

Quantum Optimisation for Logistics and Supply-Chain Finance

A benchmark-first operating framework for GCC family businesses and private-credit providers

Prepared by Matchpoint Partners

Author attribution pending CK approval

August 2026

ABSTRACT

Background. Logistics, inventory and working-capital decisions share orders, delivery, acceptance, invoice and settlement events. **Objective.** This paper develops a controlled framework for testing quantum optimisation in GCC SME and family-business logistics and in private-credit supply-chain-finance decisions. **Approach.** The analysis reviews 40 official, standards, implementation and primary research sources available through 1 August 2026. **Findings.** Current evidence supports bounded, benchmark-first hybrid experiments. It does not establish broad production quantum advantage. Strong classical baselines, event quality, constraint ownership, independent feasibility checks and component attribution are prerequisites. **Implications.** The accountable unit is a quality-adjusted accepted optimisation decision packet. Attributed revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.

Highlights

Unit. An accepted optimisation decision packet connects solver output to a feasible, reviewable action.

Baseline. Exact or bounded methods, strong classical heuristics and the current production workflow are retained.

Attribution. Decomposition, warm starts, quantum execution, repair and validation are measured separately.

Authority. Operational, legal, accounting, compliance, credit and investment decisions remain with named owners.

Economics. All worked scenarios are illustrative; commercial attribution remains zero without observed evidence.

JEL Classification: C61, C63, G21, G23, L91, M11, O33

Keywords: quantum optimisation, quantum annealing, QAOA, vehicle routing, logistics, supply-chain finance, trade finance, receivables, working capital, GCC SMEs, private credit

Research paper only. This document is not investment, legal, regulatory, accounting, tax, audit, sanctions, AML, cybersecurity, credit, valuation or technology advice. All worked inputs and thresholds are unverified illustrative management assumptions.

1. INTRODUCTION

Logistics and working capital are connected decisions. A GCC family business buys materials, holds inventory, books transport, serves customers, issues invoices, approves supplier invoices, collects receivables and draws or repays facilities. Each event changes both the physical position and the cash position. A routing plan that reduces kilometres can still damage value if it misses a delivery window, delays acceptance, blocks invoicing or concentrates service risk. A financing plan that extends buyer payment terms can still weaken the supply chain if smaller suppliers cannot fund the additional days.

The optimisation problem is therefore broader than route length. It can include vehicles, warehouses, inventory, production slots, customer priorities, delivery windows, invoice eligibility, supplier payment choices, borrowing limits, advance rates, concentration limits, expected dilution, liquidity reserves and credit authority. Many of these choices are discrete. Their combinations can grow quickly, especially when decisions are linked across locations, days and counterparties.

Quantum optimisation offers methods for selected combinatorial problems. Quantum annealing seeks low-energy states of an Ising or quadratic unconstrained binary optimisation formulation. Gate-model approaches such as the Quantum Approximate Optimization Algorithm use a hybrid quantum-classical loop to sample candidate binary solutions [20-24]. Recent logistics studies have tested vehicle routing, package delivery and automated-guided-vehicle problems using simulators, hybrid solvers and physical annealing systems [25-34]. These studies are useful evidence. Their results remain highly dependent on formulation, decomposition, constraints, classical orchestration and the tested instance.

This paper develops a controlled framework for **Quantum Optimisation for Logistics and Supply-Chain Finance**. The primary ideal customer profile is B4, GCC SME & Family-Business Owners: owners, managing directors and second-generation operators of businesses with indicative revenues of AED 10 million to AED 250 million. The secondary profile is A3, Private Credit, Direct-Lending & Special-Situations Funds: international credit and special-situations managers seeking GCC opportunities. These profiles, the title, hook, visual requirements, target length and Tech & AI service mapping were verified from Matchpoint Partners' Topic Tracker and ICP Legend on 1 August 2026.

The paper's executive conclusion is measured. Current evidence supports benchmark-first hybrid experimentation on carefully selected logistics subproblems. The reviewed evidence does not establish a broad production quantum advantage for a GCC SME's routing, inventory or supply-chain-finance decisions. Strong classical optimisation, better event data and disciplined operating controls are the immediate foundations. A quantum component deserves continuation only when it adds reproducible value against those foundations on representative workloads.

The accountable work unit is a **quality-adjusted accepted optimisation decision packet**. It contains the decision question, data snapshot, objective, hard and soft constraints, solver configuration, candidate plan, feasibility certificate, comparator result, exceptions, sensitivity analysis, named reviewer disposition and released action. Runs, shots, qubits, variables and solver time are technical operating measures. Accepted packets connect technical work to an action that a business owner, operations lead, finance lead and credit authority can reproduce and approve.

Seven propositions organise the analysis. First, a complete decision contract and reliable event data matter before the choice of solver. Second, hard feasibility outranks a lower model objective. Third, every quantum or hybrid result requires a strong classical comparator and component-level attribution. Fourth, physical and financial flows can share a model while retaining separate legal, accounting and credit authority. Fifth, decomposition, repair and post-processing are part of the method and its measured cost. Sixth, operational value requires representative shadow testing, stability and acceptance by named users. Seventh, attributed revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.




Figure 1. Before and after logistics and supply-chain-finance workflow

Source: Matchpoint operating framework.

The evidence cut-off is 1 August 2026. The paper uses official laws, regulator material, standard-setter resources, official industry definitions, primary research and implementation documentation. It contains no Matchpoint client route plan, supplier ledger, facility agreement, borrowing-base certificate, invoice file, observed pilot result or approved vendor benchmark. Worked volumes, rates, time inputs, costs, savings and thresholds are unverified illustrative management assumptions. Legal, accounting, tax, financial-crime, sanctions, cybersecurity, credit and investment determinations remain with appropriately qualified owners.

2. SCOPE, DEFINITIONS AND EVIDENCE BOUNDARIES

2.1 Physical and financial flows

A physical flow describes goods and services moving through orders, production, storage, transport, delivery, inspection, acceptance and returns. A financial flow describes purchase obligations, invoices, credit notes, receivables, payables, discounts, advances, collections, facility utilisation and settlement. A data flow records the events and documents that support both.

The optimisation model should connect these flows through verified events. A dispatched shipment does not establish customer acceptance. An issued invoice does not establish eligibility under a receivables purchase agreement. An approved supplier invoice does not establish that a finance provider will purchase it. A warehouse receipt may evidence stored goods within its legal and operational scope; ownership, control, priority and enforceability require separate analysis.

Flow	Example state	Decision relevance	Required evidence owner
Order	Confirmed quantity, price, due date and customer	Demand, priority and expected cash event	Sales and commercial owner
Inventory	SKU, lot, location, quantity, age and status	Availability, carrying cost and collateral	Operations and finance
Shipment	Vehicle, load, route, milestone and exception	Capacity, service and estimated delivery	Logistics owner
Acceptance	Delivered, inspected, accepted or disputed	Invoice trigger and dilution risk	Customer-service or contract owner
Invoice	Issuer, debtor, amount, tax, due date and status	Receivable, payable and programme eligibility	Finance and tax owner

Flow	Example state	Decision relevance	Required evidence owner
Finance	Facility, advance, rate, limit, concentration and maturity	Liquidity, cost and lender exposure	Treasury and credit authority
Settlement	Payment amount, date, allocation and exception	Cash position and exposure reduction	Treasury and reconciliation owner

2.2 Supply-chain-finance techniques

The Global Supply Chain Finance Forum defines supply-chain finance as the use of financing and risk-mitigation practices and techniques to optimise working capital and liquidity invested in supply-chain processes and transactions. The techniques are generally applied to open-account trade, triggered by supply-chain events and dependent on visibility of the underlying trade flows [1]. The Forum groups techniques across receivables purchase, loan or advance-based approaches and enabling frameworks.

Payables finance is a buyer-led programme in which participating sellers can elect to receive the discounted value of approved receivables before their due dates, commonly at a financing cost aligned with the buyer's credit risk [2]. Receivables discounting is seller-led financing based on receivables owed by one or more buyers [3]. Factoring, forfaiting, distributor finance, inventory finance, pre-shipment finance and loans against receivables have different structures, rights and risks. The model must identify the technique actually used.

