QUANTUM SIMULATION

Edge Quantum Simulators

Current-source review, evidence-bounded adoption decision, explicit limitations, reproducibility controls, and preserved source research note.

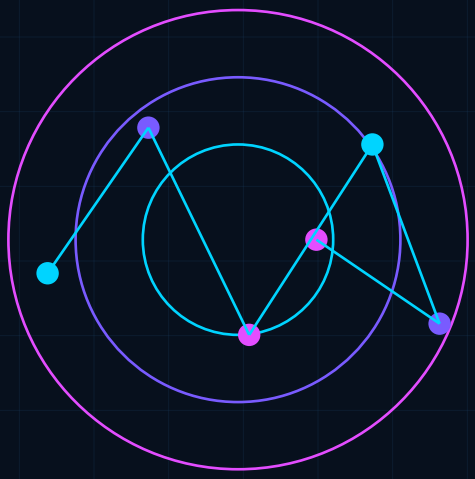

PERMANENT ID	EVIDENCE	ADOPTION	FRESHNESS
RES-021	A-	RESEARCH TOOLING ONLY	STABLE WATCH

Freshness review: 2026-09-17 / Next scheduled review: 2027-03-16

RESEARCH CONTROL SURFACE

Required Research Fields

Each field remains independently reviewable; evidence strength does not imply production authority.

EVIDENCE GRADE — **A-**

Strength of the retained source set and technical basis.

SOURCE AUTHORITY — **2 registered authorities**

Qiskit Aer Documentation; NVIDIA cuQuantum Documentation

FRESHNESS — **STABLE WATCH**

Reviewed 2026-09-17; dynamic sources require event-driven revalidation.

CONFLICTS — **VISIBLE / NOT SILENT**

Differences between specification, vendor behavior, research claims, and reproduced evidence must remain explicit.

LIMITATIONS — **EXPLICIT**

No certification, procurement approval, legal conclusion, or unbounded production claim is created by this report.

REPRODUCIBILITY — **REQUIRED**

Material claims require exact versions, representative data, failure cases, artifacts, and repeatable acceptance criteria.

ADOPTION POSTURE — **RESEARCH TOOLING ONLY**

Use edge or local quantum simulation for education, algorithm prototyping, and small reproducible experiments. Do not present classical simulation as quantum advantage.

REVIEW TRIGGER — **2027-03-16**

Scheduled at 180 days; earlier on material source, vulnerability, implementation, or contradictory-evidence change.

SOURCE AUTHORITY / FRESHNESS

Authority Register and Freshness Delta

Primary-source lineage is preserved; current-source changes are separated from unperformed implementation claims.

AUTHORITY 01

Qiskit Aer Documentation

https://qiskit.github.io/qiskit-aer/

AUTHORITY 02

NVIDIA cuQuantum Documentation

https://docs.nvidia.com/cuda/cuquantum/

2026-09-17 FRESHNESS DELTA

Primary-source register retained and freshness controls advanced to the current review date. No new product or laboratory performance claim is introduced without reproduced evidence.

Freshness policy: scheduled review + immediate review on source supersession, material release, vulnerability, retraction, or contradictory evidence.

Evidence Must Survive Contact With Reality

- Source authority is not implementation evidence. Product behavior, scale, security, interoperability, and operator fit remain environment-specific.
- Conflicting evidence is recorded rather than averaged away. Specification behavior, vendor implementation, preprint results, and production observations are separate evidence classes.
- Quantitative claims are informative until reproduced under documented hardware, software, data, topology, load, and failure conditions.
- Adoption posture is reversible. A stronger source set can coexist with a pilot-only or research-only recommendation when maturity, safety, economics, or operability remain unresolved.

ADOPTION / VALIDATION / REVIEW

Decision Gate and Validation Plan

ADOPTION POSTURE

RESEARCH TOOLING ONLY

Use edge or local quantum simulation for education, algorithm prototyping, and small reproducible experiments. Do not present classical simulation as quantum advantage.

1 / SOURCE	Pin exact versions, publication state, retrieved artifacts, and source hashes.
2 / PROTOTYPE	Demonstrate the core mechanism with representative data and instrumented execution.
3 / FAILURE	Exercise malformed input, dependency loss, stale state, overload, recovery, rollback, and misuse.
4 / BASELINE	Compare against an incumbent, classical, or simpler architecture using the same workload.
5 / ACCEPT	Record acceptance criteria, residual limitations, owner, expiration, and revalidation triggers.

EXPIRATION / REVIEW TRIGGER

Next scheduled review: 2027-03-16 (180-day cadence). Review sooner after material source revision, breaking implementation release, critical vulnerability, retraction, benchmark contradiction, changed workload, or changed legal/safety boundary.

8. Complete Source Research Note

SOURCE-LINEAGE NOTICE

The following material preserves the complete original source note, excluding only the conversational wrapper and outer Markdown fence. Legacy status labels and absolute language are retained for traceability and do not override the candidate status, limitations, or adoption decision in this edition.

