Techstrong
a Futurum company

SPECIAL REPORT

All Aboard the Quantum Train: It may be arriving and leaving at the same time

The state of quantum computing — the players, the money, the promise, what is real and what comes next.

By Alan Shimel

FOUNDER & CEO, TECHSTRONG GROUP
WITH RESEARCH DATA FROM FUTURUM RESEARCH

AUGUST 2026 TECHSTRONG RESEARCH TECHSTRONG.AI

Techstrong
a Futurum company

ABOUT THIS REPORT

The market arrived before the machine. That gap is the story.

This Techstrong Special Report examines the commercial, technical and policy state of quantum computing as of August 2026: what has genuinely shipped, what has merely been funded, what has only been promised, and what enterprise leaders should actually do about it between now and the end of the decade.

THE ARGUMENT IN FIVE FINDINGS

Finding 01

The market arrived before the machine

Enterprise end users outspent laboratories and governments for the first time in 2025, while no generally useful, fault-tolerant quantum computer runs production workloads.

Finding 02

Five milestones are being sold as one

Benchmark supremacy, below-threshold error suppression, a logical qubit, fault tolerance and utility scale are different achievements. Only the first three have been demonstrated.

Finding 03

The industry has no transistor moment

Seven credible modalities and roughly 67 QPU builders compete with no agreed equivalent of a CMOS process. Consolidation has already begun.

Finding 04

AI is already inside the machine

Learned decoders, reinforcement-learning calibration and GPU-based real-time error correction are improving quantum hardware today. Quantum for AI remains research.

Finding 05

One quantum program has a deadline

Post-quantum cryptography has finalized NIST standards, federal mandates and a migration horizon that does not depend on any vendor roadmap.

The test

Value must exceed cost

DARPA's definition of utility scale — computational value exceeding cost against the best classical alternative — is the only standard that settles the argument.

> “No enterprise can purchase a quantum server, move a material portion of its computing estate onto it and document a superior return. That is the state of quantum computing in one contradiction.”
>
> ALAN SHIMEL · ALL ABOARD THE QUANTUM TRAIN · THE OPENING

Techstrong
a Futurum company

INSIDE THIS REPORT

Contents

MASTHEAD, METHOD AND ATTRIBUTION

AUTHOR

Alan Shimel — Founder and CEO, Techstrong Group. A four-decade veteran of the technology industry and a longtime chronicler of how emerging platforms cross from laboratory to production.

RESEARCH ATTRIBUTION

This report incorporates research and analysis from **Futurum Research**, including the quantum processing unit market landscape work of **Dr. Bob Sutor**, the AI-and-quantum analysis *AI + Quantum: Hype or True Breakthroughs?*, the quantum manufacturing commentary of research director **Brendan Burke**, and **Daniel Newman**'s analysis *Quantum Fine-Tuning and the Energy Case for Quantum in AI*. Futurum figures are attributed at the point of use and in the endnotes.

METHOD

Every quantitative claim in this report is attributed to a named primary or near-primary source: peer-reviewed papers, vendor roadmaps and financial filings, government standards and policy documents, and named analyst research. Where a number is a model, a scenario or a company claim awaiting independent replication, the text says so. Numbered markers in the text correspond to the sources listed at the foot of each page and collected in *Sources and Notes*.

WHAT THIS REPORT DOES NOT DO

It does not forecast a date for useful quantum advantage, rank vendors, or treat a roadmap milestone as a delivered capability. Where the evidence stops, the report stops.

EDITORIAL NOTE

Quantum vocabulary is used precisely throughout. *Advantage*, *utility*, *below threshold*, *logical qubit* and *fault tolerance* describe different milestones and are not used interchangeably.

PUBLISHED

August 2026 · Techstrong Group · Boca Raton, Florida

DATA SPREAD

Quantum by the Numbers

Nine figures that define the market as of August 2026 — each one attributed, and each one carefully bounded.

$300M

Enterprise end-user spending on quantum computing in 2025 — exceeding research laboratories and governments combined for the first time.

WALL STREET JOURNAL

$550M

Total estimated quantum computing market in 2025, with enterprises accounting for more than half of spending.

BCG

$12.6B

Venture investment into quantum-technology startups in 2025 — more than six times the prior year.

MCKINSEY

67

Companies across 17 countries manufacturing QPUs through ten modalities at the end of 2024. A snapshot of an unconverged industry, not a stable statistic.

FUTURUM RESEARCH · BOB SUTOR

$2.013B

Proposed U.S. CHIPS Act incentives across nine quantum companies, signed as non-binding letters of intent in May 2026.

U.S. DEPT. OF COMMERCE

2.14×

Logical error suppression achieved by Google when its surface code grew from distance five to distance seven — one logical memory, not a fault-tolerant computer.

NATURE · GOOGLE QUANTUM AI

2029

IBM's delivery target for Starling: 100 million quantum operations across 200 logical qubits. An engineering target, not an appointment.

IBM QUANTUM ROADMAP

2035

The year NIST intends to disallow quantum-vulnerable public-key cryptography; deprecation begins after 2030.

NIST IR 8547

2033

DARPA's outer checkpoint for an industrially useful quantum computer — defined as one whose computational value exceeds its cost.

DARPA QBI

$80.1M

IonQ's second-quarter 2026 revenue, up 287% year on year — across computing, networking, sensing and space assets, not quantum computation alone.

IONQ · Q2 2026 RESULTS

67μs

Reaction time of a GPU-based decoder for qLDPC error correction — 32 times faster than the Helios processor's two-millisecond requirement.

QUANTINUUM · NVIDIA

40%

Share of enterprise quantum budgets directed to algorithms and software in 2024, nearly double the 21% recorded in 2022.

BCG

How to read these numbers. None of the nine figures is evidence of useful quantum advantage. Four measure money, two measure engineering progress on a single subsystem, and three are dates set by vendors or governments. Full attribution appears in *Sources and Notes*.

OPENING · THE CONTRADICTION

The Market Arrived Before the Machine

Enterprise buyers, governments and cloud providers are spending real money on quantum computing. The generally useful, fault-tolerant computer they are spending it on has not been built.

For most of its history, quantum computing has been the technology that would change everything—as soon as it arrived.

Now the money is arriving. So are the customers, government incentives, cloud services, acquisitions, roadmaps and corporate quantum teams. *The Wall Street Journal* recently reported that enterprise users spent approximately $300 million on quantum computing in 2025, more than research laboratories and governments combined for the first time.[1] Among the companies preparing are insurers, banks, professional-services firms and manufacturers that remember what happened when generative AI moved from research curiosity to boardroom priority almost overnight.

The computer they are preparing for, however, has not arrived.

No generally useful, fault-tolerant quantum computer is running mainstream production workloads today. No enterprise can purchase a quantum server, move a material portion of its computing estate onto it and document a superior return. Most commercial projects remain experiments, algorithm-development programs, cloud-access arrangements or preparations for capabilities still on vendor roadmaps. Even the U.S. government's Quantum Benchmarking Initiative defines its mission as determining whether an industrially useful quantum computer—one whose computational value exceeds its cost—can be built by 2033.[2]

That is the state of quantum computing in one contradiction. The market has arrived before the machine.

> The market has arrived before the machine. That does not make the market imaginary — but it does mean the machine still has to show up.
>
> THE CENTRAL TENSION OF 2026

The contradiction does not make the market imaginary. Boston Consulting Group estimates that quantum computing generated a market of approximately $550 million in 2025, with enterprise end users accounting for more than half of spending. Its survey found that more than 60% of participating enterprises were spending at least $1 million annually.[3] McKinsey estimates that quantum-computing companies produced more than $1 billion in worldwide revenue during 2025 and that investment in quantum-technology startups reached $12.6 billion, more than six times the prior year.[4]

Those numbers describe real economic activity. They do not prove useful quantum advantage. They include cloud access, research contracts, government programs, consulting, software, application development, integration and internal capability building. They tell us that organizations are buying tickets. They do not tell us when the train reaches its destination.

1. Isabelle Bousquette, "Businesses Are Spending Big on Quantum," *The Wall Street Journal* (CIO Journal). wsj.com/cio-journal/businesses-are-spending-big-on-quantum-c73060ad
2. DARPA, "Quantum Benchmarking Initiative (QBI)." darpa.mil/research/programs/quantum-benchmarking-initiative
3. Boston Consulting Group, "Quantum Is Getting Real. CEOs Need to Shape Where It Creates Value" (2026). bcg.com/publications/2026/ceos-need-to-shape-where-quantum-creates-value
4. McKinsey & Company, *Quantum Technology Monitor 2026: A Commercial Tipping Point*. mckinsey.com — Quantum Technology Monitor 2026

SECTION 01 · VOCABULARY

What "Arrived" Actually Means

Quantum computing suffers from a vocabulary problem as much as an engineering problem — and the vocabulary problem is now costing enterprises money.

The field has used supremacy, advantage, utility, fault tolerance and commercial readiness in ways that can make one milestone sound like all the others.

A quantum computer can perform a benchmark that is extremely difficult to reproduce with a classical computer without doing anything commercially valuable. A research team can demonstrate that adding physical qubits reduces the error rate of one encoded logical qubit without having enough logical qubits to run a useful application. A company can sell time on a quantum processor even if the customer is purchasing research access rather than superior computation. An enterprise can have a production quantum program whose production output is knowledge, intellectual property or readiness—not a quantum-generated business result.

All of these are legitimate forms of progress. They are not interchangeable.

FIVE MILESTONES, FIVE MEANINGS

Benchmark supremacy. A task that is hard to reproduce classically. No commercial value implied.