Technique	Principal trigger or asset	Common decision variable	Boundary that optimisation cannot decide
Payables finance	Buyer-approved invoice	Supplier election, funding allocation and payment timing	Legal effect of buyer undertaking and accounting treatment
Receivables discounting	Eligible receivable	Invoice selection, price and concentration	True sale, assignment, perfection and enforceability
Factoring	Transferred receivables plus servicing terms	Advance, reserve and collection allocation	Recourse, notice, priority and insolvency treatment
Inventory finance	Eligible goods and control structure	Borrowing base, advance and reserve	Ownership, inspection, warehouse control and security validity
Pre-shipment finance	Purchase order or production need	Funding amount, milestone and tenor	Performance risk and lender authority
Dynamic discounting	Buyer-funded early payment	Discount curve and supplier choice	Supplier consent, tax and accounting treatment
Trade loan	Underlying trade and borrower obligation	Draw, tenor and repayment allocation	Credit approval and product classification

The Asian Development Bank measured a global trade-finance gap of USD 2.5 trillion for 2022 using responses from 137 banks in 54 countries and 185 companies in 43 countries [4]. The result is global and historic; it does not quantify the unmet demand of the B4 profile or any specific GCC market in 2026. It supports the importance of access to trade finance and the need to measure local demand with current evidence.

2.3 Supplier-finance disclosure and legal identity

The IFRS Foundation records that the International Accounting Standards Board issued Supplier Finance Arrangements amendments to IAS 7 and IFRS 7 in May 2023. The amendments require additional disclosures about supplier-finance arrangements [5,6]. The optimisation layer should therefore preserve arrangement identity, terms, carrying amounts, due-date information and liquidity-risk data needed by reporting owners. It should not label a liability as trade payable or borrowing.

The UAE Federal Decree-Law No. 16 of 2021 defines factoring and transfer of receivables, including current and future receivables and transfers for security or sale [9]. The legal effect of a particular assignment depends on the agreement, underlying contract, notices, priority, governing law and facts. The system can assemble and test the data required by counsel and credit teams. Their conclusions remain authoritative.

UNCITRAL's Model Law on Electronic Transferable Records addresses functional equivalence, control and integrity for electronic records that correspond to transferable documents or instruments [7]. The UNCITRAL-UNIDROIT Model Law on Warehouse Receipts supports paper and electronic warehouse receipts and the use of

stored goods in commercial transactions and as collateral [8]. These model laws are reference frameworks. Their legal force depends on enactment and the applicable jurisdiction.

2.4 UAE digital invoice context

The UAE Ministry of Finance's current eInvoicing portal is the official source for the national programme. It distinguishes a structured eInvoice from PDF, Word, image, scanned and email invoices [10,11]. Version 1.1 of the guidelines describes the accredited-service-provider exchange model, the voluntary phase from 1 July 2026 and phased mandatory implementation beginning in 2027, subject to the current legislative documents [10]. The programme can create more structured invoice events for businesses within scope.

Structured data improves consistency and machine processing. It does not itself establish delivery, acceptance, absence of a dispute, ownership of a receivable, priority of an assignment or eligibility under a financing agreement. Those facts require other systems and controls. A useful optimisation data contract links the eInvoice identity to the order, shipment, acceptance, credit note, receivable, programme rule and settlement.

The Central Bank of the UAE describes the UAE as a central logistics and trading hub and emphasises the complexity of interconnected international supply chains [12]. Its current guidance distinguishes documentary trade finance, non-documentary finance and open-account trade; it also identifies trade-based financial-crime and proliferation-finance risks such as false invoices, over- or under-invoicing and misdescription of goods [12-14]. A lower objective value cannot override customer due diligence, sanctions, trade-control or suspicious-activity controls.

2.5 Quantum, quantum-inspired and hybrid labels

The term quantum optimisation can conceal materially different methods. A physical annealer uses quantum hardware designed for an Ising or related optimisation formulation. A gate-model quantum processor executes circuits, such as QAOA circuits, and measures bit strings. A classical simulation models a quantum circuit on classical hardware. A quantum-inspired solver uses ideas associated with quantum methods while running classically. A hybrid service can combine proprietary classical algorithms and quantum resources.

Label	Minimum disclosure	Excluded claim without further evidence
Physical QPU	Hardware, topology, calibration context, QPU time, access latency and samples	End-to-end speed or cost advantage
Quantum annealing	Ising/QUBO mapping, embedding, chain settings and post-processing	Exact optimisation or guaranteed global optimum
Gate-model QAOA	Encoding, circuit depth, shots, optimiser, backend and error handling	Scalable performance from a small simulation
Classical simulation	Simulator, hardware, precision and memory	Physical quantum execution
Quantum-inspired	Classical algorithm and compute resources	Quantum computation
Hybrid solver	Classical and quantum components, orchestration and attribution tests	Value caused by the quantum component

Preskill introduced the noisy intermediate-scale quantum framing for devices whose noise, depth and scale constrain computation [22]. The reviewed logistics literature remains consistent with a hybrid operating model. Every paper, vendor proposal and internal benchmark should use precise labels.

2.6 Evidence ladder and decision authority

The evidence ladder separates feasibility from production value. Theory can show that a method is defined. Simulation can test selected instances. Physical hardware can show executability. A representative benchmark can compare the complete workflow. A shadow pilot can test stability and user acceptance. Restricted production can test a controlled decision under real operating conditions. Each level answers a different question.

Level	Evidence	Question answered	Decision permitted
Q0 concept	Mathematical formulation	Can the problem be represented?	Research backlog only
Q1 simulation	Reproducible small-instance result	Does the implementation behave as expected?	Laboratory continuation

Level	Evidence	Question answered	Decision permitted
Q2 physical execution	Repeated QPU runs	Can the encoded instance execute on named hardware?	Hardware benchmark
Q3 representative benchmark	Holdout instances and strong classical baselines	Does the full method add measurable value?	Shadow-pilot decision
Q4 shadow workflow	Live inputs, no operational authority	Is performance stable and usable?	Restricted-use proposal
Q5 restricted production	Named users, limits and rollback	Does it improve a controlled decision?	Scale, retain or retire

Quantum optimisation evidence ladder




Figure 3. Evidence ladder for quantum optimisation

Source: Matchpoint synthesis of methods and studies [20-36].

3. B4 AND A3 DECISION MAP

3.1 B4 owner and operator decisions

For a GCC SME or family business, the useful optimisation question starts with a named operating pain. Examples include late deliveries, expensive rental vehicles, warehouse congestion, excess safety stock, slow invoice approval, concentration of receivables, insufficient headroom or supplier stress. The question must identify the decision frequency, planning horizon, responsible owner and action deadline.

An owner may want a daily dispatch plan that protects priority customers while limiting vehicle cost. A finance director may want to choose which invoices to discount within facility, debtor and concentration limits. A managing director may want a weekly plan that coordinates production, transport and liquidity. Each question needs its own model. One universal supply-chain optimiser would hide incompatible objectives and authorities.

B4 decision	Horizon	Hard constraints	Soft objectives	Release owner
Daily dispatch	Hours to one day	Vehicle, driver, capacity, access and legal time	Distance, lateness, rentals and workload balance	Logistics head
Warehouse movement	Minutes to shift	Collision avoidance, aisle rules and equipment	Travel time, queue and throughput	Warehouse manager
Production and delivery	Days to weeks	Material, machine, labour and customer commitments	Margin, changeover, inventory and service	Operations director
Receivables selection	Daily to weekly	Eligibility, limit, concentration and approval	Liquidity, funding cost and reserve	Finance director

B4 decision	Horizon	Hard constraints	Soft objectives	Release owner
Supplier payment	Daily to monthly	Approved obligation, cash floor and programme rule	Discount, supplier resilience and liquidity	CFO or owner
Integrated plan	Weekly to quarterly	All applicable hard constraints and authorities	Cash conversion, margin, service and resilience	Managing director

The planning model should expose trade-offs. A high-priority delivery can justify a rental vehicle. A lower-cost route can be rejected if its late-delivery risk delays acceptance. An early supplier payment can be valuable when it protects a critical source of materials. These judgements belong in approved objective weights, lexicographic priorities or scenario comparison.

3.2 A3 credit and special-situations decisions

A private-credit or special-situations fund sees the same operating system through a financing and downside lens. It needs verified receivables, payables, inventory, route, customer and cash data to assess borrowing-base quality, liquidity, working-capital volatility, concentration, dilution, disputes and covenant headroom. It may finance the business, purchase receivables, fund inventory or structure a rescue and turnaround facility.

