

NVIDIA Announces Open Physical AI Data Factory Blueprint to Accelerate Robotics, Vision AI Agents and Autonomous Vehicle Development

News Summary:

- Blueprint enables massive-scale data processing and curation, synthetic data generation, reinforcement learning and evaluation of physical AI models for vision AI agents, robotics and autonomous vehicles.
- Cloud service providers including Microsoft Azure and Nebius provide the blueprint to transform world-scale compute into agent-driven turnkey data production engines.
- Leading physical AI developers FieldAI, Hexagon Robotics, Linker Vision, Milestone Systems, Skild AI, Uber and Teradyne Robotics are using the blueprint to accelerate robotics, vision AI agents and autonomous vehicle development.

GTC--NVIDIA today announced the NVIDIA Physical AI Data Factory Blueprint, an open reference architecture that unifies and automates how training data is generated, augmented and evaluated, reducing the costs, time and complexity of training physical AI systems at scale.

The blueprint enables developers to use NVIDIA Cosmos™ open world foundation models and leading coding agents to transform limited training data into large, diverse datasets -- including rare edge cases and long-tail scenarios that are expensive, time-consuming and often impractical to capture in the real world.

NVIDIA is collaborating with [Microsoft Azure](#) and [Nebius](#) to integrate the open blueprint with their cloud infrastructure and services, enabling developers to turn accelerated computing power into high-volume training data. Leading physical AI developers [FieldAI](#), [Hexagon Robotics](#), [Linker Vision](#), [Milestone Systems](#), [RoboForce](#), Skild AI, Teradyne Robotics and Uber are using the blueprint to accelerate robotics, vision AI agents and autonomous vehicle development.

"Physical AI is the next frontier of the AI revolution, where success depends on the ability to generate massive amounts of data," said Rev Lebedev, vice president of Omniverse and simulation technologies at NVIDIA. "Together with cloud leaders, we're providing a new kind of agentic engine that transforms compute into the high-quality data required to bring the next generation of autonomous systems and robots to life. In this new era, compute is data."

A Unified Engine for Physical AI Development

Physical AI follows scaling laws: Performance improves as data, compute and model capacity grow. The Physical AI Data Factory Blueprint serves as a single reference architecture that moves teams from raw data to model-ready training sets through modular, automated workflows:

- **Curate and Search:** [NVIDIA Cosmos Curator](#) processes, refines and annotates large-scale real-world and synthetic datasets.
- **Augment and Multiply:** Cosmos Transfer exponentially expands and diversifies curated data, multiplying real and simulated inputs to better capture rare and long-tail scenarios across environments and lighting conditions.
- **Evaluate and Validate:** [NVIDIA Cosmos Evaluator](#), powered by Cosmos Reason and now available on GitHub, automatically scores, verifies and filters generated data to ensure physical accuracy and training readiness.

NVIDIA is using the Physical AI Data Factory Blueprint to train and evaluate NVIDIA Alpamayo, the world's first open reasoning-based vision language action models for long-tail autonomous driving. Skild AI is applying the blueprint to advance general-purpose robot foundation models, while Uber is using it to accelerate autonomous vehicle development.

Agent Driven Orchestration at Scale

Many robotics developers are not equipped to stand up and manage the complex AI infrastructure required to generate data at scale.

[NVIDIA OSMO](#), an open source orchestration framework, unifies and manages these workflows across compute environments, reducing manual tasks so developers can focus on building their models.

OSMO now integrates with leading coding agents such as Claude Code, OpenAI Codex and Cursor, enabling AI-native operations where agents proactively manage resources, resolve bottlenecks and accelerate model delivery at scale.

Powering the Global Physical AI Ecosystem

Cloud service providers play a critical role in providing the accelerated AI infrastructure, machine learning operations and orchestration services developers need to build and deploy physical AI at scale.

Microsoft Azure is integrating the Physical AI Data Factory Blueprint into an open physical AI toolchain, now available on GitHub. The blueprint offers integration with Azure services -- including Azure IoT Operations, Microsoft Fabric, Real-Time Intelligence and Microsoft Foundry -- to provide enterprise-grade, agent-driven workflows for training and validating physical AI systems quickly and at scale.

FieldAI, [Hexagon Robotics](#), Linker Vision and Teradyne Robotics are among the first to test the Azure physical AI toolchain for accelerating and scaling data generation, augmentation and evaluation across their perception, mobility and reinforcement learning pipelines.

Nebius has integrated OSMO into its AI Cloud, enabling developers to use the blueprint to deploy production-ready data pipelines tailored to their needs. Nebius's infrastructure powers the physical AI stack end to end, blending NVIDIA RTX PRO™ 6000 Blackwell Server Edition GPUs with ultrafast object storage, native data management and labeling, serverless execution and built-in managed inference.

Early users Milestone Systems, Voxel51 and RoboForce are harnessing the blueprint on Nebius infrastructure to accelerate model development for video analytics AI agents, autonomous vehicles and industrial humanoid robots.

The NVIDIA Physical AI Data Factory Blueprint is expected to be available on GitHub in April.

Watch the [GTC keynote](#) from NVIDIA founder and CEO Jensen Huang and explore [sessions](#).

Feature image courtesy of Milestone Systems (left), FieldAI (middle) and Hexagon Robotics and BMW (right).

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: together with cloud leaders, NVIDIA providing a new kind of agentic engine that transforms compute into the high-quality data required to bring the next generation of autonomous systems and robots to life; 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; expectations with respect to AI and related industries; 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; NVIDIA's ability to realize the potential benefits of business investments or acquisitions; 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.

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA Cosmos, and NVIDIA RTX PRO 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.

Quentin Nolibois
Corporate Communications
NVIDIA Corporation
press@nvidia.com