CHAPTER 11 / 5 MINUTE BRIEF

Exotic Hardware: Quantum and Emerging Capacity

Read the full chapter → / Read this brief online →

The main idea

"Do we have a plan for quantum?" is an awkward question to receive in a meeting about next quarter's capacity. It may mean that a research team needs access to a particular machine, that Security needs help with a cryptographic migration, or that somebody read an impressive announcement. Those are three different requests. Answering all of them with "we've got it taken care of" would be efficient only until somebody asked what we meant.

I want a capacity team to take the question seriously without making every promising technology a purchasing requirement. The useful response names what exists, what the organization might do with it, the evidence that would justify a commitment, and the work we should already be doing on ordinary infrastructure: preparing for post-quantum cryptography does not require a quantum computer.

Explore this section in the full chapter →

Quantum, honestly assessed

Quantum error correction is making measurable progress. In Google's Willow experiment, logical error rates fell as the surface code grew from distance three to five to seven: adding physical qubits made the encoded memory more reliable.[1] IBM's published roadmap targets a fault-tolerant Starling system in 2029 with 200 logical qubits and a circuit capacity

[1] Google Quantum AI and collaborators, "Quantum error correction below the surface code threshold," *Nature*, 2024, https://www.nature.com/articles/s41586-024-08449-y. Accessible author account: https://blog.google/innovation-and-ai/technology/research/google-willow-quantum-chip/. The result concerns the demonstrated error-correction experiment, not enterprise workload economics.

of 100 million gates.[2] The first is an experimental result and the second a supplier's target; both matter, in different columns of the evidence record.

Two results deepened the experimental column in October 2026. An error-corrected machine runs some operations cheaply and consumes a prepared state, called a **magic state**, to perform the rest; Google's post notes estimates that producing such states can occupy up to 95 percent of a fault-tolerant machine's capacity. On the same Willow processor, a filtering method called magic state cultivation — repeated error checks that keep and grow only the cleanest attempts — cut errors in the prepared states roughly fortyfold at a reported 99.99 percent fidelity, with the authors' own note that demonstration at larger scale remains to be done.[3] Quantinuum, whose systems are available through Azure Quantum, reported a 99.9995 percent logical fidelity on its commercial Helios machine — five errors in a million operations, ten times better than its own result the month before — and ran two error-correction codes on the same hardware at different physical-qubit costs, ten versus up to thirty-seven per logical qubit.[4]

[2] IBM, "How IBM will build the world's first large-scale, fault-tolerant quantum computer," June 10, 2025, https://www.ibm.com/quantum/blog/large-scale-ftqc. Supplier roadmap and explanation of logical versus physical qubits; dates and delivered capabilities must be checked when procuring.

[3] Google Quantum AI, "Bringing the magic: Google breaks open a bottleneck in quantum computing," October 15, 2026, https://blog.google/innovation-and-ai/models-and-research/quantum-computing/magic-state-cultivation-quantum-computing/; the paper, "Magic state cultivation on a superconducting quantum processor," is in *Science*, https://www.science.org/doi/10.1126/science.aee7961. First hardware demonstration of the technique: about a fortyfold error reduction and a reported 99.99 percent T-state logical fidelity on the Willow processor. The post notes estimates that up to 95 percent of a machine's capacity can go to producing magic states and states that demonstration at larger scale remains. A vendor's account of its own result.

[4] Quantinuum, "Quantinuum Breaks Error-Correction World Record by 10x, Validating its Leading Path to Fault-Tolerant Quantum Computing by 2029," October 19, 2026, https://www.quantinuum.com/press-releases/quantinuum-breaks-error-correction-world-record-by-10x-validating-its-leading-path-to-fault-tolerant-quantum-computing-by-2029. Vendor announcement: a 99.9995 percent logical fidelity on the commercial Helios system using a color code, ten times the company's own September 2026 result on the same hardware; the announcement states its Helix code used 10 physical qubits per logical qubit against up to 37 for the color code, and links the result to its planned Apollo system. Quantinuum appears among Azure Quantum's priced providers; the record and the roadmap claims are the company's own.

