MatchPoint Research Working Paper No. 2026/T36

Quantum Optimisation for Data-Centre Capital Stacks

A benchmark-led framework for capital tranches, draw schedules, power portfolios and evidence-gated quantum readiness

Prepared by Matchpoint Partners
Author attribution pending CK approval
August 2026

ABSTRACT

Background. Data-centre financing joins phase timing, power access, long-lead equipment, tenant ramp, liquidity, reserves and layered capital. **Objective.** This paper develops a controlled optimisation framework for B3 developers/operators and A4 infrastructure investors, with classical production baselines and quantum challengers. **Approach.** The analysis reviews 52 official, vendor, primary and peer-reviewed sources available through 2 August 2026 and defines governed model objects, mixed-integer constraints, D-Wave CQM and annealing mappings, IBM QAOA mappings and a fair benchmark protocol. **Findings.** Capital-stack choices can be represented as binary, integer, continuous and quadratic relationships. Mathematical compatibility does not establish quantum advantage. Production use requires a strong classical baseline, full resource accounting, independent feasibility validation and reproducible evidence. **Implications.** Implementation should begin with one decision, one approved instance library, one classical production model and bounded challenger tracks. Worked cases and operating inputs are unverified illustrative scenarios; attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

Highlights

Primary users. B3 AI-data-centre developers/operators and A4 infrastructure and digital-infrastructure investors.

Decision unit. A versioned project, tranche, draw, power, offtake, reserve, scenario or solver-run object with source and approval state.

Architecture. Governed inputs feed deterministic underwriting, classical MILP or CP-SAT baselines, bounded quantum challengers and independent validation.

Decision boundary. Solvers generate candidates; sponsors, advisers, lenders and investment committees retain transaction authority.

Economics. Attributed revenue, cash saving, loss reduction and alpha remain USD 0 until observed evidence is approved.

JEL Classification: C44, C53, G21, G23, G24, G31, G32, L63, L94, O33

Keywords: quantum optimisation, data centres, capital stack, project finance, quantum annealing, QAOA, D-Wave, IBM Qiskit, power contracts, mixed-integer programming

1. INTRODUCTION

Data-centre financing joins a long-duration physical asset, a fast-changing technology load, a constrained power position and a layered capital structure. The financing decision is therefore a connected set of choices rather than a single leverage calculation. A sponsor may need to select a site and phase, reserve grid capacity, sign one or more power contracts, place long-lead equipment orders, sequence equity and debt draws, satisfy lender conditions, preserve liquidity reserves and respond to changes in tenant demand. Each choice affects the feasible set available to the others.

The demand context is material. The International Energy Agency's 2025 *Energy and AI* report links data-centre expansion directly to electricity supply, grid constraints and energy security [1]. Its 2026 update reports further growth in data-centre electricity use and describes tightening bottlenecks in transformers, turbines, chips, grid connections and approvals [3]. The United States Department of Energy and Lawrence Berkeley National Laboratory also document rapid growth in electricity demand from data centres and the need for integrated generation, storage, grid and flexibility planning [4-6]. These sources establish the operating problem. They do not establish a financing outcome for any particular Gulf project.

This paper develops a decision framework for B3 AI-data-centre and digital-infrastructure developers/operators and A4 infrastructure and digital-infrastructure investors. The framework represents capital tranches, draw schedules, contracted power, reserves, construction phases and covenant headroom in a controlled optimisation model. Classical mixed-integer programming and constraint programming form the production baseline. Quantum annealing, quantum-classical hybrid constrained quadratic models and the Quantum Approximate Optimization Algorithm, or QAOA, are defined as challenger methods whose value must be demonstrated through fair, reproducible comparison.

The central conclusion is bounded. Many data-centre capital-stack choices can be expressed as binary, integer, continuous and quadratic relationships. D-Wave's Ocean documentation describes constrained quadratic models that accept those variable types and explicit constraints, while IBM and Qiskit documentation describe QUBO and Ising mappings and QAOA workflows [18-27]. Mathematical compatibility is not commercial advantage. A decision-useful claim of quantum advantage requires an approved instance set, the strongest appropriate classical baselines, equivalent problem information, complete compute-time accounting, solution-quality and feasibility metrics, repeated runs and reproducible evidence [28-35].

The paper answers five questions:

1. What decisions and constraints define a financeable data-centre capital stack?
2. Which parts belong in deterministic underwriting, scenario analysis or optimisation?
3. How should classical MIP and CP-SAT baselines be built and governed?
4. Where do quantum annealing and QAOA fit, and what mapping costs can remove an apparent advantage?
5. What evidence should be required before a sponsor or investor uses a quantum-generated recommendation?

Worked cases, amounts, probabilities, timings, solver outputs and performance values are **[Unverified illustrative scenarios]** unless explicitly attributed to a cited source. This paper does not claim observed financial benefit. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

Decision layer	Optimisation object	Retained authority
Capital structure	Tranche selection, sizing, pricing, tenor and draw order	Sponsor, investment committee and financing counterparties
Construction	Phase starts, equipment orders, contingency and completion gates	Sponsor, technical adviser and contractors
Power	Interconnection, PPA, grid, onsite generation, storage and flexibility	Sponsor, utility, offtaker and relevant authorities
Credit	DSCR, leverage, reserve, covenant and distribution tests	Lenders, credit committee and finance documents
Quantum challenger	Candidate solution or sample distribution	Model owner and independent validation; no transaction authority

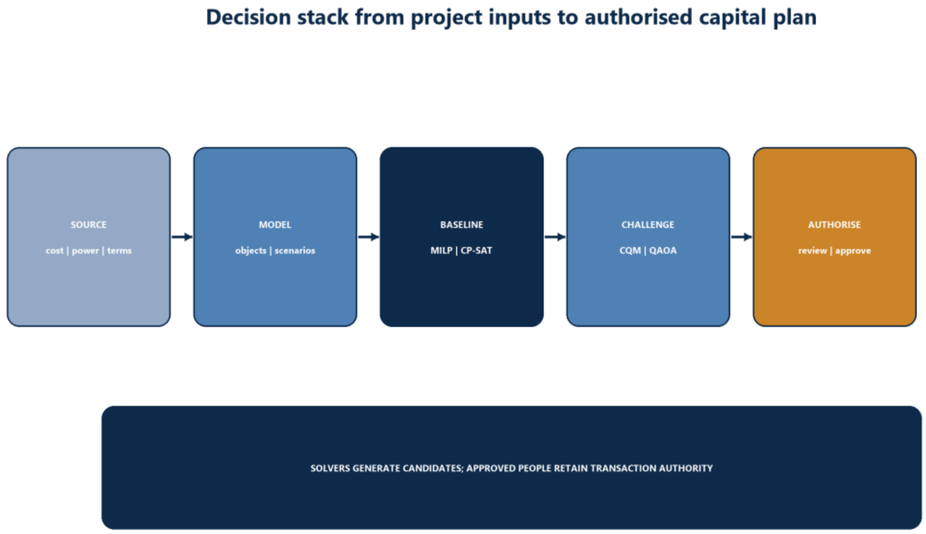

Figure 1. Decision stack from project inputs to authorised capital plan

Source: Matchpoint capital-stack framework; data-centre and governance context [1-17,49-51].

2. B3 AND A4 DECISION PERIMETER

2.1 Primary ICPs

The Topic Tracker identifies B3 as AI-data-centre and digital-infrastructure developers/operators seeking project debt and equity. It identifies A4 as infrastructure and digital-infrastructure investors providing infrastructure debt and equity. Their decision interests overlap while their duties remain distinct.

A B3 sponsor needs a financeable path from development rights and power access to phased construction and contracted capacity. Its model must expose the timing and source of every cash requirement, the conditions attached to each facility, the liquidity effect of delay, and the point at which committed demand supports additional investment. An A4 investor needs to understand downside protection, governance, completion exposure, power and offtake concentration, refinancing risk, residual value and whether the proposed financing retains resilience under adverse scenarios.

2.2 Decision rights

Role	Prepares or analyses	Approves or binds	Model boundary
Sponsor finance	Assumptions, capital plan, scenario set, draw forecast	Sponsor submissions and internal approvals	Cannot create lender commitment
Technical adviser	Programme, cost, completion and equipment evidence	Opinion within appointment	Cannot determine financing terms
Power adviser	Load, connection, tariff, PPA and flexibility analysis	Advice within mandate	Cannot grant grid capacity or execute a PPA
Debt adviser or arranger	Structure, lender mapping, term comparison	Advice and process under mandate	Cannot bind a lender without authority
Lender or private-credit fund	Credit analysis, documentation and monitoring	Commitment, waiver and consent	Must independently approve model use
Equity or infrastructure investor	Return, governance and downside analysis	Investment decision	Must independently approve model use
Quantum or classical solver	Candidate solutions, bounds and diagnostics	None	No authority over investment, credit or contract

2.3 Companion topics

The Topic Tracker identifies P109, *GPU-Backed Debt: Collateral with a Three-Year Half-Life*, and P110, *The First 20%: Financing Data-Centre Pre-Development and Power*, as companion topics. Both remained "Not started" when T36 was reconciled on 2 August 2026. T36 therefore defines interfaces to those topics without claiming that their supporting papers or live pages exist.