Below threshold. Adding physical qubits reduces the logical error rate. One subsystem, demonstrated.

Logical qubit. An encoded qubit with active error detection and classical decoding. Memory is easier than computation.

Fault tolerance. Deep, universal circuits executed at the fidelity an application requires.

Utility scale. Computational value exceeds cost against the best classical alternative. Nobody is here yet.

Why the qubit count is the wrong scoreboard

The distinction begins with the qubit. A classical bit is deliberately forced into a stable zero or one. A quantum bit can be prepared in a combination of states and manipulated through interference and entanglement. Those properties provide access to mathematical methods unavailable to conventional machines. They also make quantum information exceptionally fragile. Heat, electromagnetic noise, material defects, imperfect control and unwanted interactions can destroy or distort a computation.

That is why raw qubit count is a poor scoreboard. A system with more physical qubits may be less useful than a smaller machine with higher fidelity, better connectivity or longer coherence. Useful fault-tolerant systems will encode a logical qubit across multiple physical qubits, repeatedly detect errors and use a classical decoder to infer the correction. Depending on the hardware, algorithm and error-correcting code, one useful logical qubit could require many physical qubits plus a substantial classical-control system.

Error mitigation attempts to improve results from noisy processors without fully protecting the computation. Error correction is the more demanding process of preserving encoded quantum information as operations continue. *Below threshold* means the physical errors are sufficiently controlled that increasing the size of the error-correcting code reduces the logical error rate rather than making it worse. It is an essential threshold, not the end of the trip.

> Below threshold is an essential threshold — not the end of the trip.

Google crossed an important version of that threshold with Willow. In a peer-reviewed *Nature* paper, researchers reported that increasing a surface code from distance five to distance seven suppressed logical errors by a factor of 2.14. The 101-qubit distance-seven code reached a logical error rate of approximately 0.143% per correction cycle, and its logical memory lifetime was about 2.4 times that of its best constituent physical qubits.[5]

That was a major result. It was still one logical memory, not a fault-tolerant computer executing an economically important application. The honest description does not diminish the achievement. It shows how many distinct engineering milestones sit between a physics experiment and a production computing platform.

Bigger code, smaller error: the Willow result

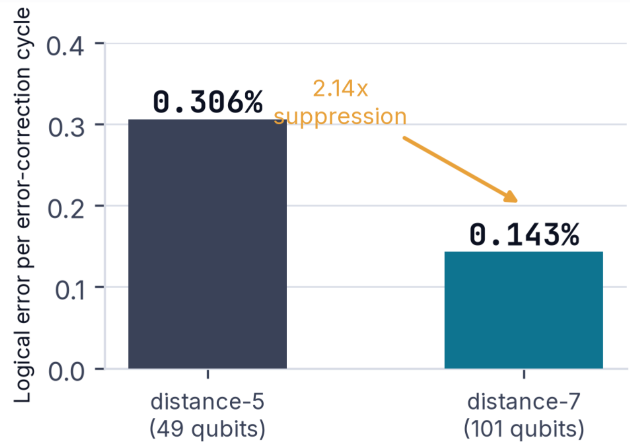

Source: Google Quantum AI, Nature (2024).

FIGURE 2 · ERROR SUPPRESSION

Adding qubits made the encoded qubit better rather than worse — the definition of operating below threshold. It did not make a computer.

Source: Google Quantum AI, *Nature* (2024). The distance-five value is derived from the reported 2.14× suppression factor.

Anatomy of one logical qubit: why a quantum computer is mostly a classical computer

LAYER 01

Physical qubits

Superconducting circuits, trapped ions, neutral atoms or photons. Fragile; degraded by heat, electromagnetic noise, material defects and crosstalk.

→

LAYER 02

Error-correcting code

Many physical qubits encode one logical qubit — a surface code or a qLDPC code. Larger code distance means stronger protection and higher overhead.

→

LAYER 03

Syndrome measurement

Repeated measurements report that *something* went wrong, without revealing the encoded state. Real hardware adds leakage, crosstalk and correlations.

LAYER 04

Classical decoder

A classical computer infers the most likely error from the syndrome stream — now increasingly a learned model or a GPU running inside a microsecond-scale budget.

→

LAYER 05

Control and calibration

Frequencies, amplitudes and phases drift continuously. Calibration has historically required stopping the computation.

→

LAYER 06

Logical qubit

One protected qubit. Useful applications need hundreds to thousands of them, executing deep universal circuits — not holding a memory.

Diagram based on the error-correction architecture described in Google Quantum AI's *Nature* (2024) below-threshold experiment and in IBM's published fault-tolerance architecture. Layer counts are conceptual; physical-qubit overhead per logical qubit varies by hardware, code and algorithm.

5. Google Quantum AI and collaborators, "Quantum error correction below the surface code threshold," *Nature* (2024). nature.com/articles/s41586-024-08449-y

SECTION 02

There is still no transistor moment

Seven credible modalities. Roughly 67 builders across 17 countries. No convergence, and no agreed equivalent of a CMOS process.

SECTION 02 · INDUSTRY STRUCTURE

From Feynman to a Market

Four decades after the theory, the industry looks less like the mature processor market than like early aviation — dozens of designs, no standard, and consolidation beginning.

The idea of using quantum mechanics to compute dates to the early 1980s. Richard Feynman argued that classical computers struggled to simulate quantum systems because nature itself was not classical. David Deutsch developed the concept of a universal quantum computer. Peter Shor changed the stakes in 1994 by showing that a sufficiently powerful quantum machine could factor large integers efficiently, threatening the public-key cryptography used throughout digital commerce. Lov Grover later showed a quadratic speedup for unstructured search.

Those algorithms established the theoretical case. The following decades became an engineering search for a physical system that could preserve, manipulate, measure and scale quantum information.

There is still no equivalent of the transistor around which the industry has converged. Superconducting circuits, trapped ions, neutral atoms, photons, silicon spins, topological approaches and quantum annealers make different trades among speed, fidelity, connectivity, manufacturability and operating complexity.

Seven bets, seven sets of unsolved problems

Superconducting processors from IBM, Google, Rigetti and others offer fast gates and build on techniques related to semiconductor fabrication. They generally require extreme cryogenic environments and carefully managed local connectivity. Trapped-ion systems from Quantinuum, IonQ and others offer high fidelity and flexible connectivity, but gates have historically been slower and scaling the optical and ion-control systems is difficult. Neutral-atom platforms from QuEra, Atom Computing and Infleqtion arrange atoms with lasers and can create large, reconfigurable arrays, although control fidelity and fault-tolerant operations remain active engineering challenges.

PsiQuantum is betting that photonics and semiconductor manufacturing can support very large fault-tolerant systems, but loss, sources, detectors and packaging remain central obstacles. Silicon-spin developers hope to benefit from established semiconductor methods, although consistent fabrication and control at scale are unresolved. Microsoft's topological program seeks qubits with inherent protection against some errors; its newest claims are striking, but the approach is still the least independently established of the leading commercial bets. D-Wave's annealers are commercially accessible and address optimization-style problems, but they are not universal gate-model quantum computers and should be evaluated in their own category.

Futurum analyst Bob Sutor counted approximately 67 companies across 17 countries manufacturing QPUs through ten modalities at the end of 2024. His conclusion was deliberately uncomfortable: too many companies were pursuing hardware while too few were investing in the networking required to combine processors into commercially useful systems.[6]

That count is a snapshot rather than a timeless market statistic. It nevertheless captures the absence of convergence.[7] Quantum looks less like the mature processor market than the early automobile, aviation or semiconductor industries, when dozens of designs competed before supply chains, standards, capital requirements and customer demand reduced the field.

6. The Futurum Group, "Is Quantum Computing Languishing? — Report Summary," analysis by Dr. Bob Sutor (December 2024). futurumgroup.com/press-release/is-quantum-computing-languishing/

7. Sutor's own later landscape work illustrates the point: a subsequent quarterly QPU market landscape from Sutor Group Intelligence and Advisory tracked 88 companies across 21 countries. drbobsutor.substack.com/p/sutor-group-quantum-processing-unit-449

> Not all 67 QPU builders need to survive for quantum computing to succeed. It is almost certain that they will not.
>
> ON CONSOLIDATION

Consolidation is already becoming part of the story. IonQ acquired Oxford Ionics and multiple networking, sensing and communications companies before agreeing in January 2026 to acquire U.S. chip manufacturer SkyWater Technology for approximately $1.8 billion — a transaction that cleared the Federal Trade Commission and closed on July 31, 2026.[8] The strategy is vertical integration: control more of the hardware, manufacturing and supply chain rather than depend on an ecosystem that does not yet exist at scale. Other companies will pursue partnerships, foundries or specialization. Not all 67 QPU builders need to survive for quantum computing to succeed. It is almost certain that they will not.

The modality matrix: seven physical bets on the same idea

MODALITY	REPRESENTATIVE BUILDERS	STRUCTURAL ADVANTAGES	UNRESOLVED ENGINEERING
Superconducting GATE MODEL	IBM, Google, Rigetti	Fast gates; fabrication techniques related to semiconductor processes	Extreme cryogenics; carefully managed local connectivity
Trapped ion GATE MODEL	Quantinuum, IonQ	High fidelity; flexible connectivity	Historically slower gates; scaling optical and ion-control systems
Neutral atom GATE MODEL	QuEra, Atom Computing, Infleqtion	Large, reconfigurable arrays assembled with lasers	Control fidelity; fault-tolerant operations
Photonic GATE MODEL	PsiQuantum	Bet on photonics plus semiconductor manufacturing at very large scale	Loss; sources; detectors; packaging
Silicon spin GATE MODEL	Diraq and others	Potential to inherit established semiconductor methods	Consistent fabrication and control at scale
Topological GATE MODEL	Microsoft	Qubits with inherent protection against some errors	Least independently established of the leading commercial bets
Quantum annealing NOT GATE MODEL	D-Wave	Commercially accessible; addresses optimization-style problems	Not a universal gate-model computer; evaluate in its own category

FIGURE 3 · MODALITY COMPARISON

Compiled from the modality assessments in this report. Builder lists are representative of the companies named in the manuscript, not exhaustive.