A **physical qubit** is part of the hardware. A **logical qubit** is encoded information protected using a particular error-correction scheme. The number of physical qubits needed for a useful logical computation depends on error rates, code, connectivity and the operations being performed. Dividing one headline by another, such as a million-qubit resource estimate by the latest logical-qubit count, therefore does not measure the remaining distance.[5]

Explore this section in the full chapter →

For an enterprise capacity decision, qubit count is only a beginning. I would ask which algorithm the workload needs, what accuracy and confidence the result requires, how many repetitions it takes, and what classical preprocessing and verification remain. Then compare the complete run with the best suitable classical approach. A benchmark designed to expose a quantum advantage can be an important scientific result without establishing a cheaper way to complete the company's work.

That is a reason to keep qualified people engaged, not to dismiss the field or to buy a machine on a roadmap's date.

Explore this section in the full chapter →

[5] IBM, "How IBM will build the world's first large-scale, fault-tolerant quantum computer," June 10, 2025, https://www.ibm.com/quantum/blog/large-scale-ftqc. Supplier roadmap and explanation of logical versus physical qubits; dates and delivered capabilities must be checked when procuring.

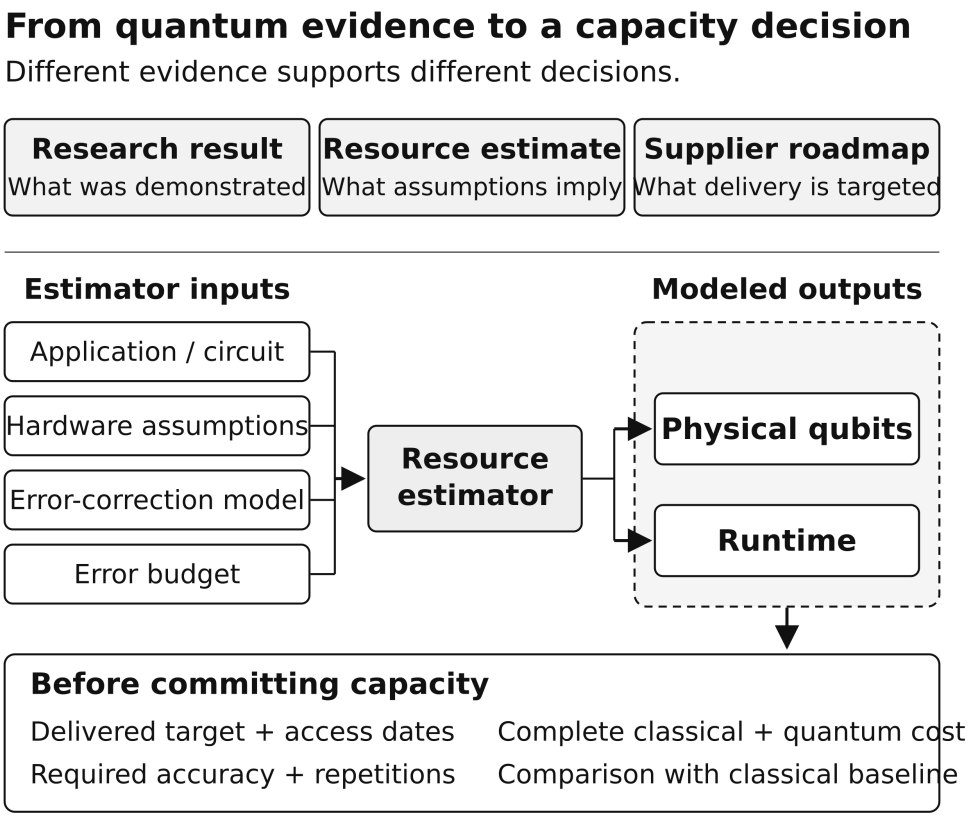


Diagram 11.1. A demonstrated result, a resource estimate and a supplier roadmap support different decisions. The estimator maps application and hardware assumptions, error correction and error budget to physical qubits and runtime. Access, complete cost and comparison with a classical baseline remain separate capacity checks.

Explore this section in the full chapter →

Post-quantum cryptography is a different project

Post-quantum cryptography replaces vulnerable cryptographic mechanisms with ones designed to resist quantum attacks. It includes key establishment and digital signatures, not just encryption. NIST finalized ML-KEM, ML-DSA and SLH-DSA in August 2024; these algorithms run on classical computers.[6] An organization can therefore have a substantial quantum-related migration while purchasing no quantum hardware.

