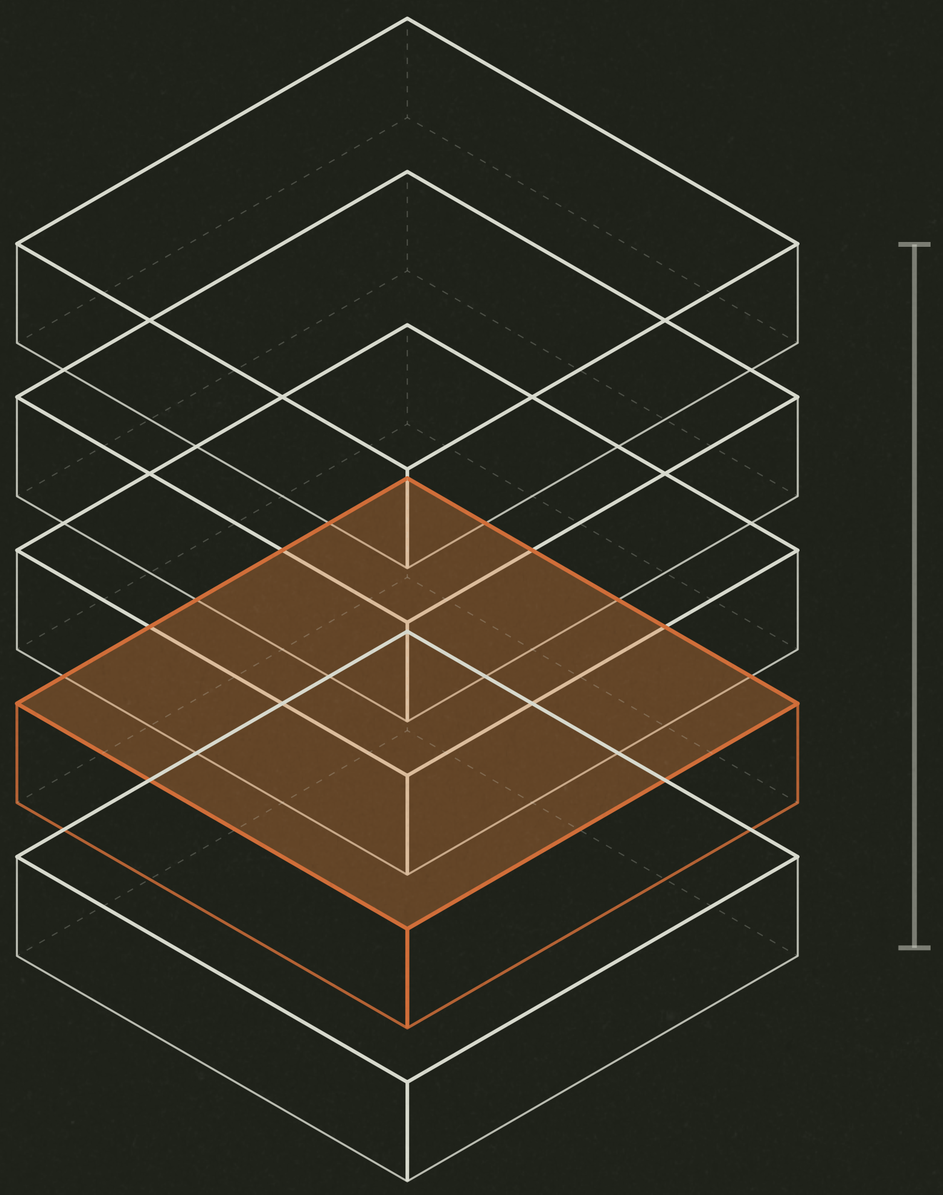
Elevate

Quantum Computing: The Market Is Mispricing the Stack



INVESTMENT THESIS

The Quantum Report

Qubit counts are the wrong scoreboard. Circuit depth, gate fidelity and error correction are what gate commercial value — and what the market is mispricing across the stack.

2026

Contents

A Note from the Editor 01

01 The Scoreboard Is Wrong 02

02 Quantum in Plain English 04

03 The Three Algorithms That Tell You Where the Money Is 14

04 Why You Should Care 22

05 The Investment Framework 33

06 Count Nines, Not Qubits 36

07 Every Revolution Has Two Numbers 38

References 40

A Note from the Editor

Quantum computing is an absolute beast of a topic. In first principles, quantum computing is built on the laws of quantum mechanics. The latter is defined as,

> “The basic physical theory that explains how matter and light behave at and below the scale of atoms.”

Enough said.

Two of our primary goals for the research desk are to entertain and to educate anyone who wants to invest their precious time. To accomplish this, we strive to articulate the subject matter in a way anyone can understand. To be entertained and educated on this one, it’s going to require some gray matter activation from you, the reader. However, we promise that if you do, you won’t be let down.

Following our research, we see quantum computing as a key technology with the potential to significantly change our lives in the not-so-distant future. As such, we think it’s worth the investment of brain power to dive in.

Brace yourself for quite the atomic performance (you’ll have to read to understand the reference). Enjoy!

the elevate research desk

01

Quantum Computing: The Market Is Mispricing the Stack

The Scoreboard Is Wrong

Every quantum investment thesis ultimately reduces to two questions: how big is the impact when each algorithm becomes viable — and what will make it viable first?

History suggests the biggest returns rarely come from the platform alone. They come from the workloads, markets, and ecosystems the platform unlocks. Railroads mattered because they moved commerce. The internet mattered because it enabled new applications. Smartphones mattered because they spawned a software ecosystem nobody priced in. Quantum computing will be no different.

The physics is fascinating. Quantum mechanics is the mathematics of nature itself, governing everything from battery chemistry to drug binding to superconductors. But capital doesn't flow to the fascinating. It flows to workloads, that is, the specific problems quantum systems will solve better than classical computers, and the companies positioned where those problems get solved.

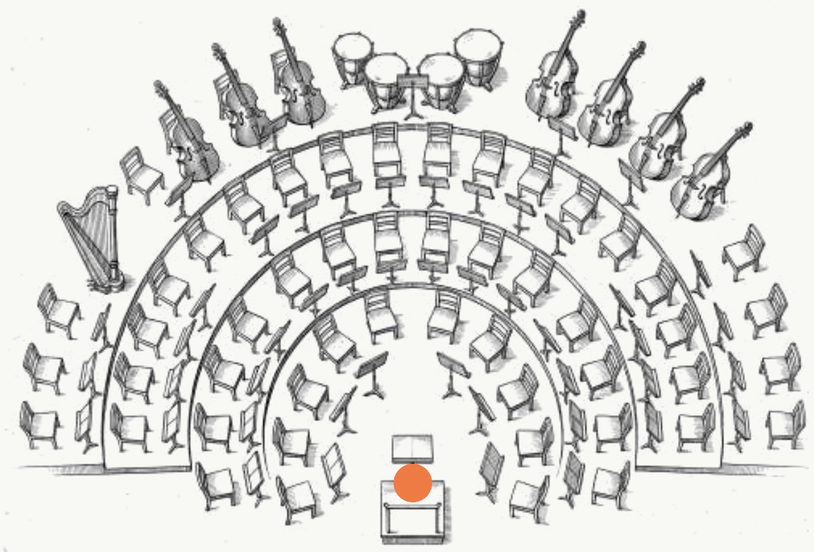
Today, billions of dollars are moving into quantum hardware, cloud infrastructure, cybersecurity, pharmaceuticals, advanced manufacturing, and national defense [12]. Nearly all of it is being allocated on a common assumption: the winners will be whoever builds the biggest quantum computer. Qubit counts are the industry's default scoreboard. They headline most every press release and anchor every valuation conversation.

We believe that the scoreboard is wrong, and that the market is mispricing the entire stack because of it.

Think of a quantum computer as an orchestra. The number of musicians on stage tells you almost nothing about whether the performance will be any good. What matters is whether they can play in tune, stay coordinated, and get through an entire score without the performance falling apart. Quantum computing is the same. The metrics that actually gate commercial value are circuit depth, gate fidelity, and error correction. In other words, how long the orchestra can play, how often each note is played correctly, and what it costs to guarantee the performance survives its own mistakes. These numbers, not qubit counts, determine whether a theoretically valuable algorithm can run in the real world, and whether a company becomes a trillion-dollar platform or an expensive science project.

They also determine which applications reach commercial relevance first. Understanding those algorithms is ultimately more important than understanding qubit counts, because they define who gets paid, when they get paid, and where value accumulates across the stack.

The rest of this piece follows that logic: from first principles, to algorithms, to commercialization, and ultimately to the investment opportunities each stage creates.



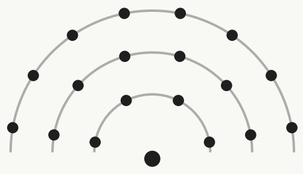
The book's master analogy arrives: an orchestra judged not by its size, but by how it plays together.

Quantum in Plain English

Before we get to the investment argument, we owe you enough intuition to follow it. Four concepts do almost all the work: **superposition, entanglement, interference, and measurement.** Decoherence, the reason all of this is so hard to build, falls out of the fourth concept. If you understand these five ideas, you can follow any serious conversation about quantum computing, read an IBM or Quantinuum roadmap without marketing gloss, and tell the difference between a real milestone and a press release. If you don’t, every algorithm and every hardware pitch will sound like the same paragraph with different nouns.



02



We said in the opening that a quantum computer is best understood as an orchestra. Here is what each player actually does.

Superposition: The Scale Unlock

A classical bit is a light switch. It is either 0 or 1. Every computer you have ever used, from your laptop to your phone to the servers that deliver websites, is made of billions of these switches flipping in sequence, billions of times a second. A classical musician, in other words, plays exactly one note at a time.

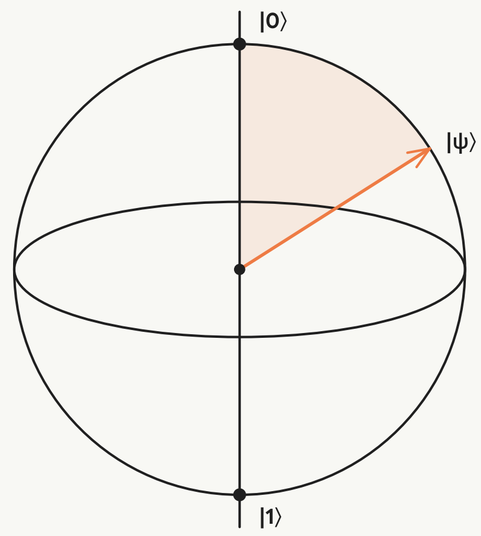
A qubit is like a spinning coin. Technically, it reads as a vector, but the coin analogy will do. While the coin spins, it is meaningfully in both states at once. It is not an unknown 0 or 1 waiting to be discovered. It is in a mathematically real superposition of both, described by complex-number amplitudes. That distinction matters. Amplitudes can be negative or complex. Probabilities cannot. This is one of the key differences between classical and quantum computing. A classical bit may be hidden from view, but it still has a definite value. A qubit is different. Before measurement, it behaves as a combination of possibilities, and only when measured does it resolve into a single answer. A qubit can be built from many physical systems: the spin of an electron, the polarization of a photon, two energy levels of a trapped ion, or the current direction in a superconducting loop. The hardware changes. The math does not.

Superposition is what puts every instrument on stage: each qubit holds its full range of notes at once, resolving to a single note only when the audience listens.

The payoff is scale. n qubits in superposition represent 2^n states simultaneously. Two qubits represent four states at once; ten represent 1,024; fifty represent more states than there are stars in the Milky Way; three hundred represent more states than there are atoms in the observable universe.

One critical caveat, and it is the sentence most popular coverage skips, is that **an exponentially large state space does not automatically give you an exponential speedup.** The speedup is algorithm-specific, and for most problems it doesn’t exist at all. A quantum computer is not a magic box that can “try every solution in parallel and return the best one.” Some problems will remain out of reach no matter how good the hardware gets.

FIG. 1 — SUPERPOSITION



One qubit holds both states at once — a point on the Bloch sphere, not a switch.

So superposition, by itself, computes nothing. It is raw material: an enormous space of possibilities waiting to be shaped. The next two ideas are what shape it.



Entanglement: The Coordination Layer

If superposition is “one qubit in two states at once,” entanglement is “two or more qubits sharing one state.”

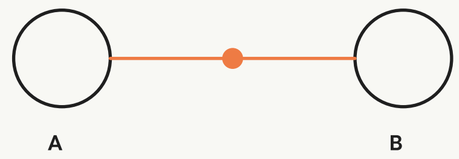
When two qubits are entangled, measuring one instantly determines the outcome of measuring the other, even if they are separated by a continent. A useful way to build intuition is to imagine a pair of quantum socks. Until you look, neither is “left” or “right.” The moment you put one on your left foot, the other one, wherever it is, becomes “right.” They didn’t decide in advance; they decided together, at the moment of measurement, and they must agree. (One clarification sophisticated readers will look for: because each measurement outcome is individually random, entanglement cannot be used to transmit information faster than light. The correlations are instantaneous; communication is not.)

Einstein called it “spooky action at a distance”, and hated it, because he thought the universe couldn’t really work this way, but it does. Bell’s 1964 inequality formalized why these correlations cannot be explained by any classical, local theory [3]. Decades of experiments have confirmed it, culminating in the 2022 Nobel Prize to Aspect, Clauser, and Zeilinger. **The maximally entangled two-qubit states are called Bell states.**

The practical point for investors: entanglement is what turns n soloists into an orchestra. It is the resource that lets an n-qubit machine behave as one coordinated system rather than n independent coins. The strings and brass have not chosen their parts in advance, but once the system is measured, their outcomes are correlated in ways no classical section could fully pre-script. That coordination is not decorative; it is central to why quantum computers can do anything distinctive at all. We will return to why that matters commercially when we meet the major algorithm families later in the piece.

But coordination alone is not enough. To understand how quantum computers move from many possible states toward a useful result, we now turn to quantum interference.

FIG. 2 — ENTANGLEMENT



Two qubits, one shared state — measure one and the other is decided at once.



> “Einstein called it **“spooky action at a distance”**, and hated it, because he thought the universe couldn’t really work this way, but it does.”



Interference: The Compute Engine

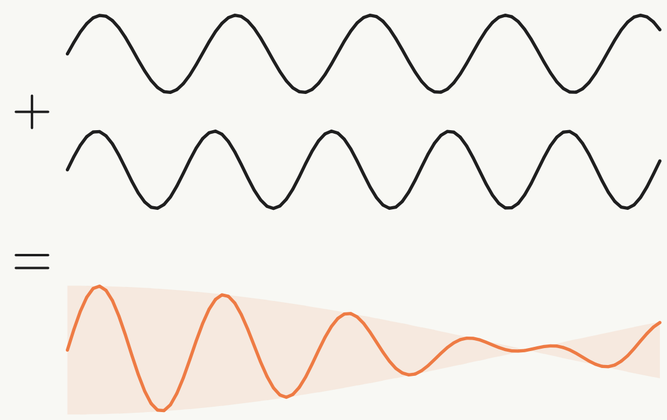
Superposition puts every instrument on stage. Entanglement lets them play in coordination. **Interference is what decides whether the result is music or noise**, and it is what actually produces useful answers.

