



16-card AI server

AMD has announced the Instinct MI350P, a PCIe accelerator aimed at enterprises that want on-premises AI inference without rebuilding their data center. The card is a dual-slot, full-height, ...

Powered by the latest NVIDIA Blackwell architecture, AMD EPYC or Intel Xeons processors, our GPU optimized AI servers deliver unbeatable performance and scalability for your most demanding ...

AMD has released a PCIe version of its flagship MI350 accelerators, the MI350P. Half of a MI350X, the card is aimed at customers who need to fit an modern AI accelerator into a traditional ...

With powerful RTX graphics and AI capabilities, NVIDIA RTX PRO Servers deliver exceptional performance for Universal Scene Description (OpenUSD)-based 3D, synthetic data generation, and ...

Targeted at agentic AI, Instinct MI350P PCIe cards are dual-slot drop-in cards for standard air-cooled servers. They are built to deploy inference on premises within customer's current data ...

AMD has introduced the Instinct MI350P PCIe GPU, a new enterprise accelerator designed for AI inference workloads in existing data center environments. The card uses a dual-slot PCIe ...

As a provider in cutting-edge AI solutions and products, Aivres continuously innovates to develop the widest array of most powerful servers on the market that support various GPUs and accelerators for ...

Scalable GPU servers for AI, Machine Learning, and HPC. Supports NVIDIA, AMD, and Intel GPUs with air or liquid cooling for faster model training.

Deploy AI Dedicated servers with low latency inference, full root access, 99.99% uptime, latest GPUs, crypto payments & 24/7 support. Best AI server hosting for GenAI workloads.

Instinct (TM) MI350 Series / MI325X. Supermicro unleashes the power of large-scale infrastructure with a server built with our proven AI building-block system and powered by 5th Ge



16-card AI server

Web: <https://safireschools.co.za>

