THEMATIC RESEARCH · QUANTUM COMPUTING

How Close Is Quantum Computing to Being Useful?

A plain-English map of where the technology actually stands, an honest estimate of its pacing, and what it means for cryptography, the financial system, Bitcoin, and the companies racing to build it.

TON618 CAPITAL RESEARCH · PREPARED UNDER CFA INSTITUTE ANALYTICAL STANDARDS · FOR INFORMATION PURPOSES ONLY ·

Not yet.

BUT THE CLOCK JUST STARTED MOVING

Quantum computing has not done a single economically useful thing that a classical computer cannot — as of mid-2026, that scoreboard still reads zero. What changed is the foundation: in 2024–25 several labs crossed the *error-correction threshold*, the precondition for ever building a reliable machine. So the honest answer runs on three different clocks. **Narrow scientific usefulness** is arriving now through the late 2020s. **Broad commercial usefulness** is a 2030s story, with full-scale machines plausibly not until 2035–2040+. And the **cryptography-breaking** milestone — the one that matters to every bank, government, and Bitcoin holder — is something experts now put at roughly **one-in-three odds within ten years**, and more likely than not within fifteen. The hype and the substance have never been further apart, and the stocks are priced as if the future already arrived.

0

Economically useful results a quantum computer has produced that a classical computer cannot — as of mid-2026.

~1 in 3

Experts' central odds that a code-breaking quantum computer exists within 10 years (≈2035) — up from ~1 in 5 a year ago.

$51B / $193M

Combined market value vs. annual revenue of the listed pure-play quantum firms — about 268× sales, none profitable.

Executive summary

Quantum computing is simultaneously the most over-hyped and one of the most genuinely important technologies in the market. Both things are true at once, and keeping them apart is the whole job of this note. **The substance:** after twenty-five years of promises, 2024–25 delivered the first results that actually matter — Google's "Willow" chip showed that adding more physical parts can make a quantum bit *more* reliable rather than less (crossing the so-called error-correction threshold), and trapped-ion machines from Quantinuum and IonQ now make errors at rates good enough to build on. That is the necessary precondition for everything else, and it is real and peer-reviewed.

The hype: none of this has yet produced a useful answer to a real-world problem. Every "quantum advantage" headline so far has either been on a contrived benchmark with no economic value, or has been matched by an ordinary computer within weeks — sometimes on a laptop. The one application with a proven, world-changing payoff (breaking the encryption that protects the internet) requires a machine roughly a thousand times larger than today's best, and it does not exist. Meanwhile the four listed pure-play quantum companies earn about $193 million between them and are valued near $52 billion.

This note explains, in plain language, what a quantum computer is and is not good at; where the technology actually stands; why "error correction" — not the qubit-count headlines — is the metric that decides everything; our best estimate of the pacing; and what it means for the cryptography that secures banks, governments, corporations, and Bitcoin. It closes with how we think about the investment opportunity and the portfolio risk. It is analysis, not advice, and contains no recommendation regarding any security or digital asset.

> **The distinction that organizes this note.** Three different milestones get blurred together as "useful." **Scientifically useful** — the machine does one narrow thing of research interest that classical computers struggle with (arriving now). **Commercially useful** — it routinely solves *profitable* problems at scale (2030s). **Cryptographically relevant** — it can break today's public-key encryption ("Q-Day"). These run on different clocks and at different odds; conflating them is the most common error in the field.

1 What a quantum computer actually is — in plain English

An ordinary computer stores information in **bits**, each a definite 0 or 1. A quantum computer uses **qubits**, which exploit two genuinely strange features of the physics of very small things.

- **Superposition.** A qubit can be placed in a blend of 0 *and* 1 at the same time. Put 50 qubits together and the system can, in a sense, hold all 2^{50} (about a quadrillion) combinations at once. This is *not* simply "trying every answer in parallel" — that is the popular myth — but it does give the machine an exponentially large space to work in.
- **Entanglement.** Qubits can be linked so that their fates are tied together; measuring one instantly tells you about the other. This lets the machine represent correlations that would be astronomically expensive to track classically.
- **Interference.** Here is the catch and the craft. When you finally *measure* the qubits, the superposition collapses to a single ordinary answer. A useful quantum algorithm is a carefully choreographed routine that makes the *wrong* answers cancel out (like noise-cancelling headphones) and the *right* answers reinforce, so the one you read out is very likely the one you wanted.

The consequence: a quantum computer is not a faster version of your laptop, and it will never be one. It is a **special-purpose instrument** that is dramatically faster for a *small* set of problems with the right hidden mathematical structure, and no better — often worse — at everything else. Email, spreadsheets, video, and the vast majority of business computing will never run on one.