Context: approximately 67 companies across 17 countries were manufacturing QPUs through ten modalities at the end of 2024 (Futurum Research / Bob Sutor).

8. Reuters, "IonQ to buy SkyWater for $1.8 billion to expand hardware capabilities" (January 26, 2026) reuters.com; IonQ, "IonQ Completes Acquisition of SkyWater Technology" (July 31, 2026) ionq.com

SECTION 03 · THE MONEY

Following the Money

The amounts entering quantum are no longer laboratory-scale — and the composition of the spending has changed more than the total.

300+
Companies worldwide engaging with quantum computing
MCKINSEY 2026

1 in 3
Large organizations that allocated more than $10M to quantum in 2025
MCKINSEY 2026

7%
Allocated more than $50M to quantum initiatives in 2025
MCKINSEY 2026

~150
Enterprise proofs of concept running at roughly 100 enterprises by end-2024
BCG

According to McKinsey's 2026 Quantum Technology Monitor, more than 300 companies worldwide are engaging with quantum computing. Of the large global organizations it analyzed, one-third allocated more than $10 million to quantum initiatives in 2025 and 7% allocated more than $50 million. Most spending went to use-case and application development, integration and internal capabilities rather than hardware or operating-system development.

BCG's figures draw the same picture from another angle. It found quantum investment among 92% of the leading finance and insurance companies it examined, 56% of top health care and biopharma organizations and 52% of leading industrial companies. At the end of 2024, roughly 100 enterprises were participating in about 150 proofs of concept, 50% more than two years earlier. The proportion of enterprise quantum budgets directed toward algorithms and software nearly doubled from 21% in 2022 to 40% in 2024.

These numbers require care. "Investing in quantum" can mean hiring several researchers, funding a university partnership, purchasing cloud credits or running a proof of concept. It does not mean 92% of banks are receiving a production quantum advantage. Still, the movement from hardware tourism toward algorithms, integration and intellectual property is meaningful. Companies are spending less time asking what a qubit is and more time asking whether one could change bond trading, risk analysis, molecular simulation or logistics.

> Companies are spending less time asking what a qubit is and more time asking whether one could change bond trading, risk analysis, molecular simulation or logistics.
>
> ON THE SHIFT FROM HARDWARE TOURISM

Where quantum investment is concentrated — and where the budget is going

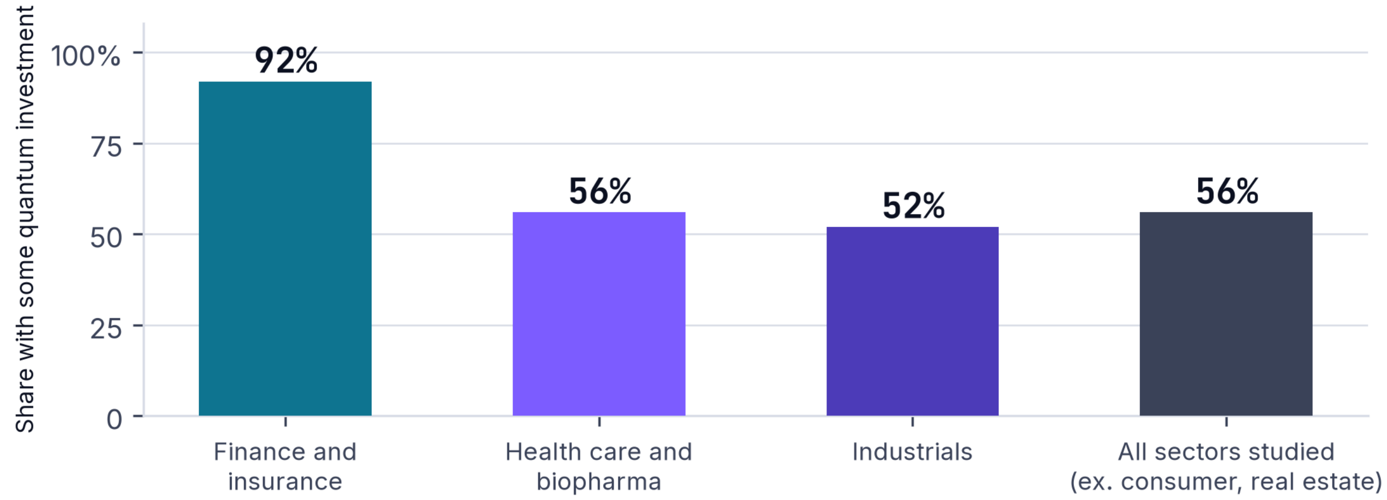

Source: BCG (2026), based on more than 200 large companies studied.

FIGURE 4 · SECTOR ADOPTION

Finance and insurance lead by a wide margin, but "investment" here is a broad category spanning research hires, university partnerships, cloud credits and proofs of concept.

Source: BCG (2026), based on more than 200 large companies studied.

Four ways to count quantum money in 2025 — and four different answers

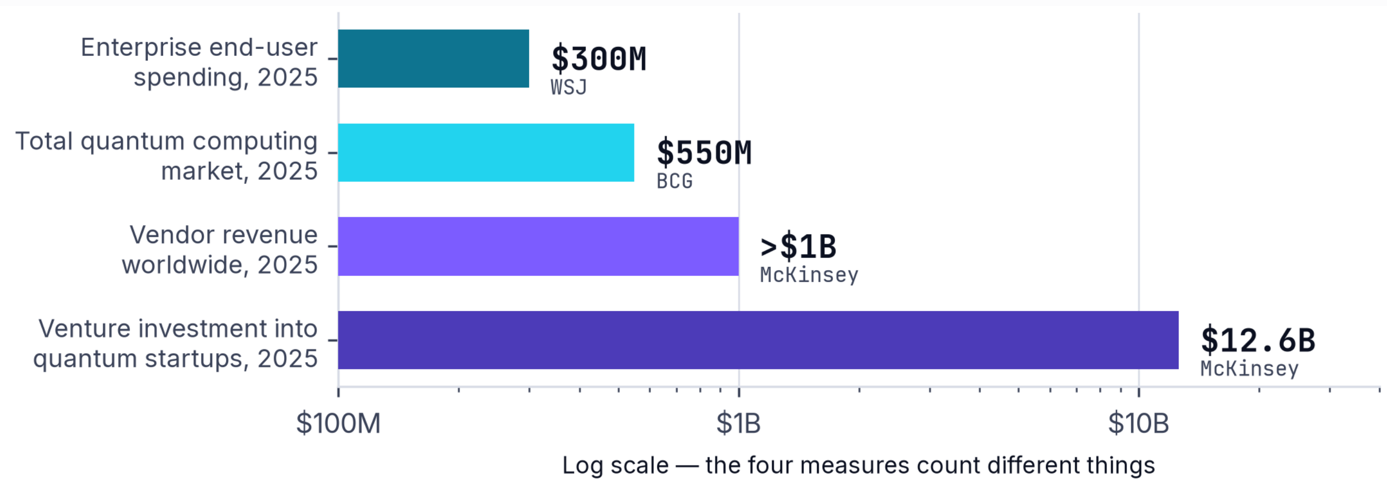

Sources: WSJ CIO Journal; BCG; McKinsey Quantum Technology Monitor 2026. Not additive.

FIGURE 1 · MARKET MEASUREMENT

The four headline figures of 2025 are frequently quoted as though they describe the same thing. They do not: one counts what enterprises spent, one counts total market size, one counts what vendors booked, and one counts capital raised by startups. Only the first two are demand signals.

Sources: *The Wall Street Journal* (enterprise spending); BCG (market size); McKinsey *Quantum Technology Monitor 2026* (vendor revenue, venture investment).

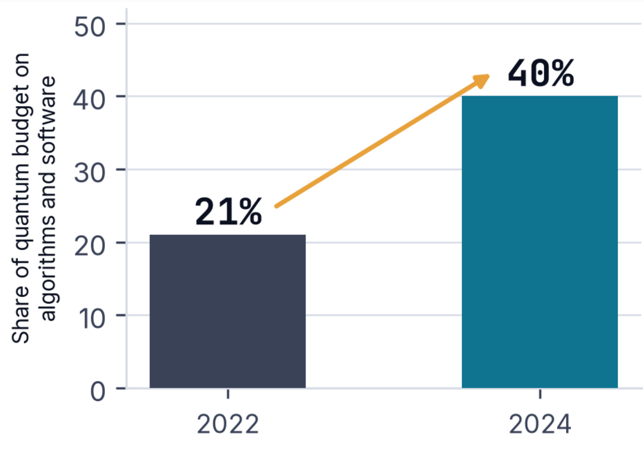

Source: BCG enterprise quantum research (2026).

FIGURE 5 · BUDGET COMPOSITION

The shift from hardware tourism to algorithms and software is the most meaningful change in enterprise quantum spending.

Source: BCG enterprise quantum research (2026).