Every quantum algorithm works by arranging amplitudes so that wrong answers cancel out (destructive interference) and the right answer reinforces itself (constructive interference), so that when you finally measure, the right answer is overwhelmingly likely. In orchestral terms: the wrong answers are dissonant voices that cancel each other out and the right answer is the chord that swells above them. Noise-cancelling headphones use the same principle in the classical-wave regime. A sound wave meets its inverted copy and the result is silence.

This is the single most important conceptual point in quantum computing, and most popular explanations skip it. Quantum algorithms aren't "faster because they try everything at once." **They're faster because they orchestrate interference so the wrong answers erase themselves.** This concept matters enough that we wanted to introduce a different analogy, especially for those that are visual thinkers. Picture a maze flooded with water: superposition and entanglement explore every path at once, then interference drains the dead ends, leaving water pooled along the one path to the exit.

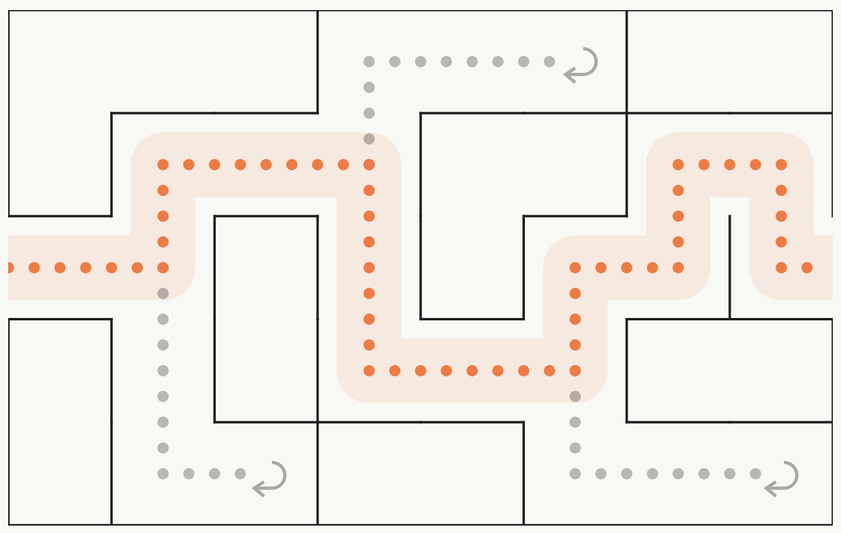
Without interference, superposition is just noise. A quantum computer without interference is a very expensive random-number generator. Which raises the obvious next question: once interference has done its work and the right answer is sitting in the machine, how do you actually get it out? You measure it, but measurement, it turns out, is both the payoff and the peril.

FIG. 3 — INTERFERENCE

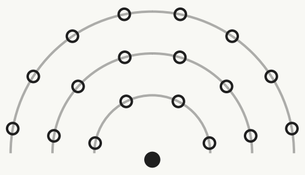


Waves add where they align and cancel where they oppose — the machine amplifies the right answer.

FIG. 4 — SEARCH BY INTERFERENCE



Every path is explored at once; interference drains the dead ends — the answer is wherever the water still pools.



Measurement: The Output Layer

Everything we have described so far takes place in a realm you can never observe directly. Measurement is the act of extracting a classical answer from a quantum system. When you measure a qubit, you do not see its superposition. Instead, you get one ordinary classical bit: either 0 or 1, with probabilities determined by the qubit's amplitudes.

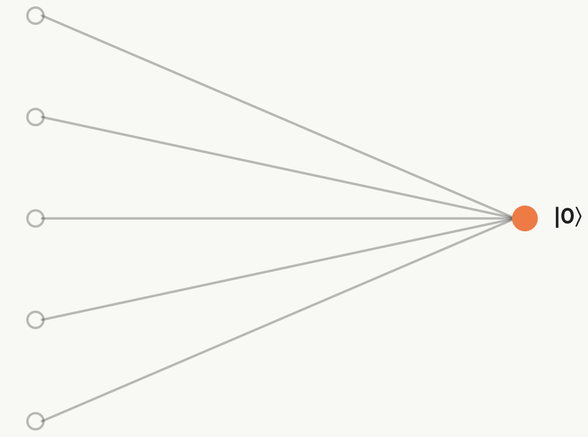
Even if you measure 300 qubits whose combined quantum state represents more possible configurations than there are atoms in the universe, the result you obtain is still just 300 ordinary bits. The exponential richness exists only while the system remains quantum. Once you measure it, that richness collapses into a classical outcome.

This is why measurement matters commercially and not just philosophically. It is the reason algorithm design is so important. The entire craft of quantum computing is engineering the interference before the measurement, so that when the collapse happens, the bits you read out are overwhelmingly likely to be the answer you wanted. Using the orchestra analogy again, you can think of measurement as the box office. However spectacular the performance, value is realized only in that final transaction.

Here is the catch, and it is one reason quantum hardware is so hard to build is that the environment does not wait for you to measure. A stray photon, a vibration, a magnetic-field fluctuation, or a tiny temperature shift can leak information about the system into the outside world. When that happens, **decoherence** kicks in and the quantum waves fall out of sync, destroying the interference needed for computation. The spinning coin from earlier has not been cleanly read out at the right moment; it has been jostled into classical-looking behavior before the computation is finished. The performance falls apart mid-phrase.

Decoherence is measurement's evil twin: not the clean readout you wanted, but an uncontrolled leak of quantum information at exactly the wrong moment.

FIG. 5 — MEASUREMENT



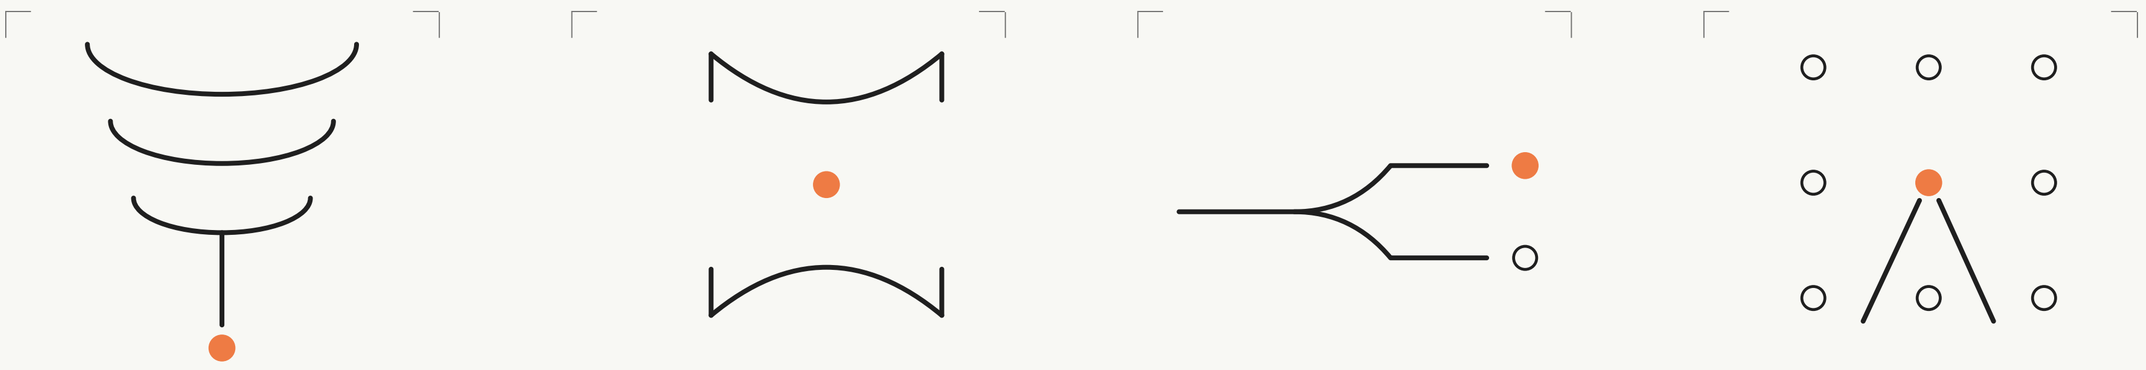
Every possibility is present until you measure — then exactly one bit comes out.

Measurement: The Output Layer

This is why IBM's superconducting qubits run colder than deep space, at roughly 15 millikelvin (about −460°F) inside dilution refrigerators the size of a vending machine. It's why Quantinuum and IonQ suspend individual ytterbium and barium ions in ultra-high vacuums, manipulated by lasers timed to the nanosecond. It's why PsiQuantum is betting on photons routed through custom silicon photonic chips, and why QuEra uses arrays of neutral rubidium atoms held in place by optical tweezers. Each modality is a different way of building a concert hall silent enough for the orchestra to hear itself. In other words, each is a different answer to the same engineering problem: ***how do you keep a quantum state alive long enough to compute with it, and then measure it cleanly when you're done?***

But that is only the first half of reliability. A quieter machine gives the computation a chance to survive. It does not guarantee the answer will be right. At commercial scale, errors will still appear. The question is whether the machine can detect them, correct them, and keep going.

FIG. 6 — FOUR WAYS TO BUILD A QUBIT



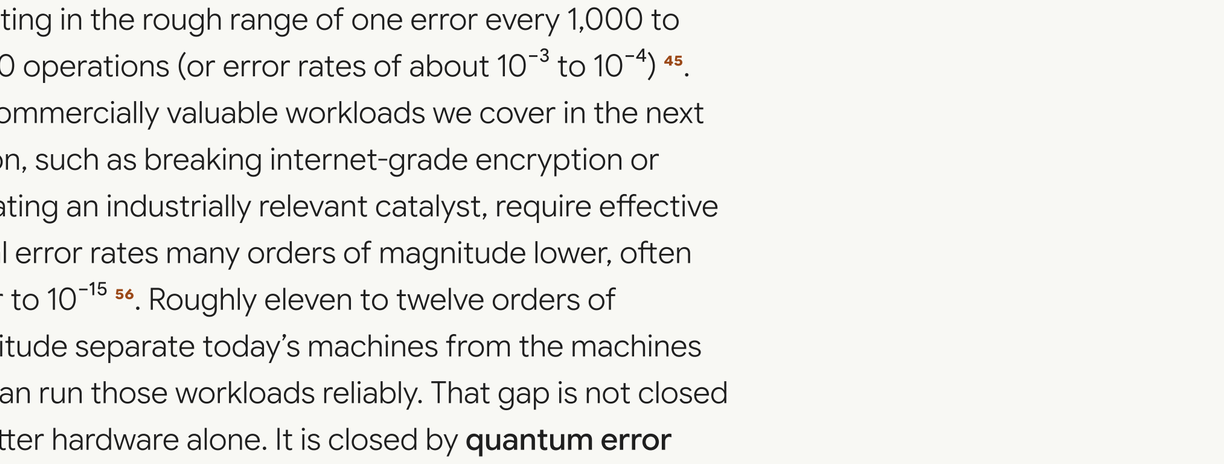
Four ways to build a hall silent enough for the orchestra.

Error Correction: The Platform Layer

That is the role of error correction. It is the layer that turns fragile physical qubits into more reliable logical qubits. And that fact, more than any other in this piece, determines when quantum computing becomes a business.

Here is why. Today's best physical qubits still make errors often enough to matter, with many leading systems operating in the rough range of one error every 1,000 to 10,000 operations (or error rates of about 10^{-3} to 10^{-4}) [45]. The commercially valuable workloads we cover in the next section, such as breaking internet-grade encryption or simulating an industrially relevant catalyst, require effective logical error rates many orders of magnitude lower, often closer to 10^{-15} [56]. Roughly eleven to twelve orders of magnitude separate today's machines from the machines that can run those workloads reliably. That gap is not closed by better hardware alone. It is closed by **quantum error correction**: encoding one reliable "logical" qubit across many unreliable physical ones, so the ensemble can detect and repair mistakes faster than they accumulate. Error correction is the main method for driving computational error probabilities dramatically lower, combining many physical qubits into a more reliable logical qubit.

FIG. 7 — MANY PLAYERS, ONE PART



Many physical players, one logical part: error correction assigns an ensemble of noisy qubits to a single reliable role.

Error Correction: The Platform Layer

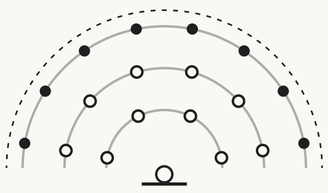
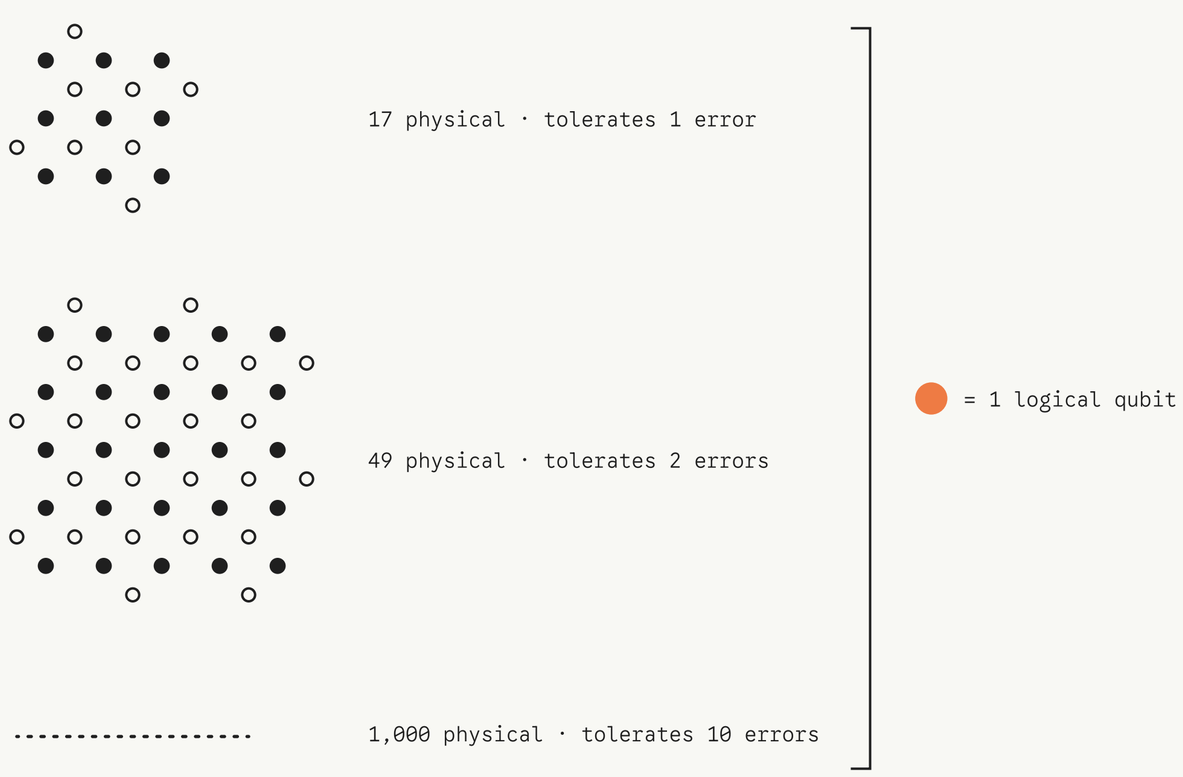


FIG. 8 — WHAT ONE LOGICAL QUBIT COSTS



Physical qubits per logical qubit, by errors tolerated. Overhead escalates fast.