An optimisation engine can prepare candidate allocations and downside scenarios. Credit authority remains with the lender's approved process. The system must preserve which constraints derive from signed documents, which are policy rules, which are management preferences and which are model assumptions.

A3 decision	Candidate optimisation	Evidence required	Human authority
Facility sizing	Draw and reserve under downside scenarios	Eligible assets, cash flows, limits and covenant definitions	Investment committee or delegated credit authority
Borrowing base	Invoice or inventory inclusion and haircut	Asset identity, ownership, ageing, dilution, concentration and control	Facility agent and lender credit team
Receivable purchase	Asset selection and price	Valid receivable, assignment, debtor, due date and dispute status	Legal and credit approval
Liquidity bridge	Funding timing against operating milestones	Uses, sources, milestone evidence and contingency	Deal team and committee
Covenant sensitivity	Headroom across demand, cost and collection shocks	Governing definitions and reconciled forecast	Finance, legal and credit owners
Workout plan	Payment, inventory and operating sequence	Current cash, claims, priorities and operational feasibility	Restructuring and legal authorities

3.3 Shared decision packet

The B4 operator and A3 lender can share a decision packet without sharing every internal objective or unrestricted record. The packet defines the approved data snapshot, rights, calculation, plan, sensitivities and exceptions. Access follows purpose and entitlement.

The packet should distinguish observed facts from forecasts. A route completion is observed after the relevant event is captured. A predicted arrival remains a forecast. An issued invoice is observed as a document event. Its expected collection date remains a forecast. Eligibility can be a system determination under encoded rules; the lender's acceptance remains an authority decision.

Packet field	B4 use	A3 use	Control
Decision question	Operational action	Credit or monitoring action	Named scope and deadline
Data snapshot	Orders, stock, routes and cash	Eligible assets, exposure and covenant data	Hash, source and timestamp
Constraint register	Operating commitments	Documents, policies and credit rules	Source and authority class
Candidate plan	Dispatch, payment or integrated schedule	Funding, reserve or monitoring response	Version and solver lineage
Feasibility result	Executable plan	Compliance with modelled limits	Independent checker

Packet field	B4 use	A3 use	Control
Sensitivity	Operational resilience	Downside and headroom	Approved scenarios
Exceptions	Manual intervention	Waiver or credit escalation	Owner and expiry
Disposition	Release, revise or reject	Approve, condition or reject	Named sign-off

3.4 Stop conditions

The optimiser should abstain when a hard constraint is incomplete, contradictory or outside the permitted authority. It should also stop when the input snapshot is stale beyond the decision's tolerance, a required system is unavailable, a candidate plan cannot be independently checked or the exception queue exceeds reviewer capacity.

Stop condition	Required response	Owner
Order, inventory or vehicle state is materially incomplete	Hold plan and restore approved manual dispatch	Operations owner
Invoice identity, debtor, amount, due date or dispute status is ambiguous	Exclude asset and route for finance review	Finance owner
Facility or covenant rule cannot be encoded from an approved source	Hold financial recommendation	Credit and legal owner
Candidate violates capacity, time, legal, safety or liquidity floor	Reject candidate regardless of objective value	Independent validator
Quantum or hybrid component fails repeatedly or exceeds decision window	Use approved classical baseline	Optimisation owner
Sanctions, AML, trade-control or fraud alert is unresolved	Block transaction and follow regulated process	Compliance owner
Model output is unstable across equivalent runs	Hold operational use and investigate	Model owner
User cannot explain or operationalise the candidate plan	Reject or revise in shadow workflow	Business owner

4. CLASSICAL LOGISTICS BASELINE

4.1 Why the baseline comes first

Vehicle routing has a long classical research history. Dantzig and Ramser formulated the truck-dispatching problem in 1959 [15]. Solomon developed influential vehicle-routing and scheduling instances with time windows in 1987 [16]. Uchoa and co-authors introduced a diverse set of capacitated vehicle-routing benchmark instances from 100 to 1,000 customers [17]. Modern open-source implementations include PyVRP, which implements hybrid genetic search and reports comparative performance in a peer-reviewed software paper [18], and Google OR-Tools, which supports capacity, time-window, resource and optional-visit constraints [19].

A quantum benchmark that omits strong classical methods produces weak evidence. The baseline should include an exact or bound-producing method on tractable instances, a strong heuristic or metaheuristic on realistic instances, and the current production or manual process. The organisation should record software versions, hardware, time limits, random seeds, warm starts, data transformations and stopping conditions.

4.2 Vehicle-routing decision model

The capacitated vehicle-routing problem assigns customer visits to vehicle routes while respecting capacity and usually requiring routes to start and end at a depot. Practical variants add multiple depots, heterogeneous fleets, time windows, pickups and deliveries, driver rules, vehicle compatibility, access restrictions, breaks, service times, traffic, optional visits and costs that vary by vehicle or time.

The business objective can be lexicographic. Safety and legal feasibility come first. Committed delivery windows and priority-customer service may come next. Cost, distance, emissions, rental use and workload balance can then be optimised within those limits. A weighted sum can be useful when weights are approved and sensitivities are shown. It can conceal unacceptable trade-offs when a large saving offsets a hard operational breach.

Model element	Example	Validation evidence
Node	Depot, customer, supplier, port or warehouse	Location master and active status
Arc	Permitted movement from one node to another	Road network, access and travel-time source
Demand	Weight, cube, pallets or service units	Order line and unit conversion
Vehicle	Capacity, type, cost, shift and compatibility	Fleet and rental master
Time window	Earliest and latest service	Customer commitment and site access
Service time	Loading, unloading or inspection	Observed distribution or approved standard
Priority	Contract, margin, criticality or recovery	Named business rule
Objective	Cost, lateness, fleet use and balance	Approved objective hierarchy

4.3 Inventory, production and warehouse decisions

Routing interacts with inventory and production. A route cannot load unavailable stock. A production plan may create stock after the dispatch cut-off. A warehouse can have the stock on hand and still lack labour, dock or equipment capacity. Returns and quarantine stock cannot be treated as freely available.

The model should identify which decisions are coupled tightly enough to optimise together. A daily route plan may accept fixed available inventory from the warehouse system. A weekly integrated plan may choose production, transfers and routes jointly. Larger models can be decomposed by time, region, warehouse, product family or decision layer, provided that shared constraints are reconciled.

Automated-guided-vehicle scheduling adds collision, path occupancy, priority and real-time response. Haba, Ohzeki and Tanaka studied reverse annealing for multi-AGV routing in a virtual plant and used a greedy solution as a starting point [26]. Nguyen Quang and co-authors later described candidate-route generation, clustering, physical quantum annealing and integration with a commercial AGV operating system; their simulations covered 1,000 AGVs, while substantial preprocessing reduced and decomposed the problem [33]. These results show the importance of orchestration and decomposition.

4.4 Data quality and event time

Optimisation uses a snapshot of a changing system. Orders arrive, vehicles break down, traffic changes, customers reschedule and invoices are disputed. Every input should carry event time, system-record time, source, status and quality. The decision contract defines a maximum acceptable age for each field.

An event ledger should preserve corrections. A quantity changed from 10 to 12 is different from a second order for two. A cancelled invoice is different from a paid invoice. A route event received late should not rewrite the earlier decision record. The released packet retains the exact snapshot used by the solver.

Data-quality test	Failure example	Treatment
Identity	Customer aliases create duplicate stops	Resolve master identity before optimisation
Completeness	Missing weight permits an overloaded route	Block or apply approved conservative default with flag
Timeliness	Old inventory creates infeasible loading plan	Refresh or hold decision
Unit consistency	Cases interpreted as individual items	Convert through controlled unit master
Status	Cancelled order remains active	Enforce state-transition rules
Referential integrity	Invoice lacks valid order or debtor	Exclude from finance optimisation
Plausibility	Negative travel time or impossible due date	Quarantine and investigate

4.5 Baseline hierarchy

The exact baseline can use mixed-integer programming, constraint programming, branch-and-cut or another method that returns an optimum or bound for the chosen formulation. Exact methods may be limited to smaller or simplified instances. Their value is a known reference for solution quality.