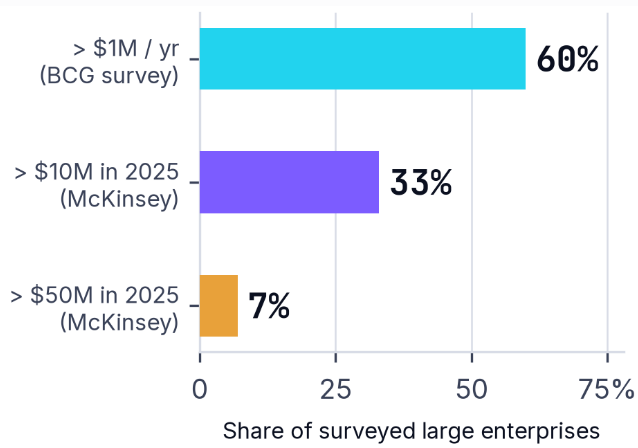

Sources: BCG; McKinsey Quantum Technology Monitor 2026.

FIGURE 6 · BUDGET SCALE

Quantum budgets are now large enough to require board-level justification at a meaningful minority of large enterprises.

Sources: BCG; McKinsey *Quantum Technology Monitor 2026*.

Public money moves from research to industrial capacity

Public investment is moving from general research toward industrial capacity. In May 2026, the U.S. Department of Commerce signed nine letters of intent proposing $2.013 billion in CHIPS Act incentives. The plan includes $1 billion for an IBM subsidiary producing quantum-grade superconducting wafers and $375 million for a GlobalFoundries facility supporting multiple modalities. Atom Computing, Diraq, D-Wave, Infleqtion, PsiQuantum, Quantinuum and Rigetti were selected for proposed awards ranging from $38 million to $100 million. In exchange, the government would receive minority, non-controlling equity stakes.[9]

These are proposed incentives, not completed awards, and the money should not be double-counted with corporate commitments. But the policy is significant. The government is not choosing one qubit. It is financing a portfolio of foundries, modalities, packaging, cryogenics, controls, readout and interconnects—the beginnings of an industrial base.

Futurum research director Brendan Burke has compared the strategy to a TSMC model for quantum. The analogy works up to a point. A neutral manufacturing layer could reduce the capital burden on individual vendors and turn repeatable processes into an ecosystem. Quantum is harder because the industry has not agreed on the equivalent of a CMOS process, and several modalities require radically different materials, operating environments and control systems. The less glamorous constraints—cryogenic components, lasers, photonic packaging, control electronics and test equipment—may decide which roadmap can actually be manufactured.

> The danger is an ecosystem optimized to finance and sell the promise to itself. The opportunity is that industrial capacity must exist before demand can scale. Both can be true.
>
> ON THE AI INFRASTRUCTURE PARALLEL

This is also where quantum starts to resemble the early AI infrastructure market. Revenue initially comes from building capacity, tools and access before end users can prove that the final applications justify the full investment. The danger is an ecosystem optimized to finance and sell the promise to itself. The opportunity is that industrial capacity must exist before demand can scale. Both can be true.

9. U.S. Department of Commerce / NIST, "Department of Commerce Announces Letters of Intent With 9 Companies for $2 Billion to Accelerate U.S. Leadership in Quantum Computing" (May 2026). nist.gov

SECTION 04

Every major vendor now has a date. None of them has an appointment.

IBM, Microsoft, Quantinuum and IonQ all cluster at the end of the decade — using four different definitions of a logical qubit.

SECTION 04 · ROADMAPS

Roadmaps Are Not Arrivals

The specificity is encouraging. It also creates false precision when engineering targets are read as delivery dates.

The leading vendors increasingly cluster around the end of the decade, but their numbers are not directly comparable.

IBM says it will deliver Starling in 2029, an error-corrected system designed to execute 100 million quantum operations across 200 logical qubits.[10] Quantinuum's roadmap points to Apollo as a universal fault-tolerant system capable of millions of gates by 2030.[11] IonQ's current roadmap projects 80,000 logical qubits and two million physical qubits by 2030.[12] Microsoft says Majorana 2 has cut its timetable for a scalable quantum computer to 2029.[13]

The specificity is encouraging. It also creates false precision if the dates are treated as appointments rather than engineering targets. Each company uses different hardware, logical-qubit definitions, error-correcting codes and performance measures. A logical qubit capable of holding a state is not necessarily a logical qubit capable of running deep, universal circuits at the required fidelity. An algorithm's value depends on gate count, connectivity, cycle time, data movement, decoding and the classical work surrounding it.

The only definition that settles the argument

DARPA's definition is more useful than any single qubit number: utility scale exists when computational value exceeds cost. That demands total time-to-solution, energy-to-solution, accuracy and price against the best classical alternative—not the classical method most convenient for the demonstration.

The industry has also learned that the classical baseline will not sit still. Better tensor-network methods, GPU simulation, approximation algorithms and AI-assisted scientific computing can erase a claimed advantage after the quantum experiment is announced. A credible result therefore needs a problem with practical meaning, a current classical comparison, transparent resource accounting, verifiable output and preferably independent replication.

The end-of-decade roadmaps could be substantially right, substantially wrong or right about the hardware while wrong about the applications. BCG identifies a risk that the industry may build impressive quantum capability without enough algorithms that solve valuable problems. That would be a peculiar arrival: a machine capable of things classical computers cannot do, but not yet capable of enough things customers will pay to have done.

The fault-tolerance roadmap: four vendors, four definitions, one decade

10. IBM, "IBM lays out clear path to fault-tolerant quantum computing" and the IBM Quantum roadmap. ibm.com/quantum/blog/large-scale-ftqc
11. Quantinuum, "Quantinuum Unveils Accelerated Roadmap to Achieve Universal Fault-Tolerant Quantum Computing by 2030." quantinuum.com
12. IonQ, "Roadmap." ionq.com/roadmap
13. Microsoft, "Majorana 2, made more reliable with Microsoft Discovery agentic AI" (June 2026). news.microsoft.com

2026	2029	2030	2033	2035
WHERE WE ARE Below-threshold logical memory demonstrated; real-time GPU decoding demonstrated; no fault-tolerant application in production.	IBM · STARLING 100 million quantum operations across 200 logical qubits. MICROSOFT Revised target for a scalable topological quantum computer.	QUANTINUUM · APOLLO Universal fault-tolerant system capable of millions of gates. IONQ 80,000 logical qubits and 2 million physical qubits.	DARPA · QBI Outer checkpoint for an industrially useful quantum computer — one whose computational value exceeds its cost.	SCENARIO HORIZON McKinsey models \$43B–\$71B of quantum computing market and up to \$2.7 trillion of economic value; BCG models \$450B–\$850B of value creation.

How to read this timeline. These are vendor engineering targets and a government checkpoint, not scheduled deliveries. The logical-qubit counts are not comparable across companies: each uses different hardware, error-correcting codes and definitions of a logical qubit. Sources: IBM Quantum roadmap; Microsoft; Quantinuum; IonQ; DARPA QBI; McKinsey *Quantum Technology Monitor 2026*; BCG (2026).

SECTION 05 · AI AND QUANTUM

AI Is the Catalyst, Not the Competition

The relationship is already symbiotic — and asymmetric. AI is improving quantum computing today, while quantum computing for AI remains primarily a research activity.

Quantum computing and AI are frequently presented as competing waves. The framing confuses a general-purpose software capability with a specialized computing architecture. AI is not waiting for quantum to fail, nor does quantum need AI demand to collapse.

The relationship is already symbiotic, but it is asymmetric. As Futurum's Sutor put it in *AI + Quantum: Hype or True Breakthroughs?*, AI is improving quantum computing today, while quantum computing for AI remains primarily a research activity.

Decoding: the first place AI got inside the machine

The most immediate example is error decoding. A quantum error-correcting code does not simply announce which qubit failed. It produces a stream of syndrome measurements from which a classical decoder must infer the most likely error. Real hardware produces leakage, crosstalk and correlations that hand-designed noise models do not capture perfectly.

Google DeepMind and Google Quantum AI trained AlphaQubit, a recurrent transformer, first on simulated errors and then on experimental data from the Sycamore processor. In the largest experiments, Google reported that AlphaQubit made 6% fewer errors than a highly accurate tensor-network decoder and 30% fewer than correlated matching. The peer-reviewed *Nature* paper found that it retained an accuracy advantage in simulations through code distance 11, representing 241 physical qubits, and generalized to experiments far longer than those on which it was trained.[14]

AlphaQubit is not yet a complete production decoder. Its speed and training demands must improve for the tightest real-time loops. But it demonstrates the principle: a learned model can capture the messy error distribution of an actual processor better than human-designed algorithms alone.

From correcting errors to steering the hardware

Google has since extended learning from correction into control. Quantum computers are analog machines

14. J. Bausch et al., "Learning high-accuracy error decoding for quantum processors," *Nature* (2024). nature.com/articles/s41586-024-08148-8; Google DeepMind, "AlphaQubit." blog.google

whose frequencies, amplitudes and phases drift. Today, computation is commonly stopped so the hardware can be recalibrated. That is incompatible with future algorithms that may need to operate continuously for days or months.

In a 2026 *Nature* experiment, Google used reinforcement learning to adjust more than 1,000 control parameters on Willow using error-detection events as the learning signal. Against deliberately injected drift, the combination of controller and decoder steering improved logical stability 3.5-fold. Even after conventional calibration and human tuning, reinforcement-learning fine-tuning suppressed the logical error rate by another 20%. Simulations extended the approach to a distance-15 code with almost 40,000 parameters and found that convergence speed did not deteriorate with system size.[15]

That is AI acting as part of the machinery, not as a workload waiting in line to replace it.