What it is genuinely good at — and what is mostly hype

The single most important fact for a non-specialist: there are really only **two** problem classes with a proven *exponential* advantage, and both are narrow. The famous "optimization," "AI," and "finance" use-cases mostly are not real — a point made forcefully by researchers at Microsoft, who showed that a merely *quadratic* speedup (the best these offer) is usually erased by the machine's enormous overheads and slow data loading.[1]

WHAT QUANTUM COMPUTERS CAN ACTUALLY DO — THE SPEEDUP, HONESTLY GRADED

APPLICATION	BEST SPEEDUP	VERDICT	PLAIN-ENGLISH STATUS
Simulating molecules & materials (catalysts, batteries, drugs)	Exponential	**Substance**	The prize. "Using quantum to model quantum." Needs a big, error-corrected machine that doesn't exist yet.
Breaking today's encryption (Shor's algorithm)	Exponential	**Substance**	Mathematically certain. Purely a hardware-scale problem — see §5.
Search / unstructured databases (Grover's algorithm)	Quadratic only	**Limited**	Real but modest; routinely eaten by clock-speed and data-loading overheads.
Optimization — logistics, scheduling, trading	Quadratic at best	**Hype**	No durable, proven edge over the best classical methods. Most "advantage" claims here do not survive scrutiny.
Machine learning / "quantum AI"	Usually none	**Hype**	Several headline speedups were "dequantized" — a classical algorithm matched them. Mostly bottlenecked by data loading.

As the complexity theorist Scott Aaronson (UT Austin) puts it, the credible hopes are "(1) simulation of quantum physics and chemistry, (2) breaking a lot of currently deployed cryptography, and (3) eventually, modest benefits for optimization [and] machine learning… but it will probably be a while." Note that even the chemistry prize is daunting: a 2021 estimate to model one important catalyst (the enzyme behind fertilizer production) required ~2.7 *million* qubits; a 2025 algorithmic breakthrough cut that to ~100,000 — "still far more than today's hardware."[2]

2 Where we actually are: real progress, zero useful output

The state of the art is best understood as a paradox. Hardware is improving rapidly and the foundational science just hit a genuine milestone — yet no quantum computer has produced an economically useful result. Both halves matter.

Start with the headline that everyone sees: **qubit counts**. They are rising, but the number is far less meaningful than it sounds — and chasing it can even be counterproductive. IBM deliberately moved its flagship *down* from 1,121 qubits (its "Condor" chip, 2023) to ~120–156 on later chips, trading raw count for quality. Different technologies are not comparable: trapped-ion machines (a few dozen qubits) are far more accurate than superconducting chips (hundreds), and neutral-atom arrays now hold thousands of atoms but cannot yet use most of them reliably.

EXHIBIT 1 · THE QUBIT-COUNT RACE — AND WHY THE NUMBER DECEIVES

Gate-based physical qubits by leading system, log scale. The dashed ceiling is what it would take to break the encryption securing the internet — about a thousand times today's best.

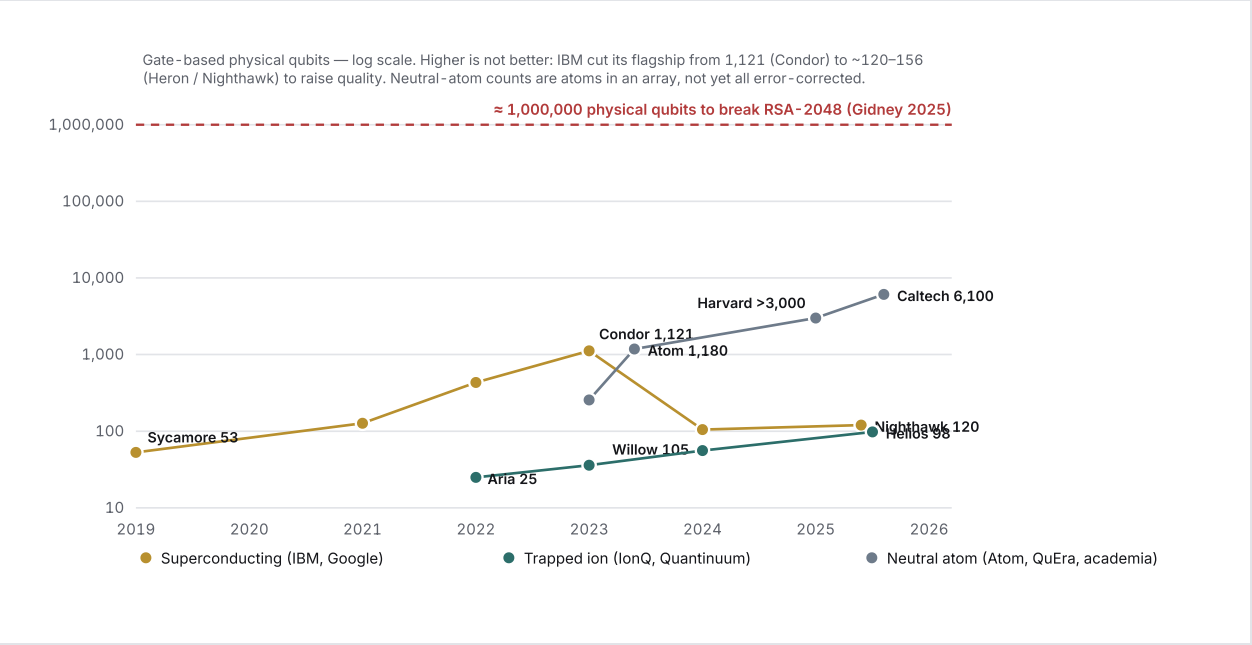

Now the substance. The right way to keep score is not "what was claimed" but "did an ordinary computer catch up?" By that test, every supremacy headline so far has an asterisk — until very recently.