Explore this section in the full chapter →

[6]NIST, FIPS 203, https://csrc.nist.gov/pubs/fips/203/final; FIPS 204, https://csrc.nist.gov/pubs/fips/204/final; FIPS 205, https://csrc.nist.gov/pubs/fips/205/final. Finalized August 13, 2024; key encapsulation and signature standards for classical implementation.

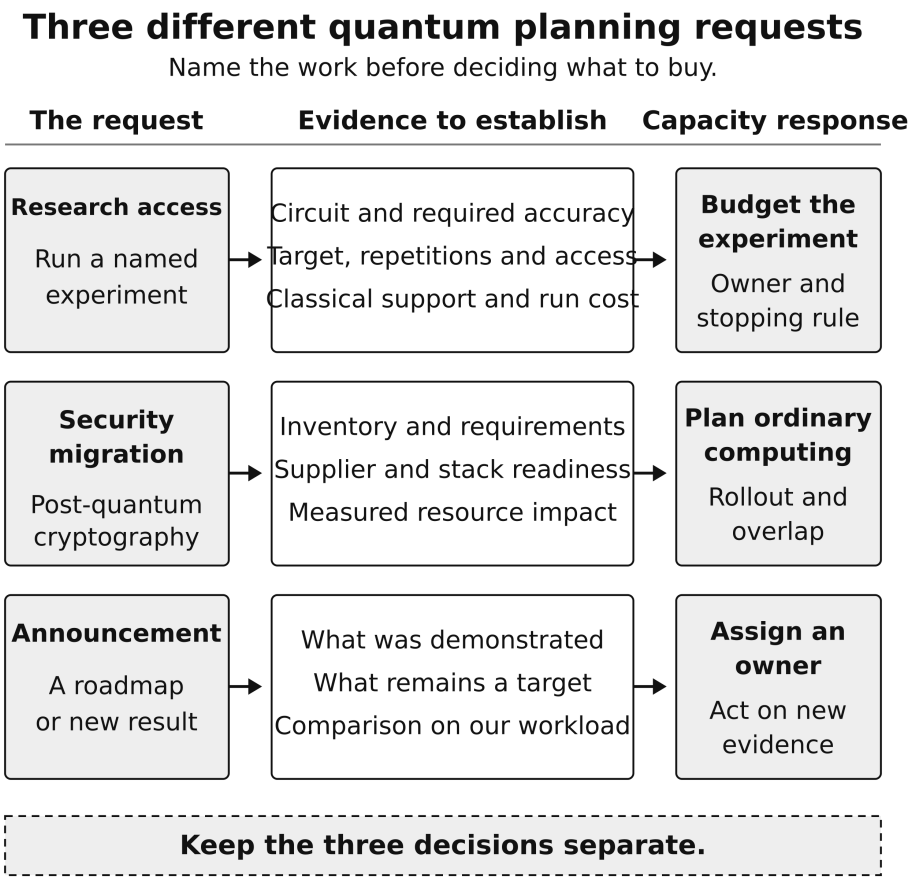


Diagram 11.2. Research access, cryptographic migration and a supplier announcement ask for different evidence and capacity responses. An experiment needs a budget and stopping rule; post-quantum migration uses ordinary computing; a roadmap needs an owner who can recognize when the evidence changes.