AI may make quantum arrive sooner while making "arrival" harder to claim.

AI can also search for fault-tolerant circuits, tune control pulses, classify readout signals and identify patterns in experiments. In 2025, researchers reported using reinforcement learning to discover fault-tolerant logical-state preparation circuits that matched and sometimes outperformed human-designed versions while requiring fewer steps and resources. The designs were demonstrated across as many as 25 physical qubits, suggesting that AI could help search a hardware-dependent design space that becomes increasingly difficult to navigate manually.[16]

Microsoft says its Discovery agentic-AI system helped researchers develop the materials stack used in Majorana 2. The company reports a 1,000-fold reliability improvement over its previous generation, a mean qubit lifetime of 20 seconds and a revised 2029 target for scalability. Those are company claims surrounding a technically controversial modality, and they need independent validation. They nevertheless illustrate how AI-driven scientific discovery could compress the experimentation loop.

The infrastructure built for AI is catalyzing quantum research as well. AWS, NVIDIA and university researchers integrated NVIDIA's cuQuantum libraries with QuTiP and reported up to a 4,000-fold speedup for selected 64-state-qubit simulations using eight H200 GPUs. That is not a quantum speedup. It is classical accelerated computing making it faster to model qubit dynamics, readout and two-qubit gates before consuming scarce hardware time.[17]

The classical system does not disappear when fault tolerance arrives. It grows more important. Quantinuum and NVIDIA demonstrated a GPU-based decoder for qLDPC error correction that reacted in 67 microseconds, 32 times faster than the Helios processor's two-millisecond requirement.[18] IBM calls its model quantum-centric supercomputing: the QPU performs the portion suited to quantum methods while CPUs, GPUs and HPC systems handle orchestration, simulation, optimization, decoding and data processing.

AI therefore pushes quantum in two directions at once. It helps design and stabilize the hardware, which could pull useful systems closer. It also improves classical chemistry, optimization, simulation and scientific discovery, which pushes the competitive threshold farther away. AI may make quantum arrive sooner while making "arrival" harder to claim.

15. Google Quantum AI and Google DeepMind, "Towards a quantum computer that learns from its errors," *Nature* (2026). research.google/blog/towards-a-quantum-computer-that-learns-from-its-errors/
16. R. Zen, J. Olle, L. Colmenarez, M. Puviani, M. Müller and F. Marquardt, "Quantum Circuit Discovery for Fault-Tolerant Logical State Preparation with Reinforcement Learning," *Physical Review Research* / APS. link.aps.org/doi/10.1103/gqpr-dgz7
17. AWS Quantum Computing Blog, "Accelerating the Quantum Toolkit for Python (QuTiP) with cuQuantum on AWS." aws.amazon.com
18. NVIDIA, "World's Leading Scientific Supercomputing Centers Adopt NVIDIA NVQLink..." — Quantinuum's Helios processor demonstrating scalable real-time decoding for quantum error correction. investor.nvidia.com

What AI has already delivered inside the quantum stack

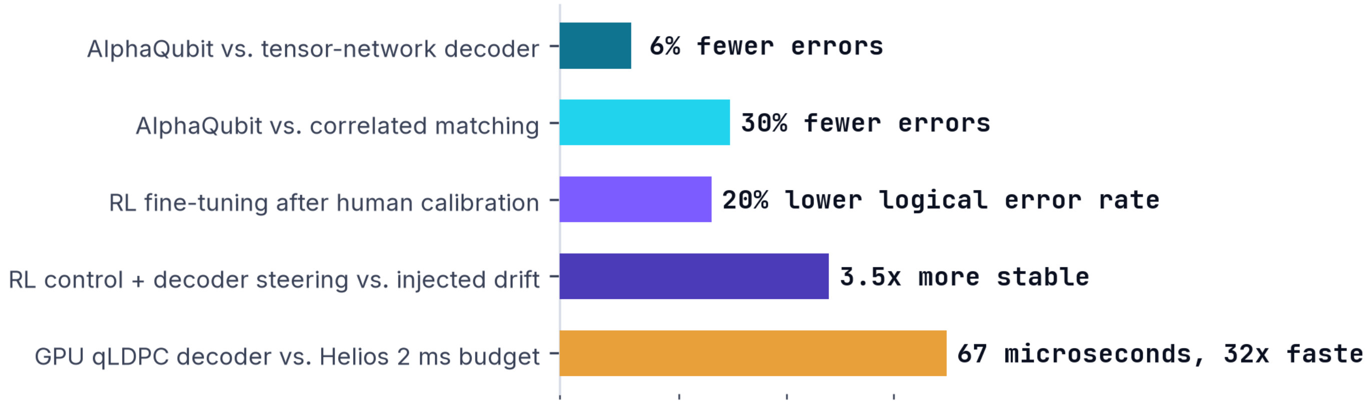

Bar length is illustrative of relative improvement only; each row is a different metric.

Sources: Google DeepMind / Google Quantum AI, Nature (2024) and Nature (2026); Quantinuum and NVIDIA (2025).

FIGURE 7 · AI INSIDE THE MACHINE

Five separate reported results, each measuring a different quantity. Bar lengths indicate relative magnitude of improvement only and are not comparable across rows.

Sources: Google DeepMind / Google Quantum AI, *Nature* (2024) and *Nature* (2026); Quantinuum and NVIDIA real-time qLDPC decoding on Helios via NVQLink (2025).

SECTION 06 · QUANTUM FOR AI

Can Quantum Return the Favor?

The reverse relationship — quantum accelerating AI — is intriguing and immature. The energy argument is the right question, not yet the answer.

Current quantum computers cannot conveniently load enormous classical datasets, copy unknown quantum states or run the depth of reliable operations associated with large neural networks. A small quantum classifier that performs well on a curated dataset is not evidence that a QPU can train a frontier model more efficiently than a GPU cluster. Many "quantum AI" results are produced on classical simulations of quantum circuits or compared against modest classical baselines.

That does not make the entire category empty. It makes the burden of proof higher.

A May 2026 preprint from IonQ, QuantumBasel and university researchers reported an end-to-end hybrid pipeline in which a parameterized quantum circuit served as a fine-tuning head for a foundation model. The researchers directly measured power consumption on an IonQ Forte Enterprise trapped-ion system. Their best quantum-tuned model reduced classification error by about 24% against the best classical fine-tuning baseline considered, and their scaling analysis projected an energy break-even against classical simulation near 34 qubits.[19]

Daniel Newman's Futurum analysis, *Quantum Fine-Tuning and the Energy Case for Quantum in AI*, argues that energy-to-solution may ultimately matter more than the qubit-count contest. That is the right commercial question. It is not yet a commercial answer.[20]

READ THE FINE PRINT

Four limits on the energy result

Scale. The measured experiments covered 12 to 18 qubits. The 34-qubit crossover was extrapolated.

Workload. The task was sentiment classification, not a frontier model.

Baseline. The energy comparison was between QPU execution and classical *simulation of the same quantum circuit* — not between the QPU and the best production GPU method for the underlying business task.

Status. A preprint submitted to conference review. A credible research signal, not a demonstrated energy advantage for mainstream AI.

The likely sequence is narrower than the hype. AI and quantum components will first share workflows. A classical model may prepare data, identify a difficult subproblem or optimize parameters; a QPU may sample, simulate or transform a carefully structured portion; classical systems will interpret and validate the result. If quantum eventually accelerates AI, it will probably enter as a specialized coprocessor long before anyone speaks sensibly of a quantum replacement for the AI factory.

> If quantum eventually accelerates AI, it will enter as a specialized coprocessor long before anyone speaks sensibly of a quantum replacement for the AI factory.

ON THE LIKELY SEQUENCE

19. O. Knitter et al., "Measuring Accuracy and Energy-to-Solution of Quantum Fine-Tuning of Foundational AI Models," arXiv preprint (May 2026). arxiv.org/html/2605.02798
20. Daniel Newman, "Quantum Fine-Tuning and the Energy Case for Quantum in AI," Futurum Research (June 25, 2026). futurumgroup.com/insights/quantum-fine-tuning-and-the-energy-case-for-quantum-in-ai/

SECTION 07

One part of the quantum market has unambiguously arrived

Post-quantum cryptography has finalized standards, federal deadlines and a multi-year migration path. It does not require believing any vendor roadmap.

SECTION 07 · POST-QUANTUM CRYPTOGRAPHY

The Quantum Application That Cannot Wait

Preparing security for the possibility that cryptographically relevant quantum computers eventually arrive is the one quantum program with a deadline that has already passed the decision point.

One part of the quantum market has unquestionably arrived: preparing security for the possibility that cryptographically relevant quantum computers eventually do.

Shor's algorithm threatens RSA and elliptic-curve cryptography, which support key exchange, digital signatures, identity, certificates, software updates, financial transactions and much of the trust infrastructure of the internet. No known quantum computer can break properly sized implementations today. The migration must begin before one can.

The reason is not only "Q-Day." Adversaries can collect encrypted traffic now and preserve it until a future machine can decrypt it. Data that must remain confidential for years is exposed to a capability that may not exist until years from now. Migration itself is a multi-year program because vulnerable cryptography is scattered across applications, devices, protocols, hardware, certificates, vendors and long-lived infrastructure.