THE "QUANTUM ADVANTAGE" SCOREBOARD — CLAIMS VS. WHAT CLASSICAL COMPUTERS DID NEXT

DATE	CLAIM	WHAT HAPPENED
Oct 2019	Google "supremacy": 200 sec vs. "10,000 years"	IBM rebutted in 3 days; by 2024 the same task ran in ~6 seconds on a classical supercomputer. **Useless benchmark.**
Jun 2023	IBM "utility" before fault tolerance (127 qubits)	Reproduced within weeks on a *single laptop core*. **No useful output.**
Dec 2024	Google **Willow**: below-threshold error correction	The random-benchmark headline is meaningless, but the *error-correction* result is real and peer-reviewed. **Genuine milestone.**
Oct 2025	Google "Quantum Echoes": first *verifiable* advantage, ~13,000× faster	A classical rebuttal was attempted in April 2026 and *failed*. The claim stands — but the task is still a proof-of-principle, not a paying problem. **Best result to date.**

So the honest summary of mid-2026: the science is finally real, the engineering is accelerating, and the commercial payoff has *still not arrived*. Aaronson's warning is worth keeping in mind — "qombies roam the earth: undead narratives of 'quantum advantage for important business problems' detached from any serious underlying truth-claim." (One 2025 bank "trading advantage" claim turned out to come entirely from hardware *noise*; the clean simulation showed no benefit at all.)[3]

3 The metric that actually matters: error correction

If you remember one idea from this note, make it this one. Qubits are exquisitely fragile — a stray vibration, a photon, a flicker of heat, and the delicate superposition collapses and the calculation is ruined. Today's qubits make an error roughly every few hundred to few thousand operations. A useful algorithm needs *billions*. That gap — not the qubit count — is the whole game.

The solution is **quantum error correction**: spread the information of one reliable "logical" qubit across many noisy physical qubits, so that errors can be detected and fixed faster than they accumulate. The catch is a chicken-and-egg threshold. The famous *threshold theorem* says error correction only helps if your physical parts are already *below* a critical error rate (about 1% for the leading scheme). Above it, adding qubits makes things *worse*; below it, adding qubits makes the logical qubit exponentially better. Crossing that line is the gate everything else passes through — and it was finally, convincingly crossed in 2024–25.

EXHIBIT 2 · ERROR RATES HAVE DROPPED BELOW THE FAULT-TOLERANCE THRESHOLD

Best two-qubit gate error rate by technology, over time (lower is better, log scale). Both leading approaches are now below the ~1% line — the necessary condition for scalable error correction.

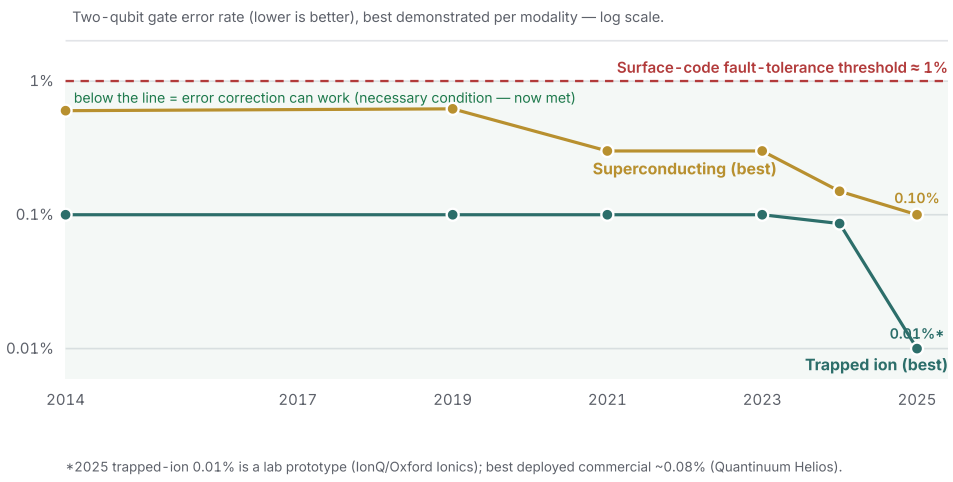

Google's Willow (December 2024) was the watershed: as they made the error-correcting code larger, the logical error rate *fell* by about half at each step — the signature of being below threshold. The peer-reviewed result, not the "ten septillion years" press headline, is what mattered. Quantinuum's "Helios" (November 2025) then ran **48 fully error-corrected logical qubits**, and academic groups at Harvard/MIT/QuEra have built integrated systems on hundreds of atoms.

But here is the sobering arithmetic. Each logical qubit can require anywhere from ~100 to ~1,000+ physical qubits. Today the entire field can muster a few dozen logical qubits. **Useful** work needs hundreds to thousands of them. That is the real distance left to travel — and it is best seen not as a single gap but as a set of *targets that are themselves moving toward us* as algorithms improve.