P109 contributes a collateral and residual-value module. It would test hardware eligibility, depreciation assumptions, refresh cycles, remarketing channels, advance rates and offtake support. P110 contributes a development-capital module. It would test land control, connection deposits, design, permits and long-lead orders before a construction lender is prepared to fund. In T36, both modules are placeholders requiring approved inputs.

2.4 Advice perimeter

An optimisation output is a model result. It is not a financing offer, valuation, legal interpretation, technical certification, grid commitment, tax opinion, accounting judgement or investment recommendation. Any transaction use requires approved source data, professional review, counterparty engagement and executed documents.

3. WHY THE CAPITAL STACK IS A CONSTRAINED OPTIMISATION PROBLEM

3.1 Connected decisions

A capital stack becomes an optimisation problem when several feasible structures exist and management wants to select among them using a declared objective. Typical objectives include minimising expected all-in funding cost, minimising expected equity requirement, maximising risk-adjusted project value, maximising liquidity headroom, minimising refinancing exposure or finding a robust compromise across several measures.

The decision set is combinatorial. A tranche can be selected or rejected. A phase can start in one of several periods. A power contract can cover a selected share of load. An equipment order can be accelerated or deferred. A reserve can be funded from equity or facility proceeds if documents permit. These discrete choices interact with continuous amounts such as debt size, equity contribution, power volume and cash reserve.

3.2 Constraint families

Constraint family	Example model statement	Required evidence
Sources and uses	Sources equal eligible uses plus required cash	Approved budget and financing terms
Funding eligibility	A facility funds only permitted cost categories	Term sheet and finance documents
Draw timing	Draw occurs after conditions precedent and before expiry	Conditions tracker and availability period
Leverage	Debt does not exceed an approved base or ratio	Valuation and lender definition
Coverage	Cash flow supports debt service under defined cases	Financial model and covenant definition
Reserve	Minimum liquidity, debt-service and completion reserves are maintained	Finance documents and policy
Construction	Phase start follows land, power, permits and procurement gates	Technical programme and approvals
Power	Contracted and available supply supports energisation plan	Utility, PPA and onsite-generation evidence
Offtake	Capacity commitments satisfy validity and concentration tests	Executed customer contracts
Distribution	Equity distributions occur only after tests are satisfied	Waterfall and covenant package

3.3 Mathematical class

Many core relationships can be written as a mixed-integer linear programme, or MILP. Binary variables represent selection and sequencing decisions; continuous variables represent money, capacity and energy; and linear constraints enforce balances, timing and thresholds. Quadratic terms arise when the objective includes variance, pairwise concentration, nonlinear penalty interactions or a direct QUBO encoding. Some relationships are nonlinear or discontinuous and may require piecewise-linear approximations, mixed-integer nonlinear programming, decomposition or simulation outside the optimiser.

The choice of mathematical form matters. A QUBO is not automatically the best representation for the classical baseline. A constrained model may be more compact and more interpretable. Fair benchmarking gives each solver the formulation that best fits its architecture while preserving equivalent information and an equivalent decision problem [28-31].

3.4 Feasibility before optimality

Every candidate solution must satisfy hard legal, contractual, technical and policy constraints before its objective value is considered. A cheaper plan that misses an energisation gate, violates a funding-eligibility rule or exhausts minimum liquidity is infeasible. Penalty formulations used for QUBO or Ising models must be calibrated and independently tested so that apparent objective improvement cannot be purchased through an unrecognised constraint breach.

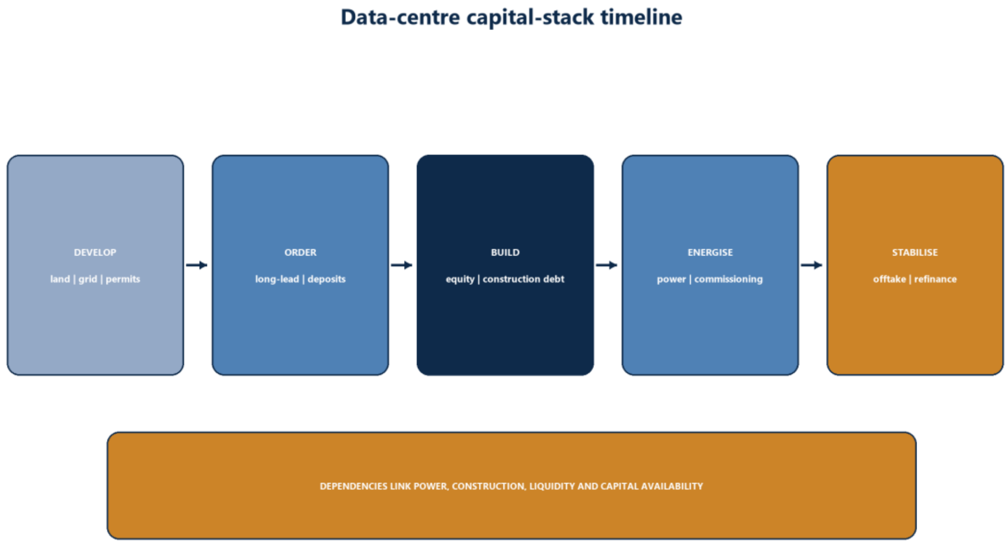

Figure 2. Data-centre capital-stack timeline and dependency graph

Source: Matchpoint framework; energy and infrastructure context [1-10,17].

4. CANONICAL DECISION MODEL

4.1 Model objects

The proposed canonical model stores each consequential assumption as a versioned object with source, owner, as-of date, unit, scenario, confidence or approval state and refresh trigger. A solver-ready matrix is derived from these objects. It is not the system of record.

Object	Minimum fields	Control purpose
Project phase	Site, building, MW, start, completion and dependency	Defines construction units
Cost item	Cost code, phase, amount, currency, timing and eligibility	Defines uses and facility access
Capital instrument	Type, commitment, price, tenor, amortisation and conditions	Defines funding choices
Draw gate	Instrument, date, evidence, approver and status	Prevents premature availability
Power source	Grid, PPA, generation or storage; volume, shape, term and price	Defines supply portfolio
Offtake contract	Customer, capacity, term, ramp, credit and termination	Defines revenue support
Reserve	Purpose, base, amount, release test and account	Protects liquidity and credit
Scenario	Named changes, probability if approved, owner and status	Separates cases from facts
Solver run	Model hash, solver, version, parameters, seed, time and result	Supports reproducibility

4.2 Time and currency

Time is represented in explicit periods, normally months for development and quarterly or annual aggregation for longer-term credit analysis. Each amount retains native currency, approved foreign-exchange source, conversion time and hedge treatment. Nominal and real values are never mixed silently. Interest conventions, day counts, compounding and fees remain explicit.

4.3 Source hierarchy

Executed contracts, lender-approved definitions, utility confirmations, board-approved budgets and validated technical schedules take precedence over marketing material, indicative terms and management estimates. Management estimates may be used in scenarios when labelled and approved. Vendor roadmaps remain vendor-supplied evidence and cannot be treated as delivered capability.

4.4 Model lineage

Every coefficient in the optimisation instance links to the approved model object from which it was derived. Transformations retain formula version, unit conversion and rounding. The system stores both the business-readable representation and the solver-ready form, enabling reviewers to trace a surprising output to the underlying term or assumption.

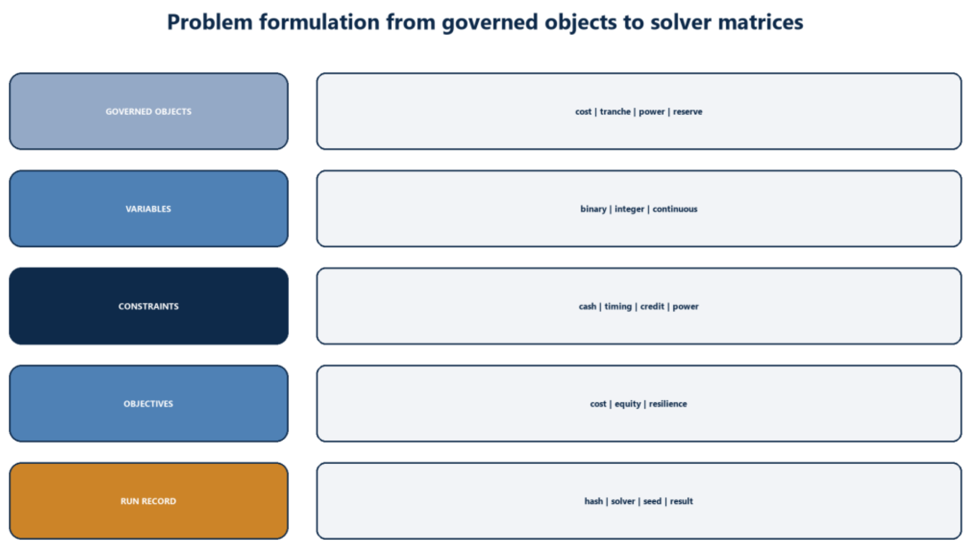

Figure 3. Problem formulation from governed objects to solver-ready matrices

Source: Matchpoint optimisation framework; modelling context [18-27,36-45,52].

