GLOBAL INDUSTRY CHALLENGE 2026

TEAM COVER PAGE

Disclaimer: *Submission must follow GIC requirements: Maximum 3 pages (excluding this cover page and references), 11-point Times New Roman, single spacing, and submitted via Aqora. File Name Requirement: TeamName__Phase2_VersionX.pdf. This official cover page template is required and may not be modified or recreated. Non-compliant submissions may be disqualified and voided.*

Challenge Name: Global Industry Challenge 2026 — U.S. DOE OTC, Energy Infrastructure (Q-Site: Quantum-Enhanced Strategic Siting of Energy Storage and Microgrids)				
Phase #: Phase 3 — Applied Execution				
Team Name: kazi				
First Name	**Last Name**	**Email**	**Aqora Username**	**Role within Team**
Edwin	George Shaji	edwingeorgeshajipadiyanickal@gmail.com	edwingeorgeshaji	Qiskit Integration and QAOA Execution (POC)
Jinsa	Mariam Thomas	jinsamariamthomas26@gmail.com	Rudy	Python Backend and Graph Modelling
Mariam	Skaria	mariamse336@gmail.com	mariam_se	Documentation and Commercial Strategy
Malavika	Krishnan	krismaalu@gmail.com	mxlavika	Classical UI and DB Architecture

Q-Site: Hybrid Quantum Resource Allocation for AI-Loaded Microgrids

Phase 3 Technical Write-Up | Team kazi | U.S. DOE Office of Technology Commercialization — Energy Infrastructure Track

---

1. Focus Area and Rationale

We address **storage siting and sizing for resilience and AI load integration**: given a power grid with a newly connected large AI data center load, which buses should receive battery energy storage systems (BESS) to best preserve voltage stability under a defined N-1 contingency set, subject to a fixed budget on the number of installable units? This use case was selected because it directly matches our team's technical strengths in quantum optimization (QAOA/QUBO) and graph-based power systems modeling, and because it maps onto a measurable, judge-reproducible planning outcome: **the number of unsafe-voltage bus incidents across a standard contingency set, before and after storage siting**.

We ground the study in a concrete, quantified problem rather than a generic claim of grid stress. Adding an 80 MW synthetic AI data center load (roughly 31% of the IEEE 14-bus system's ~259 MW base load) at Bus 9 — selected as the electrically most exposed load bus via graph centrality and hop-distance-to-generation analysis — produces real, AC-power-flow-verified voltage violations at multiple buses, and two single-line contingencies that cause full voltage collapse. This gives us a genuine planning problem with a measurable before/after resilience metric, not a synthetic or assumed one.

2. Technical Approach to Quantum Integration

2.1 Problem Encoding

We select 6 candidate storage buses: the two buses that exhibit genuine under-voltage violations across our contingency set (Buses 9 and 10), plus their four direct electrical neighbors (Buses 5, 8, 11, 13), which are plausible sites for localized voltage support. Each candidate bus i is represented by a binary decision variable $x_i \in \{0,1\}$ (1 = install storage). The siting decision is formulated as a Quadratic Unconstrained Binary Optimization (QUBO) problem:

$$\text{minimize} \quad x^T Q x = \Sigma_i c_i x_i + P\cdot(\Sigma_i x_i - K)^2 + \alpha\cdot\Sigma_{(i,j)\in E} x_i x_j$$

where c_i is a **vulnerability score** for bus i (the summed under-voltage deficit for that bus across all converged N-1 scenarios — buses that are more frequently and more severely under-voltage receive a more negative, more attractive coefficient); $P\cdot(\Sigma x_i - K)^2$ is a penalty (P=5.0) enforcing the budget constraint of exactly K=2 batteries; and the final term (α=0.05) mildly discourages selecting two electrically adjacent candidates, favoring geographic spread of coverage. This produces a 6-variable (6-qubit) QUBO — small enough for both exhaustive classical verification and near-term NISQ execution.

2.2 Modeling Assumption: Voltage-Direction-Aware Scoring

A key modeling decision, arrived at after an internal validation failure during development, is that vulnerability scoring counts **only under-voltage** deviations (v < 0.94 p.u.), not over-voltage. This is because we model installed storage as a reactive-power-injecting static generator (standard volt/VAR support behavior for grid-connected battery inverters), which can only raise local voltage — it cannot help buses that are already over-voltage, and would in fact worsen them. Our first QUBO formulation scored both directions symmetrically and selected buses that were already over-voltage, which validation against real AC power flow (Section 4) showed produced zero resilience improvement. We

consider this an important, disclosed limitation: our siting objective is directionally coupled to a specific storage control mode, and a different inverter control strategy (e.g., real-power discharge for thermal relief) would require a different scoring function.