EXHIBIT 3 · THE GAP THAT MATTERS — AND HOW IT'S CLOSING FROM BOTH ENDS

Physical qubits needed for each milestone vs. the largest machine today (log scale). Red targets have fallen sharply as the *algorithms* improved — the goalposts are moving toward the hardware.

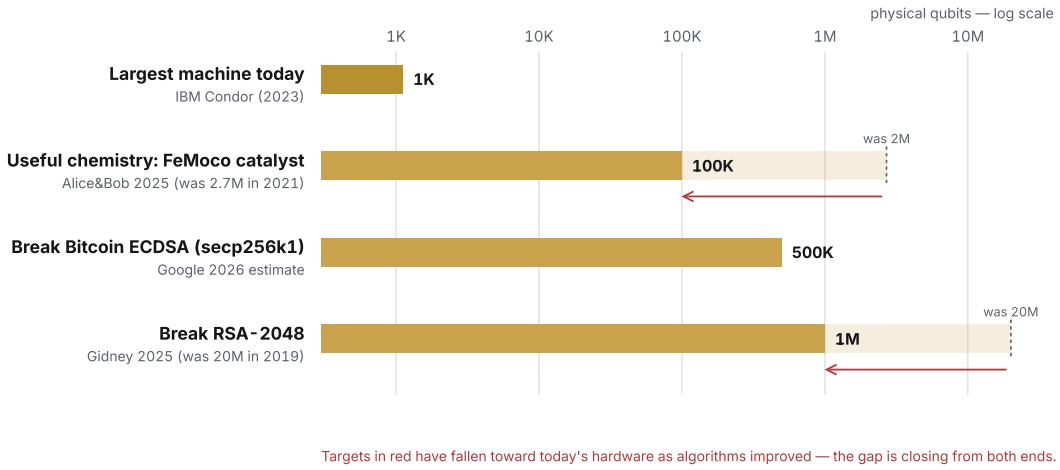

The most striking case: in 2019, breaking RSA-2048 encryption was estimated to need 20 *million* physical qubits. In May 2025 the same researcher (Google's Craig Gidney) revised that to **under 1 million** — a 20× reduction with no change in hardware, purely from cleverer algorithms. The machine still doesn't exist. But the bar for building it keeps dropping.

4 Our pacing estimate: three clocks, not one

"When will quantum computing be useful?" has no single answer because "useful" is three different milestones. Here is our read of the credible range, separating the vendors' (optimistic, historically slip-prone) roadmaps from independent and expert consensus.

EXHIBIT 4 · THE PACING ROADMAP — WHEN EACH KIND OF "USEFUL" ARRIVES

Gold bars = vendor / optimistic targets; grey = independent and expert-survey consensus. Dashed verticals mark government deadlines to retire today's encryption (§5).

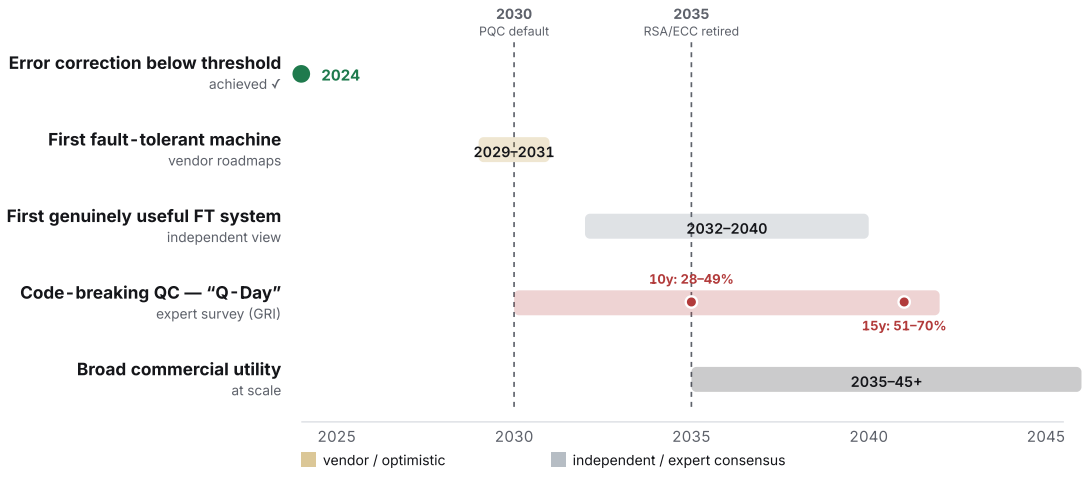

THE THREE CLOCKS — OUR SYNTHESIS OF CREDIBLE ESTIMATES

MILESTONE	WHAT IT MEANS	CREDIBLE WINDOW
Narrow scientific use	A machine does one research-interesting thing classical computers struggle with (e.g. a specific physics/chemistry simulation).	Arriving now → ~2030
Cryptographically relevant ("Q-Day")	Can break RSA-2048 / today's public-key encryption. Drives the "harvest now, decrypt later" risk in §5.	~1 in 3 by 2035; more likely than not by ~2040
Broad commercial utility	Routinely solves *profitable* problems at scale that classical machines can't.	2030s → 2040+