In orchestral terms: assign many players to each part, so one cracked note does not ruin the piece. The catch is the head count. Current schemes demand on the order of **1,000 physical qubits for every logical qubit** [67]. A logical qubit is a collection of physical qubits grouped together to reliably produce one answer; the number of physical qubits ensures accuracy. The other catch with logical qubit functionality is **surface code distance**. Surface code distance is the safety net around logical qubits, the number of physical qubits that can fail with the logical qubit still producing the right answer. In an orchestra, this would be the number of musicians who can quietly mess up without affecting the performance. Unfortunately, this number of errors scales much more slowly than the raw number of qubits. With current architectures, a logical qubit needs at least 17 physical ones to tolerate 1 error and 49 to tolerate 2. This quickly escalates, and even a 1000-qubit logical qubit can tolerate only 10 errors [4]. This perfectly highlights why every extra 9 in gate fidelity matters (more on this later). A key investable unlock in this space will be reducing the number of physical qubits required per logical qubit, enabling larger computations and creating usable computing power.

Now the “so what.” Error correction is the bottleneck between quantum science and quantum commerce. It sits between every machine anyone builds and every algorithm anyone wants to sell. It converts hardware progress into commercial capability at a brutal exchange rate, and its pace, not raw qubit counts, sets the timeline for future cash flows in the industry. That is why we believe **error correction is the platform layer:** the layer that bridges breakthrough physics and real infrastructure, and ultimately dictates where value accrues across the stack. Every serious hardware roadmap is, underneath the marketing, a bet on a specific path across the gap between today’s physical error rates and commercially useful logical error rates. How fast that path gets crossed starts with one number we are about to meet: how reliably you can execute a gate.

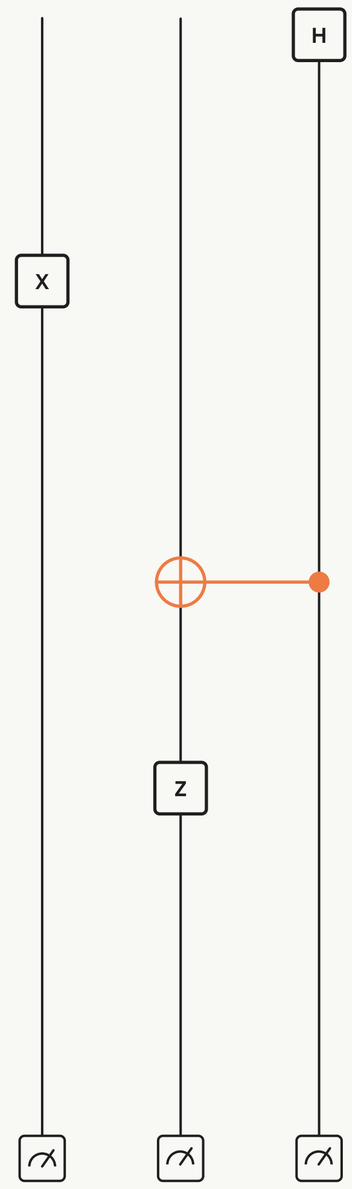
Gates & Circuits: The Performance Layer

Superposition, entanglement, and interference describe what a quantum system is. Gates describe what you do with it, and this is where the physics turns into a market. Every quantum hardware company on the planet is, underneath the physics, competing on one question: **how reliably can you execute a gate?** The companies that answer it best won't just win benchmarks. They'll determine which platforms reach commercial fault tolerance first, and on whose balance sheet that milestone lands.

We've already met the orchestra. Gates are the notes on the sheet music. Each one is a precisely controlled physical instruction (a microwave pulse on a superconducting qubit, a laser pulse on a trapped ion, a beam-splitter on a photon) that transforms the system in a specific, reversible way. Some gates direct a single player: rotating a qubit, flipping it, putting it into superposition. Others coordinate entire sections at once, creating and manipulating entanglement across qubits. **The workhorse in many circuit descriptions is a two-qubit gate called the CNOT, a standard tool for creating conditional logic and entanglement between qubits.**

A quantum circuit is the full score: a choreographed sequence of gates followed by a measurement. Every quantum algorithm, including the three families we're about to introduce, is ultimately just a different score written for the same orchestra, a recipe of gates that engineers the interference pattern needed to amplify the right answer. So when a hardware company reports a "two-qubit gate fidelity of 99.9%," decode it this way: 999 of every 1,000 notes are played correctly. That sounds impressive until you remember that commercially useful algorithms require vast numbers of notes, and every wrong one compounds. This is why gate fidelity is the metric investors should watch before they trust the qubit-count headline, and why each additional "9" is so valuable. Better gates can dramatically shrink the error-correction overhead, which means fewer backup players per part. The gap between 99.9% and 99.99% is not just an incremental improvement. It can be the difference between a science project and an actual product. That is where we think the gold is buried: not in qubit counts, which make headlines, but in fidelity curves, which make companies.

FIG. 9 — GATES & CIRCUITS

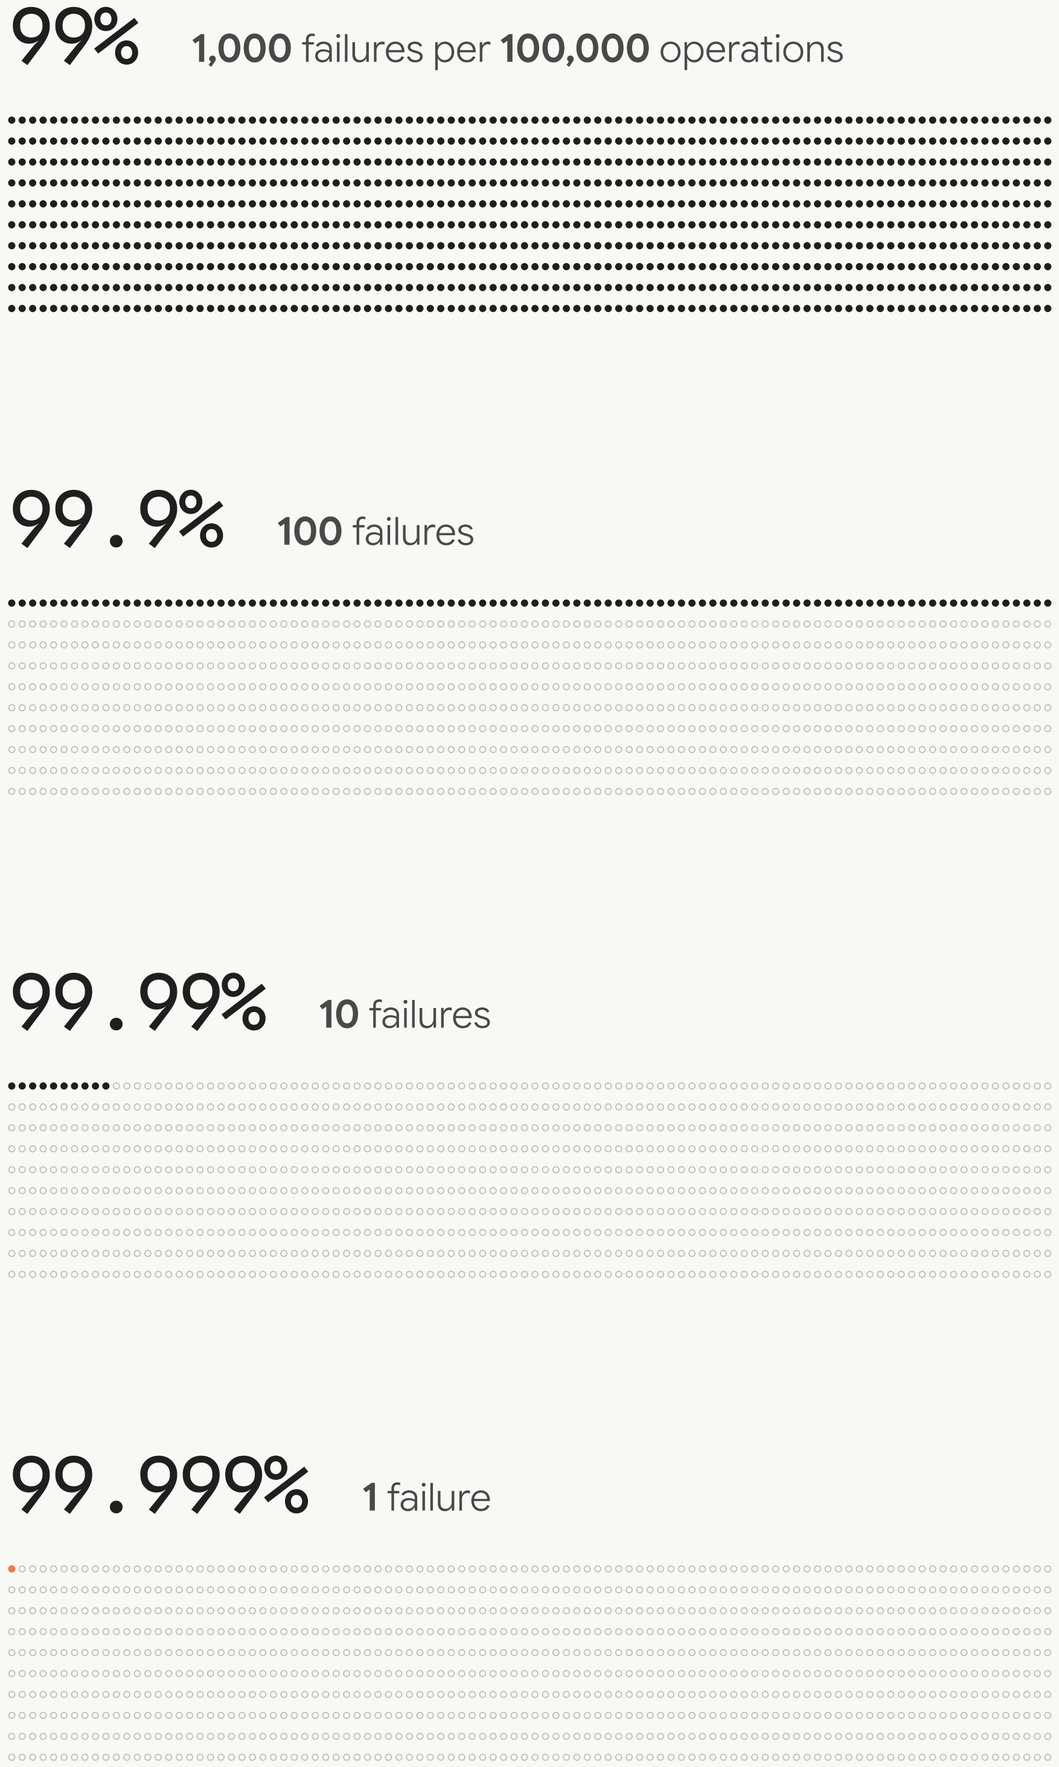


A circuit is a score, played left to right — the terracotta CNOT ties two qubits together.

So the natural question becomes: fidelity curves in service of what? A score is only worth performing if someone will pay to hear it. Three families of quantum algorithms are where the commercial value actually begins, and that is where we turn now.

Gates & Circuits: The Performance Layer

FIG. 10 — FAILURES PER 100,000 OPERATIONS



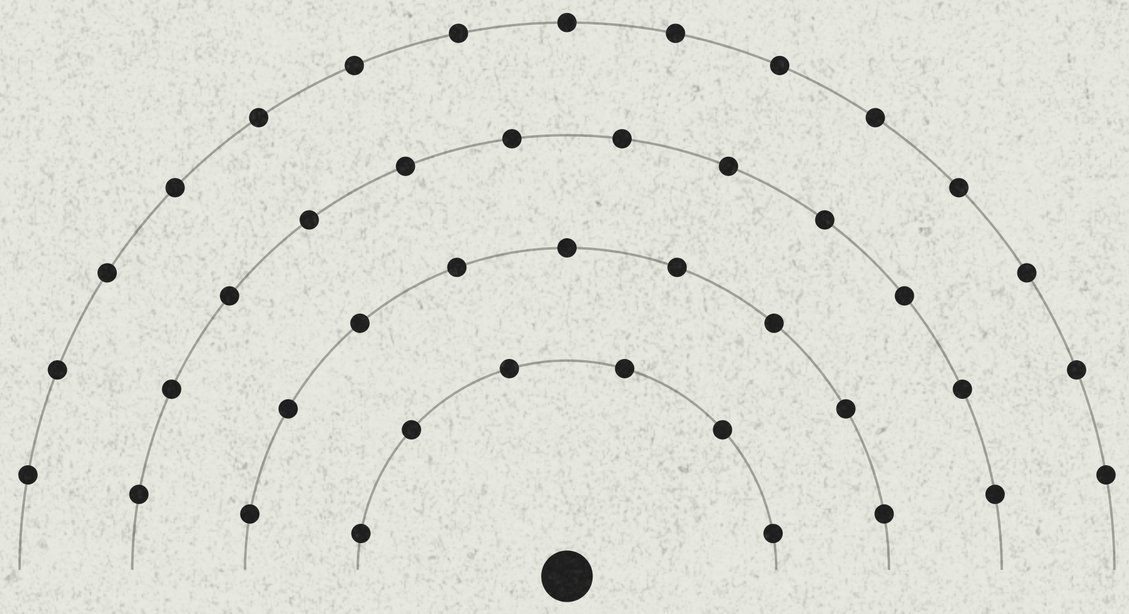
Each additional nine shrinks the error-correction overhead.

Filled: failures. Hollow: failures avoided versus the 99% baseline.

The Three Algorithms That Tell You Where the Money Is

Quantum hardware without quantum-native software is an expensive paperweight. Every algorithm in this section is, physically, a score for the same orchestra: a sequence of gates applied to qubits, the same verbs we introduced above, now doing real work.

Two numbers determine whether a score can actually be performed: **circuit depth** (how many notes the piece demands before it ends) and **gate fidelity** (how often each note is played correctly). These compound brutally. A circuit with 1,000 sequential gates at 99.9% fidelity completes correctly only about 37% of the time. At 99% fidelity, effectively never. This is the actual scoreboard the qubit-count headlines are hiding.



