


# AI Server Configuration Comparison Chart



## Overview

This comprehensive guide aims to demystify the intricacies of server hardware for AI, providing a detailed comparison of CPUs, GPUs, and RAM. The enterprise AI server market reached \$245 billion in 2025 (ABI Research) and is projected to grow at 18% CAGR through 2030. The transition from NVIDIA Hopper. Dell's AI Factory platform (e. PowerEdge XE97xx/XE9712) provides high-density rack-scale clusters (72 GPUs per rack with NVLink, ~30× LLM inference speed-up and up to 25× energy efficiency advantage over prior-gen systems ()) with both liquid- and air-cooled options. HPE's Private Cloud AI. AI infrastructure budgeting requires precise assessment of GPU performance, memory hierarchy, storage throughput, and network latency. An AI Server Cost varies depending on server configuration, interconnect type, and workload requirements. An extremely powerful professional AI and HPC solution built on next-gen Blackwell Tensor core GPUs - NVIDIA B300. Picking the right processors will jumpstart your supercomputing platform and expedite your AI-related computing.

AI Server Configuration Comparison Chart

|                                                                                     |                                                                                                                                                                                                           |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | <p>In this comprehensive guide, we have explored the key factors to consider when selecting an AI server setup, including hardware components, operating systems, storage solutions, ...</p>              |
|    | <p>A comprehensive guide to selecting the right server specifications (CPU, GPU, RAM) for AI workloads, covering deep learning, inference, and data processing."</p>                                      |
|  | <p>Our infrastructure team will analyze your AI workload requirements and recommend the optimal server configuration from our OEM partners. Receive detailed specifications, current pricing, and ...</p> |
|  | <p>AI Server configurator is a tool that enables advanced comparison and configurations of powerful HPC systems built on latest NVIDIA GPUs.</p>                                                          |
|  | <p>Compare on-prem AI infrastructure from Dell, HPE, Lenovo, Supermicro &amp; Cisco. Analyze NVIDIA GB200/GB300 NVL72 and Blackwell Ultra hardware specs, cooling, software, and performance.</p>         |



Discover expert insights on choosing CPUs and GPUs for AI servers, exploring key analysis and solutions to optimize your AI infrastructure's performance and efficiency.



Step-by-step guide to deploying AI models on GPU servers. Improve inference speed, optimize performance, and streamline your AI workflows.



Understand the factors influencing AI server price. Compare configurations and find the most cost-effective AI dedicated server for your research or business.



Understand the factors influencing AI server price. Compare configurations and find the most cost-effective AI dedicated server for your ...



Step-by-step guide to deploying AI models on GPU servers. Improve inference speed, optimize performance, and streamline your AI workflows.



Compare AI servers and traditional servers across architecture, performance, scalability, and workloads. Understand which server type fits your business needs.

## Contact Us

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

Website: <https://www.indzawo.co.za>

Email: [sales@indzawo.co.za](mailto:sales@indzawo.co.za)

Phone: +27 71 296 8473

Address: 22 Quantum Street, Midrand, 1685, Gauteng, South Africa

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