2.3 QAOA Circuit and Hybrid Loop

The QUBO is converted to an Ising Hamiltonian H (21 Pauli-Z/ZZ terms for our 6-variable problem) via the standard $x_i=(1-s_i)/2$ substitution. We build a QAOA ansatz circuit (Qiskit QAOAAnsatz, depth p=2) and optimize its variational parameters (β,γ) using a classical COBYLA optimizer in a hybrid loop: each iteration evaluates the circuit's expected energy via a quantum Estimator, and COBYLA proposes new parameters to minimize that energy. After optimization converges, we sample the final circuit (4096 shots) and select the lowest-QUBO-cost bitstring observed as our answer — standard QAOA post-processing that accounts for sampling noise.

3. Data Modeling Strategy

Base grid: IEEE 14-bus test system (via pandapower's built-in case14), a standard, citable, publicly reproducible topology with 14 buses, 20 lines/transformers, and ~259 MW of base load. **AI load overlay:** an 80 MW / 25 MVAr synthetic data center load added at Bus 9, sized within the 50–500 MW range suggested by the challenge guidance and scaled to this grid's base load so contingencies produce a realistic severity spread rather than uniform collapse. **Contingency set:** full N-1 analysis — each of the grid's 15 lines is individually removed and AC power flow (pandapower Newton-Raphson solver) is re-solved, giving 15 scenarios total. **Candidate site selection:** data-driven, not assumed — we computed under-voltage exposure for all 14 buses across all scenarios and selected the two genuinely exposed buses plus their graph-adjacent neighbors (Section 2.1). All physics results (voltages, convergence, line loading) come directly from pandapower's AC solver; no values are interpolated or hand-set.

Assumptions and simplifications, stated explicitly: we model storage as reactive-power injection only (no real-power discharge, no state-of-charge dynamics, no time-series dispatch); we use a static per-unit voltage safety band of 0.94–1.06 p.u.; N-1 contingencies are evaluated independently (no cascading or N-2 analysis); and load/generation values are the IEEE 14-bus default case values plus our single synthetic AI load addition, not a time-varying or stochastic scenario tree.

4. Results

4.1 Baseline Grid Stress (No Storage)

With the 80 MW AI load active and no storage, the base case (no line outage) already shows 2 unsafe-voltage buses. Across all 15 N-1 contingency scenarios: **13 scenarios converge with voltage violations (2–4 unsafe buses each)**, and **2 scenarios (loss of line 0–1 or line 8–9) cause total voltage collapse** (power flow fails to converge). Total unsafe-bus incidents summed across all converged scenarios: **30**.

4.2 Classical Baseline (Exhaustive Search)

Since our candidate set is small (6 buses, choose 2 → 15 combinations), we solve the QUBO exactly via brute-force enumeration, providing a guaranteed-optimal classical baseline. The optimal solution is **Bus 9 and Bus 10** (QUBO cost = −20.1027), consistent with these being the only two buses with nonzero under-voltage exposure in our scenario set. A budget-penalty sanity check confirms off-budget selections (1 or 3 batteries) always score worse than the best 2-battery combination, validating the QUBO's constraint encoding.

4.3 QAOA Results (Simulator)

Running QAOA (p=2, COBYLA, 100 optimizer iterations, Qiskit Aer statevector simulation, 4096 final shots) on the identical QUBO, the algorithm converges to **Bus 9 and Bus 10** — matching the classical global optimum exactly (optimality gap = 0.0). Table 1 summarizes key run metrics.

Metric	Value
Qubits	6
Circuit depth (QAOA reps, p)	2
Classical optimizer	COBYLA, 100 iterations
Final sampling shots	4096
QAOA result	Buses {9, 10}, cost = –20.1027
Classical optimum	Buses {9, 10}, cost = –20.1027
Optimality gap	0.0 (exact match)
Wall-clock runtime	23.5 s (statevector simulator)

Table 1. QAOA simulator run summary.

4.4 Real-Hardware Execution (IBM Quantum via qBraid)

We executed the identical QAOA circuit and QUBO from Section 4.3 on real IBM Quantum hardware via qBraid's qiskit-ibm-runtime, transpiled with generate_preset_pass_manager and run via EstimatorV2/SamplerV2. Our job was auto-assigned to the least-busy available backend, **ibm_kingston**, a 156-qubit second-generation Heron (Heron r2) processor — consistent with the challenge's supported IBM Heron r2 platform tier. Given limited free-tier runtime, we capped optimizer iterations and added a 420s hard timeout, using the best parameters found so far if reached early. Since the transpiler maps our 6 logical qubits onto 6 of the device's 156 physical qubits, measured 156-bit bitstrings are mapped back to our 6 logical variables via the transpiler's qubit layout before decoding.

