

Quantum's Grid Test: Dispatch, Routing and Capacity Planning

An Investment Framework for Comparing Quantum and Classical Utility Optimisation

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Electricity businesses considering quantum optimisation need a purchasing decision tied to a defined task and a reproducible comparison. Dispatch, maintenance routing and capacity planning present different constraints, deadlines and consequences. This paper proposes an investment framework for deciding whether a bounded quantum or hybrid evaluation deserves funding. It connects technical evidence to the cost of producing an acceptable operational recommendation and then to the cash consequences of procurement.

The evidence review considers a September 2026 US Department of Energy grid-planning challenge, three primary preprint abstracts and public benchmark-library documentation. These sources describe applications, test conditions and comparison methods. They supply no universal proof that quantum optimisation is commercially superior for a utility. The paper proposes paired tests against a maintained classical baseline, independent feasibility checks, frozen instance definitions, complete workflow timing and explicit treatment of failed or late results. A qualitative readiness sequence separates formulation, benchmark reproduction, utility review and purchasing decisions.

Three fictional financial cases illustrate the proposed appraisal. Initial expenditure is assumed at USD 420,000, annual operating costs at USD 180,000 and annual gross cash benefits at USD 220,000, USD 360,000 or USD 500,000. At a fictional 10 percent discount rate over five years, corresponding workflow net present values are approximately negative USD 268,369, positive USD 262,342 and positive USD 793,052. The framework's practical conclusion is to fund a specified comparison with capped commitments, preserve a usable conventional fallback and make expansion conditional on independently accepted evidence. The numerical cases describe assumptions rather than observed performance or supplier quotations.

JEL Classification: G31, L94, O33, C61, Q41

Keywords: quantum computing, grid optimisation, unit commitment, utility routing, capacity planning, classical baselines, investment appraisal, staged procurement

1 The investment decision

A utility considering a quantum optimisation service should define the decision that its evaluation is intended to support. That decision might be whether to pay for an offline comparison, renew a development contract or integrate a planning recommendation into an existing engineering review. Each commitment has a different evidence requirement. A prototype that helps formulate a problem can justify further research without establishing an operating capability or a procurement case.

The investment decision addressed here is approval of a bounded, offline comparison of quantum or hybrid methods for dispatch, utility crew routing and capacity planning. The comparison should establish whether an acceptable recommendation can be produced within the relevant deadline and at a complete cost that supports the proposed expenditure. Plant control, switching authorisation and autonomous operational execution sit outside this paper's proposed initial evaluation.

This scope allows an investor or sponsor to distinguish useful progress from unsupported commercial extrapolation. A supplier might demonstrate a circuit, a solution candidate or a promising decomposition. The sponsor still needs to know which constraints were included, which work occurred on classical machines, how rejected results were treated and whether the same resources could improve the incumbent process. Those questions should appear in the evaluation agreement before expenditure begins.

The paper develops a proposed decision framework. It reviews specified primary sources and provides hypothetical numerical examples. It reports no original hardware experiment, utility operating trial or verified client savings. Its contribution is an explicit connection between problem definition, comparative testing, acceptable outcomes and staged investment. A reader can use that connection to structure an evidence request and a bounded procurement decision.

2 The current research evidence

On 24 September 2026, the US Department of Energy announced the winning teams in a challenge examining quantum and hybrid approaches to energy-storage and microgrid placement and capacity. The final phase used hardware and simulation environments and compared approaches with classical baselines. The announcement provides no numerical speedup, independently usable cost series or utility cash-return estimate. It supports the relevance of comparative evaluation for a defined planning problem [1].

Gaidai and Mukherjee's July 2026 preprint abstract reports tests of hybrid methods on specified power-grid optimisation instances. Its qubit-efficient approach did not outperform uniform sampling in the reported setting. This is a boundary-specific result; it supplies a reason to examine formulation and baseline quality closely. It does not establish that every quantum method fails on every grid application [2].

Aboumrada and colleagues describe a hybrid unit-commitment method combining a variational quantum algorithm with classical procedures. Their abstract reports systems with three, ten and twenty-six generators and selected hardware demonstrations. Koretsky and colleagues describe simulated hybrid unit-commitment tests and comparisons with classical methods. Both publications should be cited as preprints with their stated experimental boundaries, without converting those boundaries into a universal commercial threshold [3, 4].

Taken together, these sources support a research question that remains application-specific. An investor should ask what the proposed test adds beyond the cited demonstration, whether the relevant constraints survive the formulation and how the complete workflow performs against credible alternatives. Published interest and a successful bounded demonstration can support further examination. Procurement requires evidence about the sponsor's actual permitted use and economic consequences.

3 Separate the utility applications

Dispatch, routing and capacity planning should have separate acceptance documents. Dispatch concerns decisions about generation and resources within a defined time horizon. A routing evaluation concerns assignments and movements of qualified crews under specified access and timing restrictions. Capacity planning concerns the configuration and size of investments under explicit demand and operating assumptions. Combining their results in one claimed performance percentage would obscure different decision requirements.