Where the disagreement is. The hardware builders cluster around *2029–2030* for their first fault-tolerant machines — IBM ("Starling," 200 logical qubits, 2029), Google (~2030), Quantinuum ("Apollo," 2029), IonQ (2030), and the most aggressive, PsiQuantum (~2027–28). But these are marketing roadmaps with a long, documented history of slipping — and indeed several 2025–26 targets already have. Independent voices are more cautious: the consulting firm BCG puts *full-scale* fault tolerance only after 2040, with a window of broad advantage opening 2030–2040. Our own view splits the difference: **genuine, narrow scientific value this decade; broad commercial value in the 2030s; and a meaningful, plan-for-it-now probability of code-breaking capability within ten to fifteen years.**

> **The one number that is rising fastest is the threat estimate.** An annual expert survey (the Global Risk Institute's Quantum Threat Timeline) saw its 10-year odds of a code-breaking machine jump from ~19–34% to **~28–49%** in a single year — the steepest move in the survey's history, driven largely by Gidney's cheaper-to-break result. The "useful for business" clock barely moved; the "dangerous to encryption" clock lurched forward.

5 The cryptography reckoning — banks, governments, and corporations first

This is where a still-nonexistent machine already changes decisions today. Almost everything that keeps digital life secure — web traffic (the padlock in your browser), VPNs, software updates, secure email, the certificates banks and governments trust, and the messaging that moves money between banks — rests on **public-key cryptography**, specifically RSA and elliptic-curve cryptography (ECC). Their security rests on math problems (factoring huge numbers; elliptic-curve discrete logs) that are hopelessly hard for classical computers.

In 1994, Peter Shor proved a quantum computer could solve *exactly those problems* efficiently. So a large enough quantum computer doesn't weaken this encryption — it **breaks it outright**. Crucially, the damage is asymmetric, and this asymmetry shapes the entire global response:

WHAT QUANTUM BREAKS — AND WHAT SURVIVES

CRYPTOGRAPHY	USED FOR	QUANTUM ATTACK	OUTCOME
RSA / ECC (public-key)	Web/TLS, VPNs, signatures, PKI, bank messaging, Bitcoin keys	Shor's algorithm (exponential)	**Broken**
AES-128 (symmetric)	Bulk data encryption	Grover's algorithm (quadratic)	**Weakened**
AES-256 (symmetric)	Bulk data encryption	Grover's algorithm	**Safe**
SHA-256 / SHA-3 (hashing)	Integrity, Bitcoin mining	Grover / partial	**Largely safe**

Why the split? Shor exploits deep *algebraic structure* unique to public-key math, yielding an exponential break. Grover, against the "unstructured" symmetric ciphers, offers only a square-root speedup — so simply doubling the key length (to AES-256) fully restores safety. **The world's entire migration effort therefore targets public-key crypto specifically.**

Why the threat is already here: "harvest now, decrypt later"

You do not need to wait for the quantum computer for the damage to begin. Adversaries can **record encrypted data today** and simply store it until a machine exists to decrypt it. Anything that must stay secret for a decade or more — state and military secrets, health and biometric records (which can never be "changed"), financial records, long-term contracts, trade secrets — is *already* exposed to a computer that won't be built for years. Security professionals capture this with **Mosca's inequality**: if the years your data must stay secret, *plus* the years it takes you to migrate, exceeds the years until Q-Day, you are *already* too late to start.

EXHIBIT 5 · "Q-DAY" — WHEN A CODE-BREAKING MACHINE ARRIVES, BY EXPERT ESTIMATE

Probability that a cryptographically-relevant quantum computer exists by a given year. Band = range of expert opinion. Dashed verticals = government deadlines to retire RSA/ECC.

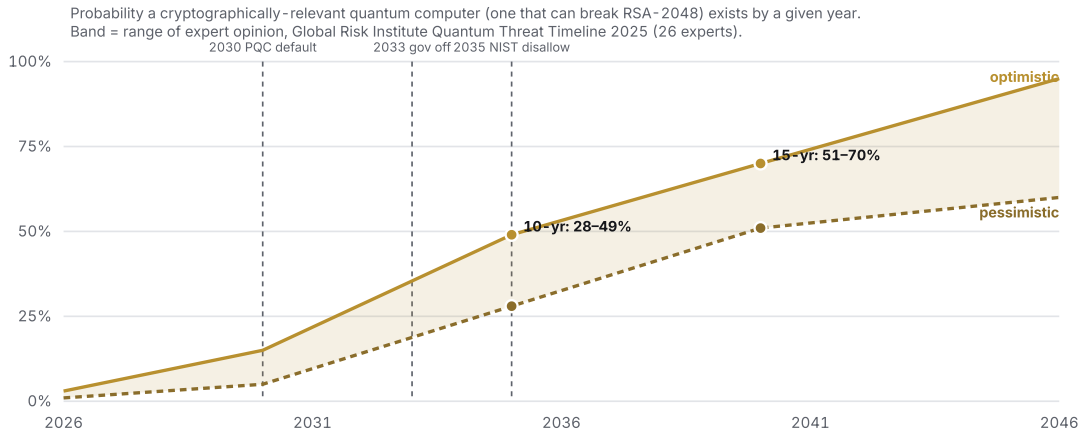

Anchors at 10-yr (2035) and 15-yr (2040) horizons are surveyed; intermediate points interpolated. One year earlier the 10-yr odds were ~19–34%.