In August 2024, NIST finalized three principal standards: FIPS 203 for ML-KEM key establishment, FIPS 204 for ML-DSA digital signatures and FIPS 205 for SLH-DSA signatures.[21] NIST now says organizations should begin applying them. Its transition plan intends to deprecate quantum-vulnerable public-key algorithms after 2030 and disallow them after 2035, with high-risk systems moving earlier.[22]

In June 2026, the White House ordered federal agencies to designate PQC migration leads, maintain cryptographic inventories and develop prioritized plans. That turns crypto-agility from a conference recommendation into an operational discipline.[23]

> "If you do nothing else, transition your cybersecurity infrastructure to one that is crypto-agile and employs the NIST PQC standards."
>
> BOB SUTOR · FUTURUM RESEARCH

Futurum's Sutor offered the most concise enterprise advice: "If you do nothing else, transition your cybersecurity infrastructure to one that is crypto-agile and employs the NIST PQC standards."

Crypto-agility means knowing where cryptography lives and being able to replace algorithms without replacing the surrounding system. The first step is inventory, not indiscriminate deployment. Organizations need to identify exposed algorithms, data-retention periods, external dependencies and systems that are difficult to update. They must test performance, interoperability, key management, certificates and hybrid transition methods. They also need procurement language requiring suppliers to disclose and update cryptographic dependencies.

Post-quantum cryptography is classical software designed to resist both classical and quantum attacks. It should not be confused with quantum key distribution, which uses quantum properties to detect interception but requires specialized infrastructure, or quantum random-number generation, which already has commercial applications. PQC is the broad enterprise migration path

21. NIST, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards" (August 13, 2024). nist.gov
22. NIST IR 8547, *Transition to Post-Quantum Cryptography Standards* (initial public draft). nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8547.ipd.pdf
23. Executive Order 14412, "Securing the Nation Against Advanced Cryptographic Attacks" (June 2026) whitehouse.gov; OMB Memorandum M-26-15, "Execution of the Migration to Post-Quantum Cryptography" (June 24, 2026) whitehouse.gov (PDF)

because it runs on existing computing and networking systems.

An organization does not need to believe a large fault-tolerant machine arrives in 2029 to act. It needs only to believe that sensitive data may outlive the migration program. That calculation has already crossed the threshold.

The migration clock: PQC is the only quantum program with published deadlines

Aug 2024	Jun 2026	2026-2030	After 2030	After 2035
NIST	WHITE HOUSE · OMB	ENTERPRISE WORK	NIST IR 8547	NIST IR 8547
FIPS 203 (ML-KEM), FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA) finalized.	Executive Order 14412 and OMB M-26-15: migration leads, cryptographic inventories and prioritized agency plans.	Inventory, data-retention analysis, supplier disclosure, interoperability and key-management testing, hybrid transition.	Quantum-vulnerable public-key algorithms deprecated.	Quantum-vulnerable public-key algorithms disallowed; high-risk systems expected to move earlier.

Sources: NIST (August 2024 standards release); NIST IR 8547 transition plan; Executive Order 14412 and OMB Memorandum M-26-15 (June 2026). PQC is classical software; it should not be confused with quantum key distribution or quantum random-number generation.

SECTION 08 · THE COMMERCIAL REALITY

What Is Commercially Real Today?

The market is real if we define it by what customers are purchasing rather than what the machines can ultimately do.

The commercial quantum market is real if we define it by what customers are purchasing rather than what the machines can ultimately do.

Customers can buy cloud access to processors from IBM, IonQ, Quantinuum, D-Wave, Rigetti, QuEra and others through provider and hyperscaler platforms. They can purchase dedicated systems or capacity, hire application-development teams, fund university partnerships, develop intellectual property, test quantum-safe security and integrate quantum jobs into HPC workflows. Governments can finance domestic capability and strategic supply chains. Researchers can perform experiments on hardware that would have been inaccessible outside a national laboratory a decade ago.

Some vendors are generating meaningful revenue, although scale and composition matter. IonQ reported $80.1 million in second-quarter 2026 revenue and raised full-year guidance to as much as $290 million, while continuing to expect a substantial adjusted EBITDA loss.[24] Its expanding portfolio now includes computing, networking, sensing and space-related assets, so rising company revenue is not a pure measure of customers paying for useful quantum computation.

The application work is becoming more serious. Banks are exploring risk, pricing, portfolio and trading problems. Chemical, pharmaceutical and industrial companies are investigating molecular simulation, catalysts, batteries and materials. Airlines, automakers and logistics providers are testing scheduling, routing and design. But a named enterprise partnership still does not establish advantage.

24. IonQ, "IonQ Announces Record Second Quarter 2026 Revenues, Growing 287% YoY" (August 2026). investors.ionq.com

> The answer today is often that the commercial product is learning. That is not worthless. A failed experiment that prevents a much larger misplaced investment has economic value.
>
> ON WHAT ENTERPRISES ARE ACTUALLY BUYING

The answer today is often that the commercial product is learning. That is not worthless. Early work can identify which problems have the mathematical structure quantum algorithms need, build scarce internal expertise and produce intellectual property. It can also reveal that a proposed use case has no credible path to advantage. A failed experiment that prevents a much larger misplaced investment has economic value.

The mistake is to call every learning program deployment and every vendor booking proof that the underlying computing advantage has arrived.

The six questions that separate a result from a press release

01

Was the result produced on real hardware?

Or on a classical simulation of a quantum circuit? Both are legitimate research. Only one is evidence about a machine.

02

What classical baseline was used?

The best current classical method — not the method most convenient for the demonstration. Tensor networks and GPU simulation keep improving.

03

Was the full cost counted?

Data loading, repeated shots, error mitigation and post-processing all consume time, energy and money.

04

Did the quantum component improve anything?

Cost, time, energy or accuracy. If the improvement cannot be named in one of those four units, there is no advantage claim to evaluate.

05

Can the result be reproduced?

Independent replication is the difference between a result and an announcement.

06

Is the vendor financeable?

A promising modality attached to a company that cannot fund the next generation is still a platform risk. Examine cash, customer concentration and government dependence.

SECTION 09 · TIMELINE

A Timeline Built From Evidence, Not Appointments

There will be no single morning when the world wakes up in the quantum era. Different capabilities will mature on different schedules.

Between now and the end of the decade, the highest-confidence developments are not universal quantum transformation. They are PQC migration, better logical qubits, tighter GPU-QPU integration, AI-assisted calibration, improved decoders, cloud access, government-backed manufacturing and consolidation among vendors. Hardware companies will attempt to hit fault-tolerance roadmaps while enterprises concentrate spending on algorithms and integration.

The 2029–2030 period contains the industry's boldest appointments: IBM Starling, Microsoft's revised scalability target, Quantinuum Apollo and aggressive IonQ scaling goals. Some will slip. Some will meet a technical definition that remains short of commercial usefulness. One or more may deliver a machine that performs valuable scientific calculations unavailable elsewhere. The evidence will matter more than the launch event.

DARPA's 2033 utility-scale horizon is a useful outer checkpoint because it asks the economic question directly. A system should create computational value exceeding its cost. By then, success should be visible through logical-qubit growth, lower logical error rates as systems scale, deeper fault-tolerant circuits, working interconnects, credible resource estimates and independently validated applications.

What would move the date in either direction

The timeline accelerates if error-correction overhead falls, manufacturing becomes repeatable, modular networking works and algorithms demonstrate large advantages on valuable problems. It slows if physical-error improvements plateau, classical methods keep erasing benchmark advantages, control infrastructure fails to scale or the industry discovers that hundreds of logical qubits are useful for fewer problems than expected.

Beyond the early 2030s, forecasts become scenarios. McKinsey estimates an internal quantum-computing market of \$43 billion to \$71 billion by 2035 and potential economic value across industries of \$1.3 trillion to \$2.7 trillion. BCG estimates long-term value creation of \$450 billion to \$850 billion. These are not revenue guarantees. They are models whose largest outputs depend on fault-tolerant hardware and applications that do not yet exist.

Three data points, two houses, one warning label

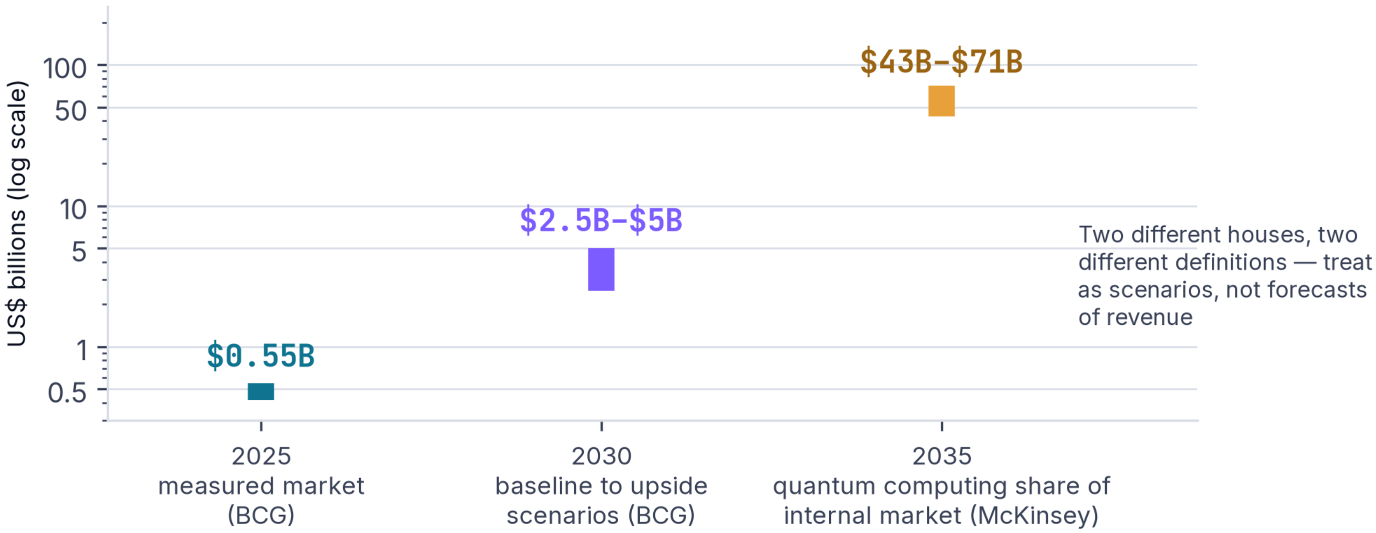

Sources: BCG (2026); McKinsey Quantum Technology Monitor 2026. 2030 and 2035 figures are scenarios.