Table 1. Proposed application boundaries for an offline evaluation

Application	Decision being evaluated	Constraints to specify
Dispatch	Candidate generation schedule	Demand, reserves, ramps, network model
Crew routing	Proposed task and crew assignments	Skills, access, shifts, precedence
Capacity planning	Candidate asset locations and sizes	Demand cases, site limits, grid studies
Benchmark reproduction	Comparable research result	Frozen instances, method, computing budget

Table 1 provides a proposed scoping checklist rather than a complete engineering specification. For each application, the sponsor should identify the organisation authorised to accept a recommendation and the information that acceptance requires. A route with the shortest travel distance can remain unusable if the assigned crew lacks the required qualification. A planning solution can remain provisional while a separate connection study is incomplete.

The contract should name exclusions clearly. An initial routing comparison might cover planned maintenance with approved task data and exclude emergency restoration. A capacity exercise might compare specified battery locations while excluding land acquisition, permitting and interconnection approval. A dispatch comparison might use a research formulation while excluding live scheduling. Results should retain these boundaries whenever they are presented to an investment committee or prospective customer.

4 Preserve the problem when changing its representation

The proposed method begins with an engineering specification that exists independently of the solver. Record the decision variables, objective, constraint definitions, input units and permitted tolerances. Name the source of each operating limit and the owner permitted to change it. This document should remain the reference when a supplier converts the problem into a representation suitable for its chosen computing method.

The conversion needs a traceable constraint map. For every original requirement, identify its representation in the transformed problem and its treatment after a candidate solution returns. Constraints may be directly represented, evaluated separately or excluded with explicit sponsor approval. A penalty term should have a documented interpretation. Its inclusion in an objective function alone provides no evidence that the final recommendation meets the original requirement.

This distinction matters for financial comparisons. Two formulations can generate different reported objective values because they contain different terms, units or tolerances. The reviewer should recalculate every accepted candidate under the common original objective. Record the feasibility result separately. An attractive objective value from a candidate that fails a mandatory condition should not enter an accepted-solution savings estimate.

Keep a versioned record of approximation choices. If the evaluation relaxes the network model, reduces the planning horizon or aggregates units, identify the effect on the decision being tested. A simplified case can be useful for development. Its interpretation should remain attached to that simplification. Expansion to a more complete problem then requires another documented comparison and the relevant engineering acceptance.

5 Establish a strong classical baseline

The sponsor should appoint an owner for the conventional comparison before requesting quantum performance claims. That owner should be able to configure the reference method appropriately, document its settings and explain its limitations. The comparison should include the existing operational method where accessible and a suitably maintained classical alternative under an agreed resource budget. A deliberately weak comparator would give the purchasing committee an unreliable account of incremental value.

Define the resources offered to each method. Record hardware access, preprocessing, tuning effort, software versions, warm-start information and the personnel used to implement the workflow. Supplier development effort and sponsor support should be visible. Equal expenditure is not always a meaningful scientific condition, but each difference should be disclosed and tested for its effect on the purchasing decision.

Separate implementation improvement from a change in computing approach. Updating data preparation, using a better starting solution or correcting a modelling error might improve the incumbent process. The sponsor should retain those changes as a separate classical case. This allows the committee to ask how much benefit remains attributable to the proposed service after conventional improvements have been considered.

Retest the baseline when the problem or software changes materially. Preserve the previously agreed result rather than rewriting its history. A new comparison can show how the candidate service performs against the current alternative. Procurement documents should specify which baseline governs payment or renewal, because an obsolete reference can support misleading claims even when the underlying calculation is accurate.

6 Use public benchmarks within their stated boundaries

Power Grid Lib's optimal-power-flow repository is maintained through an IEEE Power and Energy Society task force. Its documentation describes curated AC optimal-power-flow cases, MATPOWER-format data and reference implementations. It also cautions about using those cases for other problem variants. A proposed evaluation should record the actual selected case files and their version, alongside every approved modification [5].

The separate unit-commitment repository documents a mixed-integer formulation with specified generator, load, reserve and time-series features. Its scope should be read directly when selecting instances. A sponsor should separately identify any network or operating requirements needed for its own decision. Passing a chosen research case cannot establish acceptance against constraints that the chosen case does not contain [6].

Public libraries provide a starting point for reproducible comparison. The evaluation should freeze a dataset manifest containing filenames, checksums, licences and the exact formulation used. If a case is modified, preserve the original and the change record. The published paper or procurement report should describe modifications precisely enough for another qualified reviewer to understand the comparison without relying on a supplier presentation.

No benchmark-library experiments are performed in this paper. The recommended manifest is a proposed control for a future evaluation. A sponsor moving towards its own data should compare the research instance with the intended operating perimeter and identify additional validation work. Dataset familiarity, public availability and

mathematical reproducibility each support different parts of the evidence record; operational acceptance remains a separate engineering decision.