The defense is already being deployed

The response — **post-quantum cryptography (PQC)**, new algorithms believed safe against both classical and quantum attack — is further along than most people realize. In August 2024 the U.S. standards body NIST finalized the first three PQC standards (built on "lattice" math that even a quantum computer struggles with). Governments have set hard deadlines, and the migration is quietly happening on the live internet right now.

THE GLOBAL MIGRATION OFF TODAY'S ENCRYPTION — SELECTED MILESTONES & DEADLINES

WHEN	WHO	WHAT
Aug 2024	NIST (US)	First 3 PQC standards finalized (ML-KEM, ML-DSA, SLH-DSA)
2023–25	Apple, Google, Signal, Cloudflare	Quantum-safe encryption shipped to billions of devices. By Oct 2025, **over half of all human web traffic** through Cloudflare was already quantum-protected — up from ~2% in early 2024.
2025	Bank for Int'l Settlements + SWIFT	"Project Leap" tested quantum-safe encryption in live-like interbank payment systems; BIS issued a financial-sector migration roadmap.
2030	US gov (NSA CNSA 2.0)	Quantum-safe crypto the default for federal systems and networking
2033–35	US gov (NIST / NSA)	RSA & ECC to be *disallowed* for federal use; full quantum-resistance target. EU and UK target high-risk sectors by ~2030–31.

The takeaway for a non-specialist: the institutions with the most to lose — central banks, intelligence agencies, the companies that run the internet's plumbing — are *not* waiting for Q-Day. They have read Mosca's inequality, and they are migrating now. This is the clearest signal that the threat, while years away, is treated as real.

6 The Bitcoin angle

Because TON618 is a Bitcoin-benchmarked manager, we address Bitcoin's specific exposure directly — and put it in proportion. Bitcoin uses two kinds of cryptography with *very* different quantum exposure:

- **Digital signatures (ECDSA, the secp256k1 curve) — vulnerable to Shor.** This is what proves you own your coins. A large quantum computer could, in principle, derive a private key from a *publicly visible* public key.
- **Mining and address hashing (SHA-256, RIPEMD-160) — only weakened by Grover.** This is far more resilient. Bitcoin's mining and the basic address scheme do not face an existential quantum threat.

The practical risk centers on **exposed public keys**. A modern, never-reused Bitcoin address keeps its public key hidden behind a hash and is *not* directly Shor-exposed. But two categories are: the earliest coins (stored as raw public keys, including Satoshi's ~1.1 million BTC) and any address that has been *reused* after spending. Estimates put the share of supply in quantum-vulnerable addresses at roughly **25–33% (~4–6.9 million BTC)**. A separate, harder attack targets the brief window when a transaction is broadcast but not yet confirmed; a recent Google analysis suggested a fast future machine

might derive a key in ~9 minutes — but that requires a far more capable machine than the "at-rest" threat, and the ~10-minute block window is a tight race.

> **The good news, and the genuinely hard problem.** Bitcoin *can* migrate to quantum-safe signatures via a software upgrade (soft fork), and developers are already drafting the path: proposals like **BIP-360** (a new quantum-resistant address type) and, in April 2026, the more contentious **BIP-361** (which would *freeze* ~5.6M BTC sitting in vulnerable addresses unless owners move them). The hard problem is political, not technical: lost coins — including Satoshi's — can never be moved by their owners, forcing an eventual, divisive "freeze-or-burn" debate that pits Bitcoin's immutability against the network's defense. This is a multi-year governance question, not a 2026 emergency.

Our assessment: the quantum threat to Bitcoin is **real, specific, and manageable on the same ~10–15-year clock as the threat to banks and governments** — and Bitcoin, unlike a paper contract sitting in an archive, can upgrade. We treat it as a monitorable tail risk, not a thesis-breaker (see §7). Notably, BlackRock added a quantum-risk disclosure to its Bitcoin ETF filing in 2025 — a sign the risk is now mainstream, not that it is imminent.

7 Investment relevance & the TON618 read

Quantum computing presents an investor with a classic shape: a genuinely large prize, an honest decade-plus of uncertainty on timing, and a set of securities that have already run far ahead of the fundamentals. We separate the *opportunity* from the *portfolio risk*.

The opportunity — a barbell, not a basket

The listed pure-plays have delivered spectacular returns and now price in a great deal of success. Combined, they earn ~$193M of revenue against ~$52B of market value — about **268× sales** — and none is profitable (combined losses ~$1.3B). The *entire* global quantum industry booked roughly $1B of revenue in 2025. This is a field being valued on a 2035 outcome.