Document ID: RES-021 Version: 1.0.0 (Definitive Registry Edition) Status: Published Research Note Classification: Public Organization: Mythos Systems (MYTHOS AI) Owner: Aaron Stovall Effective Date: 2026-07-16 Applies To: Quantum software engineers, HPC architects, and AI researchers developing variational quantum algorithms (VQE, QAOA), tensor-network simulations, and hybrid classical-quantum orchestration pipelines Companion Standards: MYTHOS-QTC-0001 (Quantum Computing Benchmark & Algorithm Design), MYTHOS-HW-0001 (AI GPU Clusters), MYTHOS-AI-0015 (Federated Learning), MYTHOS-DST-0001 (Durable Workflows)

The architecture of quantum algorithm development has failed. Organizations adopt a "wait-and-see" approach, halting quantum software development until fault-tolerant quantum computers (FTQC) are commercially available. When they do experiment, they rely on metered, high-latency cloud Quantum Processing Unit (QPU) APIs to run noisy, small-scale tests. This iterative loop is economically unviable and mathematically insufficient for developing algorithms meant for future 100+ qubit fault-tolerant machines.

The Mythos Architecture mandates "Quantum Agility." We must simulate future quantum workloads on specialized classical edge hardware today to prepare the software fabric for tomorrow.

This research note defines the architectural dynamics of Edge Quantum Simulators. It dissects the mechanics of simulating variational quantum algorithms (VQE, QAOA) using State-Vector and Tensor-Network math on classical FPGAs and TPUs. By mapping the exponential memory of quantum state spaces to the massive tensor-interconnects of TPUs, and the deterministic gate operations to custom FPGA silicon, we enable rapid, offline CI/CD testing of quantum algorithms. Understanding these dynamics is a prerequisite for achieving the verifiable quantum advantage mandated by the Mythos Quantum Standards.

To simulate quantum workloads, we must understand that modern quantum algorithms (like VQE and QAOA) are not purely quantum. They are hybrid classical-quantum loops.

1.1 Variational Quantum Eigensolver (VQE)

Used in chemistry and materials science to find the ground-state energy of a molecular Hamiltonian.

- The Loop: A classical optimizer (e.g., COBYLA, Adam) proposes a set of angles (parameters) for a quantum circuit. The quantum simulator executes the circuit and measures the expectation value (energy). The classical optimizer updates the angles to minimize the energy.
- The Simulation Bottleneck: The classical optimizer must execute the quantum circuit thousands of times per iteration. If this is done via a cloud API, the latency destroys the optimization loop.

1.2 Quantum Approximate Optimization Algorithm (QAOA)

Used for combinatorial optimization (e.g., Max-Cut, routing).

- The Loop: Similar to VQE, a classical optimizer tunes quantum gate angles to maximize an objective function.
- The Simulation Bottleneck: QAOA requires deep circuits with many entangling layers, exponentially increasing the classical compute required to simulate the quantum state.

Simulating a quantum computer on a classical machine is an act of brute-force mathematical defiance. A quantum computer utilizes superposition; a classical computer must explicitly store every possible state.

2.1 The Exponential Memory Wall

An N-qubit quantum system requires a state vector of 2^N complex amplitudes.

- The Physics: Simulating 30 qubits requires 16 GB of RAM (for FP64 complex numbers). Simulating 40 qubits requires 16 Terabytes. Simulating 50 qubits requires 16 Petabytes.
- The Flaw: Standard edge hardware (CPUs, GPUs) hits the 30-qubit wall instantly. You cannot simulate a fault-tolerant 100-qubit algorithm on a classical machine using exact state-vector math.

2.2 The Tensor Contraction Overhead

To bypass the memory wall, simulators use Tensor Networks (e.g., Matrix Product States - MPS). Instead of storing 2^N amplitudes, the state is factorized into a chain of tensors.

- The Constraint: Tensor networks work brilliantly for low-entanglement circuits (like VQE for simple molecules). However, if the circuit generates high entanglement (like QAOA at high depth), the tensor network bonds must grow, and the memory savings vanish.
- The Compute Tax: Simulating a gate in a tensor network is a tensor contraction. This is heavy, sequential matrix multiplication that starves standard GPUs.

To make edge quantum simulation viable for algorithm development, the simulation must be offloaded to specialized silicon designed for tensor math or deterministic gate routing.

3.1 TPUs for Tensor-Network Simulation

Google's Tensor Processing Units (TPUs) are uniquely suited for tensor-network quantum simulation.

- The Mechanic: TPUs utilize systolic arrays and a high-speed Inter-Core Interconnect (ICI). A 2048-chip TPU pod possesses immense distributed memory and bandwidth.
- The Impact: TPUs can simulate up to 50-60 qubits for highly entangled circuits by perfectly parallelizing the tensor contractions across the pod. For an edge environment, a localized TPU v5e or v6 pod allows sovereign, offline testing of QAOA circuits that would OOM a GPU cluster.