7 Build paired tests and protected holdout cases

Every proposed comparative run should link the candidate and baseline to the same input snapshot. Give the instance a stable identifier and preserve demand assumptions, task definitions, asset limits and objective parameters. Changes after the snapshot should create a new instance. This arrangement allows a reviewer to connect differences in output to the methods actually tested and to distinguish those differences from changes in inputs.

Separate development cases from evaluation cases. Development cases can support tuning and debugging. Holdout cases should remain protected until the agreed test stage so that the result does not depend entirely on repeated optimisation of a familiar dataset. The sponsor should decide how the holdout set is selected and whether its mix represents the decisions that procurement is intended to support.

Record random seeds where applicable and report repeated runs under the agreed protocol. Preserve failures and timeouts. A best run can be reported as one descriptive result, accompanied by the full set of attempts. For commercial review, the distribution of usable outcomes matters because a sponsor must operate under the conditions encountered in practice rather than under the most favourable retained demonstration.

Give the reviewer access to the denominator. If a supplier submits ten accepted outcomes from a larger set of attempts, the report should preserve the count and disposition of all attempts. Retest rules, permitted retries and termination conditions should be agreed beforehand. This limits ambiguity when results are disappointing and makes the final investment record more useful for subsequent development or procurement decisions.

8 Examine the complete computing workflow

An investment committee needs the time and cost of producing a usable result. The proposed evaluation therefore records preparation, remote queueing, execution, repeated sampling, output reconstruction, feasibility checking and any repair. Human review required before an offline recommendation is accepted should also be measured when relevant. The report should state which stages are included in a supplier's advertised performance number.

Allocate shared work consistently. If data conversion is performed once and reused across many instances, distinguish the initial setup from recurring work. If a service requires retuning after an input change, identify the frequency assumed in the cost model. A favourable amortisation assumption needs evidence about reuse. The sponsor should test the workflow under changed inputs before adopting a recurring-cost estimate.

The QED-C application-oriented benchmark repository describes prototype reference applications for characterising quantum performance. Its documentation explicitly distinguishes those examples from performance standards. Such tools can inform technical comparison, but they should not be described as utility certification or as proof that a dispatch, routing or investment recommendation has passed the sponsor's requirements [7].

The proposed output is a connected workflow record rather than a hardware ranking. It should identify the accepted result, resources consumed and conditions under which it was obtained. Where a classical step dominates cost or time, preserve that observation. Where the quantum component adds measurable value, explain its contribution using the same evidence boundary. The purchasing decision then concerns a defined service that the sponsor can assess.

9 Create an independently checked feasibility record

Feasibility should be assessed by a checker that applies the common problem specification to the returned candidate. The sponsor should document who maintains that checker and how it is tested. Where independent engineering review is required, identify its scope and the responsible qualified team. Solver-reported success should be retained as one input to this process rather than treated as the acceptance decision itself.

The checker should record each mandatory requirement, tolerance and result. If a constraint violation is accepted under a formally approved exception, preserve the exception and its owner. A candidate that meets a relaxed development formulation should remain labelled as such in the evaluation record. This protects the interpretation of the result when it is subsequently used in contracting or investment analysis.

Repaired candidates need a separate history. Save the initial result, the repair method, resources used and final accepted output. Recalculate the common objective after repair. An investment model should use the value of the recommendation actually accepted, including the cost and timing of repair. Improvement claims derived from the original unusable candidate would otherwise disconnect commercial appraisal from engineering acceptance.

Where the checker cannot establish a requirement, the report should identify that requirement as unresolved. The committee can fund a bounded investigation, revise the evaluation scope or decline expansion. A numerical success score cannot replace a material missing constraint. The proposed record is intended to support those specific decisions and to retain the reasons when a promising method remains unsuitable for the sponsor's intended use.

10 Treat deadlines and failed attempts as economic inputs

Specify the deadline from the point at which the relevant input becomes available to the point at which an accepted recommendation is required. State whether that interval includes data preparation, computing queues, retries and review. A dispatch deadline and a multimonth investment study have different meanings. The same algorithm can merit further planning research while failing the acceptance conditions for a time-sensitive workflow.

The proposed benchmark should report the share of attempts producing accepted results by the deadline. It should also retain late accepted results, infeasible outputs and executions without a result. These categories provide different information. A late planning result might still be usable after an approved extension; a missed operational deadline might require the conventional fallback. Define the disposition before assigning a financial consequence.

The investment model should use evidence about the workflow that actually serves the decision. If an accepted result arrives only after the decision has been made, assign no claimed decision benefit unless the sponsor can show a distinct permitted use. Avoid removing the late case from the denominator solely because it complicates an average runtime or success percentage.

Stress-test timing dependencies. A remote queue, input correction or repeated review can change the usable-result rate. Record how the existing method responds under those conditions. The paper does not assign probabilities to those events. It recommends collecting the observations required to estimate their effect later. Until that evidence exists, the committee should retain the relevant uncertainty in its procurement conditions and expenditure limit.

