CERTIFICATION BOARD

Certified Quantum Solutions Architect (CQSoA)

Ten Sample Questions

These ten questions are written at the difficulty of the live exam. Six carry a figure, a data table, a configuration or a code listing; four are plain scenario questions, as on the live paper. Every one is a scenario a working architect would meet, and none of them appears in the live question bank.

The live exam is multistage adaptive: how you score on one stage selects the difficulty of the next, so no two candidates sit the same paper. It is a performance-based exam: each form can carry up to 5 to 12 performance-based questions alongside the multiple-choice questions, which means you build, sequence or complete something rather than only choosing from a list. Answers and full explanations follow each question here; the live exam does not disclose item-level answers.

Question 1

Domain 2: Hardware Platforms, Error Correction and Roadmaps — DATA

Emma Wilson must decide whether a chemistry application can run on any of three vendor roadmaps by 2030. The estimate calls for 800 logical qubits at a logical error rate of 10^-9, and magic-state factories add roughly 40 percent to the qubit budget.

Vendor	2030 physical qubits	Two-qubit error	Code distance for 10^-9	Physical per logical (2d^2)
A	100,000	10^-3	27	1,458
B	20,000	2 x 10^-4	17	578
C	1,000,000	5 x 10^-3	Above threshold	n/a

Table 1.1 Roadmap figures as published, with the distance needed at each error rate for the target logical error.

Which conclusion is correct?

- (A) Vendor C meets the application, because a million physical qubits is far more than the 800 logical qubits required.
- (B) Vendor A meets the application, because 100,000 physical qubits exceed the 800 logical qubits by two orders of magnitude.
- (C) None of the three meets it by 2030: A needs about 1.6 million qubits, B about 650,000, and C is above threshold entirely.

(D) Vendor B meets the application, because its error rate is the lowest and the code distance needed is therefore the smallest.

Correct answer: C
Resource estimation is arithmetic: logical qubits times physical-per-logical times factory overhead. 800 x 1,458 x 1.4 is about 1.63 million; 800 x 578 x 1.4 is about 647,000. A device above the code threshold cannot be error corrected regardless of size, so vendor C's count is irrelevant. The honest answer is a horizon classification with observable gates, not a vendor pick.

Question 2 Domain 3: Hybrid Architecture and Integration CODE

James Carter reviews the execution loop of a variational pilot that is taking days to converge. The optimiser needs about 4,000 iterations.

```
for step in range(4000):
    circuit = build_ansatz(params)            # rebuilt every step
    tqc = transpile(circuit, backend)          # ~1.8 s
    job = backend.run(tqc, shots=4000)         # new job, queue ~25 s
    counts = job.result().get_counts()         # exec ~1.5 s
    energy = expectation(counts)
    params = optimizer.step(energy)
# no results store; counts discarded after each step
```

Listing 2.1 The pilot's loop.

Which redesign is correct?

(A) Transpile the ansatz once and bind per step, run inside a session, batch where possible, and persist counts with provenance.

(B) Increase shots to 40,000 per step so that each energy estimate is more accurate and far fewer optimiser steps are needed.

(C) Move the pilot to a vendor with a larger qubit count so that a wider ansatz reaches convergence in fewer iterations.

(D) Replace the gradient optimiser with a random search so that the loop no longer depends on the previous step's energy.

Correct answer: A
Queue latency and repeated transpilation dominate: about 28 seconds per step, so 4,000 steps take about 31 hours. Sessions and one-time transpilation remove almost all of it. Discarding counts also destroys reproducibility; a results store with backend, calibration and mitigation metadata is a platform requirement. More qubits and more shots make the loop slower, not faster.

Question 3 Domain 4: Use-Case Evaluation and Advantage Assessment DATA

A vendor claims a 40x speed-up on a portfolio optimisation problem. Olivia Davis audits the methodology.

Item	Vendor benchmark	Audit finding
Instance size	24 assets	Exact classical solver finishes in 0.4 s
Classical baseline	Simulated annealing, single thread, default settings	Tuned commercial solver: 0.02 s
Quantum time reported	QPU execution only, 0.8 s	Queue and post-processing add 41 s
Instances	1 (best of 20 runs)	Median of 20 runs is 3x worse than reported
Solution quality	Within 2% of optimum	Exact optimum available classically

Table 3.1 Benchmark claims versus audit.

Which conclusion should the audit report state?

- (A) The 40x speed-up is confirmed, because the QPU execution time of 0.8 seconds is the figure that matters for comparison.
- (B) The claim should be reduced to a more defensible 20x, splitting the difference between the vendor's and the audit's timing.
- (C) The vendor should be asked to repeat the best-of-20 procedure on 100 more runs to strengthen the statistical evidence.
- (D) No advantage is shown: a trivial instance, an untuned baseline, partial timing and a cherry-picked run; rerun properly or not at all.

