with our broader mission.



Customise your tools and capacity



NVIDIA AI ENTERPRISE

Pre-trained models, and over 100 frameworks, including NeMo, delivered with enterprise support from NVIDIA.

SECURITY BUILT-IN

Continuous monitoring and regular releases of security patches for critical and common vulnerabilities and exposures (CVEs), with enterprise support and access to NVIDIA AI experts.

LLM TRAINING SPEED-UP

Train at scale with the NVIDIA TAO Toolkit. Create custom, production-ready AI models in hours, rather than months, by fine-tuning NVIDIA pre-trained models — without AI expertise or large training datasets.

REDUCE DEVELOPMENT TIME

Develop faster with NVIDIA AI Workflows: prepackaged reference examples for enterprises to jump-start the building of AI solutions.

DEPLOY AT SCALE

Deploy at scale with NVIDIA Triton Inference Server, which simplifies and optimises the deployment of AI models at scale and in production for both neural networks and tree-based models on GPUs.

MANAGED METAL

H100 & L40S accelerators networked at hyperscale with RDMA InfiniBand.

START-UPS AND SCALE-UPS

Cost-effective dedicated metal to experiment, train and inference GenAI workloads.

MANAGED CO-LOCATION

Enterprises, government and education teams looking to add GPU cluster capacity to existing x86 workloads, with a better TCO and time-to-use than on-premise.

AI PLATFORMS

Mid & large-scale application and platform developers who demand large-scale, high performance training and inference GPU clusters in geographically diverse locations.

WHOLESALE & HYPERSCALE

32,000+ GPUs in a single-cluster, customised to suit the bespoke requirements of enterprises and CSPs. Delivered at data centre level to expand private clusters or open new public AZs.

 **GreenFutures**
Accelerator

FOR MORE INFORMATION,
VISIT [SMC.CO/ACCELERATOR](https://smc.co/accelerator)

GreenFutures is our SMC accelerator program dedicated to propelling start-ups and scale-ups to begin and grow their AI journey on a foundation of sustainability. Access the tools, technology, and mentorship you need to grow sustainably.

Applications are open for those who meet our criteria.

End-to-end AI infrastructure

All of the frills, without the excess energy consumption.



METAL & FULL STACK

Flexible ways to use. For hyperscale, SMC provides clean, controllable metal clusters down to total ownership of the network. For enterprise, use NVIDIA AI tools and services delivered in a full-stack development environment.



RESERVED GPU CLUSTERS

Large-scale GPU clusters that are designed to NVIDIA's standards for hyperscale generative AI workloads.



SERVICE & SUPPORT

Our global operation centres are paired with a guaranteed minimum 99.50% SLA with critical support 24/7, 365 days a year — on all instances, no matter the cluster size. Receive unlimited workload migration and on-boarding support from experts.



CONNECTIVITY

A global peering network built to the needs of each customer. The SMC Secure Connect Network enables easy cross-connect to public clouds and your own data stores as part of the on-boarding process.

We're engineers and technologists, dedicated to reinventing infrastructure

OUR COMMITMENT

SMC was founded on a vision to move cloud computing towards net zero, with solutions forged through advanced technology. This vision drives everything we do.

We're here to enable people not only to build AI — but to build AI sustainably. We've pushed boundaries and pulled strings. This push and pull, through many years, led to us building a GPU cloud platform like no other — delivered through our own Sustainable AI Factories — to upend what it means to be truly efficient.

Today, nearly a decade on, with the safety of enterprise SLAs and critical support, we maintain an unwavering dedication to our planet.

Let's build AI, sustainably.





Talk to our experts

SMC is for training or inferencing. From small workloads to extra large. See how easy it is to join other innovators who are already building genuinely sustainable AI. SMC's infrastructure experts and AI solutions architects help to engage to scope and on-board new and transitioning workloads with ease.

Get in touch with us at
smc.co/enquire-now
or email contact@smc.co