03

As of this writing in mid-2026, the hardware race is not a single leaderboard. It is a set of architectural trade-offs. Quantinuum currently sets the pace on gate quality. Its Helios system, launched in late 2025, runs 98 trapped-ion qubits at 99.92% two-qubit fidelity, with all-to-all connectivity that lets any qubit interact with any other, and it hosts 48 error-corrected logical qubits [11]. Google's Willow processor made a different but more fundamental point in late 2024: logical-qubit error rates can fall as surface-code distance increases. In orchestral terms, adding more backup players actually made the performance cleaner. That is the below-threshold milestone, the point where error correction starts to look less like theoretical hope and more like an engineering path.

IBM is making another bet. Its Heron processors trail Quantinuum on fidelity, but IBM leads on qubit count, iteration cadence, and roadmap clarity. Its Starling system, targeted for 2029, is designed to run 100 million gates across 200 logical qubits. QuEra and Atom Computing are pursuing a different route again, using neutral atoms held in place by focused laser beams. Both demonstrated logical-qubit-level operations in 2024 and have been closing the fidelity gap since [12].

Which architecture ultimately wins remains an open question. What is not open is the test each one must pass. A winning system will not be judged by qubit count alone. It will be judged by whether it can run deep, accurate, commercially useful algorithms. In other words, every serious contender eventually has to perform the three scores below.

FIG. 11 — THE RECORD AND THE ROADMAP

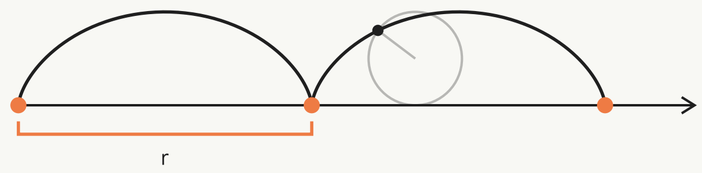


The record so far — and the roadmap ahead.

Shor’s Algorithm: The Internet’s Nightmare

Shor’s algorithm is what keeps cryptographers up at night. To see why, you need to understand one fact about modern encryption: almost all public-key cryptography on the internet, from RSA and Diffie-Hellman to elliptic curve, rests on the fact that **multiplying two large prime numbers is easy, but factoring the result back into those primes is extraordinarily hard**. When you visit a site over HTTPS, your browser and the server rely on this asymmetry to agree on a shared secret without ever transmitting it in the clear. Factoring a 2048-bit RSA number, a number more than 600 digits long, would take the best classical algorithms running on a top supercomputer billions of years. Even spread across massive parallel hardware, the task still takes millions of years [13] [14]. **That asymmetry, easy to lock and effectively impossible to unlock, is what makes the internet defensible. It is the entire security foundation of online commerce, banking, and state secrets.**

FIG. 12 — PERIOD-FINDING



The pattern repeats with period r — the quantum Fourier transform locks onto it.

Peter Shor’s 1994 breakthrough was to turn factoring into a different problem. Instead of trying to divide a large number directly, Shor’s algorithm constructs a function with a hidden repeating pattern. That pattern, called the period, is mathematically linked to the number’s factors [15]. Once you know the period, the factors can be recovered with relatively simple greatest-common-divisor calculations. The hard part is finding the period. For cryptographically large numbers, the pattern is spread across a search space so vast that no classical computer can realistically map it.

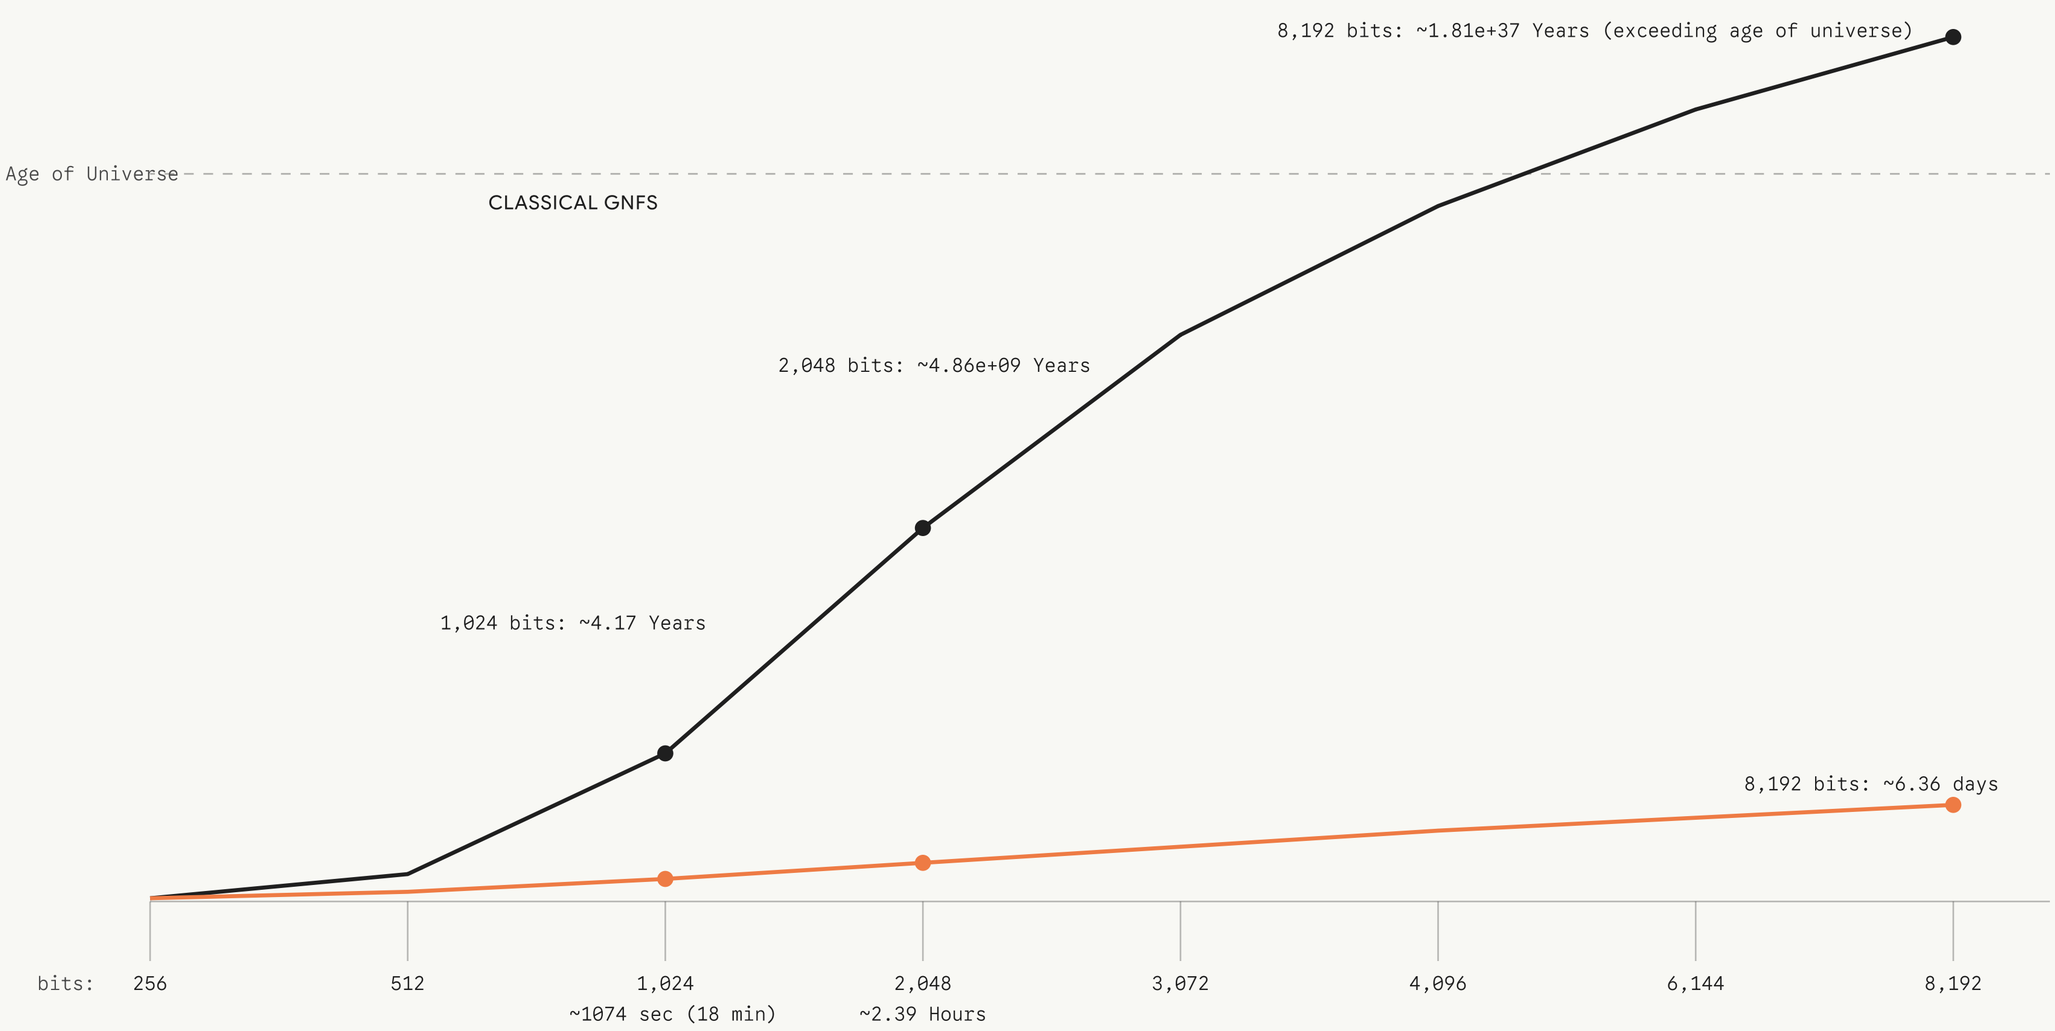
That is where the quantum part matters.Shor’s algorithm uses a quantum computer to encode many possible inputs to the periodic function at once, then applies the **Quantum Fourier Transform (QFT)**, a circuit designed to reveal periodic structure. The QFT turns the hidden rhythm of the function into measurable peaks: wrong period candidates tend to cancel through destructive interference, while information about the true period is amplified enough to survive measurement. In short, superposition creates the stage, and interference extracts the rhythm. A classical computer has no comparable way to make the wrong answers cancel out in bulk, which is why Shor’s algorithm is not merely faster than known classical factoring methods, but **exponentially faster** in the usual account of its speedup [15]. The full mechanics deserve their own piece; we’ve written that up separately here.

The net effect is stark. For cryptographically large keys, a factoring problem that would take known classical methods millions to billions of years collapses to hours or days on a sufficiently large, fault-tolerant quantum computer. That would not break the internet in one stroke, but it would break the public-key assumptions underneath much of it, including RSA, elliptic-curve cryptography, TLS handshakes, VPN authentication, code signing, and many cryptocurrency signature schemes. Shor's algorithm is not a faster crowbar for the same lock. It reveals that the lock was secured by a kind of mathematical hardness quantum computers are built to attack.

Running Shor's at cryptographic scale remains, as of mid-2026, entirely theoretical. The most credible recent resource estimate, Craig Gidney's 2025 revision of the Gidney-Ekerå 2019 estimate, puts factoring RSA-2048 at under one million noisy physical qubits running for under a week [13] [14]. The largest number actually factored on quantum hardware with Shor's or its variants remains in the low double digits of bits, many orders of magnitude below cryptographic relevance. The gap is not the algorithm. It is entirely the error-correction overhead, the roughly 1,000-to-1 penalty for turning noisy physical qubits into reliable logical ones that we met earlier.

Two investment conclusions follow. First, Shor's is the reason post-quantum cryptography (PQC) migration is happening *now*, independent of when large-scale quantum machines arrive, because the "harvest now, decrypt later" threat is already active [16] [17]. Second, Shor's is not, on its own, a reason to invest in quantum hardware today. **The cash flows from Shor's exist in the PQC migration layer**, classical software with classical margins, not in the quantum machines that may or may not eventually run it.

FIG. 13 — TIME TO FACTOR N-BIT INTEGER



Classical factoring time explodes with key size; a fault-tolerant quantum machine's does not.

Grover’s Algorithm

If Shor’s is a symphony that demands a fault-tolerant orchestra, Grover’s is a deceptively simple étude, and the most systematically overrated score in the repertoire.

Published by Lov Grover in 1996, it solves a plain-sounding problem: searching an unstructured database [18]. Given N items and a way to check whether any item is the one you want, a classical computer needs on average N/2 checks. Grover’s finds it in roughly √N. For a trillion items, that’s the difference between a trillion operations and a million. Dramatic, but crucially not exponential.

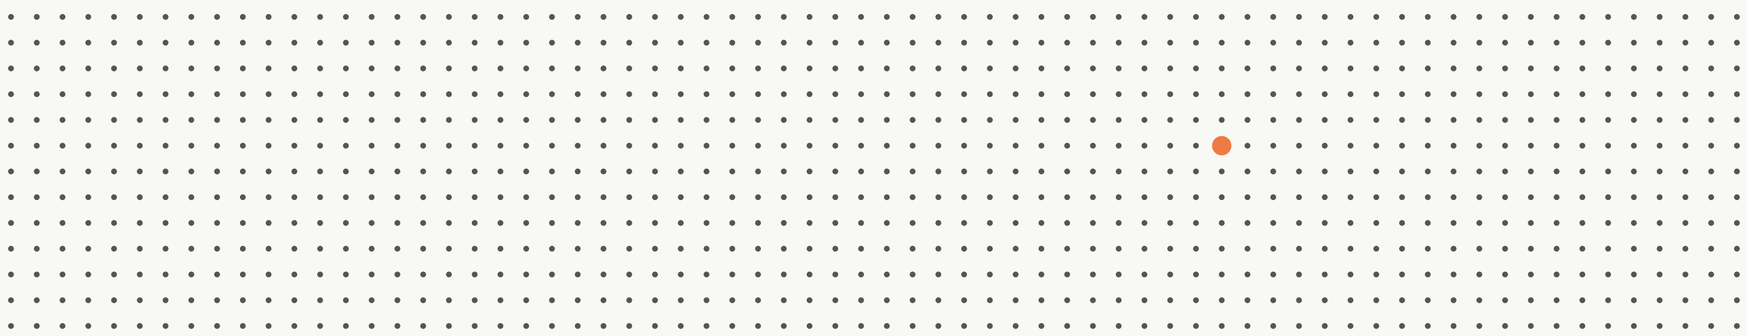
The mechanism is amplitude amplification. Start with a uniform superposition over all possible answers, then iteratively apply an “oracle” that flips the sign of the correct answer’s amplitude and a “diffusion” operator that reflects all amplitudes around their average. After roughly √N iterations, measurement returns the right answer with high probability. Picture a needle in a haystack of a million pieces: a classical search checks 500,000 pieces on average; Grover’s uses interference to amplify the needle’s amplitude and finds it in about 1,000 tries.

