



Maxis AI Cloud

Enterprise-Grade AI Compute. Delivered Locally.

AI workloads require high-performance compute, but involves high upfront investment and maintenance costs, complexity, and data risk.



Why Enterprises Struggle with AI Infrastructure



High Cost & Financial Risk

- ✗ Large upfront Capex for GPU hardware
- ✗ High Opex for maintenance costs (electricity, data centre hosting)
- ✗ Overpaying for unused GPU capacity



Data Sovereignty & Compliance Risks

- ✗ Data governed by foreign jurisdictions
- ✗ Difficulty meeting PDPA requirements
- ✗ Risk of proprietary data exposure in public large language models (LLMs)



Operational & Technical Complexity

- ✗ Hardware maintenance, cooling, and upgrades
- ✗ Rapid GPU obsolescence

What is Maxis AI Cloud


It is an extension of Maxis Right Cloud proposition that provides on-demand high performance private GPU compute for AI workloads - without upfront hardware - while keeping data and compute within Malaysia.

What It Enables



Financial Control

- ✓ Fixed MYR pricing with no hidden egress fees
- ✓ Shift from Capex to flexible Opex spending
- ✓ Pay for compute based on required workload



Performance & Speed

- ✓ Instant access to NVIDIA H200 and L40s GPUs
- ✓ Faster model training and parallel experimentation
- ✓ High-Speed Internet with Maxis Programmable Connectivity



Security & Compliance

- ✓ Offer high-performance compute within Malaysia, ensuring compliance with local data regulation, keeping sensitive data onshore
- ✓ Full compliance with PDPA requirements
- ✓ Data control via Private LLM hosting

How It Works

1 Maxis performs assessment and planning of customers computational needs for their AI use case



2 Customers then subscribe to the Maxis AI Cloud offering



3 Customers utilise the Maxis AI platform to host and run their AI applications, including private LLMs for tasks such as fine-tuning & inferencing



4 Customers are charged based on the allocated GPU resources and AI workloads processed



Why Maxis Business

We are your trusted end-to-end cloud and AI partner - delivering the right expertise, right cloud, and right network.

Build, train, and deploy AI securely with local, enterprise-grade GPU compute from Maxis.