Correct answer: D
Four classic flaws appear together: trivial instance, weak baseline, partial timing and cherry-picked runs. A correct benchmark is defined before it is run, compares against the best classical method with fair resources, counts every second the user waits, and reports distributions over many instances.

Question 4

Domain 6: Programme Leadership and Delivery — DATA

Liam Anderson inherits a quantum programme and must present a five-year budget structure to the board.

Item	Current state
Vendor contracts	5 cloud QPU contracts, 2 unused, one with an exclusivity clause
Pilots	7 pilots in 5 SDKs, no shared platform, results in notebooks
Post-quantum migration	Not started; no inventory; RSA everywhere
Use cases	No horizon classification; two "advantage" press releases drafted
Team	3 physicists, no platform engineer, no security involvement

Table 4.1 Programme as inherited.

Which structure and first-year priorities are sound?

- (A) Double the pilot budget and sign a sixth vendor to keep options open, then let the strongest pilots choose the platform.

(B) Fund PQC on a fixed timeline, a modest readiness line, pilots per gate, production contingent on milestones; fix the basics first.

(C) Cancel the programme entirely and revisit the question in 2035 when fault-tolerant hardware is expected to be widely available.

(D) Publish the two press releases first to attract executive funding, then use that funding to build the shared platform.

Correct answer: B
The programme has an unfunded certainty (PQC) and over-funded uncertainty (unstructured pilots). Budget should mirror certainty: obligation on a timeline, readiness as a stable line, experiments per gate, production as an option. Year one: inventory, consolidate vendors and remove the exclusivity clause, build the platform, withdraw the unsupported claims.

Question 5

Domain 1: Quantum Computing Foundations for Architects

DATA

A proposal claims that a 40-qubit, 600-two-qubit-gate circuit will "run today" on two candidate cloud devices. Fidelity is estimated as (1 - error) raised to the number of two-qubit gates after routing.

Device	Qubits	Two-qubit error	Connectivity	Estimated SWAP overhead for this circuit
X	1,121	9 x 10^-3	Heavy-hex lattice	+180% gates
Y	56	8 x 10^-4	All-to-all	+0%

Table 5.1 Device metrics.

Which assessment is correct?

(A) Device X runs it, because 1,121 qubits leaves ample room for the 40-qubit circuit and any routing the compiler needs.

(B) Neither runs it usefully: X ends near zero fidelity after routing, Y reaches about 62 percent, marginal and mitigation-dependent.

(C) Both run it perfectly, because 40 qubits is smaller than either device's qubit count and gate errors are below one percent.

(D) Device Y cannot run it, because 56 qubits is too close to the 40 required to leave headroom for ancilla and routing.

Correct answer: B
Qubit count beyond the circuit's width buys nothing. Post-routing gate count and per-gate error decide fidelity: 0.991^1680 is about 2 x 10^-7, while 0.9992^600 is about 0.62. An architect converts every "runs today" claim into this calculation before approving budget.

Question 6

Domain 2: Hardware Platforms, Error Correction and Roadmaps

DATA

A vendor publishes logical error rates for a surface-code experiment and asks to be named the programme's primary platform.

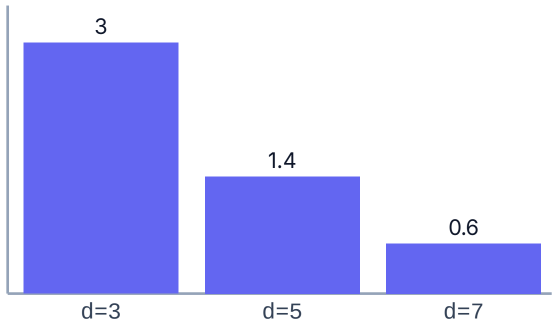

Figure 6.1 Logical error per cycle, in units of 10^-3, by code distance.

Question	Vendor answer
Independent reproduction?	Pre-print only, not yet peer reviewed
Physical two-qubit error	3 x 10^-3 average, 1 x 10^-2 worst edge
Logical gates between logical qubits demonstrated?	Not yet
Magic-state throughput	Not measured

Table 6.1 Due-diligence answers.

Which assessment is correct?

- **A** Name the vendor primary platform now, because the histogram is proof of fault tolerance and the rest is detail.
- **B** Dismiss the vendor entirely, because a physical error of 3 x 10^-3 is far too high for any practical application.
- **C** Treat it as a strong candidate: below-threshold scaling is real progress, but gate on peer review, logical gates and factories.
- **D** Ask the vendor for a larger histogram covering distances up to 25 before forming any view on the platform at all.