Explore this section in the full chapter →

The evidence standard for emerging architectures

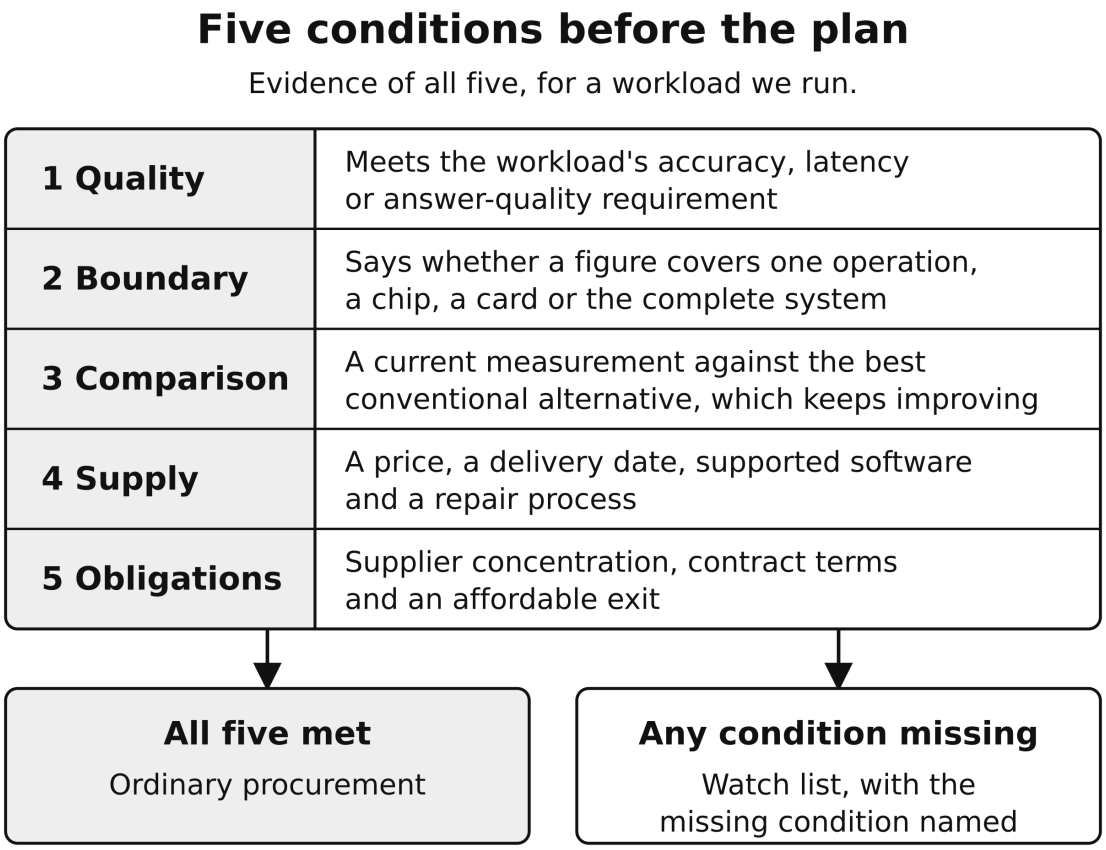


Diagram 11.3. An architecture enters the plan with evidence of all five conditions for a workload we run: quality, boundary, comparison, supply and obligations. A product that meets them all is ordinary procurement; anything short of that goes on the watch list, with the missing condition named.

