



European Ecosystem Builds World's Most Efficient AI Factories With NVIDIA DSX

NVIDIA DSX Enables European Companies to Reduce Water Consumption for Cooling by up to 100% and Generate the Same AI Output With up to 30% Less Power

News Summary:

- NVIDIA DSX reference designs use closed-loop liquid cooling and dry coolers to enable near-zero facility cooling water consumption, with potential savings of 54 million liters a day for a 2-gigawatt deployment -- equivalent to the daily water use of 180,000 households.
- AI clouds in Europe are building and operating NVIDIA DSX-based AI factories, expanding access to accelerated computing for the continent's enterprises, developers and governments.
- NVIDIA DSX MaxLPS software helps operators use up to 30% less power to generate the same number of tokens, with CoreWeave, Firmus, Google Cloud, IREN, Lambda, Nebius, Nscale and Sharon AI planning to use the technology, while DSX Flex software turns AI factories into flexible power loads to support grid reliability and enable faster grid connection.
- NVIDIA SimReady Central, a hub for simulation-ready digital models of physical equipment, launches on Hugging Face with validated power and cooling assets, helping operators build and simulate DSX AI factory digital twins faster.
- New NVIDIA DSX OS open source software automates infrastructure lifecycle and improves health automation as well as resiliency capabilities to help operators run AI factories at scale.
- An expanding ecosystem -- including technology and industrial AI leaders AVEVA, Dassault Systèmes and Siemens, and energy and cooling providers ABB, Eaton, GE Vernova, Hitachi Energy, Schneider Electric, Siemens Energy, Trane Technologies and Vertiv -- is connecting AI factory design, power, cooling and operations to support Europe's accelerated computing buildout.

GTC Berlin--NVIDIA today announced that Europe's leading ecosystem partners are adopting the [NVIDIA DSX™ AI Factory Platform](#) to build AI factories that minimize water consumption for cooling and use less power for the same AI output. The platform also helps partners [bring new capacity online faster](#) and adapt workload power consumption to support grid reliability.

NVIDIA DSX is a full-stack AI factory reference architecture, co-innovated with the world's leading infrastructure companies. Compute, networking, storage, software, power and cooling are codesigned as one system -- optimized from design to operation.

AI clouds and leading design, construction, energy and cooling providers are using DSX to build and operate infrastructure serving Europe's enterprises, developers and governments.

"Europe's technology leadership depends on how efficiently we turn power and infrastructure into intelligence," said Warren Barkley, vice president of DSX product management at NVIDIA. "NVIDIA DSX brings together the technologies and ecosystem expertise operators need to build responsibly, operate efficiently at scale and bring intelligence to the market faster."

The NVIDIA DSX platform includes [DSX MaxLPS™](#) for dynamic power allocation, DSX Flex for coordinating workload power consumption with grid conditions, DSX Reference Design for validated infrastructure architectures, [DSX Sim™](#) for simulation, [DSX OS™](#) for infrastructure management and DSX Exchange™ for connecting computing and facility systems. These technologies help operators design, deploy and run AI factories that save valuable resources.

NVIDIA DSX Rethinks Cooling

Conventional cooling towers remove heat by evaporating water, requiring a continual supply of replacement water. NVIDIA DSX reference designs connect liquid-cooled computing with dry-cooler-based heat rejection that avoids evaporation: Cooling liquid absorbs heat from computing equipment and circulates in a closed loop, while outdoor dry coolers, similar to large radiators, release that heat into the air without evaporating water.

NVIDIA Vera Rubin infrastructure supports 100% liquid cooling with coolant temperatures of up to 45 degrees Celsius. With dry-cooler-based designs, this architecture can reduce facility cooling water consumption to near zero. At 2-gigawatt scale, NVIDIA DSX can save 54 million liters of water a day compared with conventional cooling-tower-based systems -- equivalent to the daily water use of 180,000 households. Local climate determines whether supplemental mechanical cooling is needed during the hottest periods.

Delivering the Same AI Output With Less Power

Power availability defines how much capacity operators can bring online. DSX MaxLPS and [DSX Flex](#) address this constraint.

DSX MaxLPS dynamically allocates power within an AI factory's provisioned budget, helping operators generate the same number of tokens with up to 30% less power. Operators can also use this efficiency to support more GPU capacity and higher throughput within existing power budgets. CoreWeave, Firmus, Google Cloud, IREN, Lambda, Nebius, Nscale and Sharon AI are exploring the software.

In a [proof of concept](#) across NVIDIA HGX™ B200 systems with DSX MaxLPS, Lambda demonstrated 24% more cluster token throughput and 19% more GPU capacity within a fixed power budget.

Coordinating AI Factory Power Demand With the Grid

DSX Flex helps AI factories become flexible assets for the power grid, enabling operators to adapt workload power consumption to utility needs, demand-response events and energy prices. It can also coordinate electricity from the grid, onsite renewables and energy storage, helping AI factories operate as flexible power loads that support grid reliability and enable faster grid connection.

In an early trial, Emerald AI, EPRI, Nebius, the U.K. National Grid and NVIDIA demonstrated reductions of up to 40% in the cluster's power draw on an NVIDIA Blackwell Ultra GPU cluster in response to simulated grid events, while critical workloads continued running.