Grover’s quadratic speedup is real, mathematically optimal, and systematically overhyped. Against symmetric cryptography and hashing, Grover does not break the primitive so much as cut its effective brute-force strength in half. AES-256 under Grover’s attack offers roughly the security of AES-128 against classical attack, which remains a strong security level. The fix is boring but effective, simply double the key lengths. That is why NIST’s post-quantum guidance treats symmetric crypto as a sizing problem, not a replacement problem [19].

Grover's algorithm (1996) offers a quadratic speedup for searching unstructured databases, transforming N checks into roughly √N operations.

- **Mechanism:** Amplifies probability amplitudes via repeated oracle + diffusion cycles.
- **Not Exponential:** Unlike Shor's, the speedup is strictly quadratic, meaning it does not break cryptography outright.
- **Crypto Impact:** Halves effective key strength (AES-256 behaves like AES-128).
- **The Remedy:** Simply double symmetric key lengths. NIST views this strictly as a sizing adjustments problem.

FIG. 14 — THE NEEDLE IN THE HAYSTACK



A million-piece haystack: a classical search checks 500,000 pieces on average. Grover’s finds the needle in about 1,000 tries.

The subtler, more contrarian point is that Grover's speedup may not translate into a real-world speedup for many valuable problems**. Grover iterations are largely serial, and a serious attack would require fault-tolerant quantum hardware capable of surviving enormous numbers of corrected operations.** The asymptotic speedup is real, but the economic speedup is contested. Classical brute-force search parallelizes brutally well across GPUs and custom hardware, while fault-tolerant quantum search pays a heavy price in error correction, gate time, and logical-qubit overhead. For many practical targets, that overhead could erase much of the advantage. For investors, Grover's is not nothing, but it is a narrower commercial opportunity than the headlines imply, and a much less acute cryptographic threat than Shor's.

Variational Quantum Eigensolver (VQE)

The third canonical algorithm is the only one written for the orchestra we actually have, noisy, small, and badly out of tune. That makes it the most commercially relevant algorithm in the near term, and the one whose claims deserve the most skepticism.

The Variational Quantum Eigensolver, or VQE, was introduced in 2014 to estimate the ground-state energy of a quantum system, one of the central problems in quantum chemistry [21]. That number matters because the ground state anchors a molecule's real-world behavior, including how it binds, reacts, conducts, stores energy, or breaks down. Many of the hardest simulation problems in drug discovery, materials science, and energy chemistry depend on computing these quantum states accurately. Get the ground state right, or close enough, and you can screen molecules in software before spending time and money in the lab.

Variational Quantum Eigensolver (VQE)

Unlike Shor's and Grover's algorithms, which require fault-tolerant machines we do not yet have, VQE was designed for today's noisy hardware. It is the flagship algorithm of the noisy intermediate-scale quantum era, usually called NISQ. Its mechanism is a hybrid loop. A quantum computer prepares a trial molecular wavefunction and measures its energy. A classical optimizer then adjusts the parameters of that wavefunction to lower the measured energy. The process repeats until it converges. The variational principle guarantees that the lowest energy found is an upper bound on the true ground-state energy. In practical terms, the output is an approximation of the molecule's most stable configuration, with possible applications in drug binding, catalyst design, battery chemistry, and materials engineering [21].

VQE has real advantages on paper. It tolerates noise better than deep fault-tolerant algorithms, uses relatively shallow circuits, and targets problems with obvious industrial value. Its problem, and the part most coverage glosses over, is that as of **mid-2026 there is still no clean, reproducible demonstration of VQE outperforming the best classical chemistry methods on an industrially relevant molecule** [22] [23]. Classical techniques such as tensor networks, coupled-cluster methods, and neural-network wavefunctions keep improving, and they have repeatedly matched or beaten VQE results shortly after publication. Classical hardware can now emulate some VQE-class circuits in the 80-to-92-qubit range, depending on circuit structure and assumptions [22]. VQE also suffers from barren plateaus, where the optimization landscape becomes exponentially flat and defeats gradient-based training [24]. On top of that, measurement overhead can make each iteration expensive in real wall-clock time.

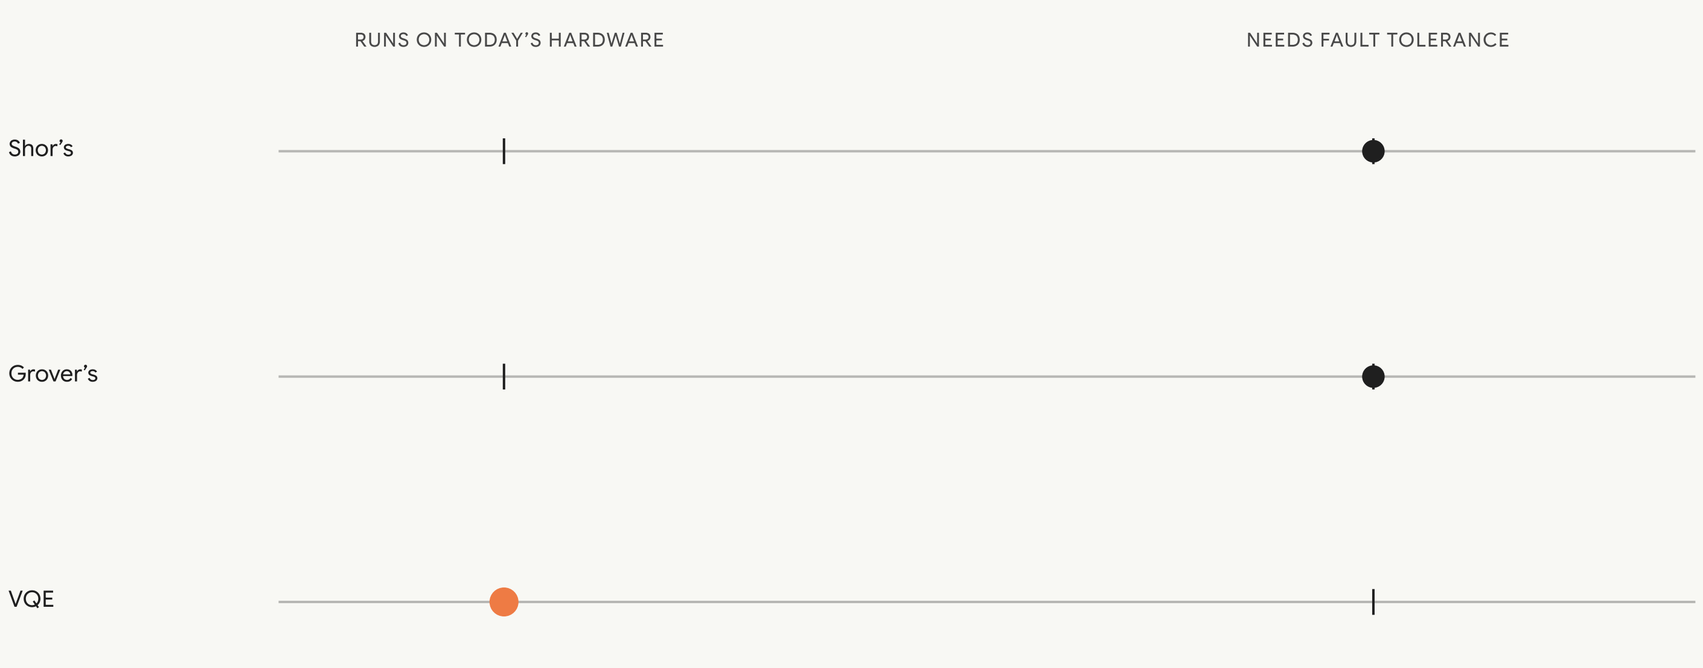
If you want one benchmark for whether NISQ-era commercial quantum computing is real, watch for this: a hybrid quantum-classical method such as VQE, ADAPT-VQE, QAOA, or a successor algorithm producing a clear advantage over the best classical methods on an industrially relevant chemistry, materials, or optimization problem. If that happens, the near-term quantum market changes. The value would not sit only in the hardware. It would spread into the software stack around it: workflow orchestration, ansatz design, error mitigation, compiler optimization, and domain-specific applications. Those layers would become real investable categories.

If it does not happen, the near-term commercial story weakens. The market will stop pricing quantum as a 2020s platform opportunity and reset around fault-tolerant machines arriving in the 2030s.

Three Scores, Three Theses, Three Timelines

Shor's gives an exponential speedup on a narrow problem, requires full fault tolerance, and creates economic urgency today via harvest-now-decrypt-later. **Grover's** gives a quadratic speedup on a broad problem and is mostly noise for investors; discount pitches built on it accordingly. **VQE** runs on the hardware we actually have, targets the most economically valuable near-term problem class, but has not yet cleared the classical bar. What these three tell us about where the money is, and when it will flow, is the subject of the rest of this piece.

FIG. 15 — THREE SCORES, THREE TRACKS



Why You Should Care

Music can be beautiful on its own, but investors care whether people will pay to listen. In quantum computing, the demand case comes from two markets where classical computing runs into hard limits: **molecular chemistry and cybersecurity.**

Chemistry is the upside case. Cybersecurity is the forced-migration case. One is a large future application market if the technology works. The other is already pulling spending forward because the risk exists before the machines do.

04

Revolutionizing Chemistry

The physical nature of quantum technology is game-changing for molecule discovery and optimization. This is where VQE comes into play and will enable humans to discover new, more effective compounds in medicine, agroscience, and materials science. This is possible with a quantum computer because nature is inherently quantum-mechanical.

The enhanced capabilities of a quantum computer can be well described to venture capitalists by looking at drug development. Traditionally, this takes years and potentially billions of dollars to produce a viable commercial drug. This is largely due to two main phases in the drug development process: early candidate discovery and human trials. Early candidate discovery is currently an incredibly manual process that takes years. Researchers are forced to test tens of thousands of potential molecules or more just to find a lead molecule to use for efficacy and safety testing, all costing money and taking time. Quantum computers will take this process and turn it completely digital, cutting down the total time from years to months. Even after that, only 10% of drugs today that go into clinical trials make it out. That number will drastically increase as most of these failures are due to off-target toxicities or efficacy issues that happen because classical supercomputers can’t take into account all the different variables of the human immune system. Quantum computers will be able to precisely simulate how a drug will interact with real cells and optimize the design for the desired effect before the drug ever leaves the lab. Lower development costs and shorter time to market will fundamentally change the math behind venture investments in therapeutics.

Quantum computing is compelling for chemistry for a simple reason. Molecules are quantum systems. Simulating them with classical computers means approximating quantum behavior on machines that were not built for it. A useful quantum computer would attack the problem more directly.

That is where VQE and its successors matter. Their goal is to estimate the ground-state energy of a molecule, one of the central quantities in quantum chemistry [21]. That energy helps determine how a molecule binds, reacts, conducts, stores energy, or breaks down. Better molecular simulation could improve drug discovery, catalyst design, battery chemistry, agroscience, and materials engineering.

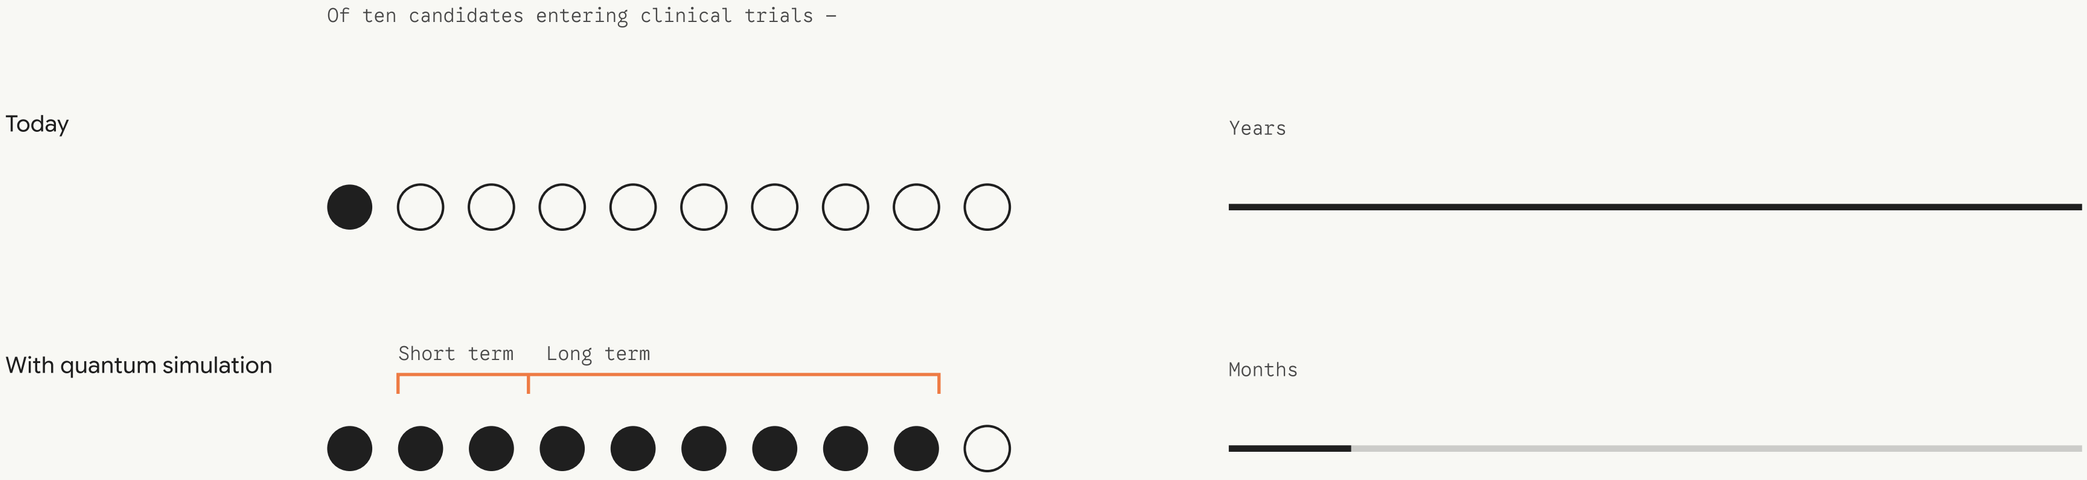
The investment case is easiest to see in drug development. Today, finding and validating a viable drug candidate can take years and cost billions. Years are spent screening large numbers of molecules, narrowing them to a few promising leads, and then discovering at the end of the process that a candidate fails on efficacy, toxicity, manufacturability, or some other constraint. Better simulation would not eliminate wet labs or clinical trials. What it could do is drastically decrease the time for candidate selection from years down to months, and increase the probability of viability, increasing successful clinical trial rates from around 10% today to around 90%.

That matters because even small improvements in candidate selection can change the underlying economics of therapeutics investment. If quantum chemistry eventually lets researchers screen molecules with meaningfully higher accuracy before synthesis and testing, it could shorten discovery cycles, lower development costs, and expand the number of commercially viable programs.

The caveat is just as important. As of mid-2026, this remains an application thesis, not a proven commercial reality. There is still no clean, reproducible VQE-class win over the best classical chemistry methods on an industrially relevant molecule [22] [23]. Classical methods keep improving. GPU clusters, tensor networks, coupled-cluster methods, and neural-network wavefunctions continue to push the benchmark higher. So chemistry is real optionality, but it should be funded against evidence, not against the calendar.