Explore this section in the full chapter →

Photonic

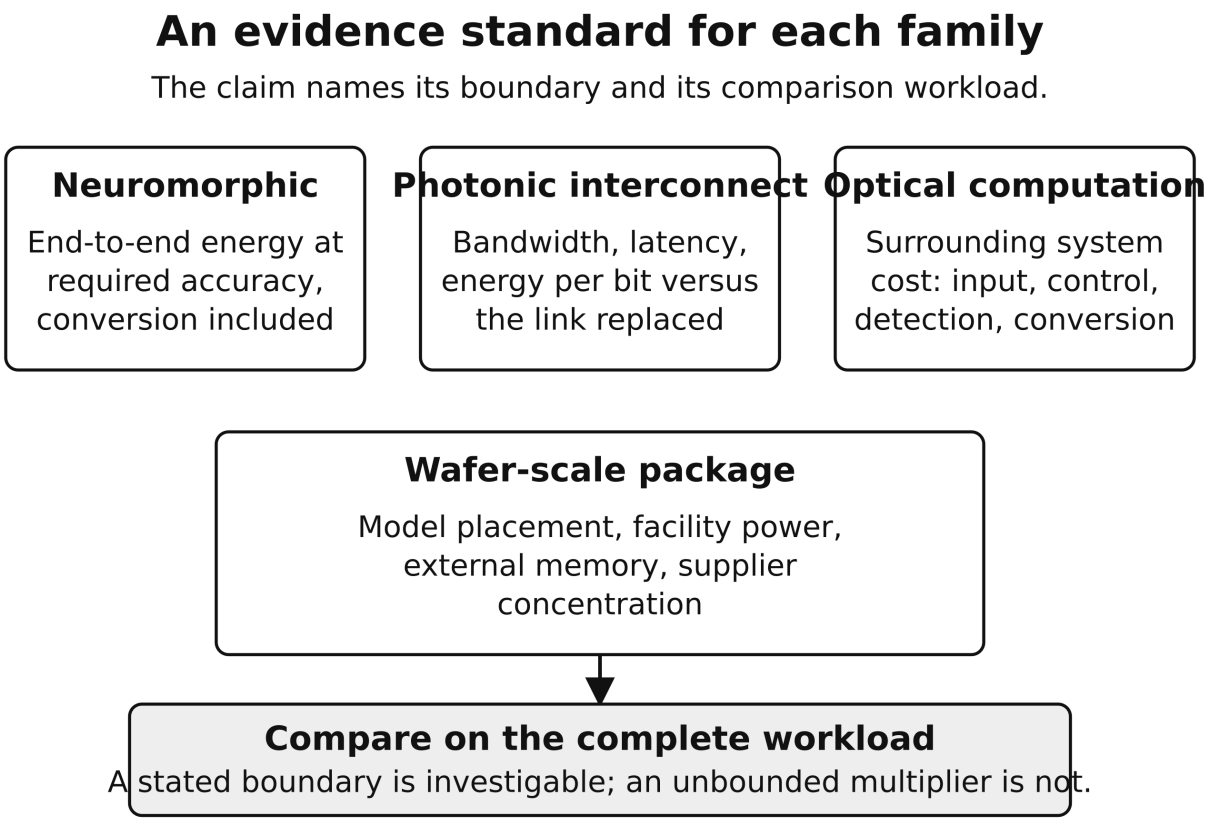


Diagram 11.4. One rule applies to every family: a claim names its boundary and its comparison workload. That means energy at the required accuracy for neuromorphic chips, energy per bit against the replaced link for optical interconnect, the surrounding system for optical computation, and placement, facility power and supplier concentration for wafer-scale systems.

Explore this section in the full chapter →

Quantum annealing

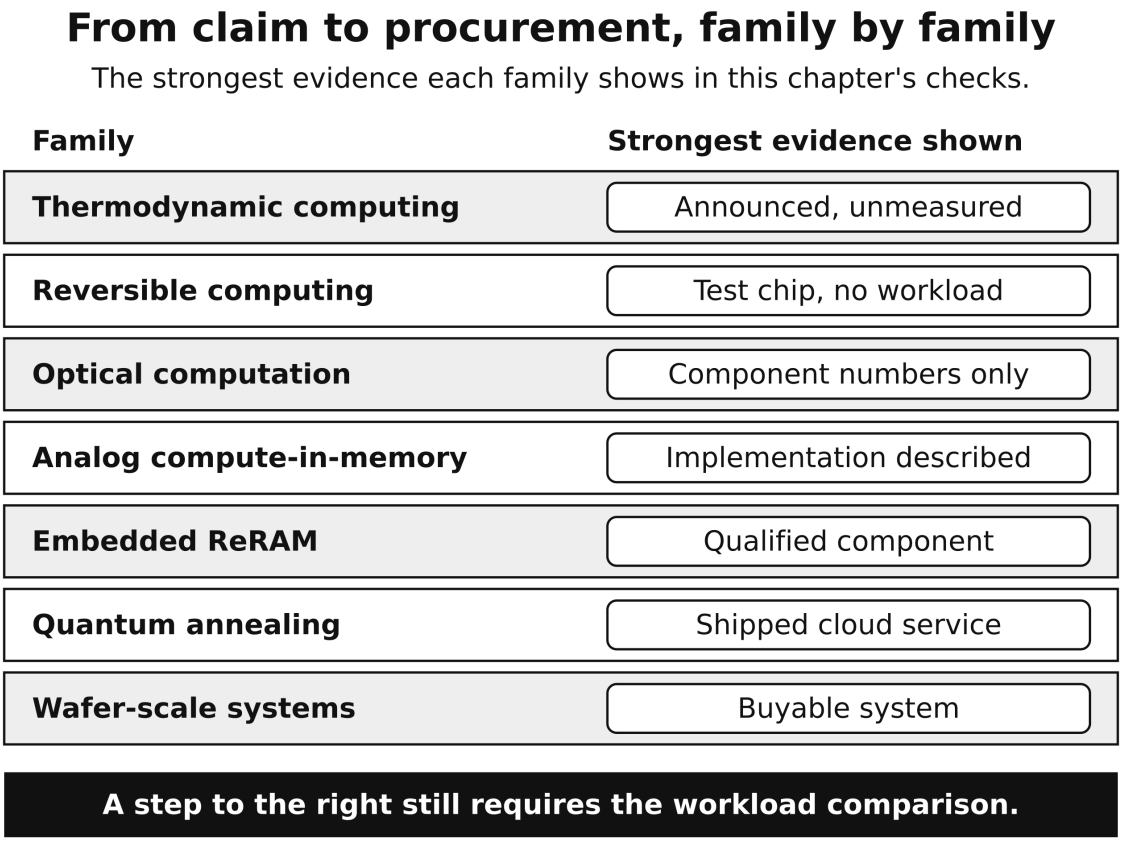


Diagram 11.5. Each hardware family in this section sits at a different evidence stage: announced designs, measured components, a qualified component you can source, or systems you can buy today. Every stage still faces the same workload comparison before it enters a capacity plan.