The heuristic baseline should be competitive and fit the problem. PyVRP, OR-Tools and specialist commercial solvers provide possible reference implementations [18,19]. A route optimiser can also use large-neighbourhood search, hybrid genetic search, local search or decomposition. The benchmark records the method rather than assuming one tool is universally strongest.

The production baseline captures what the business currently does. It can include dispatcher judgement, spreadsheet planning, routing software and manual exception handling. Its measures include planning time, feasibility, cost, service, revisions and acceptance. A technically stronger solver adds little value if the released plan arrives after trucks must leave or requires unmanageable edits.

Comparator	Purpose	Required metric	Common weakness exposed
Exact or bounded	Establish optimum or gap on tractable cases	Objective, bound, time and memory	Approximation quality
Strong heuristic	Test realistic scale	Best, median, tail and feasible rate	Instance sensitivity
Quantum-inspired	Separate algorithmic idea from QPU execution	Full classical resources and quality	Mislabelled quantum contribution
Current production	Test business relevance	End-to-end time, cost, service and edits	Integration and usability gap
Simple rule	Provide transparent floor	Feasibility and objective	Complexity without value

5. SUPPLY-CHAIN-FINANCE DECISION MODEL

5.1 Event-to-finance chain

Supply-chain financing depends on specific events. A purchase order can support a pre-shipment facility. Shipment evidence can support a milestone. Delivery and acceptance can support invoice issuance. Buyer approval can make an invoice available in a payables-finance programme. A valid receivable can be evaluated for factoring or receivables discounting. Collection repays or reduces an exposure.

The optimiser should never create a legal or accounting event by inference. It can propose that an eligible invoice be financed. It cannot make the invoice eligible. It can flag missing acceptance evidence. It cannot deem performance complete. Each transition requires a system event, document and owner appropriate to the contract and applicable law.




Figure 2. Physical and financial decision map

Source: Matchpoint framework; supply-chain-finance context [1-14].

Event	Candidate decision	Required control	Failure if assumed
Purchase order confirmed	Allocate pre-shipment funding	Counterparty, order and facility validation	Funding a cancelled or unauthorised order
Goods produced	Schedule loading and working-capital draw	Quantity, quality and title evidence	Financing unavailable or rejected goods
Shipment dispatched	Update route and cash forecast	Carrier, goods and milestone evidence	False expected collection timing
Delivery accepted	Issue or validate invoice	Contractual acceptance rule	Premature or disputed invoice
Invoice approved	Offer payables finance	Programme, amount and due-date validation	Ineligible or duplicate funding
Receivable transferred	Update ownership and exposure	Assignment, notice, priority and ledger	Double assignment or incorrect collection
Cash collected	Apply settlement and release limit	Bank match and allocation	Overstated headroom

5.2 Receivable eligibility and selection

A receivables pool can contain thousands of invoices with different debtors, tenors, currencies, dispute states, dilutions and programme rules. The selection problem can be expressed with binary variables indicating whether each invoice is financed. Constraints can cover facility size, invoice and debtor eligibility, concentration, ageing, country, currency, tenor and reserve. The objective can balance liquidity raised, financing cost, expected dilution, concentration and remaining availability.

This model is only as reliable as its data and rules. An invoice can be valid in the enterprise resource planning system and excluded by the financing agreement. A debtor can be current while its concentration limit is exhausted. A credit note can reduce the financed amount. An invoice can be assigned already. The rule engine should show which source and version supports each eligibility test.

Eligibility field	Example rule	Evidence source	Escalation
Debtor	Approved and within jurisdiction	Facility schedule and KYC master	Credit/compliance
Invoice status	Issued, undisputed and unpaid	ERP, dispute and cash systems	Finance operations
Tenor	Due within permitted range	Invoice and agreement	Facility agent
Age	Below maximum days since issue	Invoice ledger	Credit owner
Concentration	Below debtor or group cap	Exposure aggregation	Credit owner
Assignment	Available and not previously transferred	Receivables registry and internal ledger	Legal/operations
Dilution	Within threshold or reserve	Credit notes, returns and history	Credit owner
Currency	Permitted and hedged if required	Agreement and treasury record	Treasury

5.3 Payables finance and supplier resilience

In payables finance, the buyer approves invoices and the seller may choose early payment through the finance provider [2]. An optimisation model can allocate a limited programme across suppliers or invoices, subject to eligibility, facility availability and operational constraints. It can also support dynamic discounting when the buyer uses its own cash.

The objective should account for supplier resilience. Extending payment terms can improve the buyer's reported cash position while increasing the supplier's funding need. A critical small supplier can require different treatment from a diversified large supplier. The model can include criticality, substitution time, financial health and service dependency, provided that these inputs are governed and reviewed.

The IFRS supplier-finance disclosure requirements reinforce the need to preserve arrangement terms, liability amounts, payment timing and liquidity-risk information [5,6]. The accounting owner determines presentation and disclosure. Optimisation results should feed the evidence package rather than create an accounting conclusion.

5.4 Integrated liquidity forecast

The cash forecast links route and inventory events to financing decisions. A late inbound shipment can delay production, dispatch, acceptance, invoicing and collection. A faster route can bring cash forward only when contractual and operational triggers follow. The model should therefore carry event dependencies and ranges rather than one deterministic cash date.

A3 lenders can use the same structure for downside analysis. They can shock volumes, delivery times, returns, debtor payment, funding rates and facility availability. The resulting cash and covenant paths remain scenarios. Management forecasts and lender cases should be labelled separately.

Scenario input	Base treatment	Downside treatment	Authority
Demand	Approved forecast	Named volume and mix shock	Management and lender
Delivery	Observed distribution	Congestion or disruption case	Operations
Acceptance	Contract and history	Higher rejection or delay	Commercial/credit
Collection	Debtor history and terms	Late-payment curve	Finance/credit
Financing	Current eligible facilities	Lower advance, cap or availability	Treasury/lender
Cost	Contracted and observed inputs	Fuel, rental, labour and rate shock	Finance

5.5 Financial-crime and trade controls

The CBUAE guidance highlights false invoicing, over- or under-invoicing, goods misdescription and other trade-based financial-crime risks [12-14]. An optimiser that selects assets or routes at scale can propagate a weak control rapidly. Financial-crime tests should operate before funding and release.

The model can support anomaly queues by comparing quantity, price, goods, route, counterparty, invoice and shipping data. A statistical or rules-based alert remains an alert. Regulated owners decide investigation and reporting. High-risk transactions require approved customer due diligence, sanctions, trade-control and escalation processes.

6. QUANTUM METHODS AND CURRENT EVIDENCE

6.1 Ising and QUBO formulation

Many binary optimisation problems can be expressed as an objective over binary variables with linear and quadratic terms. Lucas catalogued Ising formulations for a range of NP problems and showed how constraints can be represented through penalty terms [21]. Quantum annealers commonly accept Ising or QUBO forms. Gate-model methods can also use cost Hamiltonians derived from these formulations.

Constraint conversion is consequential. A penalty that is too low can permit infeasible solutions. A penalty that is too high can compress the useful objective landscape or exceed hardware coefficient ranges. Additional variables may be needed for capacity, ordering or subtour constraints. Hardware connectivity can require minor embedding, chains and still more physical qubits.

The model record should retain the original constrained formulation, QUBO or Ising transformation, penalty logic, scaling, embedding and independent feasibility checker. The business objective should be recalculated from decoded candidates rather than inferred only from an energy value.

6.2 QAOA and gate-model methods

Farhi, Goldstone and Gutmann proposed QAOA as a hybrid method for approximate combinatorial optimisation [20]. A parameterised quantum circuit alternates cost and mixing operators. A classical optimiser updates the parameters using measured objective information. The circuit returns a distribution of bit strings.

For routing, the encoding can grow quickly. Fitzek and co-authors mapped a heterogeneous vehicle-routing problem to an Ising Hamiltonian and simulated instances using up to 21 qubits, corresponding to three customers and two trucks in their reported setup; the number of qubits in their mapping scaled quadratically with customers [31]. Leonidas and co-authors proposed qubit-efficient algorithms for vehicle routing [29]. Onah and Michielsen

proposed a coloured-permutation encoding in a 2026 preprint and reported recovery of independently verified optima on a benchmark suite; the work remains preprint evidence and is not production validation [34].