Result: the hardware run converged to Buses {9, 10} — matching the classical and simulator optimum exactly (cost = –20.1027; Table 2).

Metric	Value
Backend	ibm_kingston (156 qubits)
Transpiled circuit depth / gate count	331 / 559
Classical optimizer	COBYLA, capped 12 iter., 420s timeout
Hardware evaluations completed	11 (timeout before 12th)
Final sampling shots	2048
Hardware result	Buses {9, 10}, cost = –20.1027
Result concordance	Exact match — true optimum found
Wall-clock runtime	453.3 s

Table 2. Real IBM Quantum hardware execution summary (via qBraid).

Caveats: individual evaluations were noisier than the simulator (ranging +12.5 to –6.0 under real gate/readout noise, expected NISQ-era behavior), and our reduced iteration budget (11 vs. 100) explored a smaller parameter-space fraction; exact-optimum convergence here is a positive result but not guaranteed at larger problem sizes.

4.5 Validation Against Real AC Power Flow Physics

Critically, we do not stop at the QUBO/QAOA answer — we validate it against the same AC power flow engine used to establish the baseline problem (Section 4.1). We re-run all 15 N-1 contingency scenarios with 15 MVAr reactive-power-injecting storage installed at Bus 9 and Bus 10, and compare

unsafe-bus counts before and after (Figure 2).

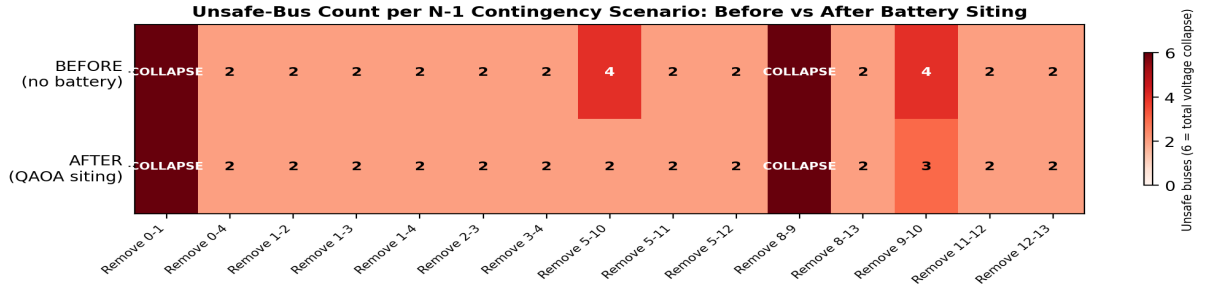


Figure 2. Unsafe-bus count per N-1 contingency scenario, before vs. after QAOA-selected battery siting. Dark cells indicate total voltage collapse.

Result: total unsafe-bus incidents fall from 30 to 27 (a **10.0% reduction**), with the two worst-hit surviving scenarios showing the clearest improvement — loss of line 5–10 improves from 4 to 2 unsafe buses, and loss of line 9–10 improves from 4 to 3. **The two total-collapse scenarios (loss of line 0–1 or line 8–9) are not resolved** by this storage configuration — 15 MVAr of local reactive support at two buses is insufficient to prevent collapse when the grid loses a structurally critical line near the slack bus. We report this honestly as a limitation: our current siting objective optimizes for aggregate under-voltage exposure, not specifically for preventing the most catastrophic (collapse) contingencies, which may require a different objective weighting or additional storage budget in future work.

5. Quantum Platform Justification and Resource Needs

Our 6-qubit problem is well within the range suitable for both exact statevector simulation (used for development, verification, and the Section 4.3 results) and near-term NISQ hardware execution. We target **IBM Quantum hardware (Heron-class processors) via qBraid**, justified as follows: (1) our QAOAAnsatz circuit at p=2 for 6 qubits has modest gate count and depth, well inside typical coherence-time budgets for Heron-generation devices without requiring analog or specialized topologies; (2) IBM Runtime's SamplerV2/EstimatorV2 primitives integrate directly with our existing Qiskit-based pipeline with no reformulation needed, minimizing translation risk; (3) qBraid provides the credentialing and job-submission layer specified by the Challenge for Phase 3. Expected resource needs: **6 qubits**, circuit depth on the order of tens of gates pre-transpilation (device-dependent post-transpilation depth, to be reported once hardware access is confirmed), **~100 classical optimizer evaluations** (each requiring one hardware circuit execution) during the COBYLA loop, and **4096 shots** for final answer sampling — a modest, cost-bounded hardware budget appropriate for a planning-stage prototype rather than a production system.

