

NVIDIA Expands Open Source CUDA-Q Platform for Fault-Tolerant Quantum Computing

News Summary:

- The NVIDIA CUDA-Q platform now includes the CUDA-Q Logical orchestration layer, giving the quantum industry an open, programmable way to design and test fault-tolerant quantum computing applications in areas such as drug discovery, financial modeling and materials development.
- Fermilab used CUDA-Q Logical to accelerate fault-tolerant architecture development from five months to three weeks -- a 7x speedup.
- QUOPS is a benchmark for the future of quantum computing, measuring the readiness of fault-tolerant quantum hardware for practical applications. Developed by Sandia National Laboratories, QUOPS is now available in NVIDIA CUDA-Q.

NVIDIA today announced an expansion of the [NVIDIA CUDA-Q™](#) open source platform with CUDA-Q Logical, an orchestration layer that provides a programmable, verifiable approach to developing useful applications for fault-tolerant quantum computers.

Fault-tolerant quantum processors with logical qubits are essential for useful quantum computing, allowing systems to overcome the errors inherent in physical qubits so they can execute the larger computations needed for practical applications such as drug discovery, financial modeling and materials development.

However, developing applications for such systems is a tedious, time-consuming codesign challenge: A change to an algorithm, error-correction code, hardware architecture or other QPU component can significantly change the resources needed to run any given application.

With CUDA-Q Logical, researchers can now design and orchestrate the many components required for fault-tolerant quantum computing applications, easily switching between options to identify optimal system configurations for performance with logical qubits.

"Quantum computing is maturing into an era of logical qubits, and researchers need an open, customizable platform capable of representing all aspects of a fault-tolerant system," said Timothy Costa, vice president and general manager of quantum at NVIDIA. "The addition of CUDA-Q Logical provides power and flexibility to explore fully integrated, co-optimized systems regardless of qubit type and architecture -- drastically shortening the timeline to useful quantum-GPU supercomputing."

CUDA-Q Logical is already being used by QPU makers and labs including Fermi National Accelerator Laboratory, [Infleqtion](#), IQM Quantum Computers, QCDesign [Quantum Motion](#) and Sandia National Laboratories.

Using CUDA-Q Logical, [Iceberg Quantum](#) modeled its fault-tolerant architecture for Diraq's qubits, showing how 1,000 logical qubits can be created with just 150,000 physical qubits, roughly 10x fewer than [Diraq's](#) previous estimates. CUDA-Q Logical gave an unprecedented way to rapidly assess potential implementations of Iceberg's architecture, helping dramatically reduce the hardware needed for useful, fault-tolerant quantum computing.

Fermilab Uses CUDA-Q Logical to Explore Fault-Tolerant Applications

In early work, Fermilab researchers used CUDA-Q Logical to validate prior results and evaluate physical qubits, runtimes and other resource requirements across different error-correction approaches and quantum hardware.

Fermilab transformed complicated fault-tolerant system designs into a repeatable and verifiable computational workflow, accelerating fault-tolerant algorithm development from five months to three weeks, a 7x speedup.

"Fault-tolerant quantum computing is the path to unlocking new scientific discovery, but getting there will require researchers to codesign algorithm, error correction, architectures and hardware together," said Anna Grassellino, chief technology officer at Fermilab and director of the Superconducting Quantum Materials and Systems Center. "Using CUDA-Q Logical, our team explored combinations of these resources in just three weeks, compared with what would have typically taken about five months of building specialized infrastructure."

Sandia Creates Benchmark for Future of Quantum Computing

QUOPS is a new, independent cross-platform benchmark developed by [Sandia National Laboratories](#) that measures the progress quantum computing systems are making toward utility-scale applications.

With Sandia's QUOPS, the industry now has access to a hardware-agnostic, open method to benchmark progress toward shared goals for fault-tolerant quantum computers. Historically, progress in quantum computing has been measured primarily through advances in physical qubits -- increasing qubit counts, improving fidelity and extending coherence.

"At Sandia, we can't wait to see fault-tolerant quantum computers helping to solve problems of national importance for the Department of Energy and the United States," said Timothy Proctor, co-director of Sandia National Laboratories' Quantum Performance Laboratory. "Our mission right now is to bring that about sooner, by accelerating our industry partners' progress. To do that, we, and other quantum computing stakeholders, have to be able to track and forecast the growth of quantum computer abilities. We created QUOPS to do exactly that, and we're excited to see it used by quantum computing vendors and customers."

Sandia shared early results in a preprint posted ahead of IEEE Quantum Week, reporting initial QUOPS benchmarks for QPUs from Google, IBM and [Quantinuum](#). A QUOPS reference implementation is available in NVIDIA CUDA-Q.

Expanding the Quantum-GPU Supercomputing Platform

The ecosystem is also adopting NVIDIA's larger lineup of quantum computing technologies for turning qubits into useful quantum-GPU supercomputers.

Diraq used NVIDIA Ising, the world's first family of open models for building and deploying AI for useful quantum computing, to calibrate its silicon-based qubit processor.

Companies are integrating their quantum processors with GPU supercomputing using [NVIDIA NVQLink™](#), the open system architecture for tightly coupling quantum processors with GPU supercomputers. With NVQLink:

- [Anyon Computing](#) developed a new quantum control system.
- [Quandela](#) developed a QPU-GPU architecture.
- Quantum Machines ran a demo of integrating state-of-the-art supercomputing resources with qubits at the Israeli Quantum Computing Center.

The NVIDIA CUDA-Q open development platform continues to be adopted across the quantum computing ecosystem:

- [BlueQubit](#) launched the Quantum Flywheel grant program for researchers, to give them access to NVIDIA accelerated computing through CUDA-Q.
- [Qedma Quantum Computing](#) and [QCentroid](#) each integrated their technology with CUDA-Q to support quantum error correction and mitigation, as well as the deployment of quantum applications.

The quantum ecosystem is also working with NVIDIA to develop new and improved quantum applications:

- IonQ announced progress on DQAOA-GPT, a quantum generative AI framework, using NVIDIA accelerated computing.
- MITRE published work on building GPU-accelerated digital twins of quantum sensors.
- [Phasecraft](#) is using NVIDIA cuQuantum to generate the largest-known variational quantum eigensolver emulated molecular database.
- [UCLA](#) and Caltech are advancing the design of control sequences needed to run quantum applications.

Availability

CUDA-Q Logical is now available through [GitHub](#). Explore the [QUOPS repository](#).

Learn more by reading these papers on [CUDA-Q Logical](#) and [QUOPS](#).

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 quantum computing; 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 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.

© 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.

Aaron Groff
Corporate Communications
NVIDIA Corporation
press@nvidia.com