

Harness the power of Artificial Intelligence

The Ultimate Learning Environment for Your AI Infrastructure

AI workloads can be deployed on cloud infrastructure leveraging vGPU, which is great for specific one-off tasks, but what if you have a regular requirement to run these workloads and therefore permanently need infrastructure in place?

The public cloud's consumption model can be costly and it's not always easy to obtain dedicated infrastructure or the exact components you need... that's where Ingenuity Cloud Services can help.

We specialise in providing dedicated infrastructure for **steady-state AI and ML workloads**.

We stock specialist scalable servers with GPU, including specific hardware (such as HGX and NVIDIA DGX), and offer this infrastructure on competitive commercial terms.

Often, having the requirement for AI infrastructure on a semi-regular basis justifies the investment in a dedicated deployment, built and optimised for your specific workload.

Cost-Effective Performance

Our bare metal is dedicated to deliver optimal performance without contention. We also offer low-cost bandwidth packages, reducing TCO, when moving large data sets on and off AI infrastructure.

Speak to us about your requirement as we can advise on specific components vs. workloads and we can offer flexible terms subject to availability.



We've Learnt the Market

The Hut Group have a close relationship with NVIDIA, having been an early adopter of the DGX internally, for our AI workloads (primarily leveraged for our eCommerce and translation tools).

Couple this expertise with having helped other customers needing infrastructure for their own AI & ML workloads, Ingenuity Cloud Services is perfectly positioned to advise customers on what infrastructure will best suit a specific requirement.

