

# NVIDIA and Storage Industry Leaders Unveil New Class of Enterprise Infrastructure for the Age of AI

## Storage Providers Build Infrastructure Featuring AI Query Agents That Tap NVIDIA Computing, Networking and Software to Reason and Rapidly Generate Accurate Responses to Complex Queries

**GTC**—NVIDIA today announced the [NVIDIA AI Data Platform](#), a customizable reference design that leading providers are using to build a new class of AI infrastructure for demanding AI inference workloads: enterprise storage platforms with AI query agents fueled by NVIDIA accelerated computing, networking and software.

Using the NVIDIA AI Data Platform, [NVIDIA-Certified Storage](#) providers can build infrastructure to speed AI reasoning workloads with specialized AI query agents. These agents help businesses generate insights from data in near real time, using [NVIDIA AI Enterprise](#) software — including [NVIDIA NIM™](#) microservices for the new [NVIDIA Llama Nemotron models](#) with reasoning capabilities — as well as the new [NVIDIA AI-Q Blueprint](#).

Storage providers can optimize their infrastructure to power these agents with [NVIDIA Blackwell GPUs](#), [NVIDIA BlueField® DPUs](#), [NVIDIA Spectrum-X™ networking](#) and the [NVIDIA Dynamo](#) open-source inference library.

Leading data platform and storage providers — including [DDN](#), [Dell Technologies](#), [Hewlett Packard Enterprise](#), [Hitachi Vantara](#), [IBM](#), [NetApp](#), [Nutanix](#), [Pure Storage](#), [VAST Data](#) and [WEKA](#) — are collaborating with NVIDIA to create customized AI data platforms that can harness enterprise data to reason and respond to complex queries.

“Data is the raw material powering industries in the age of AI,” said Jensen Huang, founder and CEO of NVIDIA. “With the world’s storage leaders, we’re building a new class of enterprise infrastructure that companies need to deploy and scale agentic AI across hybrid data centers.”

### NVIDIA AI Data Platform Adds Accelerated Computing and AI to Storage

The NVIDIA AI Data Platform brings accelerated computing and AI to the millions of businesses using enterprise storage for the data that drives their company.

NVIDIA Blackwell GPUs, BlueField DPUs and Spectrum-X networking provide an accelerated engine to speed AI query agent access to data stored on enterprise systems. BlueField DPUs deliver up to 1.6x higher performance than CPU-based storage while reducing power consumption by up to 50%, providing more than 3x higher performance per watt. Spectrum-X accelerates AI storage traffic up to 48% compared with traditional Ethernet by applying adaptive routing and congestion control.

AI Data Platform storage infrastructure uses the [NVIDIA AI-Q Blueprint](#) for developing agentic systems that can reason and connect to enterprise data. AI-Q taps into [NVIDIA NeMo Retriever™ microservices](#) to accelerate data extraction and retrieval by up to 15x on NVIDIA GPUs.

AI query agents built with the [AI-Q Blueprint](#) connect to data during inference to provide more accurate, context-aware responses. They can access large-scale data quickly and process various data types, including structured, semi-structured and unstructured data from multiple sources, including text, PDF, images and video.

### Storage Industry Leaders Building AI Data Platforms With NVIDIA

NVIDIA-Certified Storage partners are collaborating with NVIDIA to build custom AI data platforms.

- **DDN** is architecting AI Data Platform capabilities into its DDN Infinia AI platform.
- **Dell** is creating AI data platforms for its family of Dell PowerScale and Project Lightning solutions.
- **Hewlett Packard Enterprise** is infusing AI Data Platform capabilities into HPE Private Cloud for AI, HPE Data Fabric, HPE Alletra Storage MP and HPE GreenLake for File Storage.
- **Hitachi Vantara** is bringing AI Data Platform into the Hitachi IQ ecosystem, helping customers innovate with storage systems and data offerings that drive tangible AI outcomes.
- **IBM** is integrating AI Data Platform as part of its content-aware storage capability with IBM Fusion and IBM Storage Scale technology to accelerate retrieval-augmented generation applications.
- **NetApp** is advancing enterprise storage for agentic AI with the NetApp AIPod solution built with AI Data Platform.
- **Nutanix** Cloud Platform with Nutanix Unified Storage will integrate with the NVIDIA AI Data Platform and enable inferencing and agentic workflows deployed across edge, data center and public cloud.
- **Pure Storage** will deliver AI Data Platform capabilities with Pure Storage FlashBlade.
- **VAST Data** is working with AI Data Platform to curate real-time insights with VAST InsightEngine.
- **WEKA** Data Platform software integrates with NVIDIA GPUs, DPUs and networking to optimize data access for agentic

AI reasoning and insights and deliver a high-performance storage foundation that accelerates AI inference and token processing workloads.

NVIDIA-Certified Storage providers are planning to offer solutions created with the NVIDIA AI Data platform starting this month.

Learn more by watching the [NVIDIA GTC keynote](#) and [register for sessions](#) from NVIDIA and industry leaders at the show, which runs through March 21

#### **About NVIDIA**

[NVIDIA](#) (NASDAQ: NVDA) is the world leader in accelerated computing.

Certain statements in this press release including, but not limited to, statements as to: the benefits, impact, availability, and performance of NVIDIA's products, services, and technologies; NVIDIA's collaboration with third parties; third parties adopting or offering NVIDIA's products and technologies; and with the world's storage leaders, NVIDIA building a new class of enterprise infrastructure that companies need to deploy and scale agentic AI across hybrid data centers are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements above are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

© 2025 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, BlueField, NVIDIA NeMo, NVIDIA NIM and NVIDIA Spectrum-X are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Alex Shapiro  
Enterprise Networking  
1-415-608-5044  
[ashapiro@nvidia.com](mailto:ashapiro@nvidia.com)