Warm-start QAOA uses a classical relaxation to inform the initial state and mixer [23]. Warm starts can improve shallow-circuit behaviour on selected problems. They also make classical contribution central to attribution. A fair benchmark tests the warm-start or repair routine with alternative candidate generators.

6.3 Quantum annealing and reverse annealing

Quantum annealing represents a problem as an energy landscape and samples low-energy states. Reverse annealing begins from a supplied candidate and searches around it. Haba and co-authors started from a greedy multi-AGV solution and reported improvement through reverse annealing in their virtual-plant experiment [26]. Their comparison and configuration are specific to the study.

Feld and co-authors developed a hybrid approach for the capacitated vehicle-routing problem and highlighted the difficulty of obtaining feasible solutions from the quantum formulation [25]. Sinno and co-authors later tested commercial annealing platforms using CVRP benchmark data over more than 30 hours of access. Their reported QPU time was measured in microseconds, while solution error varied and quality degraded as constraint density increased [32]. QPU time is one component of end-to-end time; queueing, embedding, data movement, decoding and repair remain part of the operational measure.

6.4 Hybrid logistics evidence

Weinberg and co-authors studied a commercial supply-chain logistics problem that was too large for direct near-term embedding. They decomposed it into single-truck binary subproblems with roughly 2,500 variables and used simulated annealing and a D-Wave hybrid solver as proxies within a larger classical simulation [27]. The paper demonstrates a credible hybrid workflow and explicitly identifies the decomposition.

Osaba, Villar-Rodriguez and Asla developed Q4RPD for a real package-delivery problem supplied by a Spanish logistics company. The approach used classical computation to split the problem, manage restrictions and compose the final solution, while D-Wave's constrained hybrid solver calculated route components. The reported experiment covered six illustrative instances [30]. This is valuable real-world-oriented modelling evidence. It is not a general quantum-advantage result.

Nguyen Quang and co-authors reported a 2025 AGV route-optimisation system with candidate-route generation, problem clustering, physical annealing and a commercial AGV operating-system connection. Their simulations involved 1,000 AGVs, and preprocessing reduced variables materially before the optimisation stage [33]. The result supports the engineering value of decomposition and integration. Component-level comparison remains necessary to identify the quantum contribution.

Study	Execution label	Reported scope	Decision-relevant lesson	Limitation for T16
Feld et al. [25]	Hybrid annealing	CVRP method and small tests	Feasibility and mapping require hybrid handling	Current GCC workflow not tested
Haba et al. [26]	Physical reverse annealing plus classical start	Multi-AGV virtual plant	A good classical start can focus search	Instance, solver and timing conditions specific
Weinberg et al. [27]	Simulated annealing and D-Wave hybrid solver	Commercial-scale supply-chain model decomposed into subproblems	Decomposition makes selected subproblems tractable	Quantum hardware contribution not isolated
Osaba et al. [30]	D-Wave constrained hybrid service	Six real-world-oriented delivery instances	Real constraints need classical orchestration	Illustrative sample and proprietary hybrid composition
Fitzek et al. [31]	Gate-model simulation	Up to 21 qubits; three customers and two trucks	Encoding and optimiser trade-offs appear at small scale	Far below operating scale
Sinno et al. [32]	Commercial quantum annealing	CVRP benchmark instances and repeated runs	Constraint density affects quality	End-to-end operational comparison remains required
Nguyen Quang et al. [33]	Physical annealing, quantum-inspired and classical integration	Commercial AGV system; simulations of 1,000 AGVs	Candidate generation and clustering dominate scale management	Operational production advantage not established

Study	Execution label	Reported scope	Decision-relevant lesson	Limitation for T16
Onah and Michielsen [34]	2026 preprint	Coloured-permutation CVRP benchmark	Encoding can reduce selected auxiliary requirements	Preprint; independent production evidence unavailable

6.5 Current conclusion

The reviewed evidence establishes active progress in formulations, hybrid workflows, physical execution and commercial integration. It also shows that current systems rely on classical decomposition, strong initial solutions, hybrid services, simulation, repair and constrained test scopes. The evidence does not establish broad production quantum advantage across GCC logistics or supply-chain finance as of 1 August 2026.

BIS and ESMA reviews place optimisation within a wider financial-system opportunity and risk agenda while recording technical uncertainty and early commercial development [35,36]. Their system-level analysis supports controlled capability building and disciplined claims; it does not validate a T16 use case.

The appropriate programme is benchmark-first and component-attributed. A useful quantum experiment can still create value by clarifying the problem, improving data, establishing classical baselines, developing reproducible encodings and training a team. Any direct operating or financial return requires observed evidence from a controlled pilot.

7. INTEGRATED PHYSICAL AND FINANCIAL OPTIMISATION

7.1 Decision contract

The integrated model begins with a decision contract. It states the named decision, scope, frequency, planning horizon, deadline, owner, required systems, objective hierarchy, hard constraints, permitted approximations, scenarios and release authority. The contract also defines which layer can use quantum methods.

A practical first experiment can target one bounded subproblem, such as selecting feasible delivery clusters within a daily plan or choosing eligible receivables within a confirmed pool. Joining every physical and financial decision in the first model would increase data, constraint and authority risk. Later versions can connect modules through shared events and controlled prices or shadow values.

Contract field	Example for a shadow pilot	Required approval
Decision	Candidate daily route allocation for one depot	Logistics head
Frequency	Once each night plus named disruption rerun	Operations director
Horizon	Next operating day	Logistics head
Deadline	Candidate available before dispatch planning cut-off	Dispatcher
Hard constraints	Capacity, time, access, safety and legal rules	Operations and compliance
Objective	Lexicographic service, rentals, distance and balance	Managing director
Quantum layer	Candidate cluster or route subproblem	Model owner
Comparator	Exact on small cases, PyVRP/OR-Tools and production plan	Benchmark owner
Authority	Shadow output only	Governance committee

7.2 State and decision variables

The physical state can include orders, demand, inventory, locations, vehicles, shifts, routes and disruptions. The financial state can include invoices, payables, receivables, cash, facilities, limits, debtor groups, rates, reserves and forecast collections. Binary variables can indicate whether a vehicle traverses an arc, whether an order is served in a period, whether an invoice is financed or whether a supplier is paid early. Continuous variables can represent quantities, cash, utilisation and time.

Every model variable should map to a business concept. An opaque binary index complicates validation and audit. The decoded candidate should state which vehicle serves which order, which invoice is financed, what cash moves and which constraints bind.

7.3 Objective hierarchy

The objective can be organised in tiers. Tier 0 enforces legal, safety, sanctions, eligibility and authority gates outside the optimiser. Tier 1 enforces hard operational and financial feasibility. Tier 2 protects committed service and liquidity floors. Tier 3 optimises cost, working capital, resilience and balance. Tier 4 applies preferences whose breach is acceptable and visible.

Tier	Example	Treatment
0 external gate	Sanctions, trade control, legal eligibility and credit approval	Must pass before candidate release
1 hard feasibility	Vehicle capacity, one service assignment, facility limit and cash floor	Independent checker rejects violation
2 protected outcome	Priority delivery, critical supplier and covenant buffer	Lexicographic or high-penalty with named override
3 economic objective	Transport cost, funding cost, inventory and late-delivery cost	Optimised with sensitivity
4 preference	Driver balance, preferred carrier or payment pattern	Soft score and visible trade-off

Objective coefficients need controlled units. Distance, hours, currency and risk scores cannot be added coherently without conversion. Cost proxies should have sources and effective dates. A customer-priority score should have an approved definition. Penalties used only to enforce QUBO constraints should be separated from business costs.

7.4 Constraint register

The constraint register is the bridge between documents, systems and equations. Each row contains a plain-language rule, mathematical expression, scope, source, version, hard-or-soft class, owner, test case and change history. A signed facility agreement can define a concentration cap. A fleet record can define vehicle capacity. A customer contract can define a delivery window. A management preference should be labelled as such.

Changes require impact analysis. A new vehicle, customer, facility amendment or invoice rule can alter feasibility. The model version and released packet should identify which constraint set was active.

7.5 Decomposition and orchestration

Decomposition can separate strategic, tactical and operational horizons. A quarterly network or facility model can set capacity and financing envelopes. A weekly model can plan production, inventory and transfers. A daily model can route vehicles and select funding actions. Real-time rules can respond to disruptions within approved boundaries.