FIG. 16 — TIME TO A USEFUL DRUG

Of ten candidates entering clinical trials –



Candidate selection falls from years to months; the share of candidates surviving clinical trials rises from around 10% today to around 30% in the short term, and as high as 90% in the long term.

Source: Gadbail et al., “Quantum computing in drug discovery,” Frontiers in Drug Discovery 6 (2026).

Cybersecurity and Q-Day

FIG. 17 — HARVEST NOW, DECRYPT LATER



While quantum computers provide a wholesome boost to molecular chemistry, the change to the cybersecurity market could be devastating instead. As mentioned previously, Shor's algorithm, when implemented on a functioning quantum computer, can completely bypass classical cybersecurity systems. This is called Q-Day, the day quantum computing breaks the internet. Once this technology is unlocked, it will be a double-edged sword. Some companies will suffer due to not preparing adequately for Q-Day; meanwhile, investors who prepared properly for Q-Day will make billions, as an entirely new industry built around post-quantum cryptography (PQC) explodes as every company in the world is forced to adopt a PQC solution.

Today, internet cryptography is taken for granted. Every login, file drop, and wire transfer runs on invisible math that no one really thinks about. This invisibility is exactly what makes Q-Day so easy to underprice. It makes it hard for people to realize that, one day, a switch could be flipped and every bit of value a company has can be stolen in a heartbeat. This impending danger will create a forced migration of not only individual companies, but entire governments, entrusting their cyber existence to PQC. Such an inelastic demand paired with the technological moat of operating PQC will create massive profit margins for providers. The new cybersecurity market is already beginning to form. A service that has been inherently included in the system will now be licensed out to individual companies, providing investors with a golden opportunity to enter a market with high necessity and zero incumbent competitors.

Cybersecurity is different. It does not require quantum computers to be useful today. It only requires enough people to believe that useful quantum computers may arrive within the lifetime of sensitive data.

Shor's algorithm, running on a sufficiently large fault-tolerant quantum computer, would break the public-key assumptions behind RSA and elliptic-curve cryptography. That does not mean every part of cybersecurity disappears. Symmetric encryption can be hardened with larger keys. Hashing survives with adjustments. However, the entire public-key layer that secures identity, signatures, key exchange, software updates, VPNs, cryptocurrency wallets, and much of the internet's trust infrastructure would need to change.

That is the logic behind Q-Day. The phrase is dramatic, but the spending case is practical and the risk is real. Companies and governments cannot wait until the first cryptographically relevant quantum computer exists. Adversaries can collect encrypted traffic today and decrypt it later if the data remains valuable. That is the "harvest now, decrypt later" threat, and it turns post-quantum cryptography from a future technical upgrade into a current risk-management problem [16] [17]. Every world government and large corporation will need protection from Q-Day, making this market a necessity rather than a luxury.

This is why PQC migration is the cleanest near-term quantum market. It has regulatory pull, enterprise urgency, classical software margins, and no dependence on a single quantum computer functioning. If quantum arrives quickly, the migration was necessary. If quantum slips, the mandates and data-retention risk still drive spending. It is the rare quantum-linked market with revenue today and no science risk.

The Application & Defense Layer

Post-Quantum Cryptography vs. NISQ Application Platforms

REVENUE TODAY

ENTERPRISE PQC MIGRATION

NIST-standardized · Crypto inventory & remediation · FIPS 140-3

PQShield $65M

QuSecure $41M

CryptoNext $12M

QIZ Security $27M

QKD & CRYPTO PROTECTION

Hardware key exchange · QRNG · Secure comms

Quintessence Labs $55M

Quantum Xchange $24M

Project Eleven $26M

AWAITING FIDELITY THRESHOLD

VQE / CHEMISTRY & MATERIALS

Drug discovery · Molecular simulation — the core NISQ workload, requires gate fidelity breakthroughs per Section 4

Algorithmiq $40M · Ser B

Qunova $14M · Ser A

Also: Qubit Pharmaceuticals, 1QBit (acq'd by Quantinuum)

These platforms require error rates to drop from ~10^{-3} to ~10^{-6} before commercial scaling. The QEC layer is the unlock.

Capital urgency: PQC is the immediate software revenue play — NIST standards finalized, "Harvest Now, Decrypt Later" threat drives enterprise urgency. NISQ molecular simulation awaits the gate fidelity breakthroughs the paper identifies as the real metric.

Source: Harmonic

Why Now

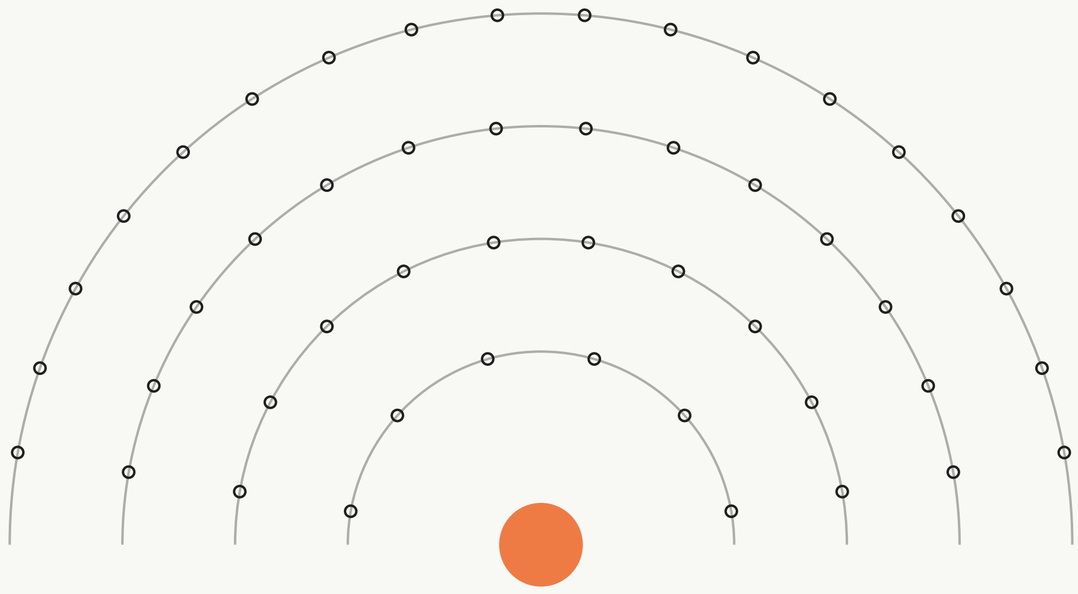
Opening night has not arrived. The machines are still too small, too noisy, and too error-prone for most commercially decisive workloads. But the venue is already under construction.

Quantum computing is moving from science project to industrial buildout. Indiana and Illinois are developing the Quantum Corridor, a regional infrastructure network designed to connect companies, research institutions, and quantum assets across Indiana and Chicago including Fermilab, one of the leading physics institutes in the world [33] [34]. The Mountain West, especially Colorado and New Mexico, is forming another major cluster around national labs, university research, photonics, atomic physics, quantum networking, and technical talent. Albany, New York brings a different but critical capability: semiconductor R&D, advanced fabrication, packaging, and metrology infrastructure that could help quantum hardware move from bespoke lab systems toward manufacturable machines.

That is why now matters. The first major cash flows in quantum will not necessarily come from the companies that eventually discover new drugs or run Shor's algorithm at cryptographic scale. They will come from the companies building the systems that make those outcomes possible.

The application layer is still waiting for proof. The infrastructure layer is already being built.

FIG. 18 — THE VENUE, UNDER CONSTRUCTION



The seats are set before opening night: infrastructure, institutions, and capital assembling ahead of the first performance.

Behind the Scenes

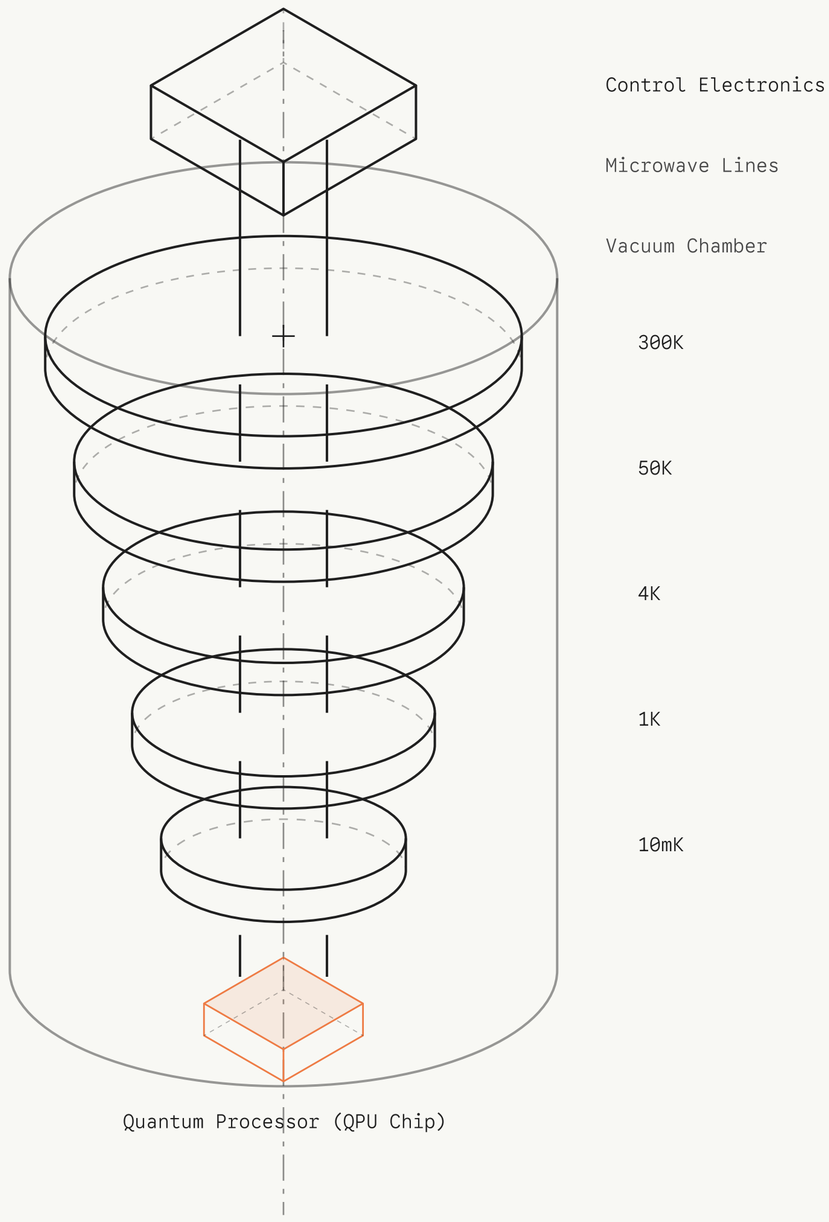
In quantum computing, the performance depends on what happens offstage.

The visible machine gets the attention: the qubit count, the processor photo, the lab demo, the benchmark headline. But the machine only works if a much larger stack works around it. Cryogenics keep superconducting qubits cold enough to operate. Lasers and photonics control trapped ions, neutral atoms, and optical systems. Microwave electronics, vacuum systems, test equipment, control software, compilers, chip fabrication, packaging, cloud access, and error correction all sit between the scientific breakthrough and the commercial product.

This is where the early market is forming. The supply chain resembles existing advanced-technology markets, but with different bottlenecks and different winners. Companies such as Bluefors and Oxford Instruments in cryogenics, Keysight in measurement, Rosenberger and Amphenol in specialized electronics, and Linde and Air Liquide in industrial gases and chemical supply are already positioned near parts of the quantum buildout.

The more open opportunity is in control software and error correction. This layer is not just supporting infrastructure. It is the translation layer between noisy physical hardware and useful logical computation. The commercial technology is still developing, but funding is already moving toward it. Riverlane's $75 million Series C at a reported $400 million-plus valuation is one example of capital recognizing that error correction may become one of the central platforms in the industry.

FIG. 19 — COMPONENTS OF A QUANTUM COMPUTER



One machine, five temperature stages: the quantum chip lives at the coldest point.

Quantum cloud infrastructure will be harder for startups to attack directly. AWS, Azure, and other large cloud providers already have distribution, enterprise relationships, security certifications, and the ability to integrate quantum access into existing compute workflows. Smaller companies can still win in orchestration, middleware, workflow tools, developer interfaces, and domain-specific software. But competing head-on with hyperscale cloud distribution will be difficult.

Chips and quantum processors remain the most technically important and most capital-intensive layer. Architecture, fabrication, packaging, control integration, and manufacturing yield could create deep moats for whichever hardware approaches scale successfully. But this is also where science risk is highest. That is why we separate hardware exposure from the bottleneck layers around it. The full-stack winners may eventually acquire the best component, control, software, and error-correction companies rather than build everything themselves.

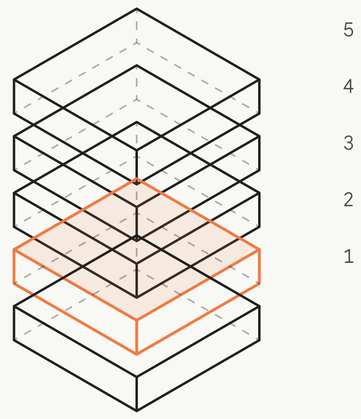
The full-stack developers are the companies that bring all these technologies together. These are also the companies that are providing exit liquidity to venture investors of quantum component companies. IonQ has spent over $3B in cash and stock on acquisitions since the start of 2025, being the largest of the early movers in this space. Quantinuum also had an IPO in June of 2026 at a $17.6B valuation. Even though the technology is not yet viable, the investment is.

That is the point of the backstage metaphor. The audience sees the processor, the qubit count, and the benchmark headline. Investors should be watching the less visible layers: the cryogenics, lasers, control systems, packaging, compilers, error correction, and workflow software that turn a physics experiment into a product. Without those layers, there is no performance. With them, the most valuable companies in quantum may not be the ones standing at center stage.

Behind the Scenes

The Quantum Tech Stack

Five layers of the quantum value chain — structured by where capital should flow



1 SUPPLY CHAIN

"Picks & Shovels" — profits regardless of which modality wins

CRYOGENICS & GAS

Bluefors

Oxford Instruments

Linde
NYSE: LIN

Air Liquide
EPA: AI

ELECTRONICS & HARDWARE

Keysight
NYSE: KEYS

Quantum Machines

SEEQC

Amphenol
NYSE: APH

Rosenberger

2 ERROR CORRECTION

Section 4's "Real Scoreboard"

Riverlane
Ser C · $75M · $400M val.

Q-CTRL
Ser B · $133M

QEDMA
Ser A · $31M

Iceberg Quantum
Seed · $7.5M

Whoever owns the error correction stack bridges the $10^{-3} \rightarrow 10^{-15}$ error rate gap and controls the value chain. Riverlane's Deltaflow chip decodes billions of errors/sec in real-time.

