

Palantir and NVIDIA Team Up to Operationalize AI — Turning Enterprise Data Into Dynamic Decision Intelligence

News Summary

- Palantir is integrating NVIDIA accelerated computing, NVIDIA CUDA-X libraries and open-source NVIDIA Nemotron models into its Ontology framework at the core of the Palantir AI Platform.
- Lowe's is pioneering operational AI for its supply chain logistics with Palantir and NVIDIA.

GTC Washington, D.C.—NVIDIA today announced a collaboration with [Palantir](#) Technologies Inc. to build a first-of-its-kind integrated technology stack for operational AI — including analytics capabilities, reference workflows, automation features and customizable, specialized AI agents — to accelerate and optimize complex enterprise and government systems.

[Palantir Ontology](#), at the core of the [Palantir AI Platform](#) (AIP), will integrate NVIDIA GPU-accelerated data processing and route optimization libraries, open models and accelerated computing. This combination of Ontology and NVIDIA AI will support customers by providing the advanced, context-aware reasoning necessary for operational AI.

Enterprises using the customizable technology stack will be able to tap into their data to power domain-specific automations and AI agents for the sophisticated operating environments of retailers, healthcare providers, financial services and the public sector.

“Palantir and NVIDIA share a vision: to put AI into action, turning enterprise data into decision intelligence,” said Jensen Huang, founder and CEO of NVIDIA. “By combining Palantir’s powerful AI-driven platform with NVIDIA CUDA-X accelerated computing and Nemotron open AI models, we’re creating a next-generation engine to fuel AI-specialized applications and agents that run the world’s most complex industrial and operational pipelines.”

“Palantir is focused on deploying AI that delivers immediate, asymmetric value to our customers,” said Alex Karp, cofounder and CEO of Palantir Technologies. “We are proud to partner with NVIDIA to fuse our AI-driven decision intelligence systems with the world’s most advanced AI infrastructure.”

Lowe’s Pioneers AI-Driven Logistics With Palantir and NVIDIA

Lowe’s, among the first to tap this integrated technology stack from Palantir and NVIDIA, is creating a digital replica of its global supply chain network to enable dynamic and continuous AI optimization. This technology can support supply chain agility while boosting cost savings and customer satisfaction.

“Modern supply chains are incredibly complex, dynamic systems, and AI will be critical to helping Lowe’s adapt and optimize quickly amid constantly changing conditions,” said Seemantini Godbole, executive vice president and chief digital and information officer at Lowe’s. “Even small shifts in demand can create ripple effects across the global network. By combining Palantir technologies with NVIDIA AI, Lowe’s is reimagining retail logistics, enabling us to serve customers better every day.”

Advancing Operational Intelligence

Palantir AIP workloads run in the most complex compliance domains and require the highest standards of privacy and data security. The Ontology at the heart of AIP creates a digital replica of an organization by organizing complex data and logic into interconnected virtual objects, links and actions that represent real-world concepts and their relationships.

Together, this provides enterprises with an intelligent, AI-enabled operating system that drives efficiency through business process automation.

To advance enterprise intelligence for the era of AI, NVIDIA data processing, AI software, open models and accelerated computing are now natively integrated with and available through Ontology and AIP. Customers can use [NVIDIA CUDA-X](#)[™] data science libraries for data processing, paired with NVIDIA accelerated computing, via Ontology to drive real-time, AI-driven decision-making for complex, business-critical workflows.

The [NVIDIA AI Enterprise](#) platform, including [NVIDIA cuOpt](#)[™] decision optimization software, will enable enterprises to use AI for dynamic supply-chain management.

[NVIDIA Nemotron](#)[™] reasoning and [NVIDIA NeMo Retriever](#)[™] open models will enable enterprises to rapidly build AI agents informed by Ontology.

NVIDIA and Palantir are also working to bring the [NVIDIA Blackwell](#) architecture to Palantir AIP. This will accelerate the end-

to-end AI pipeline, from data processing and analytics to model development and fine-tuning to production AI, using long-thinking, reasoning agents. Enterprises will be able to run AIP in NVIDIA AI factories for optimized acceleration.

Palantir AIP will also be supported in the new [NVIDIA AI Factory for Government reference design](#), announced separately today.

At NVIDIA GTC Washington, D.C., attendees can register to join Palantir and NVIDIA for a hands-on workshop on operationalizing AI, taking place Wednesday, Oct. 29, from 3:15–5:00 p.m. ET.

Learn more about how NVIDIA and partners are advancing AI innovation in the U.S. by watching the [NVIDIA GTC Washington, D.C., keynote by Huang](#).

About NVIDIA

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

Certain statements in this press release including, but not limited to, statements as to: by combining Palantir’s powerful AI-driven platform with NVIDIA CUDA-X accelerated computing and Nemotron open AI models, NVIDIA creating a next-generation engine to fuel AI specialized applications and agents that run the world’s most complex industrial and operational pipelines; the benefits, impact, performance, and availability of NVIDIA’s products, services, and technologies; expectations with respect to NVIDIA’s third party arrangements, including with its collaborators and partners; expectations with respect to technology developments; and other statements that are not historical facts are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which are subject to the “safe harbor” created by those sections based on management’s beliefs and assumptions and on information currently available to management and 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 and political conditions; NVIDIA’s reliance on third parties to manufacture, assemble, package and test NVIDIA’s products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA’s existing product and technologies; market acceptance of NVIDIA’s products or NVIDIA’s partners’ products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA’s products or technologies when integrated into systems; and changes in applicable laws and regulations, 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, CUDA-X, cuOpt, NVIDIA NeMo Retriever and NVIDIA Nemotron 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.

Shannon McPhee
NVIDIA Corporation
+1-310-920-9642
smcphee@nvidia.com