11 Structure a dispatch evaluation

A proposed dispatch evaluation should begin by specifying whether it concerns unit commitment, dispatch conditional on commitment or a different optimisation problem. Record the horizon, time resolution and demand inputs. Name the reserve and ramping requirements included. The physical network representation and treatment of outages should be explicit so that an engineering reviewer can identify what the proposed calculation can establish.

For a unit-commitment comparison, preserve each tested schedule and its accompanying power allocation. Check startup, shutdown and temporal restrictions under the chosen formulation. Recompute costs using the same parameters for every method. If a decomposition creates candidates in one stage and evaluates them in another, preserve the costs and outputs of both stages. The complete recommendation remains the relevant unit of comparison.

The research abstracts reviewed here provide bounded examples of hybrid approaches to this application [3, 4]. The proposed commercial test should extend those questions to its own defined cases without claiming that the cited studies establish the sponsor's result. A technical team can investigate whether the approach remains useful when additional periods, network requirements or operating restrictions are included.

The initial sponsor decision should remain offline. A qualified team can compare returned schedules with the existing process and identify the tests required before any further use. Investment approval should state that boundary. Passing an offline benchmark authorises the next specified evaluation step only. Any proposal affecting live operations needs its own technical, security and operating acceptance through the utility's established approval procedures.

12 Structure a utility crew routing evaluation

A crew-routing comparison should identify the tasks being assigned and the period covered. Record approved access windows, crew qualifications, shift limits, travel assumptions and task precedence. Define how tasks that require multiple crews or specialised equipment are handled. A shortest route can be a useful intermediate calculation, but the acceptance decision concerns a workable set of assignments within the specified service perimeter.

The proposed initial use is planned work using an approved information extract. Emergency restoration is excluded from this example because its information changes and operating consequences require a separate specification. The sponsor should preserve that exclusion when communicating results. If it later evaluates emergency work, create a new protocol describing how updates, uncertain access and priorities enter the decision.

Agree the objective before implementation. Travel time, task completion, overtime and customer effects can enter a model in different ways. Establish units and permitted priorities, then compare candidate assignments under the same objective. Where the sponsor cannot assign a defensible monetary weight, retain the measure separately. A composite score with arbitrary weights should not become a cash-benefit estimate.

The evaluation should include an independent assignment check and a review by the responsible scheduling team. Record corrections, rejected routes and time spent inspecting proposals. The paper provides no demonstrated quantum routing performance. Its proposed framework allows the sponsor to discover whether any candidate method adds useful value after operational restrictions and review effort are accounted for. The accepted assignment and its observed consequence should govern later commercial claims.

13 Structure capacity planning under explicit scenarios

Capacity planning should define the assets, locations and decisions included in the comparison. A proposed storage exercise might specify candidate sites, permitted sizes, demand cases and the investment horizon. Identify engineering assessments that remain outside the optimiser. A recommended configuration should retain those dependencies so that an investment committee can distinguish a model output from an approved development decision.


DOE's September 2026 challenge provides a contemporary example of quantum and hybrid evaluation for energy-storage and microgrid siting and capacity [1]. The announcement does not disclose a cash-benefit series suitable for this paper's investment model. The proposed framework therefore uses its application boundary as research context while keeping all later financial assumptions separate and fictional.

Scenario selection should be the sponsor's explicit responsibility. Preserve each demand case, outage assumption, asset-cost input and operating policy. Compare solutions on the same scenarios. A candidate tuned to one forecast might behave differently when demand or availability changes. Report that behaviour directly and explain which conditions govern the purchasing decision. Avoid combining cases into an expected value without a documented basis for probabilities.

Inspect the stability of the recommendation. If modest approved input changes move the preferred site or size materially, the committee should understand why and what further information might resolve the sensitivity. A faster solver may support more exploration, but that advantage becomes valuable only through a defined use. The sponsor should document the resulting investment decision and avoid assigning invented development revenue to computational progress alone.

14 Interpret readiness as a sequence of evidence

The proposed readiness curve expresses the order of evidence collection. It does not measure a technology's maturity or predict a hardware development date. Begin with an agreed formulation, reproduce an appropriate benchmark, run paired tests, obtain utility review and then assess bounded procurement. Figure 1 shows these stages as a qualitative sequence. The vertical positions have no empirical or probabilistic interpretation.

Figure 1. Proposed evidence sequence for a utility quantum evaluation

Each stage should produce a specific record. Formulation supplies the problem and constraint map. Reproduction supplies an implementation and comparable result. Paired testing supplies complete outcomes under the agreed protocol. Utility review supplies acceptance findings and unresolved requirements. Procurement review connects that evidence to permitted use, cost, contractual obligations and a stated expenditure decision.

A failure should identify the next decision rather than disappear into a general readiness score. If the representation omits a material constraint, revise it before comparison. If results cannot be independently checked, resolve access to the evidence. If a test is technically acceptable but economically unattractive, retain that conclusion and reconsider the service or spending level. These actions preserve the interpretation of earlier work.