5. OBJECTIVE FUNCTION AND TRADE-OFFS

5.1 Single-objective formulation

A simple formulation minimises the present value of funding cost and expected penalties:

minimise funding cost + commitment fees + hedging cost + reserve carry + expected delay cost + refinancing penalty + concentration penalty

The components must use approved definitions. Expected values require an approved scenario probability. When probabilities are unavailable, the model should report separate scenarios or robust worst-case measures instead of inventing a probability distribution.

5.2 Multi-objective formulation

Sponsors and investors rarely have one objective. A weighted formulation may combine cost, equity requirement, liquidity headroom and resilience. Weights are governance decisions, not solver facts. The preferred approach is to produce a Pareto frontier or a set of constrained alternatives:

- minimum cost subject to a required liquidity floor;
- minimum equity subject to leverage and coverage limits;
- maximum resilience subject to an approved cost cap;
- minimum refinancing exposure subject to a target return;
- minimum power concentration subject to budget and schedule constraints.

5.3 Lexicographic priority

Hard feasibility comes first. The next priority may be completion and liquidity resilience, followed by cost and return. A lexicographic solve preserves that order. It avoids the governance problem created when a small reduction in price mathematically offsets a material reduction in resilience because of an arbitrary weight.

5.4 Risk measures

The model can use scenario loss, maximum shortfall, conditional value at risk or a robust uncertainty set. Each method has data and governance requirements. Historical distributions from unrelated markets do not become project probabilities without justification. Scenario design should separate observed inputs, contractual cases, management estimates and deliberately severe stresses.

5.5 Tie-breakers

Two solutions can have the same reported objective within tolerance while differing in complexity. The model should use declared tie-breakers such as fewer facilities, fewer intercreditor dependencies, longer committed availability, higher liquidity, lower refinancing concentration or greater source diversity. The rationale remains visible to the investment or credit committee.

6. CAPITAL-INSTRUMENT VARIABLES AND CONSTRAINTS

6.1 Instrument library

The instrument library may include sponsor equity, preferred equity, shareholder loans, senior construction debt, delayed-draw term debt, bridge facilities, mezzanine or subordinated debt, equipment finance, vendor finance, working-capital facilities, letters of credit, hedging lines, project bonds, sukuk and asset-level or portfolio-level refinancing. Availability depends on jurisdiction, project facts and market engagement.

6.2 Selection and sizing

For each instrument *k*, a binary variable *y-k* records whether it is selected. A continuous variable *d-k* records commitment or drawn amount. Linking constraints enforce *0 <= d-k <= y-k x approved-limit-k*. Additional constraints represent minimum ticket, currency, cost eligibility, tenor, amortisation, collateral and cross-default relationships.

6.3 Draw order

Draw order can be contractual, economic or both. A senior facility may require a minimum equity contribution before first draw. A bridge may be repaid from construction debt after specified conditions. Equipment finance may fund only accepted hardware. The optimiser enforces each rule by period and cost category.

6.4 Fees and all-in cost

All-in funding cost includes margins, base rates, original issue discount, arrangement fees, commitment fees, undrawn fees, agency costs, hedging, legal and diligence costs where approved, reserve carry and refinancing assumptions. Taxes and accounting treatment require specialist confirmation. Indicative market terms remain unverified until supported by current correspondence.

6.5 Intercreditor and structural constraints

Structural complexity creates constraints that may not appear in a simple weighted average cost of capital. These include collateral sharing, payment blockage, cure rights, cash sweeps, permitted debt, guarantor coverage, account control, transfer restrictions and consent requirements. A structure that cannot be documented on the assumed timetable is not a valid low-cost solution.

7. DRAW-SCHEDULE OPTIMISATION

7.1 Monthly cash balance

For every period, opening cash plus permitted funding plus operating receipts equals eligible uses, debt service, reserve movements and closing cash. The minimum cash balance and designated reserves remain separate. A facility cannot fund a use before the use becomes eligible or after its availability expires.

7.2 Construction dependencies

The construction schedule is represented as a directed dependency graph. Land control, design, permits, grid studies, interconnection deposits, substation work, long-lead electrical equipment, shell construction, mechanical and electrical completion, commissioning and tenant acceptance may each gate later activities. The model must not compress professional schedules without a validated alternative programme.

7.3 Conditions precedent

Each draw gate has evidence, owner and status. The model may test a dated assumption that a condition becomes satisfied. It cannot mark a legal or technical condition satisfied. For transaction use, the current conditions tracker controls availability.

7.4 Interest during construction

Interest during construction depends on the timing of debt draws, capitalisation rules and rates. The financing need and financing cost are therefore recursive. A robust formulation may iterate or include linearised balances. The exact implementation must reconcile to the approved financial model.

7.5 Delay and liquidity

Delay changes interest, fees, contingency use, reserve duration, offtake commencement and refinancing. The model evaluates explicit delay cases. It should identify the first period and constraint that cause infeasibility, enabling management to distinguish a pricing problem from a liquidity or timing problem.

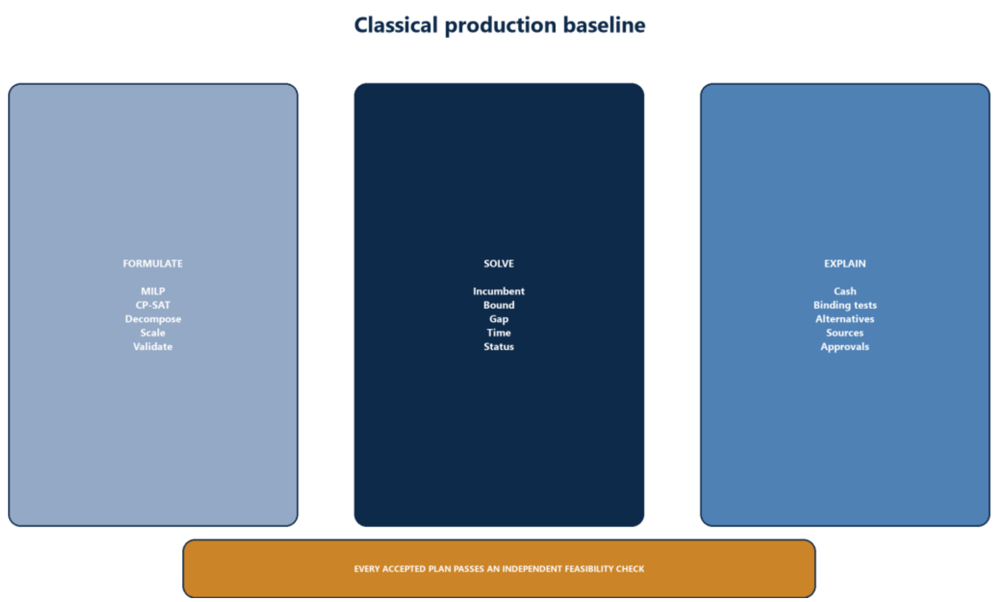

Figure 4. Classical optimisation baseline and audit trail

Source: Matchpoint framework; classical solver guidance [36-40,49-51].

8. POWER-CONTRACT PORTFOLIO

8.1 Power as a financing dependency

IEA and DOE sources describe electricity supply, grid connection and equipment availability as binding data-centre constraints [1-6]. A capital plan that assumes unavailable or uncontracted power is not financeable. The model therefore represents power as a dated portfolio of enforceable or explicitly contingent sources.

8.2 Source types

Potential sources include grid supply, utility tariff arrangements, physical or financial PPAs, onsite generation, batteries, backup generation and demand flexibility. Each source has capacity, energy, shape, availability, price, tenor, credit, emissions attributes, development risk and legal constraints. The model does not treat different products as economically identical megawatts.

8.3 Contract-selection variables

Binary variables select contracts and options. Continuous variables allocate demand by source and period. Constraints enforce contracted capacity, energy balance, ramping, outage or derating assumptions, minimum purchase, maximum volume, start date, tenor, settlement and concentration. Approved engineering studies define technical feasibility.

8.4 Shape and coincidence

Annual energy is insufficient for a high-load facility. The model must represent peak demand, ramp, maintenance, cooling load, utilisation and temporal coincidence between supply and load at a resolution supported by evidence. Coarse optimisation may screen strategies; detailed dispatch and grid analysis remain separate specialist work.

8.5 Financing interaction

Power choices affect capex, opex, credit support, deposits, letters of credit, completion date, operating risk and lender diligence. A longer PPA may support price certainty while adding termination exposure. Onsite assets may add construction and fuel risk. Flexibility may improve grid integration where technically and contractually permitted. Each effect enters through explicit model objects.

9. UNCERTAINTY, SCENARIOS AND ROBUSTNESS

9.1 Scenario register

The scenario register separates base assumptions from stresses. Examples include delayed interconnection, equipment-price change, tenant ramp delay, lower utilisation, higher base rates, refinancing margin widening, construction overrun, exchange-rate movement, power-price change and a failed funding source. Each scenario states its source, owner, approval status and affected variables.