EXHIBIT 6 · THE VALUATION GAP — MARKET VALUE VS. REVENUE OF THE LISTED PURE-PLAYS

Market capitalization vs. trailing revenue (log scale). The revenue bars are the thin dark stubs. Prices as of 24 June 2026.

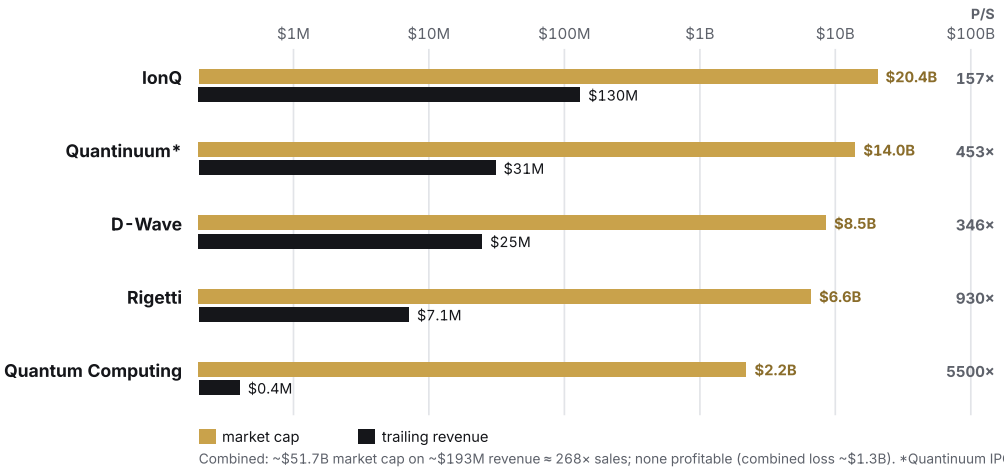

Combined: ~$51.7B market cap on ~$193M revenue ≈ 268× sales; none profitable (combined loss ~$1.3B). *Quantinuum IPO'd Jun 2026.

THE QUANTUM-EXPOSED UNIVERSE — LIVE SNAPSHOT, 24 JUNE 2026

COMPANY	PRICE	MKT CAP	52-WK RANGE	QUANTUM EXPOSURE
Pure-plays				
IonQ (IONQ)	$54.65	$20.4B	$25.89–84.64	Trapped-ion; the largest pure-play
D-Wave (QBTS)	$23.24	$8.5B	$12.75–46.75	Annealing (a different, non-universal paradigm)
Rigetti (RGTI)	$19.86	$6.6B	$10.30–58.15	Superconducting; no dated fault-tolerance target
Quantum Computing (QUBT)	$9.92	$2.2B	$6.18–25.84	Photonics; negligible revenue
Arqit (ARQQ)	$28.25	$0.5B	$11.52–62.00	Quantum-safe encryption (software)
Embedded optionality				
IBM	$260.00	$244B	$212–332	Clearest big-cap roadmap (Starling 2029); profitable
Alphabet (GOOGL)	$344.90	$2.0T	$162–409	"Willow," the QEC leader; rounding error on the P&L
Microsoft (MSFT)	$365.50	$2.7T	$356–555	Topological bet (contested); Azure Quantum cloud
Honeywell (HON)	$225.68	$143B	$187–248	Majority owner of Quantinuum (the fidelity leader)
Nvidia (NVDA)	$200.55	$4.86T	$142–237	"Picks & shovels" — quantum-classical software, control systems

How we frame it. The cleanest way to own the theme is a *barbell*: the prize is overwhelmingly likely to be captured inside trillion-dollar platforms (Google, IBM, Microsoft) and their suppliers (Nvidia, control-electronics and cryogenics vendors), where quantum is cheap, diversified, *profitable* optionality you are barely paying for — paired, if at all, with small, deliberately-sized positions in the pure-plays as high-beta lottery tickets, sized for the real risks of dilution (these firms fund losses by issuing stock) and roadmap slippage. We would be wary of the pure-play basket as a core holding at 268× sales; the more durable trade may be the *defensive* side — the post-quantum-security vendors and the consultancies running the migration, whose revenue is driven by deadlines (2030–2035) that are far more certain than Q-Day itself.

The portfolio risk — and why we are not losing sleep

- **To the fund's Bitcoin mandate,** quantum is a genuine but distant tail risk (§6), on the same ~10–15-year clock as the threat to the entire banking and government system — and Bitcoin can upgrade its cryptography, which a paper record in an archive cannot. We monitor it: the live signals are the BIP-360/361 process, the share of supply in vulnerable addresses, and any sharp jump in *logical*-qubit counts (not the physical-qubit headlines).
- **The asymmetry cuts the right way for us.** A credible acceleration of the threat timeline would damage *every* incumbent financial and technology asset that relies on RSA/ECC — arguably more than it damages a Bitcoin network that is actively building its own fix.
- **What would change our mind:** a sustained jump in logical-qubit counts (from dozens toward hundreds with low error), a fault-tolerant machine shipping *ahead* of the 2029 vendor targets, or a credible public demonstration of Shor's algorithm breaking a non-trivial key. Project Eleven's "Q-Day Prize" — awarded in April 2026 for breaking a *15-bit* key — is a useful tripwire, but 15 bits is a universe away from Bitcoin's 256.