FIGURE 8 · SCENARIOS, NOT FORECASTS

The 2025 figure is measured. The 2030 and 2035 ranges are models built on different definitions, and their largest outputs depend on hardware and applications that do not yet exist.

Sources: BCG (2026) — \$550M in 2025, \$2.5B baseline and \$5B upside by 2030; McKinsey *Quantum Technology Monitor 2026* — \$43B–\$71B quantum computing share of a \$60B–\$100B internal quantum market by 2035.

HIGHEST-CONFIDENCE DEVELOPMENTS THROUGH 2030

PQC migration · better logical qubits · tighter GPU-QPU integration · AI-assisted calibration · improved decoders · broader cloud access · government-backed manufacturing · vendor consolidation.

Not on that list: universal quantum transformation of enterprise computing.

SECTION 10 · THE PLAYBOOK

What Enterprises Should Do Now

The sensible strategy lies between building a quantum laboratory in every company and dismissing the field until a perfect machine appears.

Every organization with long-lived sensitive information should begin cryptographic discovery and PQC planning now. That is not speculative quantum investment; it is cybersecurity lifecycle management based on finalized standards and known migration time.

Companies in finance, insurance, chemistry, pharmaceuticals, materials, energy, logistics, aerospace and advanced manufacturing have a stronger case for targeted experimentation. They should begin with business constraints rather than quantum demonstrations. Define the problem and the best classical baseline. Determine what improvement would be economically meaningful. Then test whether a quantum or quantum-inspired method offers a credible path to it.

A pilot should disclose the hardware, modality, number and quality of qubits, circuit depth, mitigation, number of shots, classical computation and total elapsed time. Success criteria should include cost, accuracy, energy and repeatability. If the work runs on a simulator, say so. If the result is useful only when classical preprocessing removes most of the problem, include that cost. If a vendor compares itself with a weak classical method, rerun the comparison.

> Quantum readiness should create options, not blank checks.

Most companies do not need to own a QPU. Cloud access reduces commitment and lets teams compare modalities. Partnerships with universities, national laboratories and vendors can build literacy without forcing an early architectural bet. Enterprises should retain ownership of their problem formulation, data, evaluation method and intellectual property rather than allowing a hardware provider to define success.

Vendor viability belongs in the technical assessment. Quantum development consumes capital long before broad product revenue arrives. Buyers should examine cash, customer concentration, government dependence, roadmap history, manufacturing access and acquisition strategy. A technically promising modality attached to a company that cannot finance the next generation is still a platform risk.

Executives should also resist the fear of being "late like we were to AI" as a substitute for a business case. The lesson of AI is not that every organization should fund every emerging technology early. It is that organizations need the data, infrastructure, talent and governance to move when capability becomes useful. Quantum readiness should create options, not blank checks.

The disclosure standard: what every quantum pilot report should state

Hardware

Machine, modality and access path

Which processor, which modality, and whether the run was on hardware, on a simulator, or a hybrid of both.

Qubits

Number and quality

Count is not enough: fidelity, connectivity and coherence determine what the run could actually do.

Circuit

Depth, shots and mitigation

Circuit depth, number of repeated shots, and every error-mitigation technique applied to the result.

Classical

All surrounding computation

Preprocessing, decoding, post-processing and orchestration — including the cost of classical work that simplified the problem.

Baseline

The best current classical method

Not the most convenient comparison. If the vendor supplied the baseline, rerun it.

Outcome

Cost, accuracy, energy, repeatability

Total elapsed time and total cost, stated against the economically meaningful threshold defined before the pilot started.

CLOSING · ARRIVALS AND DEPARTURES

Arriving, Departing and Still Under Construction

Quantum computing is no longer only a physics experiment. It is also not yet what decades of headlines promised.

Quantum computing is no longer only a physics experiment. It is a commercial market, a national industrial policy, an investment category, a cybersecurity deadline and an emerging part of the heterogeneous computing stack.

It is also not yet what decades of headlines promised.

The industry has demonstrated below-threshold error correction, longer-lived logical information, higher-fidelity systems, real-time decoding and AI-assisted calibration. Companies are spending hundreds of millions of dollars on access, algorithms and readiness. Governments are investing billions in foundries and engineering. Revenue is growing. Those are arrivals.

> The future is not a quantum computer replacing the data center. It is a data center that knows when to call one.
>
> ALAN SHIMEL

Broad, repeatable quantum advantage on production problems remains ahead. Fault-tolerant roadmaps are still roadmaps. Competing modalities have not converged. Error-correction overhead remains formidable. Useful algorithms lag hardware ambition. Classical computing does not wait politely while quantum catches up. Those are departures pushed farther down the line.

AI makes the picture more complicated and more hopeful. It is not quantum's competitor. It is becoming part of the control system, decoder, simulator, design environment and scientific process required to build quantum machines. If quantum eventually accelerates AI, it will almost certainly do so as one specialized component inside an overwhelmingly classical workflow.

The future is not a quantum computer replacing the data center. It is a data center that knows when to call one.

That leaves enterprises with a more useful question than “Has quantum arrived?” Ask which part has arrived for you. For security teams, the answer is migration planning now. For scientists and organizations with unusually valuable simulation or optimization problems, it may be disciplined experimentation. For most others, it is building literacy, monitoring logical performance and preserving the option to move.

The quantum train is in the station, but not every carriage is going to the same place—and several are still being built while passengers board. The organizations most likely to benefit will not be those that declare quantum either inevitable or impossible. They will be the ones that know what evidence they need before they call it useful.

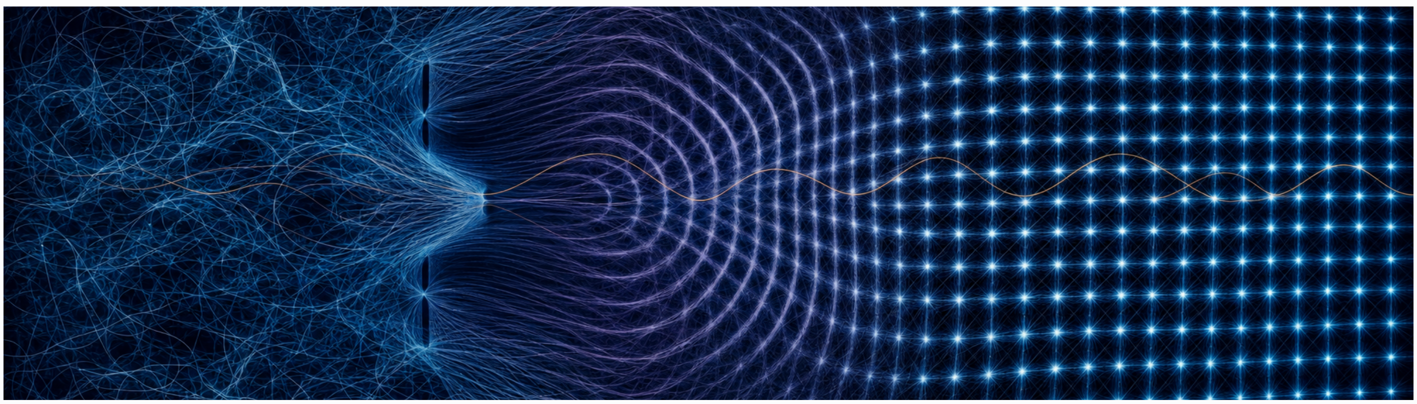

CODA · ORIGINAL ILLUSTRATION

Interference resolving into order: the field on the left is noise, the lattice on the right is structure. Quantum engineering is the work of moving from one to the other — and in August 2026 most of the industry is still somewhere in the middle.

ASK WHICH PART HAS ARRIVED FOR YOU

Security teams

The answer is migration planning now — cryptographic inventory, data-retention analysis and the NIST PQC standards.

Valuable simulation or optimization problems

For scientists and those organizations, it may be disciplined experimentation against a named classical baseline.

Most others

Building literacy, monitoring logical performance and preserving the option to move.

FROM “ARRIVING, DEPARTING AND STILL UNDER CONSTRUCTION” · ALAN SHIMEL

Alan Shimel

FOUNDER & CEO, TECHSTRONG GROUP
BOCA RATON, FLORIDA · AUGUST 2026

REFERENCE

Sources and Notes

Every quantitative claim in this report is attributed below, numbered in the order in which it appears in the text. Numbers marked as scenarios or models are projections, not measured results; numbers marked as company claims await independent replication. Research and analysis attributed to **Futurum Research** — including the QPU market landscape work of Dr. Bob Sutor, *AI + Quantum: Hype or True Breakthroughs?*, commentary from research director Brendan Burke, and Daniel Newman's *Quantum Fine-Tuning and the Energy Case for Quantum in AI* — is used with attribution at the point of use.