9.2 Stochastic formulation

A two-stage formulation chooses commitments before uncertainty and draws or operating actions after a scenario is realised. This can model the value of committed liquidity or flexible contracts. Scenario probabilities must be approved. If probabilities are management estimates, the full output is labelled accordingly.

9.3 Robust formulation

Robust optimisation seeks a plan that remains feasible across an approved uncertainty set. The uncertainty set may bound cost, delay, rate and demand changes. Wider sets improve stated resilience while usually increasing cost. Management chooses the set; the solver exposes the trade-off.

9.4 Stress and reverse stress

Stress testing asks whether the plan survives a defined shock. Reverse stress testing asks which combination of changes first breaches liquidity, coverage, leverage or completion. The latter can reveal hidden dependence on one facility, one power date or one tenant ramp.

9.5 Correlation

Power delay, construction delay, capex escalation and tenant commencement can be correlated. Treating them as independent may understate tail exposure. Correlation assumptions require evidence or explicit management approval. The model should provide correlated scenario bundles rather than false precision.

10. CLASSICAL PRODUCTION BASELINE

10.1 Baseline requirement

Quantum evaluation begins after a well-engineered classical baseline exists. Google OR-Tools documents MIP, CP-SAT, branch-and-bound and SAT-based approaches; COIN-OR documents the CBC branch-and-cut solver [36-40]. These tools demonstrate that the relevant classical method depends on model structure. Commercial solvers may also be included under approved licences.

10.2 Formulation variants

The baseline suite should include:

1. a direct MILP for continuous amounts, binary choices and linear constraints;
2. CP-SAT for predominantly Boolean or integer scheduling and logical structure;
3. decomposition where master capital choices and subproblem schedules separate naturally;
4. scenario decomposition for many uncertainty cases;
5. local search or other heuristics for rapid incumbent generation;
6. deterministic rules for calculations that do not require optimisation.

10.3 Solver evidence

For exact or bound-producing solvers, the run record stores incumbent objective, best bound, optimality gap, feasibility tolerance, elapsed time, node count and termination reason. For heuristics, it stores solution quality, constraint residuals, repeated-run distribution and time-to-target. A result labelled optimal must use the solver's precise status and approved tolerances.

10.4 Numerical controls

Scaling and numerical precision matter. OR-Tools guidance notes that floating-point tolerances influence feasibility and optimality claims [39]. Money, capacity and energy units should avoid extreme coefficient ranges. Big-M constants require documented upper bounds. Penalties require sensitivity testing. Every accepted solution is rechecked by an independent deterministic validator.

10.5 Production output

The production baseline produces a business-readable plan, sources-and-uses table, draw schedule, covenant and reserve tests, binding constraints, alternative solutions, scenario results, source lineage and approval state. The solver file alone is insufficient.

11. QUANTUM ANNEALING AND D-WAVE CQM

11.1 Quadratic models

D-Wave's Ocean documentation describes Ising models, QUBO, binary quadratic models and constrained quadratic models [18-22]. A QUBO expresses a binary objective in linear and pairwise terms. Hard constraints are commonly converted into penalties when using an unconstrained representation. D-Wave's CQM abstraction supports binary, integer and real variables with linear and quadratic objectives and constraints [18-20]. These are vendor-supplied descriptions of supported software interfaces.

11.2 Annealing workflow

A direct quantum-annealing workflow prepares an Ising or QUBO problem, embeds its logical graph onto the QPU topology, samples many anneals, unembeds the results and validates feasibility. D-Wave documentation explains that minor embedding can represent one logical variable by a chain of physical qubits and that broken chains can occur [21-24]. Embedding overhead and chain-strength selection therefore belong inside the measured compute and quality boundary.

11.3 Hybrid CQM workflow

The Leap hybrid CQM solver accepts a constrained model and uses a quantum-classical service [19-20]. The user does not control every internal allocation between classical and quantum components. Any benchmark report must identify the product, version, stated properties, submission size, service time and returned metadata. A favourable hybrid result establishes performance of that delivered hybrid service on the tested instances. It does not isolate QPU advantage.

11.4 Capital-stack mapping

Capital-instrument selection, phase activation and contract selection map naturally to binary variables. Amounts, MW and cash balances may remain continuous in a CQM, or they may be discretised into binary encodings for QUBO. Discretisation adds variables and approximation error. Equality constraints, path dependencies and risk measures add couplings. Dense couplings increase direct embedding burden.

11.5 Penalty calibration

For QUBO, a hard constraint becomes a penalty. A penalty that is too small can make an infeasible solution look attractive. A penalty that is too large can compress useful objective differences or exceed practical coefficient precision. The benchmark must report penalty derivation, coefficient scaling, feasibility rate and sensitivity. Independent validation rejects every infeasible sample.

11.6 Claim boundary

This paper does not verify that D-Wave quantum annealing outperforms a classical solver on a data-centre capital-stack instance. D-Wave documentation establishes available representations and workflows. Peer-reviewed research reports problem-specific quantum-annealing results, including both favourable scaling evidence and limitations [32-35]. Those results cannot be transferred to T36 without a reproduced benchmark.

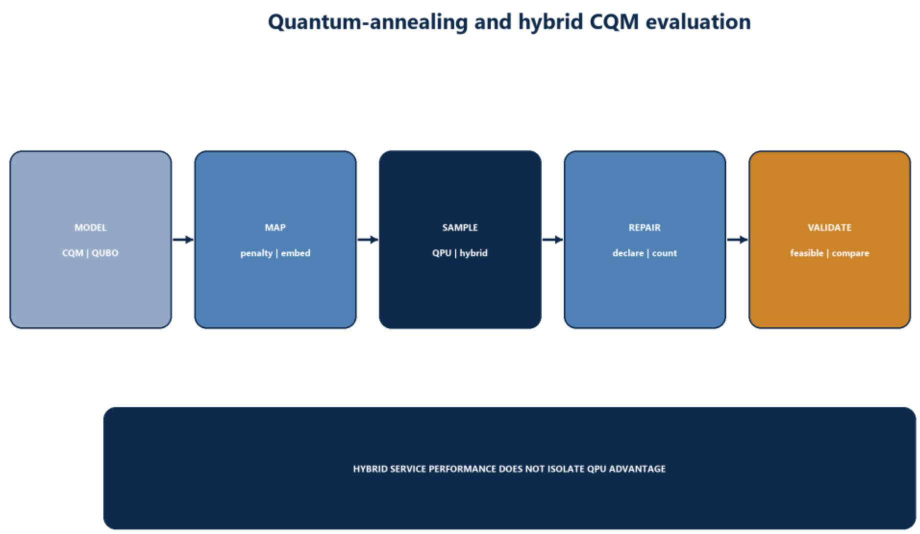

Figure 5. Quantum-annealing and hybrid CQM evaluation pipeline
Source: Matchpoint evaluation framework; D-Wave documentation and peer-reviewed research [18-24,28-35].

12. QAOA AND IBM QISKIT

12.1 Algorithm role

QAOA is a hybrid variational algorithm for combinatorial optimisation. It alternates parameterised cost and mixing operations, measures candidate bitstrings and uses a classical optimiser to update circuit parameters [25-27,41-44]. The algorithm is approximate; performance depends on problem encoding, circuit depth, parameter training, hardware noise, shot budget and classical post-processing.

12.2 QUBO and Ising mapping

The capital-stack instance must first be reduced to binary variables and a cost Hamiltonian. Continuous debt or power quantities require discretisation or decomposition. Constraints require a feasible-subspace mixer or penalties. The number of logical variables can grow with the number of instruments, periods, scenarios and precision bits. Mapping the resulting interactions to hardware connectivity adds gates and depth.

12.3 Qiskit workflow

IBM Quantum and Qiskit documentation provide tutorials for constructing QUBO problems, transpiling QAOA circuits and executing hybrid quantum-classical workflows [25-27]. Current Qiskit community materials also state the support status of specific optimisation packages and direct users to newer mapping components [27]. The implementation must pin versions and archive source, environment and transpilation outputs.

12.4 Warm starts

A classical relaxation or incumbent can initialise a quantum challenger. Qiskit documentation includes a warm-start QAOA workflow [26]. Peer-reviewed work also studies adaptive-bias and warm-start variants [45]. In a fair benchmark, the classical effort used to create the warm start belongs inside total resource accounting and is also available to a comparable classical improvement method.

12.5 Evidence from the literature

Peer-reviewed studies provide mixed, problem-specific evidence. PRX and Physical Review papers describe QAOA mechanisms, possible scaling windows and limitations [41-45]. A 2025 Physical Review Research paper concludes that QAOA's ability to deliver quantum advantage on the tested higher-order constraint-satisfaction problems remained unclear and found a mean-field classical method at least as strong on average in the examined regime [43]. A 2024 PRX Quantum paper reports a favourable scaling result for a specific random 8-SAT setting beyond a stated depth relative to the classical solver it tested [44]. Neither result is evidence for a finance problem.