An expansion proposal should explain which evidence remains reusable. Adding sites, changing the horizon or using a different computing backend can affect the comparison. The sponsor should approve the additional tests and funds explicitly. This approach provides a proposed commercialisation discipline for the defined application. It supplies no universal deployment timetable, sector adoption forecast or assured conversion from a technical evaluation to a paying contract.

15 Define comparable performance measures

The proposed performance record has several distinct measures. Feasibility describes compliance with the specified requirements. Objective value describes the agreed mathematical or operational target. Elapsed time describes the complete interval used to obtain an accepted recommendation. Cost describes the resources required. Each measure should retain its unit and boundary when results are presented to a procurement team.

Table 2. Proposed metrics and evidence boundaries

Measure	Evaluation boundary	Evidence to retain
Feasibility	Common original constraints	Independent checking report
Objective	Accepted recommendation	Recalculation under shared inputs
Elapsed time	Input availability to acceptance	Stage timestamps and retries
Failure rate	Every agreed attempt	Full outcome denominator
Review effort	Required human acceptance	Recorded work and corrections
Full cost	Setup and recurring service	Resource and payment schedule

Table 2 is a proposed measurement plan. It should be completed with the sponsor's actual protocol before testing. Record whether a reported objective is compared with an established bound, a best-known solution or an incumbent result. These reference points have different meanings. A percentage improvement should name the denominator and the treatment of failed attempts so that it can be independently interpreted.


Use distributions when one average obscures acceptance. A workflow with occasional long delays may require more conventional capacity than its average runtime suggests. Report the agreed summaries alongside individual records. If a measure cannot be established, retain it as unresolved and explain its relevance to expenditure approval. The committee should receive the evidence needed for a decision without being asked to reconstruct missing definitions from a performance chart.

16 A fictional end to end timing benchmark

Consider a deliberately fictional timing example for one offline task. A baseline classical workflow uses two minutes for preparation, twelve for execution and four for repair and checking, giving eighteen minutes in total. A tuned classical workflow uses two, seven and three minutes respectively, giving twelve minutes. A hypothetical hybrid workflow uses five minutes for preparation, six for queueing, two for execution and seven for repair and checking, giving twenty minutes.

For this comparison only, assume that every final result has the same accepted quality and constraint coverage. The assumption exists to isolate timing; it would need independent proof in a real evaluation. Set a fictional fifteen-minute acceptance deadline. The tuned conventional workflow meets that deadline. The baseline and hybrid examples exceed it. Figure 2 displays the full components, including the remote queue assumed for the hybrid case.

Figure 2. Fictional complete workflow timing against a fifteen minute deadline



The two-minute hybrid execution component cannot be used as the complete delivery time. Its hypothetical twenty-minute total includes work needed to produce the assumed acceptable output. A proposed improvement in queueing or repair can be examined as a new case, with the changed assumptions stated. A procurement claim should follow the result of that new test rather than omit the stages that currently determine acceptance.

These figures are author assumptions, not measurements of a named solver, quantum device or utility. Their purpose is to show how a complete-workflow comparison can change an investment interpretation. Real appraisal should replace them with recorded timestamps and independent acceptance findings. It should also test timing variability, failures and the cost of retaining the fallback required when a recommendation arrives late.

17 Convert accepted outcomes into cash evidence

Cash-benefit estimation should begin with an accepted operational or planning outcome. Identify the payment, expenditure or receipt affected and the person responsible for verifying it. A lower objective value may provide a technical indication worth examining, but its cash consequence depends on the decision and contractual setting. Document that bridge explicitly before assigning revenue or avoided cost to the proposed computing service.

For dispatch, a sponsor might examine cost differences under the same approved assumptions and then reconcile them with the applicable settlement or procurement arrangements. For routing, it might investigate actual overtime or contracted travel expenditure. For capacity planning, it might examine an approved change in investment scope. These are proposed evidence paths; the paper reports no realised cash benefit in any of them.

Separate cash savings from released staff time. Time can be useful without changing payroll or contracted expenditure. If the sponsor assumes that released capacity produces additional revenue, specify the mechanism, delivery capacity and demand evidence. Retain unmonetised measures where no defensible cash bridge exists. This gives the committee a clearer view of what the financial model includes and what remains a separate operational consideration.

Avoid counting one consequence repeatedly. If an investment decision reduces both a modelled cost and a payment already captured in the same cash series, reconcile the overlap. Maintain an evidence ledger linking each claimed benefit to its baseline, measurement period and approved calculation. That ledger should survive procurement renewal so that a later reviewer can distinguish realised value from an earlier forecast or an illustrative scenario.

18 Establish the full investment cost boundary

The proposed cost model separates initial expenditure from recurring commitments. Initial work might include formulation, data preparation, integration design and acceptance testing. Recurring work might include computing access, maintenance, independent checking, review and the conventional fallback. These categories are proposed diligence questions. Their actual amounts, timing and treatment should come from approved budgets, quotations and contractual terms before an investment decision.