1. Isabelle Bousquette, "Businesses Are Spending Big on Quantum," *The Wall Street Journal* (CIO Journal). https://www.wsj.com/cio-journal/businesses-are-spending-big-on-quantum-c73060ad
2. DARPA, "Quantum Benchmarking Initiative (QBI)." https://www.darpa.mil/research/programs/quantum-benchmarking-initiative
3. Boston Consulting Group, "Quantum Is Getting Real. CEOs Need to Shape Where It Creates Value" (2026). https://www.bcg.com/publications/2026/ceos-need-to-shape-where-quantum-creates-value
4. McKinsey & Company, *Quantum Technology Monitor 2026: A Commercial Tipping Point*. https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/mckinsey-quantum-technology-monitor-2026-a-commercial-tipping-point
5. Google Quantum AI and collaborators, "Quantum error correction below the surface code threshold," *Nature* (2024). https://www.nature.com/articles/s41586-024-08449-y
6. The Futurum Group, "Is Quantum Computing Languishing? — Report Summary," analysis by Dr. Bob Sutor (December 2024). https://futurumgroup.com/press-release/is-quantum-computing-languishing/
7. Sutor's own later landscape work illustrates the point: a subsequent quarterly QPU market landscape from Sutor Group Intelligence and Advisory tracked 88 companies across 21 countries. https://drbobsutor.substack.com/p/sutor-group-quantum-processing-unit-449
8. Reuters, "IonQ to buy SkyWater for $1.8 billion to expand hardware capabilities" (January 26, 2026) https://www.reuters.com/technology/ionq-buy-skywater-18-billion-expand-hardware-capabilities-2026-01-26/; IonQ, "IonQ Completes Acquisition of SkyWater Technology" (July 31, 2026) https://www.ionq.com/news/ionq-completes-acquisition-of-skywater-technology
9. U.S. Department of Commerce / NIST, "Department of Commerce Announces Letters of Intent With 9 Companies for $2 Billion to Accelerate U.S. Leadership in Quantum Computing" (May 2026). https://www.nist.gov/news-events/news/2026/05/department-commerce-announces-letters-intent-9-companies-2-billion
10. IBM, "IBM lays out clear path to fault-tolerant quantum computing" and the IBM Quantum roadmap. https://www.ibm.com/quantum/blog/large-scale-ftqc
11. Quantinuum, "Quantinuum Unveils Accelerated Roadmap to Achieve Universal Fault-Tolerant Quantum Computing by 2030." https://www.quantinuum.com/press-releases/quantinuum-unveils-accelerated-roadmap-to-achieve-universal-fault-tolerant-quantum-computing-by-2030
12. IonQ, "Roadmap." https://www.ionq.com/roadmap
13. Microsoft, "Majorana 2, made more reliable with Microsoft Discovery agentic AI" (June 2026). https://news.microsoft.com/source/features/innovation/majorana-2-microsoft-discovery-agentic-ai/
14. J. Bausch et al., "Learning high-accuracy error decoding for quantum processors," *Nature* (2024). https://www.nature.com/articles/s41586-024-08148-8; Google DeepMind, "AlphaQubit." https://blog.google/innovation-and-ai/models-and-research/google-deepmind/alphaqubit-quantum-error-correction/
15. Google Quantum AI and Google DeepMind, "Towards a quantum computer that learns from its errors," *Nature* (2026). https://research.google/blog/towards-a-quantum-computer-that-learns-from-its-errors/
16. R. Zen, J. Olle, L. Colmenarez, M. Puviani, M. Müller and F. Marquardt, "Quantum Circuit Discovery for Fault-Tolerant Logical State Preparation with Reinforcement Learning," *Physical Review Research* / APS. https://link.aps.org/doi/10.1103/gqpr-dgz7
17. AWS Quantum Computing Blog, "Accelerating the Quantum Toolkit for Python (QuTiP) with cuQuantum on AWS." https://aws.amazon.com/blogs/quantum-computing/accelerating-the-quantum-toolkit-for-python-qutip-with-cuquantum-on-aws/
18. NVIDIA, "World's Leading Scientific Supercomputing Centers Adopt NVIDIA NVQLink..." — Quantinuum's Helios processor demonstrating scalable real-time decoding for quantum error correction. https://investor.nvidia.com/news/press-release-details/2025/Worlds-Leading-Scientific-Supercomputing-Centers-Adopt-NVIDIA-NVQLink-to-Integrate-Grace-Blackwell-Platform-With-Quantum-Processors/
19. O. Knitter et al., "Measuring Accuracy and Energy-to-Solution of Quantum Fine-Tuning of Foundational AI Models," arXiv preprint (May 2026). https://arxiv.org/html/2605.02798
20. Daniel Newman, "Quantum Fine-Tuning and the Energy Case for Quantum in AI," Futurum Research (June 25, 2026). https://futurumgroup.com/insights/quantum-fine-tuning-and-the-energy-case-for-quantum-in-ai/
21. NIST, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards" (August 13, 2024). https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards
22. NIST IR 8547, *Transition to Post-Quantum Cryptography Standards* (initial public draft). https://nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8547.ipd.pdf
23. Executive Order 14412, "Securing the Nation Against Advanced Cryptographic Attacks" (June 2026) https://www.whitehouse.gov/presidential-actions/2026/06/securing-the-nation-against-advanced-cryptographic-attacks/; OMB Memorandum M-26-15, "Execution of the Migration to Post-Quantum Cryptography" (June 24, 2026) https://www.whitehouse.gov/wp-content/uploads/2026/06/M-26-15-Execution-of-the-Migration-to-Post-Quantum-Cryptography.pdf

24. IonQ, "IonQ Announces Record Second Quarter 2026 Revenues, Growing 287% YoY" (August 2026). https://investors.ionq.com/news/news-details/2026/IonQ-Announces-Record-Second-Quarter-2026-Revenues-Growing-287-YoY/default.aspx

ABOUT THIS PUBLICATION

TECHSTRONG GROUP

Techstrong Group is a technology media and research organization covering DevOps, cybersecurity, cloud-native engineering and AI for practitioners and executives. Techstrong Special Reports examine markets where engineering reality and commercial narrative have diverged.

FUTURUM RESEARCH

Research and analysis from Futurum Research appears throughout this report, including the QPU market landscape work of Dr. Bob Sutor, *AI + Quantum: Hype or True Breakthroughs?*, quantum manufacturing commentary from research director Brendan Burke, and Daniel Newman's *Quantum Fine-Tuning and the Energy Case for Quantum in AI*. Futurum figures are attributed at the point of use.

STANDARD OF EVIDENCE

Peer-reviewed results are identified as such. Vendor roadmaps are identified as engineering targets. Analyst projections are identified as models. Company claims awaiting independent replication are identified as company claims.

CORRECTIONS

Figures in this report reflect the sources listed above as of August 2026. Quantum hardware counts, vendor roadmaps and revenue figures change quickly; readers should confirm current values against the primary sources linked here.

GLOSSARY: THE WORDS THIS REPORT USES PRECISELY

QUBIT

A quantum bit. Unlike a classical bit forced into a stable zero or one, it can be prepared in a combination of states and manipulated through interference and entanglement — which is also what makes it fragile.

PHYSICAL VS. LOGICAL QUBIT

A physical qubit is one piece of hardware. A logical qubit is an encoded qubit protected by an error-correcting code spread across many physical qubits, with repeated error detection and classical decoding.

ERROR MITIGATION

Improving results from a noisy processor without fully protecting the computation.

ERROR CORRECTION

Preserving encoded quantum information as operations continue — a far more demanding requirement than mitigation.

BELOW THRESHOLD

Physical errors are controlled well enough that increasing the size of the error-correcting code reduces the logical error rate instead of increasing it.

CODE DISTANCE

A measure of an error-correcting code's size and protective strength. Larger distance means stronger protection and higher physical-qubit overhead.

SYNDROME AND DECODER

Repeated measurements that report an error occurred without revealing the encoded state, and the classical computation that infers which error most likely happened.

FAULT TOLERANCE

The ability to execute deep, universal circuits at the fidelity an application requires — not merely to hold a protected memory.

UTILITY SCALE

DARPA's economic test: computational value exceeds cost, measured against the best classical alternative rather than the most convenient comparison.

QUANTUM ANNEALING

A non-universal approach aimed at optimization-style problems. Commercially accessible, but it is not a gate-model quantum computer and belongs in its own category.

POST-QUANTUM CRYPTOGRAPHY (PQC)

Classical algorithms designed to resist both classical and quantum attack, standardized by NIST in FIPS 203, 204 and 205. PQC runs on existing computing and networking systems.

CRYPTO-AGILITY

Knowing where cryptography lives in an estate and being able to replace algorithms without replacing the systems around them.

QKD AND QRNG

Quantum key distribution detects interception but requires specialized infrastructure; quantum random-number generation already has commercial applications. Neither is PQC.

QUANTUM-CENTRIC SUPERCOMPUTING

IBM's term for an architecture in which a QPU handles the quantum-suited portion of a workload while CPUs, GPUs and HPC systems handle everything around it.

Techstrong
a Futurum company

SPECIAL REPORT

“The future is not a quantum computer replacing the data center. It is a data center that knows when to call one.”

ALAN SHIMEL · FOUNDER & CEO, TECHSTRONG GROUP

ALL ABOARD THE QUANTUM TRAIN: IT MAY BE ARRIVING AND LEAVING AT THE SAME TIME

A TECHSTRONG SPECIAL REPORT · AUGUST 2026 · BY ALAN SHIMEL

RESEARCH DATA FROM FUTURUM RESEARCH