The quantum component can sit inside this orchestration as a candidate generator for a bounded binary subproblem. The classical controller validates inputs, builds the formulation, selects the solver, decodes output, repairs candidates where allowed, runs the independent feasibility checker and returns alternatives. The released plan comes from the governed workflow.

Layer	Decision	Typical method	Quantum experiment candidate
Strategic	Network, depot, facility and funding design	Scenario MIP and simulation	Facility-location subproblem
Tactical	Production, inventory, fleet and liquidity envelope	MIP, constraint programming and heuristics	Selected binary allocation subproblem
Operational	Daily routes, loading and invoice selection	Routing heuristic, MIP and rules	Route cluster or asset-selection candidate
Real time	Disruption reroute and exception	Fast heuristic and policy	Restricted local search if latency permits

7.6 Feasibility certificate and repair

The solver's output should be checked by an independent deterministic validator. The validator confirms every hard rule against the decoded candidate and the exact input snapshot. It also recalculates the business objective.

Repair can convert an infeasible sample into a feasible plan by dropping, reassigning or rescheduling decisions. Repair is valuable operationally. It can also account for most of the improvement. Benchmarks should report raw feasible rate, repaired feasible rate, repair time and objective change. The classical repair routine should be tested with classical candidate generators.

8. CONTROLLED ARCHITECTURE AND TOOL STACK

8.1 Architecture principles

The architecture separates source systems, the event and rights layer, the model registry, solver adapters, validation, decision packets and operational release. Systems of record retain authority for orders, inventory, routes, invoices, cash and facilities. The optimisation platform reads controlled snapshots and returns candidate actions. It does not silently rewrite source records.




Figure 4. Controlled architecture and tool stack

Source: Matchpoint architecture; governance and security context [37-40].

The architecture should remain solver-portable. A common problem contract can be translated to an exact solver, heuristic, simulator, quantum-inspired method, quantum annealer or gate-model backend. Solver-specific transformations are versioned. This design supports fair comparison and avoids embedding one vendor's representation in the business definition.

8.2 Source and event layer

The source layer can include enterprise resource planning, warehouse management, transport management, fleet telemetry, customer orders, invoice exchange, treasury, banking and facility records. Each snapshot carries source identity, extraction time, event time, schema version and quality result.

The UAE eInvoicing programme creates a useful structured-invoice interface for transactions within scope [10,11]. The optimisation event model can ingest the approved invoice identifier and status while keeping the Ministry or accredited provider record distinct from internal acceptance, dispute, financing and settlement events.

8.3 Rights, policy and model registry

The rights layer enforces user, business, matter, counterparty and purpose entitlements. NIST SP 800-207 frames zero trust around explicit, least-privilege decisions for resources [38]. A logistics planner may see route data and

masked finance indicators. A treasury user may see cash and invoice data. A lender may receive only the agreed monitoring package.

The model registry stores the decision contract, variable dictionary, constraints, objective, transformations, code, dependencies, tests, owner, approved use and retirement state. NIST's AI Risk Management Framework offers a voluntary govern, map, measure and manage structure that can inform model and decision governance [37]. The framework does not validate a quantum optimiser by itself.

8.4 Solver adapters and experiment record

Each solver adapter should return a common record: candidate, raw status, objective, bound if available, feasibility, run time, resource use, seed, configuration and logs. For a physical QPU it also records backend, hardware context, queue time, QPU time, samples, embedding and post-processing. For a hybrid service it records the disclosed classical and quantum components and any unavailable detail.

The experiment record should support rerunning a claim. It retains training or tuning instances separately from holdout instances. It records failed runs and infeasible samples. Selective reporting of the best run is unsuitable for a probabilistic method.

Record block	Minimum fields	Review question
Problem	Contract, snapshot, variables, constraints and objective	Is this the same business problem?
Transformation	Scaling, penalties, reductions and encoding	What changed for the solver?
Execution	Hardware, software, seed, shots/samples and limits	Can the run be reproduced?
Post-processing	Decode, repair, local search and selection	Where did quality originate?
Result	Feasibility, objective, gap, time and resources	Is the result decision-grade?
Disposition	Reviewer, exception and release	Who accepted the action?

8.5 Release and rollback

An approved candidate becomes an operational plan through a release adapter. The adapter can create a draft route, funding proposal or exception queue. A named user accepts it before source-system action during restricted use. Automated execution requires a later authority decision and stronger controls.

Rollback uses the approved classical or manual process. The system should preserve the last known feasible plan, decision deadline and recovery owner. Solver or vendor unavailability should not prevent dispatch, payment or compliance activity.

9. BENCHMARK AND VALIDATION PROTOCOL

9.1 Falsifiable hypothesis

A benchmark begins with a claim that can fail. An example is: for a specified family of daily capacitated routes from one depot, a named hybrid quantum-classical method will produce independently feasible candidates whose median objective is within an approved gap of the strongest classical comparator, inside the dispatch decision window and at an acceptable total resource cost. The hypothesis defines the problem family, metrics, threshold, comparator and test period.

A supply-chain-finance hypothesis can target invoice selection: for a fixed eligible pool and signed constraint set, the candidate method will match or improve the classical objective while producing the same independently verified eligibility result. Legal and credit acceptance remain outside the technical hypothesis.

9.2 Instance design

The suite should contain public benchmarks, synthetic stress cases and representative private instances. Public benchmarks support comparability. Synthetic cases isolate capacity, time-window, concentration or disruption effects. Private cases test the business distribution. Private data should be minimised, protected and retained under policy.

Training or tuning instances are separated from holdouts. The suite includes easy, medium and hard cases, along with cases expected to fail. Repeated runs measure variability. The organisation freezes the benchmark before reviewing the holdout result.

9.3 Metric stack

Feasibility is measured first. Solution quality can then be expressed as objective gap from an optimum, best-known result or strongest comparator. Reliability includes feasible rate, dispersion and tail behaviour. Time includes preprocessing, queueing, solver, QPU, decoding, repair, validation and human review. Resource cost includes compute, cloud, licences, data engineering and control work.

Business metrics include accepted-plan rate, manual edits, planning time, on-time service, vehicle or facility use, working-capital measures and exception burden. These metrics require a causal attribution design before they become evidence of value.




Figure 5. Benchmark and validation gates

Source: Matchpoint benchmark framework.

Metric family	Core measures	Required disclosure
Feasibility	Raw, repaired and released feasible rate	Independent checker and rule version
Quality	Best, median, tail, gap and bound	Comparator and objective units
Reliability	Run variance, failure rate and stability	Repetitions, seeds and missing results
Time	Full wall clock and component time	Queue, QPU, classical and review time
Resources	Compute, energy if measured, licence and labour	Measurement boundary
Operations	Acceptance, edits, overrides and rollback	Named users and decision window
Economics	Observed cost, revenue or loss attribution	Approved causal method
Controls	Exceptions, access, incidents and audit completion	Owner and severity

9.4 Component attribution

Hybrid methods require ablation tests. The benchmark can compare the complete hybrid workflow, the same workflow without the QPU, the same classical preprocessor with a classical sampler, the same repair routine with random or heuristic candidates, and the strongest independent classical solver. This reveals the source of quality and cost.

QPU time should never be presented as end-to-end latency. A microsecond sampling interval can coexist with long queueing, embedding, orchestration or validation. Total time should be measured from accepted input snapshot to validated candidate available to the user.

9.5 Gates

The programme advances through explicit gates. Gate 1 confirms a valuable and bounded decision. Gate 2 confirms data and constraint quality. Gate 3 confirms a strong classical baseline. Gate 4 confirms reproducible quantum execution. Gate 5 confirms representative holdout performance. Gate 6 confirms shadow-user acceptance and control capacity. Gate 7 confirms restricted production value and rollback.

Failure can lead to repair, redesign, retention as research or retirement. A stopped experiment still creates a reusable formulation and benchmark.

10. ILLUSTRATIVE PRODUCTIVITY AND ECONOMICS

10.1 Measurement boundary

Productivity should be measured per accepted decision packet. Solver runs or plans generated can increase while accepted operational value declines. The baseline and assisted workflow should process the same representative decision family with the same service, safety, credit and control requirements.