Allocate costs according to the intended use. A pilot can share equipment or personnel with an existing research programme, but the report should identify that support and its continued availability. A commercial service may need different resilience, documentation or review. The committee should see how the cost boundary changes when moving from a bounded comparison to recurring use.

Include obligations that arise when performance disappoints. Retaining another solver, rerunning an analysis or correcting outputs can create costs that a simple computing-access quote excludes. Define those obligations and investigate their financial effect. The paper does not estimate incident losses or assign an operating-failure probability. It recommends preserving such exposures as unresolved where the evidence required to quantify them is unavailable.

Review payment timing and flexibility. A fixed minimum access commitment creates a different cash profile from pay-per-use evaluation. A supplier-funded arrangement needs its own payment schedule, funding charges and allocation of responsibilities. Compare those structures using the same accepted-use assumptions. A financing structure cannot itself demonstrate that the proposed optimisation produces usable results or that the forecast benefit has an evidential basis.

19 Three fictional investment cases

The following model evaluates a hypothetical workflow investment. Initial expenditure is USD 420,000 at the start. Annual recurring expenditure is USD 180,000, comprising fictional access cost of USD 70,000, integration maintenance of USD 50,000 and validation and review of USD 60,000. Annual gross cash benefits are assumed at USD 220,000, USD 360,000 or USD 500,000. All amounts are author scenarios rather than supplier prices or utility outcomes.

Table 3. Fictional annual cash and five year investment results in USD thousands

Case	Annual gross benefit	Annual net cash	Five year NPV
Lower benefit	220	40	-268.369
Middle benefit	360	180	262.342
Higher benefit	500	320	793.052

Annual net cash is gross benefit less USD 180,000 recurring expenditure. Assume five equal year-end cash flows discounted at a fictional 10 percent annually. Their annuity factor is approximately 3.79078677. Net present value equals annual net cash multiplied by that factor, less the USD 420,000 initial expenditure. Table 3 reports thousands of dollars, including the negative sign for the lower case.

The model excludes tax, debt, working capital, inflation, implementation delay, residual value and incident losses. It assumes accepted outcomes support the stated cash benefits every year. No probabilities are assigned.

These workflow values describe the proposed expenditure under stated assumptions; they are not enterprise valuations, evidence of quantum advantage or forecasts of a technology company's worth. A committee should replace the inputs and extend the boundary before using the model for a real commitment.

20 Test the break even requirement

Under the fictional model, the annual net cash needed to recover the initial expenditure at the stated discount rate is USD 420,000 divided by 3.79078677, or approximately USD 110,795. Adding recurring expenditure gives a gross annual cash-benefit threshold of approximately USD 290,795. That threshold depends on the exact model boundary. It contains no supplier margin, financing charge or estimate of the price available in a market.

Figure 3. Fictional workflow investment values and proposed commercialisation decisions




Figure 3 places the hypothetical investment values beside a proposed commercial sequence. The lower case falls below the break-even requirement and has a negative workflow value. The middle and higher cases exceed it under their stated assumptions. A technically accepted result with no cash consequence would produce zero gross benefit and an unfavourable investment result under this cost boundary.

Sensitivity should address each material assumption separately. A shorter useful period, greater validation effort or delayed commencement changes the result. A supplier payment that rises with activity changes the recurring cost function. The sponsor should examine those variations using documented inputs rather than choose an attractive case without establishing how its benefit would occur.

The committee can use the threshold as an evidence request. Ask what observations support accepted-use frequency, gross cash benefit and the complete cost estimate. If the required benefit cannot be supported, cap further commitment or reconsider the scope. A positive fictional case provides a calculation to reproduce. It supplies no technical acceptance, implementation certainty or assurance that a buyer will pay for the service.

21 Contract around accepted evidence

An evaluation agreement should identify the deliverables required for each payment. The proposed initial deliverables are a frozen formulation, constraint map, dataset manifest, baseline configuration and complete run record. Later payments can depend on independently checked comparative results and a specified review report. The sponsor should obtain legal advice on enforceability, intellectual-property rights and the permitted use of results in its jurisdiction.

Define the evidence needed when a test fails. The supplier should retain attempts, errors and unresolved findings rather than deliver only selected demonstrations. Agree whether the sponsor can reproduce results and inspect the software or interfaces needed for acceptance. Confidentiality can be handled through approved access arrangements, with material restrictions disclosed to the purchasing committee before commitment.

Payment conditions should distinguish delivery of research work from achievement of an accepted performance requirement. A well-executed comparison can produce an economically negative result. The agreement should

state how that outcome affects payment and expansion. This supports accurate reporting without requiring the research team to present an unsuccessful comparison as a commercially successful deployment.

Preserve sponsor flexibility at decision points. A bounded contract can specify a maximum authorised spend, the next review and the conditions required before extension. Renewal should use the agreed current baseline and complete cost boundary. Any promise about future hardware or method performance should be evaluated as a contractual statement with its own evidence and remedy. The paper supplies a proposed structure rather than legal terms ready for signature.