12.6 Current hardware roadmap

IBM's 2026 quantum roadmap sets goals for quantum advantage, deeper circuits and later fault-tolerant systems [46-48]. IBM explicitly states that roadmap information represents current intent and may change [46]. The T36 readiness plan therefore treats roadmap dates and capabilities as **vendor-supplied goals**, not delivered project assumptions.

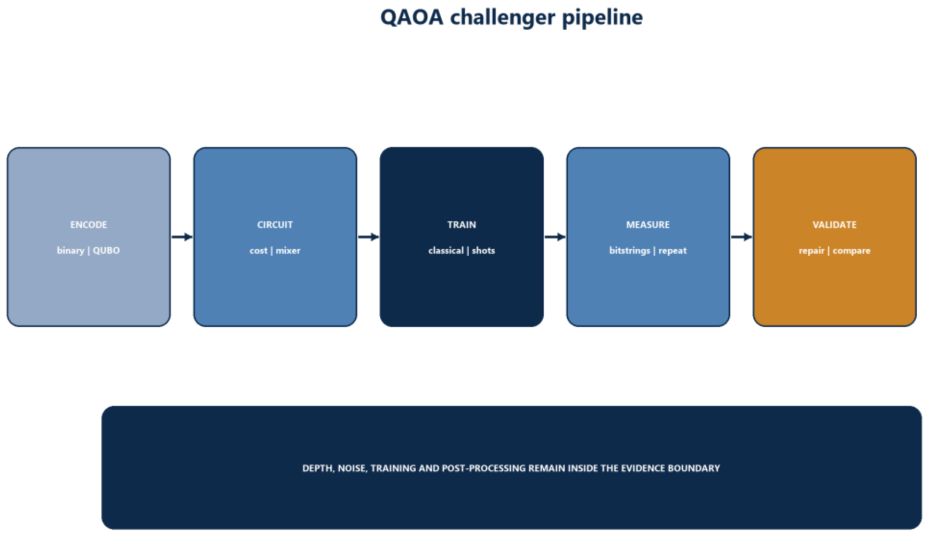

Figure 6. QAOA challenger from QUBO encoding to validated candidate plan

Source: Matchpoint evaluation framework; IBM, Qiskit and peer-reviewed research [25-31,41-48].

13. D-WAVE AND IBM LANDSCAPE

13.1 Architecture comparison

Dimension	D-Wave annealing / hybrid CQM	IBM gate-model / QAOA	Classical production baseline
Native abstraction	Ising, QUBO, BQM; CQM via hybrid service	Quantum circuits and cost Hamiltonians	MILP, CP-SAT, decomposition and heuristics
Variable treatment	Binary directly; integer and real in CQM service	Binary encoding for QAOA instance	Binary, integer and continuous directly
Constraint treatment	Explicit in CQM; penalties or specialised methods for QUBO	Penalties or constraint-preserving design	Explicit algebraic and logical constraints
Hardware mapping	Minor embedding and chains for direct QPU	Transpilation, connectivity, gates and shots	Presolve, cuts, branching and numerical tolerances
Output	Sample set or hybrid candidate set	Measured bitstring distribution	Incumbent, bound, gap and diagnostics
Key evidence gap	Isolated QPU contribution and fair baseline	Depth, noise, training and scaling	Formulation quality and instance hardness

13.2 What can be tested now

Small and medium research instances can be generated from a controlled capital-stack model. Classical exact or bound-producing solvers establish reference results where tractable. D-Wave CQM can test the constrained service representation. Direct annealing can test selected binary subproblems. IBM simulators can test QAOA formulation and parameter behaviour before approved hardware runs. Current hardware tests remain experiments; transaction decisions continue to use validated classical or professionally reviewed plans.

13.3 What should wait

Large fault-tolerant algorithm claims, deep circuit requirements and hardware capabilities outside delivered service should remain in a research backlog. A sponsor should not make a financing commitment, power contract or construction schedule depend on a vendor roadmap milestone.

13.4 Procurement neutrality

The benchmark specification should remain vendor-neutral. It defines inputs, outputs, metrics, compute boundary and acceptance tests. A supplier may propose annealing, gate-model, tensor-network, GPU, CPU, hybrid or specialist classical methods. The best validated solution within the approved service, risk and cost envelope advances.

14. ENCODING AND DECOMPOSITION

14.1 Decompose by decision horizon

The full data-centre model may be too large or too dense for a single challenger formulation. A controlled decomposition can separate strategic facility and power-contract selection from detailed monthly draws, dispatch and covenant calculation. The master problem selects a structure. Deterministic subproblems validate schedule, cash and technical feasibility.

14.2 Decompose by asset or phase

Campus phases may be modelled separately with shared constraints for grid capacity, common infrastructure, portfolio facilities and tenant concentration. Coordination variables allocate scarce capacity and liquidity. Any decomposition must converge to a globally consistent plan or state its approximation.

14.3 Binary precision

A continuous amount represented with binary bits requires a range and increment. Finer increments increase logical variables. Coarser increments create rounding and may falsely remove feasible structures. The model records precision by variable and reoptimises or validates the rounded result in the original continuous model.

14.4 Constraint-preserving design

Some constraints can be encoded through variable construction or mixers rather than penalties. Others are easier to validate after sampling. The selected method must disclose which constraints are guaranteed by representation, encouraged by penalty or tested after solution.

14.5 Repair

A classical repair heuristic may convert a near-feasible quantum sample into a feasible plan. The repair algorithm, time and objective change are part of the challenger. The benchmark reports raw feasibility and repaired feasibility separately. A repaired result cannot be presented as the direct quantum output.

15. CLASSICAL-VERSUS-QUANTUM BENCHMARK

15.1 Evidence state

No T36 benchmark was executed against an approved data-centre project dataset for this paper. The benchmark in this section is a design specification. All instance sizes, target times and acceptance thresholds require approval before execution.

15.2 Benchmark question

The question is: for a declared family of capital-stack instances, does a challenger produce feasible solutions of better quality, or reach an accepted quality sooner, than the strongest appropriate classical baseline when both receive fair formulation and tuning effort?

15.3 Instance families

The suite should contain synthetic controlled instances and permissioned real or realistically transformed instances. It varies instruments, periods, phases, power contracts, scenarios, density, precision and constraint tightness. Easy, medium and hard labels are assigned retrospectively from classical evidence, not from management expectation.

Family	Primary difficulty	Required controls
F1 tranche choice	Binary selection and fixed costs	Known optimum at small size
F2 draw schedule	Timing, eligibility and liquidity	Exact cash validator
F3 power portfolio	Contract selection and concentration	Energy-balance validator
F4 integrated stack	Capital, schedule and power couplings	Full deterministic validator
F5 scenario robust	Shared first-stage choices across cases	Scenario-by-scenario feasibility
F6 stress instances	Tight or conflicting constraints	Infeasibility certificate or diagnosis

15.4 Baselines

Each family receives a tuned MILP or CP-SAT formulation, relevant decomposition, a constructive heuristic and a local-improvement method. Commercial solvers may be included if licensed. Baseline versions, hardware, threads, memory, parameters, warm starts and tuning budget are fixed and recorded.

15.5 Challenger accounting

Quantum challenger time includes model conversion, encoding, embedding or transpilation, queue where the chosen metric requires wall-clock service, quantum execution, shots or reads, parameter optimisation, unembedding, repair, validation and repeated runs. Both access cost and energy information are recorded where available and comparable. Missing vendor data is reported as unavailable.

15.6 Metrics

Metric	Definition	Decision use
Feasibility rate	Valid solutions divided by returned samples or runs	Tests constraint control
Objective gap	Difference from best verified reference using declared denominator	Tests solution quality
Time to target	Total measured time to first accepted-quality feasible result	Tests operational utility
Time to solution	Time adjusted for success probability under declared method	Tests stochastic solvers
Stability	Distribution across seeds, runs and hardware calibrations	Tests repeatability
Scaling	Metric trend across controlled instance size and hardness	Tests growth behaviour
Total cost	Service, compute and implementation cost within declared boundary	Tests commercial relevance
Explainability	Binding constraints, lineage and reproducible run record	Tests governance readiness

15.7 Fairness rules

The robust-benchmarking literature recommends application-appropriate formulations, real and hard instances, holistic metrics and equitable hyperparameter effort [28-31]. T36 adds four finance controls: identical approved business inputs, independent feasibility validation, no exclusion of failed challenger runs, and separate reporting for raw and repaired outputs.

15.8 Advancement gate

A challenger advances from research only if it meets pre-approved thresholds across a representative suite, survives independent replication, adds decision value after full cost and integration, and preserves the audit trail. A single favourable instance does not satisfy the gate.