3 CLOUD & EXECUTION

How end-users consume quantum compute

AWS
Braket

Microsoft
Azure Quantum

4 CHIPS / QPU MODALITIES

- SUPERCONDUCTING
  Rigetti
- TRAPPED ION
  IonQ
  + Oxford Ionics acq.
- NEUTRAL ATOM
  QuEra
  Atom Computing
- PHOTONIC
  PsiQuantum

5 FULL STACK / INTEGRATED SYSTEMS

Hardware + Software + Cloud in-house

Quantinuum
NASDAQ: QNT · $17.6B

Google
Quantum AI

IBM
Quantum

IonQ
+ Oxford Ionics

Rigetti
RGTI

D-Wave
QBTS

Section 4 thesis: Capital shouldn't just chase hardware OEMs. It needs to flow to the enabling supply chain (picks-and-shovels), the QEC/middleware layer (the true platform play), and PQC (the immediate cash-flow play).

Source: Harmonic

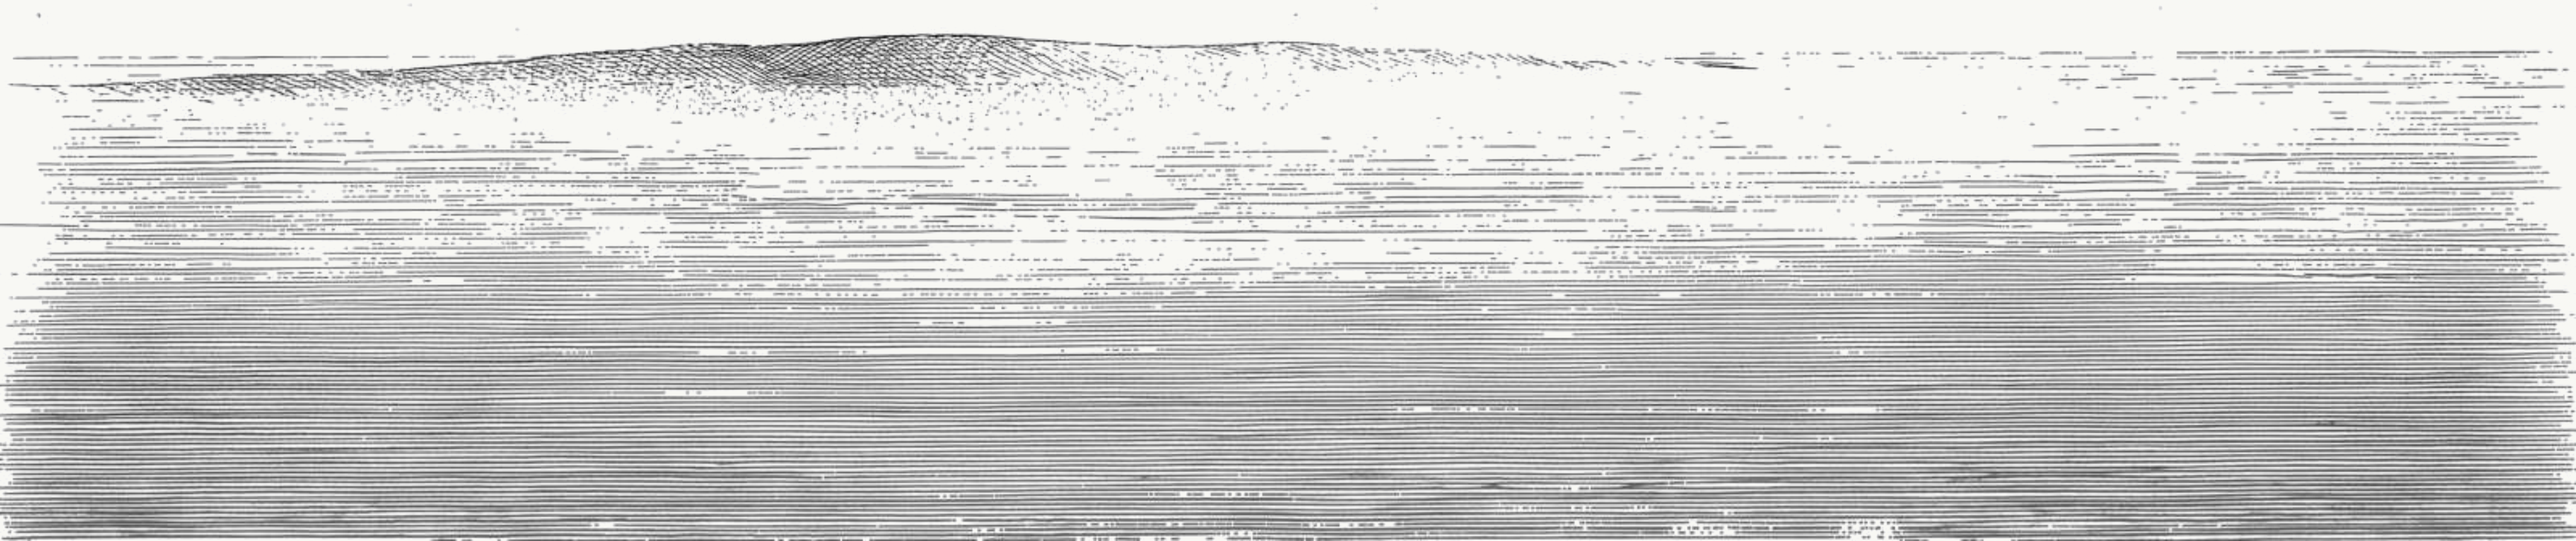
Getting the Investment Right

Available liquidity for exits means the time is right for venture funding to flow into quantum science. Not only have specialized companies been acquired, but Quantinuum, a full-stack developer, went public in June of 2026 with a $17.6B valuation. Many investors are still wary of investing in the category, and for good reason. Quantum is technically complex, timelines are uncertain, and speculative investing in new technologies is hard to get right.

In our 2026 Venture Report, “Crossing the Tilt,” we described the core mistake investors make during technological revolutions:

“The most common mistake investors make during secular technological revolutions isn’t a lack of intelligence or convictions. It’s misjudging where the wave actually is. When a tsunami forms far out at sea, there’s very little to see from shore. The surface looks calm, signals are faint, and early warnings are easy to dismiss as noise. As a result, capital often rushes toward whatever looks active or visible in the moment, rather than the deeper forces quietly building beneath it. Human behavior tends to reinforce this pattern. Markets are wired to anchor on what’s observable and consensual, taking surface-level motion at face value instead of looking farther out to understand what’s actually forming beneath the surface.”

“Crossing the Tilt” — the 2026 Elevate Venture Report. Scan to read.



This is exactly why a fundamental understanding of the science behind quantum computing is important for investors. The first step is understanding that real progress is being made. Deeper forces are at work, and the wave is forming. The second, and more important, step is differentiating between which headlines are meaningful and which are just noise. This step is the key to an investor's success. As meaningful breakthroughs are made in gate fidelity and logical qubit scaling, it will be important for venture investors to get in the game before being squeezed out by high valuations and early acquisitions by full-stack companies. At the same time, it is equally important for the same investors not to be fooled by pitches parading higher qubit counts as their value differentiator.

The show is still not ready to go on, but the build-up is in full effect. The audience is buzzing with excitement as advertisements spread the word about how incredible the show will be. Even before the final show is ready, money is flowing behind the scenes. Now is the perfect time for venture investors to get involved. The market has proven it will pay for material improvements in the quantum space, and still, the biggest innovations are around the corner. The market is new but not fundamentally different from past technological emergences. Each sector carries distinct risks and payoffs, and the indicators of success vary in technological and financial sophistication. The burden on investors is to discern which indicators are real and which aren't; those who succeed will be on the leading edge of transforming the world.

Our framework leads to three areas where the risk-reward looks most favorable today.

Getting the Investment Right

The Investment Framework

Bet One: PQC Migration, Now

Bet Two: The Layers That Win Regardless of Who Wins

Bet Three: Staged Optionality on Applications

What We Avoid

What Could Kill the Thesis

05

The Investment Framework

1

Bet One: PQC Migration, Now

PQC migration is the only quantum-driven market with mandated demand, current revenue, and minimal science risk. Governments are pushing migration. Enterprises need crypto inventories, implementation support, key-management upgrades, vendor audits, and compliance tooling.

The harvest-now-decrypt-later threat makes the demand immediate even if large-scale quantum hardware is years away [16] [17] [26].

This is classical software and services sold into a quantum-driven budget cycle. That is why it is attractive.

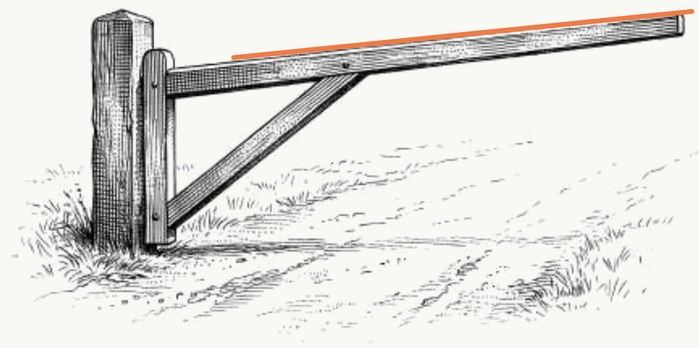
2

Bet Two: The Layers That Win Regardless of Who Wins

The market loves to debate superconducting versus trapped ion versus neutral atom versus photonic systems. That debate matters scientifically, but it is not the only way to invest. Every credible architecture needs supply-chain components, control systems, calibration software, compilers, error mitigation, and ultimately error correction.

Our conviction pick is the error-correction stack. It is modality-agnostic, sits between every useful algorithm and every imperfect machine, and becomes non-discretionary the moment anyone tries to scale. Whoever owns the abstraction layer that turns unreliable hardware into reliable logical qubits can collect a toll on the entire industry.

In private markets, this layer is still often priced like infrastructure. We think it is closer to the platform.



Bet Two — the layer every architecture must pass through.

3

Bet Three: Staged Optionality on Applications

The molecular-discovery story is real, but it is not yet proven. We would not underwrite it as if the benchmark has already been crossed. The right approach is to stage capital against evidence.

The benchmark is straightforward. Watch for VQE, ADAPT-VQE, QAOA, or another hybrid quantum-classical method to beat the best classical baseline on an industrially relevant chemistry or optimization problem [22] [23]. If that happens cleanly and reproducibly, the application layer becomes investable very quickly. The winners will likely emerge in workflow software, domain-specific algorithms, ansatz design, error mitigation, and integration with existing chemistry and materials platforms.

Until then, applications are option value. Infrastructure is the base case.

What We Avoid

We would systematically avoid three kinds of pitches.

First, anything whose headline metric is simply qubit count. More qubits matter only if they can be controlled, corrected, and used in circuits deep enough to perform valuable work.

Second, anything whose commercial case rests primarily on Grover's algorithm. The quadratic speedup is real, but the wall-clock economics are much less compelling once serial iterations, error-correction overhead, and classical parallelism are included [20].

Third, public hardware platforms priced as if scientific uncertainty has already cleared. Hardware may produce enormous winners eventually, but at today's stage many of those investments still combine platform multiples with science-project risk.

What Could Kill the Thesis

The risks are real.

Classical computing keeps moving the bar. GPU clusters and improved algorithms can emulate larger quantum circuits, match published quantum chemistry results, and reduce the practical advantage claimed by noisy quantum machines [20] [22] [23]. Error-correction overhead may remain stubbornly high. Logical-qubit scaling may take longer than expected. Fault tolerance may arrive after the market loses patience.

That is why position sizing matters. It is also why our first two bets are structured to work in a slow-quantum world. PQC migration can pay before quantum hardware succeeds. Bottleneck infrastructure can be valuable across multiple architectures and timelines. The application layer should wait for proof.

Count Nines, Not Qubits

By now, the pattern should be clear. Every concept we have discussed, from algorithms and gate fidelity to error correction and commercialization, points toward the same conclusion: **the market is watching the wrong scoreboard.**

Qubit counts make headlines. Reliability, error correction, and the software stack determine where durable value is created.

06

The distinction does not just change how we think about quantum computing. It changes how we allocate capital. We are not trying to predict which hardware architecture wins. We are asking a simpler question: **what does every successful quantum computer inevitably require?**

Answer that correctly, and much of the investment landscape comes into focus.

Count Nines, Not Qubits

The scoreboard we watch is simple:

1. **Two-qubit gate fidelity,** because every additional 9 changes the economics of error correction.
2. **Logical-qubit counts,** because physical qubits are raw material, while logical qubits are what useful algorithms actually consume.
3. **Error-corrected performance**, especially whether systems can move from memory to sustained computation.
4. **VQE-class results versus the best classical baselines,** because chemistry is the leading indicator for whether NISQ-era commercial quantum is real.
5. **PQC budget conversion,** because fear is not revenue. Signed migration work is. Watch for crypto inventories, key-management upgrades, vendor attestations, procurement language, and repeatable contracts as quantum readiness becomes an operational requirement rather than a future technology discussion

Five numbers. None of them is the headline qubit count. Together, they are the investor's scoreboard: not how large the machine looks, but how close it is to doing valuable work.

Every Revolution Has Two Numbers

A bigger orchestra does not guarantee better music. If the instruments are out of tune and the players miss their notes, adding more musicians only makes the problem louder. Quantum computing works the same way. More qubits make headlines, but reliability determines whether the machine can actually perform the score.

07

Every technological revolution has two numbers. One is the number the crowd watches. The other is the number that rations progress.