Explore this section in the full chapter →

Wafer-scale

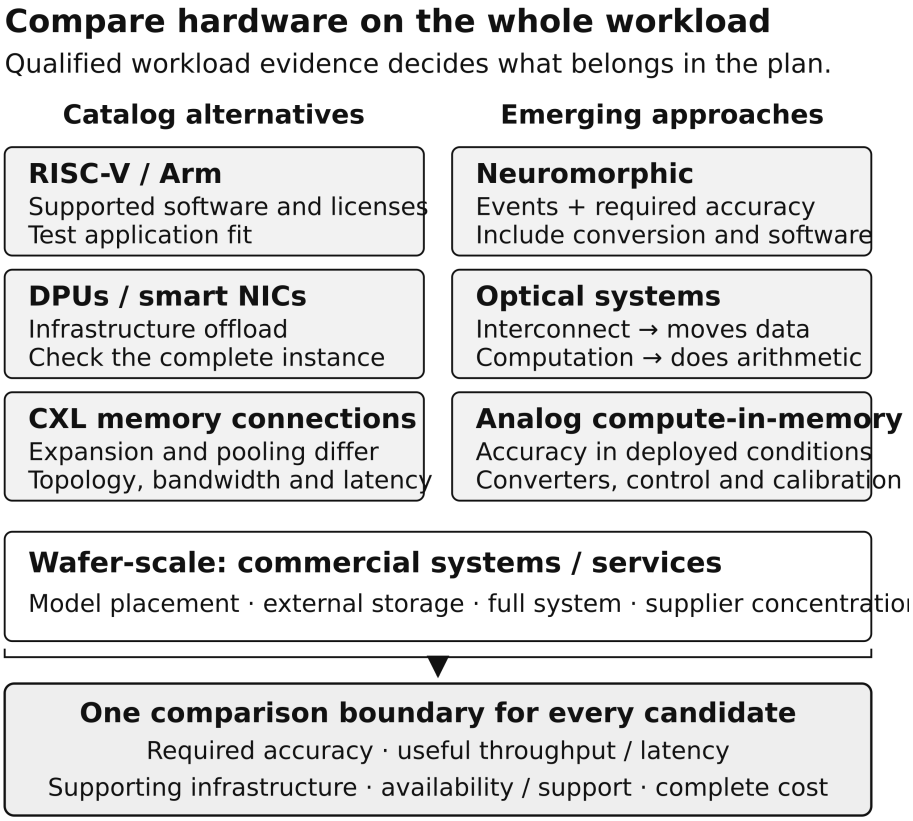


Diagram 11.6. Catalog alternatives, emerging approaches and commercial wafer-scale systems change different parts of a computing system. Every candidate is compared on the complete workload, including support, surrounding infrastructure and cost.

Explore this section in the full chapter →

The watch list

The watch list should say what would change a decision and who will recognize the change. “Quantum is interesting” is too vague to review. “Run a funded comparison when a supplier can execute this circuit at this accuracy within this experimental budget” gives the specialist something to evaluate.

Entry and owner	Evidence to keep	Decision it could change
Quantum experiment (research lead)	Circuit, accuracy, repetitions, classical baseline, access and complete run cost	Whether to fund the next research stage
Emerging accelerator or priced specialist service (platform owner)	Quality and latency on a representative workload, full-system power, price and terms, supported software	Whether to pay for a comparison or procure for a named workload
New memory or interconnect (hardware specialists)	Supported configuration, measured bottleneck, placement and failure behavior	Whether the next pool should use it
Cryptographic migration (Security, with capacity)	Required algorithm and stack, interoperability results, measured resource impact, deadline	When to fund rollout capacity

The last entry belongs to an active migration program when requirements are established; it should not be left waiting for a quantum breakthrough. The others can remain conditional without becoming forgotten.

Explore this section in the full chapter →