Commercial attribution remains USD 0 in this paper. No approved observed Matchpoint client result was supplied. The worked examples below are unverified illustrative management assumptions designed to show the arithmetic and the measurement discipline.

10.2 Illustrative planning-time bridge

Assume, solely for illustration, that a business prepares 25 daily route and financing packets in a month. The baseline requires 3.2 planner hours per packet, or 80 hours. An assisted workflow requires 1.5 hours of preparation and review per packet, or 37.5 hours, plus 14 hours of monthly data, operations and exception work, for 51.5 hours. The illustrative gross time reduction is 28.5 hours.




Figure 6. Illustrative productivity and cost bridge

Source: Unverified illustrative management assumptions; Matchpoint calculation.

Input	Illustrative assumption	Status
Accepted packets per month	25	Unverified management assumption
Baseline hours per packet	3.2	Unverified management assumption

Input	Illustrative assumption	Status
Assisted preparation and review hours per packet	1.5	Unverified management assumption
Assisted monthly data and exception overhead	14.0	Unverified management assumption
Baseline monthly hours	80.0	Calculated from assumptions
Assisted monthly hours	51.5	Calculated from assumptions
Illustrative gross hours released	28.5	Calculated; attributed value remains USD 0

The time released is not automatically a saving. It becomes an observed cost reduction only if the organisation measures the baseline, confirms comparable quality, identifies a real change in paid labour or capacity and receives approval for the attribution. It can instead become capacity for customer service, collections or planning.

10.3 Illustrative decision economics

An experiment can compare the full cost of the classical and quantum-assisted workflows. Costs include data engineering, model development, solver licences, cloud or QPU access, integration, security, validation, user training, review and maintenance. Benefits can include lower route cost, fewer rentals, reduced planning effort, earlier accepted delivery, lower funding cost, reduced idle cash or lower losses. Each benefit requires an observed causal measure.

Economic line	Measurement method	Attribution gate
Route cost	Released-plan cost compared with controlled baseline	Same demand and service; observed execution
Rental avoidance	Actual rental vehicles avoided	Verified capacity and no service degradation
Working-capital timing	Observed cash-date change from accepted event	Contractual trigger and collection verified
Funding cost	Actual financed amount, rate and days	Facility ledger and counterfactual approved
Planning capacity	Hours measured through time study	Workload and quality comparable
Loss reduction	Approved prevented-loss methodology	Risk owner and finance sign-off
Revenue	Collected incremental revenue	Customer, channel and counterfactual evidence

10.4 Break-even formula

Let annual verified benefit equal approved route-cost reduction plus approved funding-cost reduction plus approved labour-cost reduction plus approved loss reduction plus approved collected incremental revenue. Let annual full cost equal recurring platform, solver, data, security, review and support costs plus amortised implementation cost. Break-even occurs when verified benefit is at least full cost.

This formula is simple. Its evidence requirements are demanding. Management should maintain a claims register that records the source, counterfactual, owner, approval and period for every benefit. Unverified scenario values remain separate from management accounts and investment-committee materials.

11. GOVERNANCE, SECURITY AND PROFESSIONAL BOUNDARIES

The operating owner approves the physical plan. Treasury approves cash actions. The lender's credit authority approves financing. Legal counsel determines assignment, perfection, enforceability and contractual effect. Accounting and tax owners determine classification, recognition, disclosure and tax treatment. Compliance owners retain sanctions, AML and trade-control authority.

Data access should follow source rights and purpose. NIST SP 800-207 supports explicit, least-privilege access to resources [38]. ISO/IEC 27001:2022 provides information-security-management requirements [40]. ISO 28000:2022 addresses security management systems relevant to supply chains [39]. Use of these sources does not establish certification.

Quantum and classical vendors should be assessed for data location, retention, training use, encryption, access, subprocessors, incident response, availability, export controls and exit. Sensitive business data should be minimised.

Synthetic or masked instances can support early testing; representative private data enters only under approved controls.

Boundary	Required control	Release evidence
Legal rights	Source-linked agreement and qualified review	Approved legal position or explicit hold
Credit	Delegated authority, limits and conditions	Named approval and expiry
Accounting and tax	Controlled ledger mapping and professional judgement	Reconciled reporting package
Financial crime	KYC, sanctions, trade and escalation workflow	Cleared status or blocked transaction
Cybersecurity	Identity, least privilege, encryption and monitoring	Access and security test
Model	Version, validation, limits and change control	Model approval and use restriction

12. GATED ADOPTION ROADMAP

The roadmap begins with one decision whose value, data and deadline are clear. The team then builds the event contract, constraint register and classical benchmark before translating a bounded subproblem. Laboratory results remain separate from operational authority. Shadow use follows representative holdout success. Restricted production follows stable user acceptance, control capacity and rollback tests.

Adoption progresses through evidence gates




Figure 7. Gated adoption roadmap

Source: Matchpoint adoption framework.

Phase	Deliverable	Exit gate
0 Charter	Decision contract, owner and value hypothesis	Named sponsor and falsifiable question
1 Data	Snapshot, quality report and rights map	Critical fields meet approved threshold
2 Baseline	Exact, heuristic and production comparators	Reproducible classical benchmark
3 Formulation	Constrained model, QUBO/encoding and validator	Independent feasibility tests pass
4 Laboratory	Repeated simulator and QPU results	Holdout threshold and attribution pass
5 Shadow	Live-input candidate packets	Stable acceptance and decision-window fit
6 Restricted	Named users, limits and rollback	Observed value and controls approved
7 Scale or retire	Controlled rollout or archived benchmark	Governance decision and retained evidence

13. LIMITATIONS, RESEARCH AGENDA AND CONCLUSION

This paper is a framework based on public evidence. It does not report a GCC operating pilot, lender book, production quantum deployment or approved financial return. Logistics studies use different formulations, hardware, simulators, hybrid services, datasets and timing boundaries. Results cannot be transferred to a new business without representative testing.

Further research should publish complete classical baselines, end-to-end timing, failed runs, raw and repaired feasibility, resource cost and component ablations. Finance-aware logistics benchmarks should link physical events to synthetic receivables and facilities while preserving legal and accounting boundaries. Representative GCC cases should be built only with authorised data and named owners.

The present operating recommendation is clear. B4 owners should invest first in event quality, constraint ownership and strong classical optimisation. A3 lenders should invest first in verified collateral, eligibility, concentration, liquidity and covenant data. A bounded quantum experiment becomes credible after those foundations exist. It advances only when repeated holdout evidence shows feasible, timely and attributable improvement over the strongest classical workflow.

DECLARATIONS

Author. Matchpoint Partners. Named-person author attribution remains pending CK approval.

Evidence disclosure. Quantum evidence includes theory, classical simulation, proprietary hybrid services and physical quantum-annealing studies. These categories are labelled separately. Broad production quantum advantage for GCC logistics or supply-chain finance was not established in the reviewed evidence.

Scenario disclosure. All worked volumes, hours, costs, rates, savings and thresholds are unverified illustrative management assumptions. Attributed revenue, cost reduction and loss reduction remain USD 0.

Advice boundary. This paper is not investment, legal, regulatory, accounting, tax, audit, sanctions, AML, cybersecurity, credit, valuation or technology advice.

