


Our research found Dell has:



Broadest Generative AI portfolio*



Per MLPerf® results, the Dell PowerEdge XE8640 server with four NVIDIA H100 SXM5 GPUs achieved the highest AI throughput among all four-GPU submissions in nine different categories



A broader portfolio of professional services tailored for AI

Meeting the challenges of AI workloads with the Dell AI portfolio

A comparison of the Dell AI portfolio vs. similar offerings from HPE

AI adoption presents new challenges for businesses and their data center IT staff, including:

- Addressing the existing skill gaps in their current staff through training or external hiring
- Understanding the data preparation needs of AI, including the quality, quantity, location, and current state of the business’s data
- Assessing business goals to determine which AI models and implementations will provide benefits
- Assessing the computational, networking, and storage needs of the planned AI systems and acquiring these systems

With AI-ready portfolios, infrastructure vendors such as Dell Technologies and HPE offer integrated solutions that encompass the entire AI lifecycle. We researched publicly available information about both the Dell and HPE AI portfolios, and examined benchmark testing from MLPerf®. We found that Dell is poised to help businesses embrace AI with a portfolio comprising high-performing server and storage options, validated solutions, and professional services that guide the process from planning to production.

AI model benchmark performance: MLPerf result comparison

Publicly available MLPerf® Benchmark testing shows that offerings in the Dell AI portfolio offer consistent, strong performance for AI workloads. MLPerf® tests performance for both training and inferencing on several AI models. The data referenced in this report is based on MLPerf® v3.1 Inference Datacenter results published on the MLCommons® website from November 2023.¹ We compared both 8- and 4-GPU servers and include one result from each category in this summary; to see them all, read the full report at <https://facts.pt/zPmSx4c>.

*Based on Dell analysis, August 2023. Dell Technologies offers solutions engineered to support AI workloads from Workstations PCs (mobile and fixed) to Servers for High-performance Computing, Data Storage, Cloud Native Software-Defined Infrastructure, Networking Switches, Data Protection, HCI and Services. Data from: "Robert McNeal, "Dell, VMware and NVIDIA Bring AI to Your Data," accessed January 17, 2024, <https://www.dell.com/en-us/blog/dell-vmware-and-nvidia-bring-ai-to-your-data/>.

1. MLCommons, "MLPerf Inference: Datacenter Benchmark Suite Results," accessed December 12, 2023, <https://mlcommons.org/en/inference-datacenter-31/>.

Comparing MLPerf performance for eight-GPU servers

In the MLPerf® v3.1 results for eight-GPU servers, the Dell PowerEdge XE9680 with NVIDIA SXM5 H100 GPUs outperformed the HPE ProLiant XL675d Gen10 Plus with NVIDIA SXM4 A100 GPUs by up to 4.25x (see Figure 1).




Figure 1: Published MLPerf® results for the Dell PowerEdge XE9680 and HPE ProLiant XL675d Gen10 Plus as of 11/29/23. The Dell system uses the NVIDIA H100 GPU, while the GPUs in the HPE system are one generation older. Source: Principled Technologies using data from MLCommons®.*

Comparing MLPerf performance for four-GPU servers

The Dell PowerEdge XE8640 server with four NVIDIA H100 SXM5 GPUs achieved the highest AI throughput among all four-GPU submissions in nine different categories. As Figure 2 shows, compared to the HPE ProLiant DL380a server, it scored up to 2.07 times as high on the MLPerf® benchmark.




Figure 2: Published MLPerf® results for the Dell PowerEdge XE8640 and HPE ProLiant DL380a Gen11 as of 11/29/23. The Dell system uses the NVIDIA H100 SXM form factor, while the HPE system uses the less powerful PCIe form factor. Source: Principled Technologies using data from MLCommons®.*

In brief: Additional contributors to AI success

AI models require more than high performance servers for success. You must also consider storage for unstructured data and professional services to plan, prepare, deploy, and manage your AI solution. The Dell portfolio offers storage for AI datasets with the PowerScale series for file storage and Elastic Cloud Storage or ObjectScale storage for object storage.

Organizations can reap the advantages of Dell's professional and consulting services for AI, which offers some services that HPE does not, such as data preparation. The Dell portfolio also includes Validated Designs for AI, which takes the guesswork out of designing and deploying AI.

*Verified MLPerf. score of v3.1 Inference Closed. Retrieved from <https://mlcommons.org/benchmarks/inference-datacenter/> 5 December 2023, entries 3.1-0069, 3.1-0085, 3.1-0067, and 3.1-0084. The MLPerf. name and logo are registered and unregistered trademarks of MLCommons. Association in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited. See www.mlcommons.org for more information.

Read the report at <https://facts.pt/zPmSx4c>



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the report.