

4-card AI server



Overview

Performance-optimised single-socket AMD EPYC 4th generation based servers with up to four double-slot NVIDIA cards and a wide range of configuration options in space saving 2U form factor. They deliver superior performance, exceptional computational speed, scalability, and cost-efficiency, making them the industry standard for AI research and deployment, enabling faster model training. For Machine Learning a new kind of server is needed, the multi GPU AIME R410 takes on the task for delivering maximum Deep Learning training and inference performance. Cards are adaptable to customer preferences, the AI server provides the processing power required for 8, 16 or 64 GPU RAM models. Unlock exceptional performance and efficiency with PowerEdge accelerated compute servers.



Article Content

Cards Against AI Server | openai/openai-apps-sdk-examples | DeepWiki

The Cards Against AI Node.js server is the most sophisticated example in the repository. It demonstrates a complex,

PowerEdge AI Servers with GPU Acceleration | Dell USA

Drive faster results with servers equipped with the latest multi-core processors. Leverage CPUs to deliver rapid data processing and

AI server – configure your AI system with NVIDIA Enterprise GPUs

Performance-optimised servers with four GPUs and a wide range of configuration options. The most powerful servers that can

4x or 8x RTX 3090 local AI servers: still worth building in 2026?

A practical guide to 4x and 8x RTX 3090 local AI servers, covering VRAM, EPYC, Threadripper Pro, power, cooling,

AI Server 4 GPU Configurator ??up to 4 NVIDIA GPUs in 2U

Performance-optimised single-socket AMD EPYC 4th generation based servers with up to four double-slot NVIDIA cards and a wide

AIME R410

Built to perform 24/7 at your inhouse data center or co-location. Configurable with 4 high end deep learning GPUs (NVIDIA RTX

4U AI Server

Designed with versatility in mind, this 4U server allows up to four Neural Processing Unit (NPU) cards for exceptional computing

GitHub

NOTE: Our KuhnPoker server implementation supports 4-cards variant of the game with the extra A card in the game. First, the

AIME R410

Aime R410 - Machine Learning Server
Epyc CPU Performance
Up to 4TB High-Speed SSD Storage
High Connectivity and Management Interface
Well Balanced Components
Developed For Machine Learning Applications
The AIME R410 was first designed for our own machine learning server needs and evolved in years of experience in deep learning frameworks and customized PC hardware building. Our machines come with preinstalled Linux OS configured with latest drivers and frameworks like Tensorflow, Keras, PyTorch and Mxnet. Just login and start right away with your...
See more on aime
Network: 2 x 10 GBit LAN
1 x IPMI LAN
Type: Rack Server
6HE, 45cm depth
RAM: 64 / 128 / 256 GB ECC memory
AIME

AIME R500 - 4 GPU Rack Server

For Deep Learning a new kind of server is needed, the multi GPU AIME R500

AIME R500

For Deep Learning a new kind of server is needed, the multi GPU AIME R500 takes on the task for delivering maximum Deep

Building a 4-Card V100 Server for Under RMB10K: Complete Tutorial

This tutorial demonstrates how to build a professional-grade AI computing server with extremely high cost

Contact Us

For more information, pricing, or custom solutions, please contact us:

Website: <https://benniedesignstudio.co.za>

Email: sales@benniedesignstudio.co.za

Phone: +27 82 391 4765

Address: The Towers, 1 St Georges Mall, Cape Town, 8001, South Africa

This document is for informational purposes only. Specifications subject to change without notice.

