AI Server Components Price List



Overview

Explore the real costs of deploying AI-ready infrastructure, from GPU servers to advanced cooling and power delivery. Learn how to plan and optimize AI server data center costs for 2025.

Key Takeaways

AI server costs vary widely, from $150,000 for entry-level 4x GPU servers to over $3 million for high-end multi-GPU configurations, with GPUs being the primary cost driver.

GP

U CostsHigh-performance GPUs dominate AI server pricing, accounting for 60–75% of total cost. Common options include: NVIDIA A100: Supports up to 312 TFLOPS, suitable for mid-range AI workloads.

NVIDI

A H100: Scales to 1,000+ TFLOPS for mixed-precision tensor operations, ideal for large-scale training. GB200 NVL72: Used in top-tier servers, pushing total system costs to $3M.

GP

U selection impacts memory requirements, interconnects, and overall server efficiency, with NVLink 4 and NVSwitch improving multi-GPU communication.

© 2026 ENSAA OPTICAL NETWORKS - All rights reserved

Server Platform Costs

The server platform, including CPU, RAM, storage, and chassis, typically represents 15–25% of total cost. Examples of configurations: Supermicro SYS-421GE-TNRT: ~$165,000, suitable for inference and smaller training runs. Supermicro SYS-421GU-TNX: ~$340,000, supports models up to 70B parameters.

Supermicro SYS-421GU-TNXR: ~$400,000, designed for long-context inference and multi-modal AI.

CP

U generation, PCIe/NVLink interconnects, and cooling redundancy are key factors in platform pricing.

Networking and Storage

Networking and storage contribute 5–10% of total cost for multi-node clusters. High-bandwidth interconnects, low-latency networking, and dense storage are essential for AI workloads, especially in GPU-intensive servers.

Additional Cost Considerations

Power and Cooling: AI servers consume 5–7 kW per node, requiring advanced cooling and power distribution, which can double or triple per-rack costs. Rack Density: High-performance nodes in 5U configurations reduce the number of servers per rack, increasing infrastructure costs. Operational Costs: Maintenance, backup storage (100–500 GB), and DDoS protection add to total cost of ownership.

Vendor Options

Top A

I server manufacturers include Supermicro, Dell, and HPE, offering modular configurations, liquid cooling, and professional services for deployment and lifecycle management. Prices vary depending on GPU count, server platform, and additional services. SummaryEntry-level AI server (4x H100 GPUs): ~$150,000Mid-range server (8x H100 GPUs): ~$340,000–$400,000High-end server (GB200 NVL72): ~$3,000,000Primary cost driver: GPUs (60–75% of total)Other costs: CPU, RAM, storage, networking, cooling, and power infrastructure For daily updated component prices across multiple vendors, AI Hardware Index provides a comprehensive tracker ranging from $109 to $500k+ for GPUs, servers, and accelerators.

This breakdown helps in budgeting AI infrastructure, balancing GPU performance, memory, storage, and network throughput to optimize both CAPEX and OPEX.

© 2026 ENSAA OPTICAL NETWORKS - All rights reserved

Article Content

AI Server Data Center Cost Breakdown: 2025 Infrastructure Guide

Explore the real costs of deploying AI-ready infrastructure, from GPU servers to advanced cooling and power delivery. Learn how to plan and optimize AI server data center costs for 2025.

AI Hardware Index

Track AI hardware prices across 23+ vendors. Daily updated pricing for GPU servers, workstations, and accelerators from $109 to $500k+.

Cost of AI Server

Discover the cost of AI server across on-prem, data centers, and hyperscalers. Explore CAPEX, OPEX, and TCO to find the best AI solution for your business.

AI Server Price | Servers

This bundle combines powerful GPUs, cutting-edge server components, and advanced software to deliver optimal performance for artificial intelligence applications.

GPU Server Pricing Guide: What Does an AI Training Server Cost in

Buying a GPU server for AI isn''t like buying a workstation. Prices range from $150,000 to $3 million depending on GPU type, count, and configuration. This guide breaks down what you''ll

AI Server Cost Analysis – Memory Is The Biggest Loser

As datacenters move to accelerated computing, the cost allocated to various components shifts heavily. Nvidia sells DGX H100''s at around $270,000. The cost breakdown below

AI Server | TrendForce

Our comprehensive investigation spans critical components like AI chips, HBM, power systems, and cutting-edge cooling solutions. Partner with us in shaping the future of AI technology.

High-Speed AI Dedicated Servers | Top AI Servers Hosting

Are you looking for a high-performance AI dedicated server that is both capable and affordable? Our servers support top vendors and are offered in multiple form factors, such as rack and tower.

Tight DRAM Supply Gives Suppliers Greater Pricing Power in HBM,

© 2026 ENSAA OPTICAL NETWORKS - All rights reserved

This ensures that HBM remains the core memory component supporting AI training and inference infrastructure. The dynamic is also expected to drive broader demand across the AI

Think | IBM

Experience an integrated media property for tech workers—latest news, explainers and market insights to help stay ahead of the curve.

Welcome to Channel Dive | Channel Dive

Welcome to Channel Dive. We''re Informa TechTarget''s new publication, focused on delivering daily news and analysis for executives at North

Collaborative Content Management, and Secure File Sharing

Create, share, and govern trusted knowledge with Microsoft SharePoint—powering collaboration, communication, automation, and AI experiences across Microsoft 365.

ITPro Today, Network Computing, IoT World Today combine

ITPro Today, Network Computing and IoT World Today have combined with TechTarget . The page you are looking for may no longer exist.

WebProcure

WebProcure offers best-in-class functionality, reaching end-to-end from requester to procurement buyer to merchant, and all the way back! Designed specifically for the public sector.

TechTarget

TechTarget provides purchase intent insight-powered solutions to identify, influence, and engage active buyers in the tech market.

Top 400+ daily C2C jobs quick overview and apply now

For recruiters send your daily C2C Jobs requirements to jobs rptocorp @gmail with valid contact details, without valid contact details Job posting is rejected automatically first time and

Choose Your Best Authentication Solutions in Azure App Service

In App Service, we can leverage on following features to make outgoing authentication. Securely connect to Azure services and databases from Azure App Service Managed Identity

Ars Technica

© 2026 ENSAA OPTICAL NETWORKS - All rights reserved

News and reviews, covering IT, AI, science, space, health, gaming, cybersecurity, tech policy, computers, mobile devices, and operating systems.

Contact Us

Continue your research online:

Website: https://ensaa.es
Technical inquiry: https://ensaa.es/contact.html

This document is for preliminary research. Verify specifications against current standards, product documentation and project requirements.

© 2026 ENSAA OPTICAL NETWORKS - All rights reserved