Railroads were scored in miles of track laid, but the fortunes were decided by freight economics. The internet was scored in eyeballs, but durable value accrued to the infrastructure that turned attention into transactions. AI is scored in model releases, but much of the value has pooled in the hardware, data centers, and tooling required to train and deploy them. **In every case, the visible metric inflates the excitement, while the rationing constraint quietly decides who survives it.**

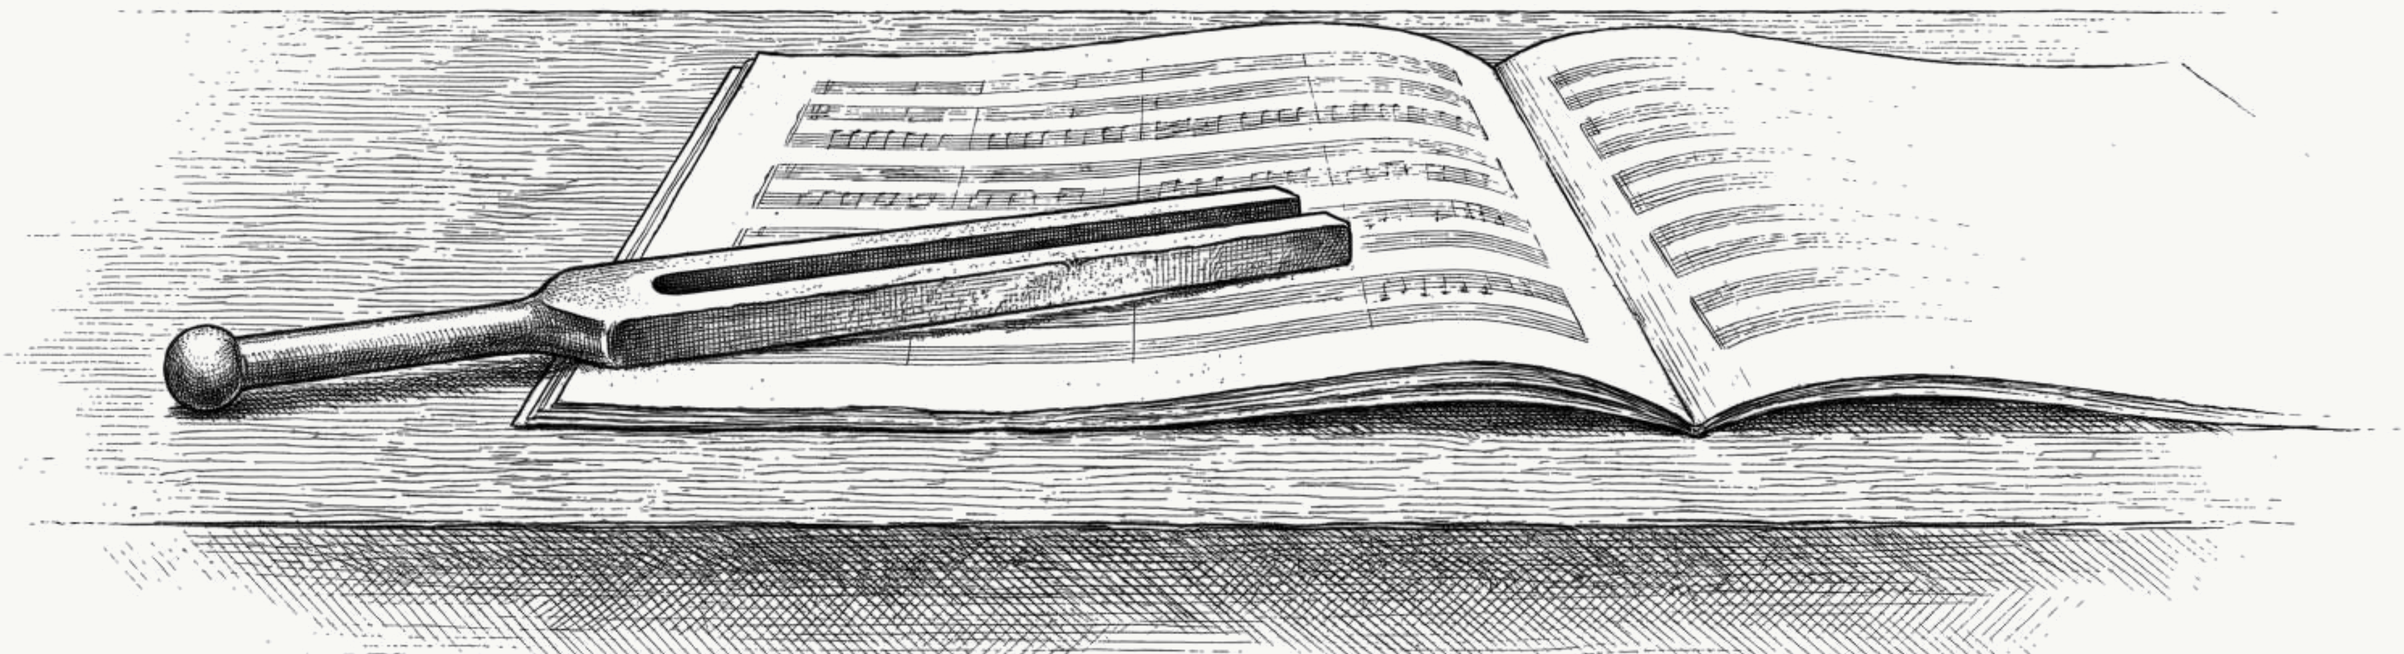
Quantum is the same pattern in a new technical language. The visible number is qubit count. The rationing constraint is fault tolerance.

That is the mental model we hope investors take from this piece. Do not ask only how big the machines are getting. Ask what scarce, unglamorous capability stands between the demo and the invoice. Ask who is being paid to build that capability now, and who will own it once the market scales.

In quantum computing, that capability is the error-corrected gap between a machine that works in the lab and a platform that works in the market. The scoreboard is circuit depth, gate fidelity, usable logical qubits, physical qubits per logical qubit, and evidence that error correction can support sustained computation rather than isolated demonstrations.

Count nines, not qubits.

The market prices the miracle. The returns go to whoever prices the bottleneck.



References

MARKET & FUNDING

1. Crunchbase News – *“Quantum Computing Startup Investment Slows In 2026 While Public Markets Hold Strong”*

2. PitchBook – *“Q2 2026 Bit by Qubit: Global Quantum Computing Funding Hits New Records”*

3. The Nobel Prize in Physics 2022 (Aspect, Clauser, Zeilinger)

10. The Quantum Insider – *“Quantinuum to Debut on Nasdaq After Raising $1.68 Billion in IPO”* (June 2026)

11. CNBC – *“Quantinuum closes flat in Nasdaq debut, after upsized offering”* (June 2026)

29. CNBC – *“Nvidia becomes first company to reach $5 trillion valuation”* (October 2025)

30. CNBC – *“IonQ buys UK quantum startup Oxford Ionics for more than $1 billion”* (June 2025)

31. The Quantum Insider – *“IonQ Completes Acquisition of Oxford Ionics”* (September 2025; ~$1.59B total consideration at closing)

32. TechTimes – *“Quantinuum IPO Prices at $60 on Nasdaq Debut: $14B Valuation Meets $31M Revenue”* (June 2026)

33. Indiana Business Journal – *“Indiana plays increasing role in quantum future”* (July 2026)

HARDWARE & FUNDAMENTALS

8. Quantum Computing Report – *“Quantinuum Launches Helios Quantum Computer with Industry-Leading Fidelity”*

9. *“A 98-qubit trapped-ion quantum computer with all-to-all connectivity,”* Nature (2026)

12. Quantum Zeitgeist – *“Quantum Logical Qubit Leaderboard 2026”*

20. *“Leapfrogging Sycamore: harnessing 1432 GPUs for 7x faster quantum random circuit sampling,”* National Science Review

34. Nature – “Top 10 institutions in ‘big science’ physics and astronomy (December 2019)

ERROR CORRECTION

4. Google Quantum AI – *“Quantum error correction below the surface code threshold,”* Nature (2024)

5. Riverlane – *“Quantum Error Correction: 2025 trends and 2026 predictions”*

6. IBM Quantum – *“IBM lays out clear path to fault-tolerant quantum computing”* (Starling roadmap)

27. Riverlane – *“Riverlane raises $75m to meet surging global demand for quantum error correction technology”*

28. Sifted – *“Building error-free quantum computers: UK startup Riverlane raises $75m Series C”*

ALGORITHMS & COMPUTATION

13. C. Gidney – *“How to factor 2048 bit RSA integers with less than a million noisy qubits,”* arXiv (2025)

14. C. Gidney & M. Ekerå – *“How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits,”* arXiv (2019)

15. P. Shor – *“Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,”* arXiv (1995)

18. L. Grover – *“A fast quantum mechanical algorithm for database search,”* arXiv (1996)

21. A. Peruzzo et al. – *“A variational eigenvalue solver on a photonic quantum processor,”* Nature Communications (2014)

22. *“Towards Quantum Advantage in Chemistry,”* arXiv (2025)

23. *“Parallel iQCC Enables 200 Qubit Scale Quantum Chemistry on Accelerated Computing Platforms Surpassing Classical Benchmarks,”* arXiv (2026)

24. J. McClean et al. – *“Barren plateaus in quantum neural network training landscapes,”* Nature Communications (2018)

25. D. Sun et al. – *“Why 90% of clinical drug development fails and how to improve it?”* Acta Pharmaceutica Sinica B (2022)

CRYPTOGRAPHY & Q-DAY

7. NIST NCCoE – *“Migration to Post-Quantum Cryptography”* (background on CRQC resource requirements)

16. Cloudflare – *“The White House’s post-quantum executive order is an important milestone”* (2026)

17. FedTech – *“Harvest Now, Decrypt Later: A Federal Quantum Threat”* (July 2026)

19. NIST – *“NIST Releases First 3 Finalized Post-Quantum Encryption Standards”* (August 2024)

26. AxelSpire – *“CNSA 2.0 & NIST PQC Deadlines 2026–2035”*

Colophon

Disclaimer

The Quantum Report is published by the Elevate Ventures Research Desk for informational and educational purposes only. It is not investment, legal, accounting, or tax advice, and nothing in it is an offer to sell or a solicitation of an offer to buy any security or other financial instrument. The views expressed are those of the Research Desk as of the date of publication and are subject to change without notice. Certain information is drawn from third-party sources believed to be reliable; Elevate Ventures makes no representation or warranty, express or implied, as to its accuracy or completeness, and accepts no liability for any loss arising from its use. Statements about future events, technologies, or market outcomes are forward-looking and inherently uncertain; actual results may differ materially. Elevate Ventures and its affiliates may hold, or may in the future hold, interests in companies referenced in this report. Nothing herein should be the basis of an investment decision; readers should conduct their own diligence and consult their own advisers.

The Quantum Report is set in Google Sans Flex, with data and annotations in Google Sans Code. Figures, diagrams, and engravings are original to the Elevate Ventures Research Desk. Market data: Harmonic. Saddle-stitched, with a soft-touch cover.

Google Sans Flex

Aa Bb Gg Qq fi fl ffi

Regular · Medium · **SemiBold** · *Italic*

Google Sans Code

|0⟩ |1⟩ 2^{n} √N 10^{-1} {0123456789}

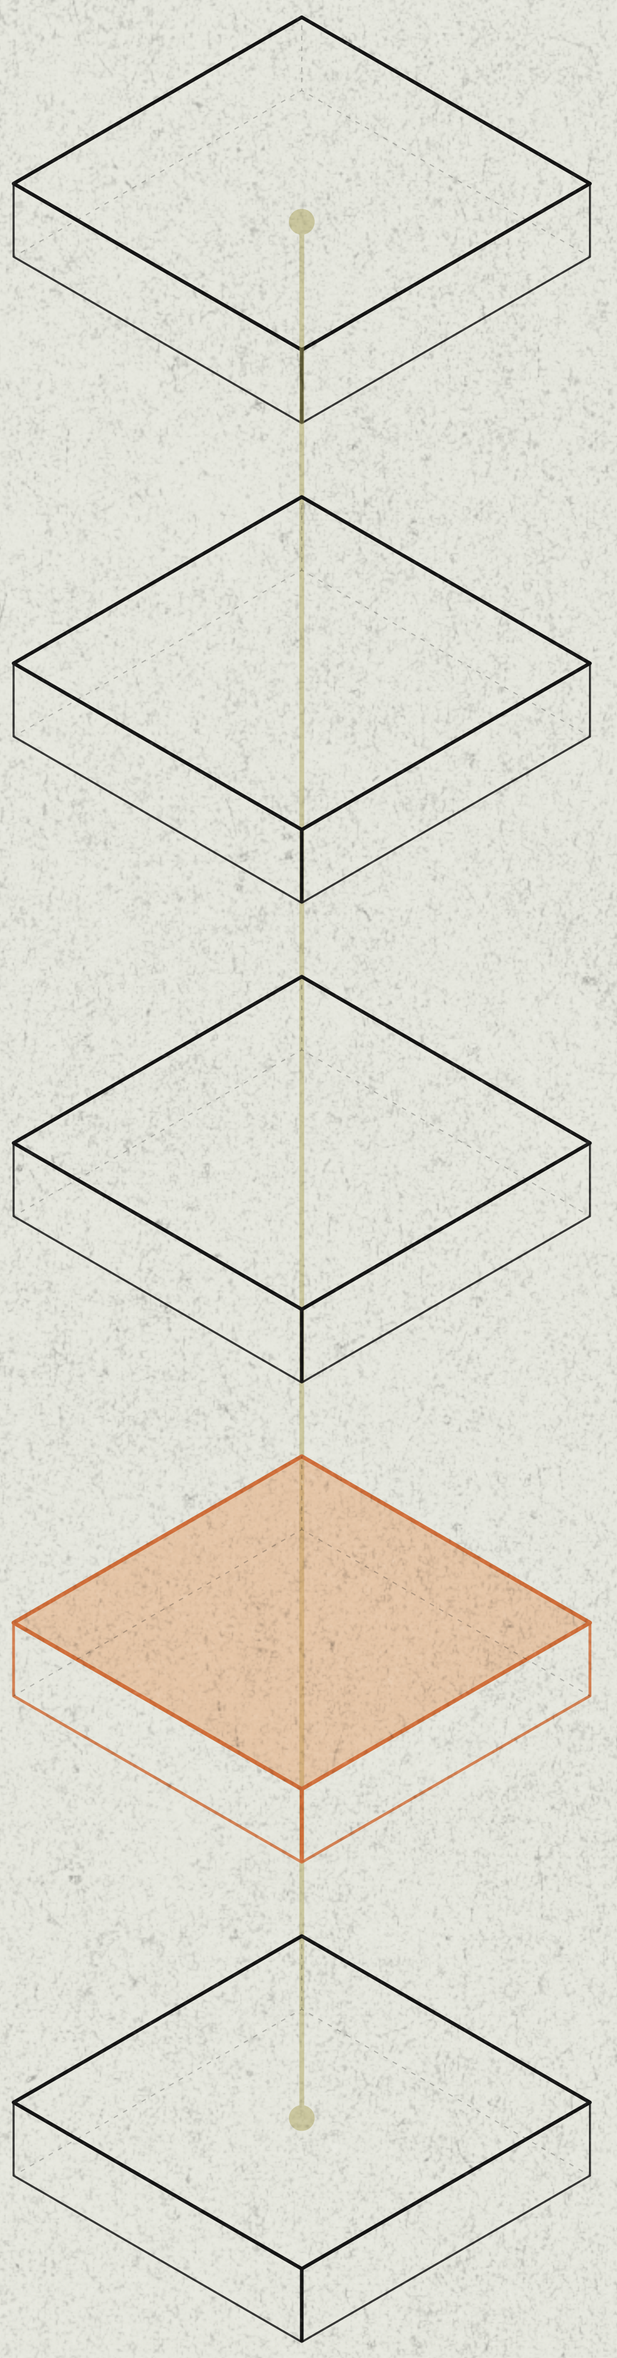
Regular · Medium · **Bold**

Both typefaces by Google. Set at 19/150 for reading, with data and annotation in Google Sans Code.

elevateventures.com

Elevate

THE SCOREBOARD IS WRONG



01
Logical-qubit counts

02
PQC budget conversion

03
VQE-class results versus the best classical baselines

04
Error-corrected performance

05
Two-qubit gate fidelity

The market prices the miracle.
The returns go to whoever prices the bottleneck.