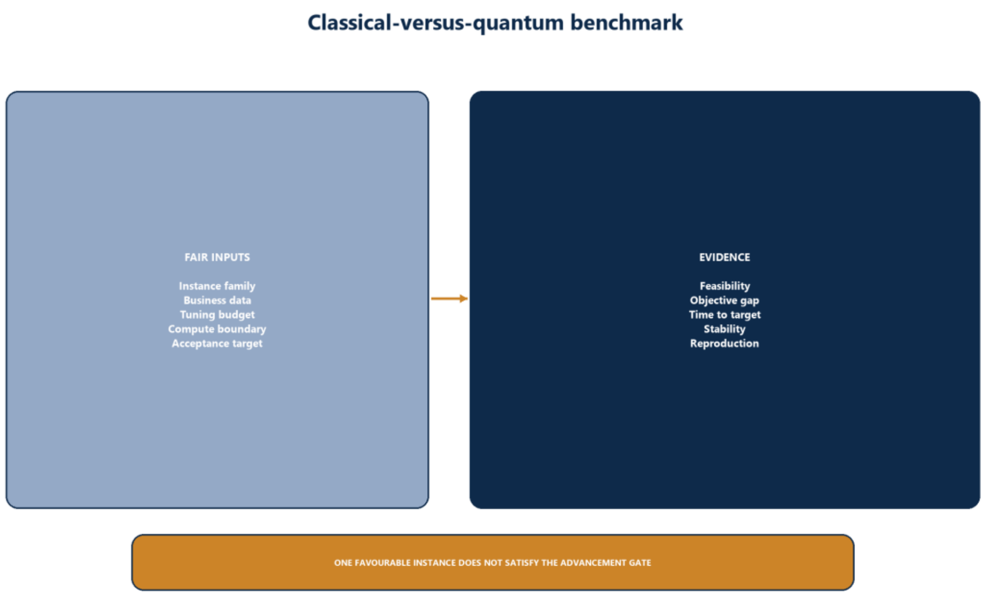

Figure 7. Classical-versus-quantum benchmark matrix and evidence gates

Source: Matchpoint benchmark framework; benchmarking research [28-31].

16. ILLUSTRATIVE CAPITAL-STACK CASE

16.1 Scenario status

This entire section is an **[Unverified illustrative scenario]**. The project, parties, prices, amounts, probabilities and outputs are invented to demonstrate the model. They are not a client case, market quote, financing offer or observed benchmark.

16.2 Project perimeter

An illustrative sponsor evaluates a three-phase Gulf data-centre campus. Phase 1 is 24 MW, Phase 2 adds 32 MW and Phase 3 adds 48 MW. The sponsor considers equity, a pre-development bridge, senior construction debt, equipment finance and a later term refinancing. Grid power, a renewable PPA, onsite generation and battery support are candidate power components.

16.3 Illustrative funding choices

Item	Illustrative value	Status
Total project uses	USD 1.20bn	Unverified illustrative
Sponsor equity availability	USD 360m	Unverified illustrative
Pre-development bridge cap	USD 90m	Unverified illustrative
Construction facility cap	USD 720m	Unverified illustrative
Equipment finance cap	USD 240m	Unverified illustrative
Minimum unrestricted liquidity	USD 45m	Unverified illustrative
Construction delay stress	9 months	Unverified illustrative

The model minimises discounted all-in funding cost subject to completion, cash, leverage, debt-service, reserve, availability, power and offtake constraints. Equity must fund specified early costs. The construction facility requires approved gates. Equipment finance funds eligible accepted equipment. Refinancing cannot occur before an illustrative stabilisation test.

16.4 Illustrative alternatives

Alternative A uses lower initial equity and heavier bridge reliance. It is cheaper in the base case under the invented inputs and becomes infeasible under the nine-month delay because the bridge expires before construction debt availability. Alternative B commits more equity and liquidity earlier. It has a higher base-case carrying cost and remains feasible in the stated delay. Alternative C adds a larger committed facility and a power-contract option; it has the highest stated fees and the lowest reliance on one energisation date.

No numeric solver advantage is stated. A classical solver would establish feasibility, objective and bound for the model. Quantum challengers would receive controlled encodings of selected instance families. The business decision remains an approved choice among transparent alternatives.

16.5 Binding constraints

The useful output is often the constraint diagnosis. In the illustrative case, the binding items may include equity-before-debt, interconnection deposit timing, equipment advance rate, minimum liquidity and bridge maturity. A committee can challenge these items and request real counterparty alternatives. The optimiser cannot create commercial terms.

16.6 Sensitivity

The paper's scenario chart displays qualitative movement across cost, liquidity, completion and refinancing exposure. It contains no observed client data. Before transaction use, every axis requires approved units, source and tolerance.

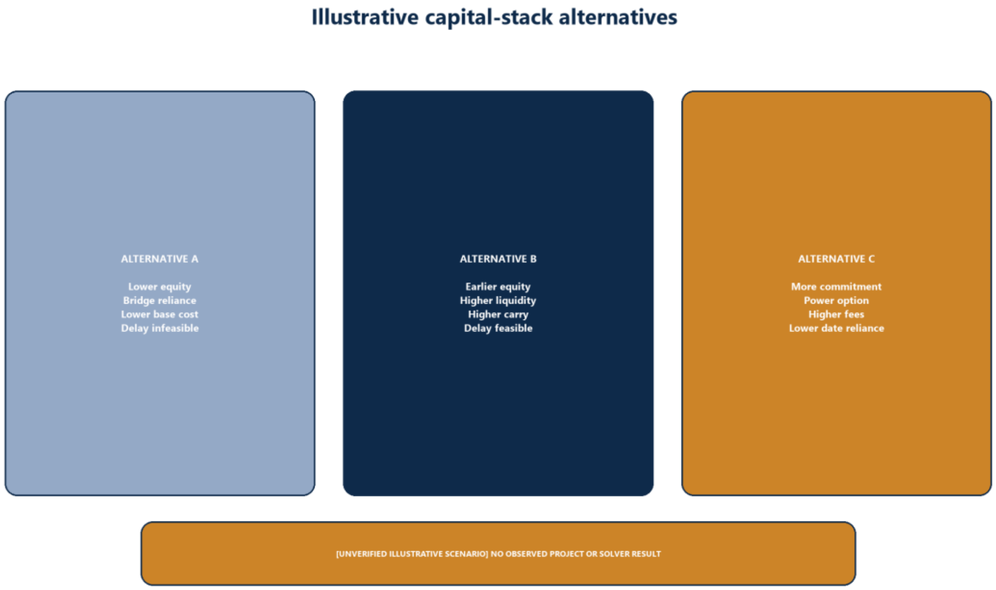

Figure 8. Illustrative capital-stack alternatives and binding constraints

Source: Matchpoint [Unverified illustrative scenario]; no observed project or solver result.

17. GOVERNANCE, VALIDATION AND SECURITY

17.1 Three lines of control

The first line owns the business model and operating use. Independent model validation challenges data, formulation, code, numerical behaviour, benchmark design and limitations. Internal audit or another independent assurance function evaluates governance and control design. The exact organisational allocation depends on approved policy.

17.2 Model-risk standard

Federal Reserve and OCC model-risk guidance emphasises robust development, effective challenge and sound governance [49]. NIST's AI Risk Management Framework provides a broader governance, mapping, measurement and management structure [50]. These sources are applied as control references; legal applicability requires specialist confirmation.

17.3 Validation tests

Validation includes unit tests, balance tests, dimensional checks, known-optimum instances, infeasible cases, perturbation tests, dual or bound reconciliation where available, solver cross-checks, seed repetition, penalty sensitivity, encoding precision, repair analysis and independent business review. Every material defect receives an owner and resolution state.

17.4 Quantum-specific controls

Quantum tests record provider, backend, calibration context if available, queue and execution timestamps, circuit or embedding, number of shots or reads, parameters, seed, post-processing and raw result reference. Vendor-managed hybrid services record all metadata the service exposes and explicitly state what cannot be observed.

17.5 Data protection and confidentiality

Financing terms, tenant contracts, grid correspondence, equipment pricing and credit information may be confidential. Data minimisation, access control, encryption, regional processing, retention, supplier terms and incident response require approval. A public or shared quantum service receives no project data unless the approved security and legal process permits it.

17.6 Software supply chain

Open-source packages, solver binaries, cloud SDKs and generated environments are pinned and scanned. NIST's Secure Software Development Framework provides control context [51]. The production repository stores source, lock files, build records and checksums. Secrets remain outside code and reports.

17.7 Decision record

The approved decision record contains the considered alternatives, sources, model version, solver runs, validation status, binding constraints, sensitivities, professional advice, conflicts, approvals and reasons for any management override. A solver recommendation never replaces that record.

18. QUANTUM-READY OPERATING MODEL

18.1 Stage 0: model hygiene

Create canonical objects, identifiers, units, version controls, deterministic calculations and a reconciled classical model. Establish data permissions and decision rights. No quantum work proceeds with unreconciled sources-and-uses or undefined constraints.

18.2 Stage 1: research instances

Generate small instances with known optima and progressively harder synthetic instances. Test QUBO, CQM and QAOA mappings on simulators and approved services. Report failures and full resource accounting.

18.3 Stage 2: shadow benchmark

Run classical and quantum challengers on permissioned historical or realistic transformed inputs. Outputs remain outside the live transaction process. Independent validators reproduce a sample and assess decision relevance.

18.4 Stage 3: analyst support

A validated challenger may propose alternatives to an analyst while the classical production model remains authoritative. The analyst sees lineage, feasibility, objective difference, uncertainty and limitations. All accepted changes are re-solved and validated in the production model.

