



NVIDIA DSX Gives Infrastructure Builders the Playbook for AI Factories

News Summary:

- Engineered from the ground up for AI factories, the NVIDIA DSX platform defines how next-generation infrastructure is designed, built and operated -- driving lowest token cost and accelerating time to first production across NVIDIA chips, systems, software, facilities and partner technologies.
- New DSX MaxLPS software enables AI infrastructure and factories to deliver the lowest token cost by maximizing token performance per megawatt.
- Open source, modular DSX OS software brings together lifecycle management, runtime consistency and health automation, resiliency, multi-tenant AI factory operations and platform services.
- Industry-leading manufacturers are building NVIDIA DSX-ready systems supporting the buildout of AI factories with extreme codesign.
- Growing DSX partnerships across every layer of the stack accelerate the design, deployment and operation of AI factories worldwide.

NVIDIA GTC Taipei--NVIDIA today announced the [NVIDIA DSX™](#) platform, which gives infrastructure builders a complete playbook to create AI factories.

NVIDIA DSX brings together open source, modular software libraries, application programming interfaces, reference designs, NVIDIA accelerated computing platforms and partner technologies into a common, codesigned platform for AI factory design, deployment and operations.

NVIDIA is the only company that builds [the full AI factory](#). By aligning every layer of the stack across compute, software, facilities and partner technologies, DSX provides infrastructure builders with a proven framework to design, deploy and operate AI factories at scale.

The integrated platform accelerates deployment, improves operational reliability and resiliency at scale and enables a broad ecosystem of solutions designed to turn every megawatt into more intelligence at the lowest token cost.

"We're not just shipping chips -- we're giving every infrastructure builder a complete playbook to build AI factories," said Jensen Huang, founder and CEO of NVIDIA. "With the DSX platform, you can simulate the entire factory before you spend a dollar, validate performance before a single rack is installed and operate with the kind of reliability that production AI demands."

DSX Platform Elements

DSX now spans the full stack, from silicon and systems to infrastructure software, facilities and partner technologies. The latest additions to the platform include new open source software:

- **DSX MaxLPS™**: A suite of technologies to maximize token performance per megawatt within a fixed power budget, enabling lowest token cost for AI factories. Combining 45-degrees-Celsius liquid cooling with in-rack technologies that optimize performance per watt, DSX MaxLPS lets operators run up to 40% more GPUs at their most energy-efficient operating point with minimal impact on workload performance.
- **DSX OS™**: Open source, modular software purpose-built for AI factory operations, providing lifecycle management, intelligence scheduling, runtime consistency, health automation, resiliency, multi-tenant operations and platform services.

DSX MaxLPS and [DSX OS](#) join an existing set of features under the DSX platform:

- **DSX Reference Design**: Generation-specific, validated AI factory architectures covering compute, networking, storage, hardware cluster design and facilities infrastructure -- including power, cooling and controls, as well as civil, structural and architectural design.
- **DSX Sim™**: High-fidelity simulation layer for the AI factory lifecycle, helping NVIDIA, partners and customers to model, validate and optimize infrastructure decisions from planning and design through deployment and operations.
- **DSX Flex**: Connects AI factories to power-grid services, enabling dynamic workload adaptation to grid signals such as load shedding, demand response and pricing events, and orchestrating renewable and hybrid power across utility, onsite renewables and storage.
- **DSX Exchange™**: Enables scalable, secure integration of compute, network, energy, power and cooling plant signals between IT, operational technology and operations agents.

Growing DSX Ecosystem

NVIDIA is partnering with industry-leading Taiwan system manufacturers to [expand the DSX ecosystem](#), supporting the buildout of AI factories with extreme codesign at their core.

NVIDIA cloud partners CoreWeave, Crusoe, Firmus, IREN, Lambda, Nebius, Nscale and Yotta Data Services are deploying core components of the DSX platform stack -- DSX Sim, DSX MaxLPS and DSX OS -- to reduce risk, improve GPU utilization and bring AI cloud capacity online faster.

Dell Technologies, HPE, Lenovo and Supermicro together with [ASUS](#), Foxconn, GIGABYTE, Pegatron, Quanta Cloud Technology (QCT), Wistron and Wiwynn are building NVIDIA DSX-ready systems and contributing simulation-ready assets that enable customers to deploy complete, full-stack AI factory solutions at global scale.

Within the ecosystem, model-based systems engineering serves as the bridge between rack design to facility deployment, for an AI infrastructure optimized for token performance per megawatt. Quanta Cloud Technology (QCT) and Pegatron are working with [Dassault Systèmes](#) to create a live AI factory digital twin configurator to automate rack-to-facility design with increased quality and reduced workload. The adoption of DSX Sim by system manufacturers expands the [NVIDIA Omniverse DSX Blueprint](#) ecosystem, deepening integration with software partners [Cadence](#), [PTC](#) and [Siemens](#).

DSX Flex is powering a commercial, multi-megawatt pilot with [Emerald AI](#) and Silicon Valley Power to demonstrate grid-responsive AI factories that can dynamically adjust power consumption in response to utility signals while protecting AI workload performance, helping safeguard grid reliability and affordability for customers while unlocking additional power capacity to support AI growth.

Partners are adopting various DSX OS software components for lifecycle management, multi-tenancy, security, health automation, resilience and platform services. Ecosystem partners adopting DSX OS components include [Aible](#), BeyondAI, Bhashini, DCAI, Mirantis, OpenNebula Systems, Rafay, Red Hat, Sarvam, Simplismart, Spectro Cloud, Supermicro, vCluster and Vultr.

Watch Huang's [keynote](#) and learn more at [NVIDIA GTC Taipei](#).

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: expectations with respect to growth, performance, availability, and benefits of NVIDIA's products, services and technologies, and related trends and drivers; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; expectations with respect to technology developments, and related trends and drivers; projected market growth and trends; 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 products 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, DSX Exchange, DSX OS, DSX Sim, DSX MaxLPS and NVIDIA DSX are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and/or other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Kristin Uchiyama
Corporate Communications

NVIDIA Corporation
press@nvidia.com