3.2 FPGAs for State-Vector and Gate Routing

For smaller (20-30 qubit) but ultra-low-latency simulations, Field Programmable Gate Arrays (FPGAs) are unmatched.

- The Mechanic: A quantum gate (e.g., CNOT, Hadamard) is a deterministic matrix multiplication. On an FPGA, these gates can be compiled directly into the hardware logic (LUTs and DSP slices).
- The Impact: An FPGA can execute a quantum gate in nanoseconds, with zero instruction-fetch overhead. This allows the classical VQE optimizer (running on the CPU) to request thousands of circuit evaluations per second, compressing a multi-day cloud QPU optimization loop into minutes on the edge.

The fundamental shift of edge simulation is the decoupling of software development from hardware availability.

- Current Paradigm: Algorithms are developed for NISQ (Noisy Intermediate-Scale Quantum) hardware, which is too error-prone to run deep circuits. Therefore, the algorithms are weak.
- Simulation Paradigm: Engineers develop algorithms assuming fault-tolerant, logical qubits. They simulate the circuit (including the error-correction overhead) on an edge TPU/FPGA. They optimize the mathematical logic before the physical FTQC hardware exists.

The Scaling Law: By utilizing edge simulators, organizations build a library of verified, fault-tolerant-ready quantum algorithms. The moment a physical FTQC coprocessor is deployed, the organization drops the simulation backend and swaps in the real QPU, instantly achieving quantum advantage while competitors are still learning to write the code.

To deploy edge quantum simulators in production, the Mythos Architecture enforces strict orchestration and hardware-aware compilation boundaries.

5.1 The Durable Hybrid Orchestration (Clio)

The hybrid classical-quantum loop (VQE/QAOA) is inherently fragile. If the classical optimizer crashes, the quantum state context is lost.

- The Mitigation: The classical optimization loop MUST be wrapped in a Durable Execution Runtime (Temporal/Clio). The optimizer proposes parameters, which are checkpointed. The FPGA/TPU simulator executes the circuit and returns the expectation value. If the process crashes, the runtime resumes exactly at the last parameter update.

5.2 Hardware-Aware Simulation Transpilation

A simulation compiled for a TPU (tensor contraction) is structurally different from a simulation compiled for an FPGA (state-vector).

- The Mitigation: The Mythos orchestration layer MUST utilize a unified intermediate representation (e.g., OpenQASM 3). The control plane dynamically transpiles the quantum circuit into either TPU tensor operations or FPGA gate logic based on the available edge hardware, ensuring optimal simulation performance.

5.3 The QPU Abstraction Boundary (The "Drop-In" Mandate)

The simulator MUST be architecturally indistinguishable from a physical QPU to the application layer.

- The Mitigation: The classical optimizer communicates with the "quantum device" via a standardized gRPC API. Today, that API routes to the TPU/FPGA simulator. Tomorrow, when a fault-tolerant IonQ or IBM QPU is plugged into the network, the routing table changes, and the real QPU executes the exact same workload. Zero application code is rewritten.

Waiting for fault-tolerant quantum hardware before developing quantum software is architectural procrastination. By leveraging specialized classical edge hardware—TPUs for massive tensor entanglement and FPGAs for ultra-low-latency gate execution—we can simulate future quantum workloads today. Edge quantum simulation transforms quantum software development from a metered, cloud-dependent experiment into a rapid, sovereign, CI/CD-driven engineering discipline, guaranteeing absolute readiness for the fault-tolerant future.

```
A cloud QPU is a metered bottleneck; an edge simulator is a sandbox.
A CPU simulates sequentially; an FPGA simulates in silicon.
A GPU hits the memory wall; a TPU shatters the tensor network.

The algorithm must be ready before the hardware arrives.

> Quantum hardware may be the future.
> Classical simulation must be absolute today.
```

End of Research Note RES-021

9. Source Register

Source	Authority and Status	Freshness Control
Qiskit Aer Documentation https://qiskit.github.io/qiskit-aer/	Primary project documentation Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.
NVIDIA cuQuantum Documentation https://docs.nvidia.com/cuda/cuquantum/	Primary vendor documentation Living Published: Living	Validated: 2026-07-22 Review on material revision, withdrawal, supersession, or scheduled report review.

10. Lineage, Review, and Publication Record

Record	Value
Original source file	RES-021_Edge_Quantum_Simulators.md
Original source words	1384
Upgrade type	Enterprise research report; source and freshness validation; limitations; adoption decision; reproducibility plan
Generated on	2026-07-22
Current status	Candidate Research Report
Publication authority	Formal Mythos approval not yet recorded
Next review	180 days or event-driven trigger, whichever occurs first