18.5 Stage 4: controlled production component

Only a narrow component with approved evidence, fallback, service-level terms and monitoring can enter production. Examples may include candidate generation for a bounded binary subproblem. A deterministic validator and classical fallback remain mandatory.

18.6 Exit and pause criteria

Research pauses if data cannot support the formulation, classical methods solve all material instances within the required time, quantum mapping overhead dominates, feasibility is unreliable, provider metadata is insufficient, confidentiality cannot be approved or expected decision value does not justify cost.

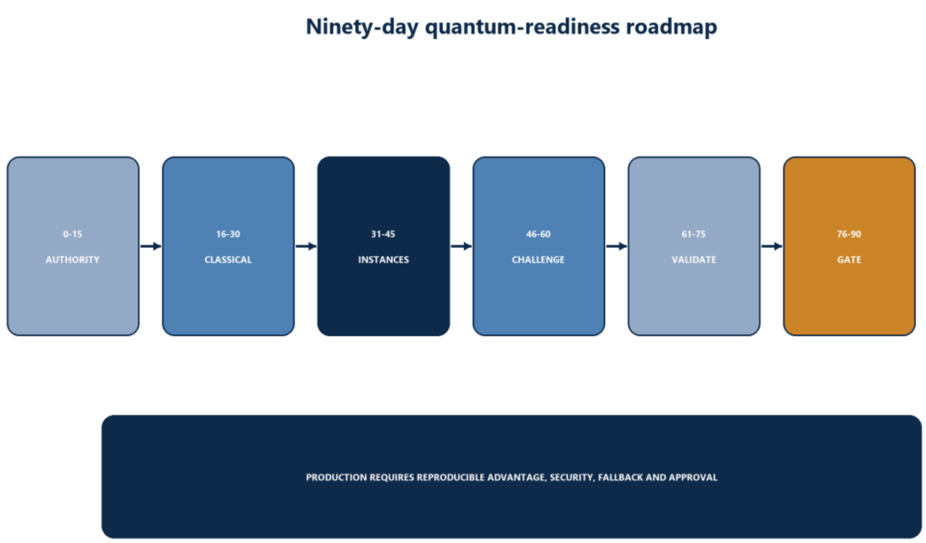

Figure 9. Quantum-readiness roadmap from classical baseline to controlled component

Source: Matchpoint implementation framework; governance and roadmap context [31,46-51].

19. ECONOMICS, CAPACITY AND IMPLEMENTATION

19.1 Evidence boundary

No approved Matchpoint or client dataset establishes time saved, cost reduction, financing improvement, prevented loss, additional revenue or alpha from the T36 framework. Attributed values remain USD 0. Vendor service prices, solver licences and internal effort require current quotes and approved records.

19.2 Measurement model

Measure	Definition	Evidence source
Model preparation time	Accepted analyst time from source freeze to solver-ready instance	Approved time record
Feasible alternative yield	Distinct accepted plans within the decision window	Validated run register
Time to accepted result	Total controlled time to the first feasible result meeting target	Benchmark telemetry
Review effort	Specialist time to reproduce and approve a result	Review log
Decision change	Approved change attributable to new validated evidence	Decision record
Provider cost	Measured service and data-transfer cost	Invoice or provider record
Prevented exposure	Approved exposure avoided because a valid control blocked a decision	Approved incident evidence

19.3 Ninety-day research plan

Days 0 to 15: authority and data. Confirm ICP use cases, source owners, confidentiality, decision rights, approved classical tools and model-risk governance.

Days 16 to 30: canonical classical model. Build and reconcile a small MILP or CP-SAT instance with deterministic validators, known-optimum tests and business-readable outputs.

Days 31 to 45: instance library. Create controlled tranche, draw, power and integrated instance families. Archive generators, hashes and reference solutions.

Days 46 to 60: D-Wave and QAOA experiments. Build CQM, QUBO and QAOA challengers, starting with simulators and approved test services. Measure the complete resource boundary.

Days 61 to 75: independent validation. Reproduce selected runs, challenge fairness, test penalty and precision sensitivity, and assess whether outputs add decision value.

Days 76 to 90: management gate. Stop, extend research or approve a tightly scoped shadow component. Record evidence, limitations, cost and next acceptance criteria.

19.4 Minimum gates

Gate	Required evidence
Business	Named decision, owner, service window and acceptance target
Data	Permissioned sources, lineage, freshness and reconciliation
Classical	Strong baseline, validation and reproducible run records
Quantum	Declared mapping, hardware or service, total accounting and raw outputs
Fairness	Equivalent information and equitable tuning budget
Security	Approved provider, access, retention and incident controls
Validation	Independent feasibility and replication evidence
Economics	Observed cost and approved decision-value evidence

20. LIMITATIONS AND CONCLUSION

20.1 Limitations

This paper is a research and operating-design document. It does not provide legal, regulatory, engineering, power-market, environmental, accounting, tax, valuation, credit, investment or quantum-technology advice. Data-centre projects differ by jurisdiction, grid, customer, design, equipment, financing and contract.

The proposed model simplifies construction, power dispatch, tenant credit, tax, accounting, collateral value, refinancing and market execution. Specialist models may need to remain separate and exchange approved boundary values. A compact QUBO may omit detail that is material to underwriting. A high-quality candidate can still be commercially unavailable.

The quantum evidence is evolving and problem-specific. D-Wave and IBM materials describe current products, documentation and roadmaps; capability and roadmap claims from vendors remain vendor-supplied. Peer-reviewed results on MaxCut, SAT, spin models or other benchmarks do not establish advantage on capital-stack optimisation. Hardware noise, embedding, transpilation, parameter training, discretisation, queueing, repair and classical post-processing can alter measured performance.

No approved T36 project dataset or live quantum benchmark was available for this paper. The worked case and all operating values are unverified illustrative scenarios. P109 and P110 were not started at reconciliation and are not represented as published companion research.

20.2 Conclusion

Data-centre capital-stack design is a suitable subject for disciplined optimisation because capital instruments, construction phases, power contracts, reserves and covenant tests form a connected decision system. The business model should begin with controlled source data, explicit authority, deterministic calculations and a well-engineered classical baseline.

Quantum annealing, hybrid CQM and QAOA offer research pathways for selected combinatorial components. Their mathematical fit is testable now. Their transaction value remains unverified for this use case. A credible programme compares them with the strongest relevant classical methods, counts the full computation and repair boundary, validates every candidate independently and advances only on pre-approved evidence.

The recommended first implementation is one capital-stack decision, one approved instance library, one classical production model and two bounded challenger tracks. The sponsor and investors retain authority. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

REFERENCES

[1] International Energy Agency. Energy and AI. 10 April 2025. https://www.iea.org/reports/energy-and-ai/
[2] International Energy Agency. Electricity 2025. 2025. https://www.iea.org/reports/electricity-2025
[3] International Energy Agency. Data centre electricity use surged in 2025, even with tightening bottlenecks driving a scramble for solutions. 16 April 2026. https://www.iea.org/news/data-centre-electricity-use-surged-in-2025-even-with-tightening-bottlenecks-driving-a-scramble-for-solutions
[4] United States Department of Energy. DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers. 20 December 2024. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers
[5] Shehabi, A. et al. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory, 2024. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf
[6] United States Department of Energy. Recommendations on Powering Artificial Intelligence and Data Center Infrastructure. 2024. https://www.energy.gov/sites/default/files/2024-11/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf
[7] World Bank. Infrastructure Monitor 2024: Blended Finance and Guarantees in Infrastructure. 2025. https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099042225161531787
[8] World Bank Group. Building Data Infrastructure for AI Readiness. 6 May 2026. https://www.worldbank.org/en/results/2026/05/06/data-infrastructure-for-ai
[9] International Finance Corporation. Data Centers. Sector overview. https://www.ifc.org/en/what-we-do/sector-expertise/telecommunications-media-and-technology/data-centers
[10] International Organization for Standardization. ISO/IEC 22237-1:2021 Data centre facilities and infrastructures, Part 1: General concepts. https://www.iso.org/standard/78550.html
[11] Basel Committee on Banking Supervision. Principles for effective risk data aggregation and risk reporting. BIS, 2013. https://www.bis.org/publ/bcbs239.htm
[12] Principles for Responsible Investment. Infrastructure. Investor resources. https://www.unpri.org/infrastructure
[13] Project Management Institute. The Standard for Risk Management in Portfolios, Programs, and Projects. https://www.pmi.org/standards/risk-management
[14] International Energy Agency. World Energy Outlook 2025: Executive summary. https://www.iea.org/reports/world-energy-outlook-2025/executive-summary
[15] International Energy Agency. Global Energy Review 2026: Global trends. https://www.iea.org/reports/global-energy-review-2026/global-trends
[16] International Energy Agency. Electricity Mid-Year Update 2026. https://www.iea.org/reports/electricity-mid-year-update-2026
[17] World Bank. Guidance on PPP Contractual Provisions. https://ppp.worldbank.org/public-private-partnership/library/guidance-ppp-contractual-provisions-2019
[18] D-Wave Quantum. Models. Ocean documentation. https://docs.dwavequantum.com/en/latest/concepts/models.html
[19] D-Wave Quantum. CQM Solver Parameters. Industrial optimisation documentation. https://docs.dwavequantum.com/en/latest/industrial_optimization/solver_cqm_parameters.html
[20] D-Wave Quantum. Samplers: LeapHybridCQMSampler. Ocean documentation. https://docs.dwavequantum.com/en/latest/ocean/api_ref_system/samplers.html
[21] D-Wave Quantum. Ocean Software Stack. https://docs.dwavequantum.com/en/latest/ocean/stack.html
[22] D-Wave Quantum. Concepts. https://docs.dwavequantum.com/en/latest/concepts/index.html
[23] D-Wave Quantum. Minor Embedding. https://docs.dwavequantum.com/en/latest/quantum_research/embedding_intro.html
[24] D-Wave Quantum. Minor-Embedding: Best Practices. https://docs.dwavequantum.com/en/latest/quantum_research/embedding_guidance.html
[25] IBM Quantum. Quantum approximate optimization algorithm. IBM Quantum documentation. https://qiskit.qotlabs.org/docs/tutorials/quantum-approximate-optimization-algorithm
[26] Qiskit Community. Warm-starting quantum optimization. Qiskit Optimization 0.7.0. https://qiskit-community.github.io/qiskit-optimization/tutorials/10_warm_start_qaoa.html