22 Assess suppliers and transaction claims

A financing or acquisition review should connect a supplier's technology claims to reproducible application evidence. Identify what the supplier owns, what it licenses and what depends on another computing provider. Record access commitments, software dependencies and the rights required to reproduce results. A product demonstration should remain associated with its tested boundary when it is included in diligence materials.

Revenue analysis needs verified commercial records. An investor should distinguish a research award, unpaid evaluation, contracted pilot and recurring paid use. Inspect the relevant agreements and receipts rather than infer demand from technical interest. Where a supplier forecasts expansion, ask which acceptance conditions remain and what expenditure the customer has actually authorised. This paper provides no estimate of market demand or mandate conversion.

For acquisition appraisal, maintain separate models for the tested workflow and the target company. A positive hypothetical workflow NPV cannot determine an enterprise value. Company appraisal needs its own verified revenue, costs, liabilities, intellectual-property position and assumptions about future cash flows. The proposed technical evidence framework can inform those assumptions, with unresolved dependencies retained in the diligence record.

Debt funding requires a further evidence boundary. A lender considering customer payments or contracted service revenue should inspect the actual contract, acceptance provisions and counterparty obligations. Hardware progress does not establish receivable collectability. Structure any financing appraisal around verified rights and cash flows, supported by qualified legal and technical review. The proposed paper supplies questions for transaction diligence without asserting that a particular quantum business is bankable.

23 Govern commercialisation and expansion

The proposed commercialisation roadmap begins with sponsor problem definition and ends with a bounded procurement review. Allocate expenditure to the actual evidence sought at each stage. Name the owner responsible for technical acceptance, the person responsible for verifying the financial bridge and the committee authorised to approve extension. Record the decision before moving to a broader application or a larger recurring commitment.

Expansion should describe the new operating perimeter. Increasing instance size, adding locations or changing from offline planning to time-sensitive use can introduce additional constraints and dependencies. The sponsor should specify the evidence required for those changes. Retain the earlier result as a result for its original scope and assess the wider proposal under a new versioned protocol.

The Department of Energy's grid-modelling programme provides broader context for research concerned with useful power-system analysis [8]. A supplier should still demonstrate relevance to the sponsor's specific task. The proposed roadmap assigns no readiness date or expected revenue to public research activity. Its purpose is to connect a named application with a bounded evidence collection and purchasing process.

Publish the accepted result with its limits. Describe the formulation, baseline, attempted cases, complete workflow measures and unresolved requirements. A negative or mixed comparison can inform subsequent investment by identifying what remains expensive or technically incomplete. Commercialisation then follows the evidence actually obtained. The sponsor can continue research, revise the application or decline further expenditure without rewriting the comparison or assuming a future improvement.

24 The proposed investment conclusion

A defensible quantum-grid investment case should connect one permitted decision to a reproducible comparison and an independently accepted outcome. The framework developed here proposes that connection for dispatch, crew routing and capacity planning. It retains task-specific constraints, complete workflow measures, strong conventional alternatives and a documented cash bridge. Each component should be inspectable before a purchasing committee relies on the result.

The reviewed primary sources establish active research and bounded demonstrations. They do not supply a universal commercial advantage or the financial inputs needed for a utility investment. The fictional cases show how accepted cash benefit and full recurring cost determine a workflow appraisal. Their break-even requirement is specific to five annual cash flows, the assumed initial expenditure and the fictional discount rate.

The recommended sponsor action is a capped offline evaluation with agreed deliverables, independent checking and explicit expansion conditions. Retain a usable conventional fallback and the complete attempt record. Require qualified engineering acceptance for any later operating application. The committee should receive a specific recommendation tied to those results, including a clear reason to stop if the comparison is unusable or the financial bridge remains unsupported.

Further research should test the framework using actual paired instances, recorded resource consumption and independently reviewed outcomes. It should examine how formulation changes, backend availability and review effort affect useful results across the specified applications. Those empirical tests remain future work. This paper offers the proposed evaluation structure and reproducible fictional arithmetic required to frame them without claiming an experiment that has not been performed.

Frequently Asked Questions

Does current research establish a commercial quantum advantage for utility optimisation?

The eight sources reviewed here provide bounded research and benchmark context. They do not establish a universal commercial advantage for dispatch, crew routing or capacity planning. A purchasing decision should depend on the sponsor's defined problem, credible classical alternatives, independently checked outputs, complete workflow costs and evidence of usable outcomes. The paper proposes a method for obtaining that evidence and reports no original utility experiment.

Which application should a utility evaluate first?

Select an application with a documented problem definition, accessible approved data, a maintained comparator and a qualified acceptance owner. An offline capacity or planned-work comparison can provide a bounded evaluation, subject to its own engineering requirements. This paper assigns no universal priority or maturity ranking. The proposed initial scope excludes live plant control and autonomous operational execution. The sponsor should record its specific decision and expenditure limit.