Nscale Advances AI Factory Efficiency With NVIDIA DSX

Using DSX MaxLPS, DSX Flex and DSX Sim, Nscale is demonstrating how the NVIDIA DSX platform connects power optimization, grid-responsive operations and design.

Based on an evaluation run in Nscale AI factory in Iceland, using DSX MaxLPS across four NVIDIA GB300 NVL72 racks, Nscale demonstrated 49% higher inference throughput and 37% more managed GPU capacity within the same 264-kilowatt provisioned power budget. The results demonstrate how dynamic power allocation can unlock additional capacity from existing infrastructure.

Nscale is also exploring DSX Flex to align workload power consumption with grid conditions.

"Europe's AI ambitions demand that we make every megawatt count," said Josh Payne, founder and CEO of Nscale. "With NVIDIA DSX, we're building AI factories that deliver more compute from available power, respond to grid conditions and move faster from design to deployment."

For its facilities in Portugal and Iceland, Nscale is using DSX Sim to create [digital twins](#) that help evaluate power, cooling and infrastructure configurations before deployment. By validating designs virtually, Nscale can identify potential issues earlier, reduce deployment risk and accelerate the path from planning to productive capacity.

Accelerating AI Factory Design With DSX Sim and SimReady Central Library

NVIDIA is launching [SimReady Central Library](#) with validated power and cooling equipment assets from partners like Vertiv, accelerating the creation of DSX AI factory digital twins.

Used with DSX Sim, ecosystem partners including [PTC](#) are helping create these simulation-ready assets with computer-aided design to enable SimReady workflows that let builders model infrastructure configurations and evaluate power and cooling decisions before equipment is installed. Operators can validate designs earlier, reduce deployment risk and move faster from planning to construction.

Driving Production Readiness With NVIDIA DSX OS

To help operators turn installed infrastructure into productive capacity, NVIDIA is expanding its modular, open source DSX OS portfolio with three components in developer preview:

- NVIDIA Config Manager automates AI factory network configuration.
- NVIDIA Rack Management Service provides rack-level power, inventory, firmware and switch management.
- NVIDIA Cluster Readiness Engine runs distributed workloads to validate GPU clusters before production workloads arrive and isolate failures to specific nodes.

Together with NVIDIA Infrastructure Controller, NVIDIA AI Cluster Runtime and NVSentinel, these components help operators automate infrastructure lifecycle management, maintain consistent runtime environments and strengthen health validation, fault attribution and recovery.

Cloud partners including Aleria, Cassava, CoreWeave, Crusoe, Firmus, Firebird, G42, GMI Cloud, IREN, Lambda, Nebius, Nscale, Radiant, Sharon AI and Volta are exploring the latest components of NVIDIA DSX OS to automate infrastructure lifecycle management, and improve resiliency and reliability.

Connecting Design, Power and Cooling Across Europe's AI Factory Ecosystem

NVIDIA's industrial software, power and cooling partners are advancing Europe's AI factory buildout with DSX reference designs, lifecycle digital twins and qualified equipment that connect facility infrastructure with accelerated computing.

- Schneider Electric has developed power and cooling reference designs for NVIDIA Vera Rubin infrastructure and is working with AVEVA on [digital twins spanning](#) infrastructure design, construction and operations.
- Procore is helping manage AI factory builds with its construction management platform, integrated with DSX Sim.
- [Siemens](#) is using DSX Sim to simulate power, cooling and acoustics, while [Dassault Systèmes](#) and Vertiv are developing digital twin models of Vertiv's OneCore infrastructure, enabling simulation and optimization through model-based systems engineering.
- ABB is developing SimReady assets with NVIDIA DSX Sim, enabling engineers to evaluate source-to-rack architectures before construction while advancing 800 VDC development through its newly announced Infinitus Direct Current portfolio.

Partners are also delivering coordinated power and cooling infrastructure aligned with DSX reference designs. Trane Technologies and Eaton recently collaborated with NVIDIA to help boost data center efficiency and cut installation costs with an integrated design built for the NVIDIA Vera Rubin DSX AI Factory Reference Design, bringing together Eaton's Beam platform and power infrastructure with Trane Continuum cooling and coordinated controls in a grid-to-chip architecture.

Trane Technologies has published seven thermal management system reference designs, including three 250-megawatt DSX-based cooling reference designs to support thermal infrastructure planning at scale.

Through the [NVIDIA DSX Ready](#) program, builders can select equipment qualified under applicable DSX categories. Trane Technologies' LiquidStack GigaModular and Vertiv's 2.3-megawatt CoolChip -- both coolant distribution units -- have qualified under their applicable categories, while Hitachi Energy's battery energy storage system has earned NVIDIA DSX Ready BESS qualification.

For energy supply and grid integration, GE Vernova, Hitachi Energy and Siemens Energy are collaborating with NVIDIA on power infrastructure and digital twin initiatives that connect energy-system engineering and grid planning based on the facility's requirements. Hitachi Energy is developing an integration between its energy management system and DSX Flex to optimize grid, generation and storage resources with workload demand.

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

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, and related deployments and timing; 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 NVIDIA'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, and other NVIDIA product and service names are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries and regions. Other company,

product, and service names may be trademarks of the respective companies with which they are associated.

Kristin Uchiyama
Corporate Communications
NVIDIA Corporation
press@nvidia.com