REFERENCES

- [1] Global Supply Chain Finance Forum (2016). Standard Definitions for Techniques of Supply Chain Finance. <https://cdn.iccwbo.org/content/uploads/sites/3/2017/01/ICC-Standard-Definitions-for-Techniques-of-Supply-Chain-Finance-Global-SCF-Forum-2016.pdf>
- [2] Global Supply Chain Finance Forum (2020). Market Practices in Supply Chain Finance: Payables Finance. <https://supplychainfinanceforum.org/GSCFF-APF-Techniques.pdf>
- [3] Global Supply Chain Finance Forum (2019). Market Practices in Supply Chain Finance: Receivables Discounting. <https://supplychainfinanceforum.org/GSCFF-Receivables-Discounting-Common-Practices.pdf>
- [4] Asian Development Bank (2023). Trade Finance Gaps, Growth, and Jobs Survey. <https://doi.org/10.22617/BRF230334-2>
- [5] IFRS Foundation. IAS 7 Statement of Cash Flows. <https://www.ifrs.org/issued-standards/list-of-standards/ias-7-statement-of-cash-flows.html>
- [6] IFRS Foundation. IFRS 7 Financial Instruments: Disclosures. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-7-financial-instruments-disclosures/>
- [7] UNCITRAL (2017). Model Law on Electronic Transferable Records. https://uncitral.un.org/en/texts/e-commerce/modellaw/electronic_transferable_records
- [8] UNCITRAL and UNIDROIT (2024). Model Law on Warehouse Receipts. <https://uncitral.un.org/en/mlwr>
- [9] United Arab Emirates (2021). Federal Decree-Law No. 16 of 2021 on Factoring and Transfer of Receivables. <https://uaelegislation.gov.ae/en/legislations/1515>
- [10] UAE Ministry of Finance (2026). UAE Electronic Invoicing Guidelines, version 1.1. https://mof.gov.ae/wp-content/uploads/2026/06/UAE-Electronic-Invoicing-Guidelines_V-1.1-01June2026.pdf
- [11] UAE Ministry of Finance. eInvoicing official portal. <https://mof.gov.ae/en/about-us/initiatives/einvoicing/>
- [12] Central Bank of the UAE (2025). Background of International Trade System and Trade Finance. <https://rulebook.centralbank.ae/en/rulebook/21-background-international-trade-system-and-trade-finance>
- [13] Central Bank of the UAE (2025). Non-Documentary Trade Finance and Open Account Trade. <https://rulebook.centralbank.ae/en/rulebook/212-non-documentary-trade-finance-and-open-account-trade>
- [14] Central Bank of the UAE (2025). Trade Finance. <https://rulebook.centralbank.ae/en/rulebook/321-trade-finance>
- [15] Dantzig, G. B. and Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, 6(1), 80-91. <https://doi.org/10.1287/mnsc.6.1.80>
- [16] Solomon, M. M. (1987). Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints. *Operations Research*, 35(2), 254-265. <https://doi.org/10.1287/opre.35.2.254>
- [17] Uchoa, E. et al. (2017). New benchmark instances for the Capacitated Vehicle Routing Problem. *European Journal of Operational Research*, 257(3), 845-858. <https://doi.org/10.1016/j.ejor.2016.08.012>
- [18] Wouda, N. A., Lan, L. and Kool, W. (2024). PyVRP: A High-Performance VRP Solver Package. *INFORMS Journal on Computing*, 36(4), 943-955. <https://doi.org/10.1287/ijoc.2023.0055>
- [19] Google. OR-Tools Vehicle Routing documentation. <https://developers.google.com/optimization/routing>
- [20] Farhi, E., Goldstone, J. and Gutmann, S. (2014). A Quantum Approximate Optimization Algorithm. *arXiv:1411.4028*. <https://arxiv.org/abs/1411.4028>
- [21] Lucas, A. (2014). Ising formulations of many NP problems. *Frontiers in Physics*, 2, 5. <https://doi.org/10.3389/fphy.2014.00005>
- [22] Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- [23] Egger, D. J., Marecek, J. and Woerner, S. (2021). Warm-starting quantum optimization. *Quantum*, 5, 479. <https://doi.org/10.22331/q-2021-06-17-479>
- [24] Weidenfeller, J. et al. (2022). Scaling of the quantum approximate optimization algorithm on superconducting qubit based hardware. *Quantum*, 6, 870. <https://doi.org/10.22331/q-2022-12-07-870>

- [25] Feld, S. et al. (2019). A Hybrid Solution Method for the Capacitated Vehicle Routing Problem Using a Quantum Annealer. *Frontiers in ICT*, 6, 13. <https://doi.org/10.3389/fict.2019.00013>
- [26] Haba, R., Ohzeki, M. and Tanaka, K. (2022). Travel time optimization on multi-AGV routing by reverse annealing. *Scientific Reports*, 12, 17753. <https://doi.org/10.1038/s41598-022-22704-0>
- [27] Weinberg, S. J. et al. (2023). Supply chain logistics with quantum and classical annealing algorithms. *Scientific Reports*, 13, 4770. <https://doi.org/10.1038/s41598-023-31765-8>
- [28] Bentley, C. D. et al. (2022). Quantum computing for transport optimization. *arXiv:2206.07313*. <https://arxiv.org/abs/2206.07313>
- [29] Leonidas, I. D. et al. (2023). Qubit efficient quantum algorithms for the vehicle routing problem on quantum computers of the NISQ era. *arXiv:2306.08507*. <https://arxiv.org/abs/2306.08507>
- [30] Osaba, E., Villar-Rodriguez, E. and Asla, A. (2024). Solving a real-world package delivery routing problem using quantum annealers. *Scientific Reports*, 14, 24791. <https://doi.org/10.1038/s41598-024-75572-1>
- [31] Fitzek, D. et al. (2024). Applying quantum approximate optimization to the heterogeneous vehicle routing problem. *Scientific Reports*, 14, 25415. <https://doi.org/10.1038/s41598-024-76967-w>
- [32] Sinno, S. et al. (2025). Analyzing performance of commercial quantum annealing solvers for the capacitated vehicle routing problem. *AIP Advances*, 15, 125003. <https://doi.org/10.1063/5.0277110>
- [33] Nguyen Quang, T. et al. (2025). Quantum annealing-based route optimization for commercial AGV operating systems in large-scale logistics warehouses. *Scientific Reports*, 15, 44047. <https://doi.org/10.1038/s41598-025-28481-w>
- [34] Onah, C. and Michielsen, K. (2026). Optimal, Qubit-Efficient Quantum Vehicle Routing via Colored-Permutations. *arXiv:2604.04570*. Preprint. <https://arxiv.org/abs/2604.04570>
- [35] Auer, R. et al. (2024). Quantum computing and the financial system: opportunities and risks. *BIS Papers No. 149*. <https://www.bis.org/publ/bppdf/bispap149.htm>
- [36] European Securities and Markets Authority (2026). Quantum computing in financial markets: applications, investments and prospects. https://www.esma.europa.eu/sites/default/files/2026-05/ESMA50-481369926-33801_TRV_Risk_Analysis_Quantum_Computing_in_Financial_Markets.pdf
- [37] NIST (2023). Artificial Intelligence Risk Management Framework 1.0. NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
- [38] NIST (2020). Zero Trust Architecture. NIST SP 800-207. <https://doi.org/10.6028/NIST.SP.800-207>
- [39] ISO (2022). ISO 28000:2022, Security and resilience; Security management systems; Requirements. <https://www.iso.org/standard/79612.html>
- [40] ISO/IEC (2022). ISO/IEC 27001:2022, Information security, cybersecurity and privacy protection; Information security management systems; Requirements. <https://www.iso.org/standard/27001>

APPENDIX A. ILLUSTRATIVE BINARY SUBPROBLEMS

For route selection, let a binary variable equal one when a candidate route is selected. The objective minimises approved route cost plus penalties for uncovered orders and soft service breaches. Hard constraints require each mandatory order to appear exactly once, vehicle and depot limits to hold and incompatible routes to remain mutually exclusive. A separate checker validates the decoded plan.

For receivable selection, let a binary variable equal one when an eligible invoice is proposed for finance. The objective can maximise liquidity net of approved funding, concentration and dilution costs. Constraints cover facility availability, debtor caps, tenor, currency and asset-level eligibility. Legal and credit owners confirm the rules and approve the transaction.

APPENDIX B. MINIMUM BENCHMARK RECORD

Block	Minimum content
Identity	Experiment, owner, date, code and approved use
Problem	Snapshot hash, instance family, horizon and decision deadline
Model	Variables, constraints, objective, units and version
Solver	Method, backend, configuration, limits, seeds and repetitions
Hybrid attribution	Preprocessing, quantum step, decoding, repair and local search
Results	Raw and repaired feasibility, objective, gap, time and resources
Operations	Acceptance, edits, exceptions, rollback and user disposition

APPENDIX C. MANAGEMENT QUESTIONS

1. Which recurring decision has a measurable deadline, owner and economic consequence?
2. Which data fields and constraints currently make the production plan unreliable?
3. What is the strongest classical comparator, and can the team reproduce it?
4. Which bounded binary subproblem is suitable for a transparent experiment?
5. What result would stop the programme?
6. Who can approve operational, financial, legal and credit actions?
7. How will full end-to-end time, cost and component contribution be measured?
8. What manual process remains available when the platform or vendor is unavailable?