6. Solution Visualization

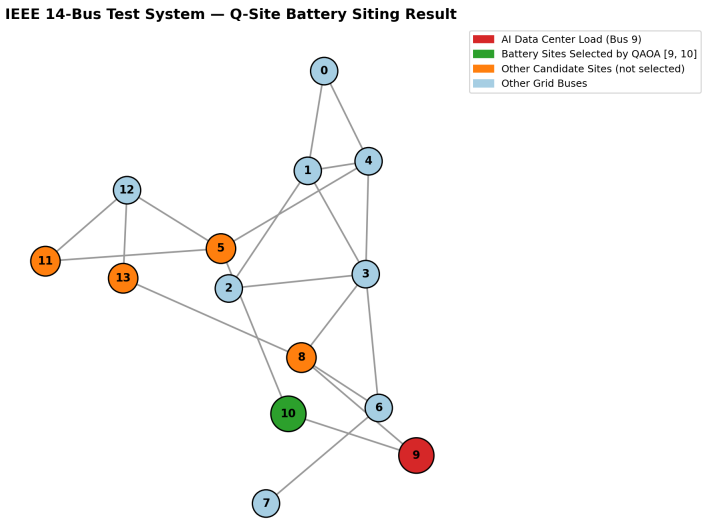


Figure 1. IEEE 14-bus topology showing the AI data center load bus (red), QAOA-selected battery sites (green), and other evaluated candidates (orange).

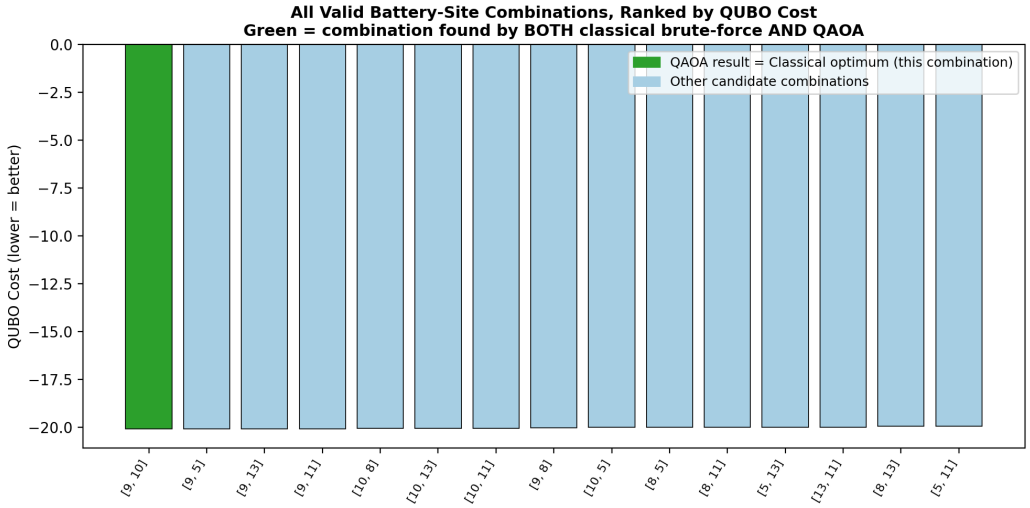


Figure 3. QUBO cost for all 15 valid 2-battery combinations. The combination found by both QAOA and exhaustive classical search (green) is the unique global minimum.

7. Stakeholder Relevance

Utility and ISO/RTO planners making multi-year capital storage-siting decisions benefit directly from a tool that identifies globally-optimal (not merely heuristic) siting under explicit contingency and new-load stress, with runtime fast enough to test many load-growth and outage scenarios per planning cycle. **DOE and state energy offices** evaluating resilience investments gain a reproducible, quantum-benchmarked methodology for comparing storage-siting approaches as AI and industrial load growth accelerates interconnection demand nationally. **Communities and large-load customers** (data center operators negotiating interconnection) benefit indirectly from siting decisions that reduce voltage violation risk introduced by their own load additions. For a planner to adopt this specific method operationally, near-term hardware would need lower per-shot latency and larger practical qubit counts with maintained fidelity, to scale from our 6-candidate prototype to realistic candidate sets of dozens of buses on transmission-scale systems.