Correct answer: C
Error suppression with distance (roughly 2x per step) is the milestone that matters and deserves attention. It is still one of four milestones on the path to application scale, and an architect turns each into a gate with evidence requirements rather than a platform commitment.

Question 7 Domain 5: Security, Post-Quantum and Governance

A quantum programme wants to use a new cloud QPU provider that is ISO 27001 certified, documents tenant isolation, retains circuits for 90 days, grants itself a licence to reuse submitted circuits "to improve services", and processes only in a region outside the jurisdiction required for regulated data. Three workloads are planned: synthetic benchmarks, proprietary pricing models, and circuits built from regulated customer data. Which decision set is correct?

- **A** All three workloads may proceed, because ISO 27001 certification and documented tenant isolation cover the material risks.
- **B** Synthetic benchmarks may proceed; pricing models need the reuse licence removed first; regulated data is blocked on residency.

- (C) No workload may proceed until the provider agrees to delete every retained circuit and result immediately after each job.
- (D) Pricing models and regulated data may proceed if the team uses post-quantum transport; synthetic benchmarks need no review.

Correct answer: B
Assurance certificates do not override contract terms or residency. A licence to reuse circuits is a licence to reuse the model they encode. Residency is a legal constraint that transport encryption does not change. Synthetic workloads carry no such risk. Proportionate, per-workload decisions are what governance looks like in practice.

Question 8

Domain 3: Hybrid Architecture and Integration

An internal quantum platform used by four teams exposes a single function that accepts one vendor's circuit type, runs it on one named device, and returns raw counts with no metadata. Teams import that vendor's SDK to call it. Which redesign addresses the architectural problems?

- (A) Add a second vendor's SDK import next to the first, so that teams can choose whichever backend they prefer per job.
- (B) Accept an interchange format, route through backend adapters that isolate each SDK, and return results with full provenance.
- (C) Rename the function to make its purpose clearer and document which device it targets in the platform's README file.
- (D) Hard-code a second device name as a fallback, so that jobs still run when the primary device is offline for calibration.

Correct answer: B
The platform has bound every team to one vendor's types and one device, and it throws away the metadata that makes results comparable. Portability and reproducibility are properties of the interface layer: an interchange format, adapters and a results schema fix both. Adding a second import spreads the lock-in to two vendors.

Question 9

Domain 4: Use-Case Evaluation and Advantage Assessment

A programme office scores five candidate use cases: (A) a catalyst ground-state simulation needing fault-tolerant hardware, (B) fleet routing via QAOA on instances an exact solver handles, (C) fraud detection via quantum machine learning over 100 million rows, (D) Monte Carlo pricing via amplitude estimation, and (E) a quantum-inspired tensor solver already beating the incumbent scheduler. Which portfolio decision is correct?

- (A) Fund C first because fraud has the highest business urgency, and fund B for production because routing is a known quantum use.
- (B) Fund none of them until a fault-tolerant quantum computer exists, since every candidate depends on hardware that is not ready.
- (C) Treat A and D as long-horizon targets with gates, B as learning at most, C as deprioritised, and adopt E as classical software.
- (D) Fund B and C for production now because both have clear business owners, and defer A, D and E until budgets allow.

Correct answer: C
Business value is necessary but not sufficient. Grounding and resource horizon decide when a use case can pay off; baseline strength and data loading decide whether it ever will. A quantum-inspired classical solver that works now should be adopted as software and reported honestly as such.

Question 10

Domain 6: Programme Leadership and Delivery

Two months before a shareholder meeting, the CEO asks the architect to endorse four statements: "we have achieved quantum advantage in supply-chain optimisation" (one pilot matched but did not beat a tuned solver), "our quantum platform is production ready" (used for pilots only), "we are fully protected against quantum attacks" (hybrid TLS on 60 percent of long-lived flows, PKI scheduled for 2028), and "we partner with leading quantum hardware providers" (two active contracts). Which response is correct?

- (A) Endorse all four statements, because the shareholder meeting is important and each claim is at least partly supported.
- (B) Endorse the first and third statements because they are close enough, and decline the second as clearly premature.
- (C) Decline the first three as written, supply accurate alternatives for each, and let the fourth stand because it is factual.
- (D) Refuse to comment on any of the four statements and refer the CEO to the communications team for a final decision.

Correct answer: C
Overclaiming is the most common way quantum programmes lose credibility and, when security is involved, it can also mislead regulators. The architect's job is to supply statements that are both true and useful: learning pilots with published methodology, a shared platform for research, and a post-quantum migration under way on the NIST timeline.

Published by the CertLabz Certification Board, a division of CertLabz LLC. Companion to objectives document CQSOA-001. https://certlabz.com/certifications/certified-quantum-solutions-architect.html