Why should timing include preprocessing and repair?

The purchasing decision concerns an accepted recommendation. Data preparation, queueing, repeated execution, reconstruction, checking and repair can all consume resources before that recommendation exists. Record their times and costs under one agreed boundary. A component runtime remains useful technical evidence, accompanied by the complete workflow result. The fictional example shows a two-minute execution stage inside a twenty-minute workflow; those figures are assumptions rather than hardware measurements.

Can public power grid benchmarks prove operating readiness?

Public benchmark libraries can support reproducible mathematical comparisons within their documented formulations. The sponsor should preserve case versions and modifications, map the library's constraints to the intended application and identify additional requirements. Its own engineering acceptance remains separate. This paper does not run the cited libraries or claim operating readiness from them. A future test should publish enough information to make its precise boundary independently understandable.

How should failed and late results enter the evaluation?

Retain every agreed attempt and its disposition, including infeasible outputs, timeouts and late accepted results. Define retry rules and deadlines before execution. For commercial appraisal, connect useful outcomes to the sponsor's permitted decision and record the fallback required when acceptance fails. A selected best result can be shown alongside the full denominator. The paper assigns no observed failure probability and recommends collecting the evidence needed for such an estimate.

What does the fictional break even calculation show?

With USD 420,000 initial expenditure, USD 180,000 annual recurring cost, five equal year-end cash flows and a 10 percent discount rate, the gross annual cash-benefit threshold is approximately USD 290,795. That calculation is reproducible under the stated assumptions. It excludes taxes, debt, delay and other listed items. It is neither a supplier quotation nor a forecast of utility savings. A real investment model requires supported inputs and an appropriate cost boundary.

How can an evaluation agreement support accurate reporting?

Specify the formulation, dataset manifest, baseline, complete run record and independent acceptance report as deliverables. State how unsuccessful comparisons affect payment and expansion. Preserve rights needed to inspect or reproduce evidence under agreed confidentiality arrangements. Qualified counsel should review the actual terms and jurisdictional requirements. A paid research deliverable and a demonstrated operating advantage are different contractual outcomes, and each should retain its defined evidence requirement.

Can a workflow investment value be used to value a quantum company?

The hypothetical workflow values in this paper describe one proposed expenditure under stated assumptions. Enterprise appraisal requires a separate analysis of verified company cash flows, customer contracts, costs, liabilities, intellectual property and dependencies. Technical comparison can inform that analysis when its scope is preserved. The paper supplies no company valuation, lending recommendation or market-demand forecast. Its proposed diligence questions help identify evidence that a transaction team would need.

References

1. US Department of Energy, Office of Technology Commercialization. DOE Announces Winning Teams in Quantum Grid Planning Competition. 24 September 2026. <https://www.energy.gov/technologycommercialization/articles/doe-announces-winning-teams-quantum-grid-planning-competition> Accessed 4 October 2026.
2. Gaidai, I. and Mukherjee, R. Benchmarking Hybrid Quantum-Classical Algorithms for Power Grid Optimization Problems. arXiv preprint 2607.15543, 17 July 2026. Abstract and stated test boundaries reviewed. <https://arxiv.org/abs/2607.15543> Accessed 4 October 2026.
3. Aboumrad, W., Marthi, P. R. V., Debnath, S., Roetteler, M. and Epifanovsky, E. A New Hybrid Quantum-Classical Algorithm for Solving the Unit Commitment Problem. arXiv preprint 2505.00145, 30 April 2025. Abstract reviewed. <https://arxiv.org/abs/2505.00145> Accessed 4 October 2026.
4. Koretsky, S., Gokhale, P., Baker, J. M., Viszlai, J., Zheng, H., Gurung, N., Burg, R., Paaso, E. A., Khodaei, A., Eskandarpour, R. and Chong, F. T. Adapting Quantum Approximation Optimization Algorithm QAOA for Unit Commitment. arXiv preprint 2110.12624, 25 October 2021. Abstract reviewed. <https://arxiv.org/abs/2110.12624> Accessed 4 October 2026.
5. IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms. Power Grid Lib Optimal Power Flow. Repository README and formulation boundaries. <https://github.com/power-grid-lib/pglib-opf> Accessed 4 October 2026.
6. IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms. Power Grid Lib Unit Commitment. Repository README and model features. <https://github.com/power-grid-lib/pglib-uc> Accessed 4 October 2026.
7. Quantum Economic Development Consortium, Standards and Performance Metrics Technical Advisory Committee. Application-Oriented Performance Benchmarks for Quantum Computing. SRI International repository README; prototype applications, not performance standards. <https://github.com/SRI-International/QC-App-Oriented-Benchmarks> Accessed 4 October 2026.
8. US Department of Energy, Office of Electricity. Grid Modeling. Programme description. <https://www.energy.gov/oe/grid-modeling> Accessed 4 October 2026.

About the Author

Authored by Chennakeshav Adya
Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.