[27] Qiskit Community. Qiskit Optimization. Source repository and support notice. https://github.com/qiskit-community/qiskit-optimization
[28] Bucher, D. et al. Towards Robust Benchmarking of Quantum Optimization Algorithms. arXiv:2405.07624, 2024. https://arxiv.org/abs/2405.07624
[29] Lubinski, T. et al. Application-Oriented Performance Benchmarks for Quantum Computing. IEEE Transactions on Quantum Engineering 4 (2023). https://doi.org/10.1109/TQE.2023.3253761
[30] Amico, M. et al. Defining standard strategies for quantum benchmarks. arXiv:2303.02108, 2023. https://arxiv.org/abs/2303.02108
[31] DARPA. Quantum Benchmarking Initiative. https://www.darpa.mil/research/programs/quantum-benchmarking-initiative
[32] Pelofske, E., Baertschi, A. and Eidenbenz, S. Short-depth QAOA circuits and quantum annealing on higher-order Ising models. npj Quantum Information 10, 30 (2024). https://doi.org/10.1038/s41534-024-00825-w
[33] Munoz-Bauza, H. and Lidar, D. Scaling Advantage in Approximate Optimization with Quantum Annealing. Physical Review Letters 134, 160601 (2025). https://doi.org/10.1103/PhysRevLett.134.160601
[34] Katzgraber, H. G. et al. Unraveling Quantum Annealers using Classical Hardness. Scientific Reports 5, 15324 (2015). https://doi.org/10.1038/srep15324
[35] Kadowaki, T. and Nishimori, H. Quantum annealing in the transverse Ising model. Physical Review E 58, 5355 (1998). https://doi.org/10.1103/PhysRevE.58.5355
[36] Google. OR-Tools: Integer Optimization. https://developers.google.com/optimization/mip
[37] Google. OR-Tools. https://developers.google.com/optimization
[38] Google. Solving a MIP Problem. https://developers.google.com/optimization/mip/mip_example
[39] Google. Advanced LP Solving. https://developers.google.com/optimization/lp/lp_advanced
[40] COIN-OR. CBC: Introduction. https://coin-or.github.io/Cbc/intro.html
[41] Farhi, E., Goldstone, J. and Gutmann, S. A Quantum Approximate Optimization Algorithm. arXiv:1411.4028, 2014. https://arxiv.org/abs/1411.4028
[42] Zhou, L. et al. Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. Physical Review X 10, 021067 (2020). https://doi.org/10.1103/PhysRevX.10.021067
[43] Mueller, T. et al. Limitations of quantum approximate optimization in solving generic higher-order constraint-satisfaction problems. Physical Review Research 7, 023165 (2025). https://doi.org/10.1103/PhysRevResearch.7.023165
[44] Boulebnane, S. and Montanaro, A. Solving Boolean Satisfiability Problems With the Quantum Approximate Optimization Algorithm. PRX Quantum 5, 030348 (2024). https://doi.org/10.1103/PRXQuantum.5.030348
[45] Yu, Y. et al. Warm-start adaptive-bias quantum approximate optimization algorithm. Physical Review A 112, 012422 (2025). https://doi.org/10.1103/PhysRevA.112.012422
[46] IBM. IBM Quantum Roadmap. 2026. Vendor roadmap. https://www.ibm.com/roadmaps/quantum/
[47] IBM Quantum. Scaling for quantum advantage and beyond. 2025. Vendor update. https://www.ibm.com/quantum/blog/qdc-2025
[48] IBM Quantum. How IBM will build the world's first large-scale, fault-tolerant quantum computer. 10 June 2025. Vendor roadmap update. https://www.ibm.com/quantum/blog/large-scale-ftqc
[49] Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency. Supervisory Guidance on Model Risk Management, SR 11-7 / OCC 2011-12. https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm
[50] Tabassi, E. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1, 2023. https://doi.org/10.6028/NIST.AI.100-1
[51] National Institute of Standards and Technology. Secure Software Development Framework (SSDF) Version 1.1. NIST SP 800-218. https://doi.org/10.6028/NIST.SP.800-218
[52] Lucas, A. Ising formulations of many NP problems. Frontiers in Physics 2, 5 (2014). https://doi.org/10.3389/fphy.2014.00005

APPENDIX A. CAPITAL-STACK MODEL OBJECT

Field	Type	Required control
model_id	Stable string	Unique and immutable
project_id	Reference	Approved project perimeter
as_of	Timestamp	Source cut-off and time zone
currency	Code	Native and reporting currency
periods	Ordered list	Start, end and granularity
phases	Reference list	Approved technical programme
uses	Object list	Cost, timing, eligibility and source
instruments	Object list	Terms, conditions, pricing and availability
power_contracts	Object list	Capacity, energy, shape, price and credit
offtake	Object list	Capacity, ramp, term and counterparty state
reserves	Object list	Base, amount, account and release test
scenarios	Object list	Source, approval and affected variables
objective	Versioned expression	Approved measure and priorities
constraints	Versioned list	Source, owner, hard or soft treatment
validation_state	Enum	Draft, tested, challenged or approved

APPENDIX B. BENCHMARK RUN MANIFEST

Field	Description
run_id	Unique execution identifier
instance_hash	Immutable benchmark instance hash
business_model_version	Source model and transformation version
solver	Product, algorithm and version
execution_target	CPU, GPU, simulator, QPU or hybrid service
hardware	Processor, memory, threads, backend or QPU identity if available
formulation	MILP, CP-SAT, CQM, QUBO, Ising or QAOA
encoding	Precision, penalties, embedding or circuit mapping
parameters	Full approved solver or algorithm configuration
seed	Random seed or documented absence
preparation_time	Measured formulation and mapping time
execution_time	Measured solve, sample or circuit time
post_time	Unembedding, repair and validation time
queue_time	Recorded separately where available
result_hash	Immutable raw output reference
feasibility	Raw and repaired status
objective	Validated business objective
bound_gap	Best bound and declared gap where available
reviewer	Independent review identity and state

APPENDIX C. CAPITAL-STACK VALIDATION CHECKLIST

- [ ] Project, currency, period and phase identifiers are approved.
- [ ] Sources and uses reconcile by period and in total.
- [ ] Every funding use is eligible under the stated source.
- [ ] Draw gates and facility availability dates are current.
- [ ] Interest, fees, hedging and reserve carry reconcile to the financial model.
- [ ] Power capacity, energy, shape and connection dates are professionally reviewed.
- [ ] Offtake assumptions link to executed or explicitly labelled scenario evidence.
- [ ] Leverage, coverage, reserve and distribution tests use approved definitions.
- [ ] Every candidate passes the independent feasibility validator.
- [ ] Classical baselines received appropriate formulation and tuning effort.
- [ ] Quantum preparation, mapping, execution, repair and validation are fully counted.
- [ ] Vendor roadmap statements are labelled vendor-supplied.
- [ ] Raw and repaired challenger outputs are reported separately.
- [ ] Worked examples and management estimates are labelled unverified.
- [ ] Decision authority and required professional reviews are recorded.
- [ ] Attributed Matchpoint or client economic benefit remains USD 0 absent approved evidence.

APPENDIX D. BENCHMARK ACCEPTANCE RECORD

Question	Required answer
What decision is being improved?	Named capital-stack decision and owner
What is the service window?	Approved elapsed-time requirement
What is the reference quality?	Verified optimum, bound or accepted target
Which classical methods were used?	Versions, formulations and tuning budget
Which quantum method was used?	Encoding, backend, parameters and resource boundary
Were all outputs feasible?	Raw rate, repair rate and residuals
Was the result reproduced?	Independent run and evidence reference
Did it change a decision?	Approved decision record or no observed change
What did it cost?	Approved measured cost or unavailable
May it enter production?	Named approval, scope, fallback and expiry