This note expresses no price target on any security or on Bitcoin. It is a thematic assessment intended as one input to the fund's understanding of a long-horizon technology and the risks and opportunities around it, consistent with our mandate to outperform spot BTC over multi-year horizons.

DISCLOSURES & CFA-STANDARD NOTES

As-of date. All hardware milestones, expert estimates, and prices are as of 24 June 2026 and will drift. Quantum hardware and the related equities move quickly; this is a point-in-time read. **Separation of fact and opinion.** Qubit counts, error rates, the dates of demonstrations, the published resource estimates, the NIST/government deadlines, and the live prices are facts drawn from the cited primary sources; the pacing windows ("three clocks"), the barbell framing, the hype/substance grades, and all forward-looking judgments are TON618 Capital opinion and may prove wrong.

Method & data. Synthesized from primary sources (peer-reviewed papers in *Nature* and *Physical Review Letters*; arXiv preprints; NIST, NSA, CISA and BIS publications; the Global Risk Institute expert survey; company roadmaps and newsrooms) and reputable secondary coverage, gathered June 2026. Equity quotes: Schwab market-data API, 24 June 2026 (intraday). Pure-play revenue/loss figures are most-recent fiscal-year company reports as compiled in secondary roundups; market caps are live and move daily. Charts are illustrative: Exhibit 5's intermediate points are interpolated between surveyed 10- and 15-year horizons; modality series in Exhibit 1 are not directly comparable (see chart notes).

Key uncertainties. Vendor roadmaps have a documented history of slippage and are treated as optimistic. The physical-to-logical qubit ratio, the true difficulty of fault-tolerant scaling, and the date of any "Q-Day" are genuinely unknown; credible experts disagree by a decade or more. Microsoft's topological-qubit ("Majorana 1") claim remains scientifically contested as of mid-2026 and is treated here as unproven. Several March–April 2026 items (Google's Bitcoin-ECC estimate, the Project Eleven award, BIP-361) are recent and partly sourced from reputable secondary outlets.

Conflicts. TON618 Capital holds long digital-asset exposure (Bitcoin and Bitcoin-linked instruments). The fund may hold positions in securities mentioned. This note is research for information purposes only — it is not an offer, solicitation, or recommendation to buy or sell any security, fund interest, or digital asset, and it is not investment, legal, or tax advice. It contains forward-looking estimates that may prove wrong. Past performance is not indicative of future results. Digital assets and early-stage technology equities are highly volatile and may result in total loss of capital. Prepared in keeping with CFA Institute Standards of Professional Conduct (diligence, reasonable basis, fair representation, and disclosure of conflicts).

SELECTED SOURCES

1. Babbush / Troyer / Hoefler et al., "Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage," *Communications of the ACM* / arXiv:2307.00523 (2023) — quadratic speedups insufficient for practical advantage.
2. Quantum-chemistry resource estimates: Google/collaborators FeMoco estimate (2021, ~2.7M qubits); Alice & Bob algorithmic reduction to ~100k (2025), via *Chemical & Engineering News*, Nov 2025.
3. S. Aaronson, "Shtetl-Optimized" (scottaaronson.blog) — credible application classes; "qombies" and the noise-driven false-advantage critique.
4. Google Quantum AI, "Willow" / below-threshold error correction: *Nature* 638, 920 (2024), s41586-024-08449-y; blog.google. "Quantum Echoes" verifiable advantage (Oct 2025); failed classical rebuttal arXiv:2604.15427 (Apr 2026).
5. Quantinuum "Helios" (98 physical / 48 logical, 99.921% two-qubit fidelity), Nov 2025; IonQ 99.99% two-qubit prototype, Oct 2025; Harvard/MIT/QuEra integrated FTQC, *Nature* s41586-025-09848-5 (Nov 2025).
6. Hardware roadmaps: IBM "Starling" 2029 / "Blue Jay" 2033 (IBM Newsroom, 10 Jun 2025); Google Quantum AI roadmap; Quantinuum "Apollo" 2029; IonQ 2030; PsiQuantum ~2027–28.
7. Shor's resource estimates: Gidney & Ekerå, "...20 million noisy qubits," arXiv:1905.09749 (2019); Gidney, "...less than a million noisy qubits," arXiv:2505.15917 (2025). ECC/secp256k1: Roetteler et al., ASIACRYPT 2017 (arXiv:1706.06752); Google Quantum AI Bitcoin-ECDLP estimate (2026).
8. PQC standards & mandates: NIST FIPS 203/204/205 (13 Aug 2024); NIST IR 8547 (deprecate RSA/ECC after 2030, disallow after 2035); NSA CNSA 2.0; US NSM-10 (2022); EU PQC roadmap (2025); UK NCSC migration timelines (2025). Industry: Cloudflare PQC report (2025), Apple iMessage PQ3 (2024), BIS Project Leap / BIS Papers No. 158 (2025).
9. Q-Day timeline: Global Risk Institute / evolutionQ, "Quantum Threat Timeline Report 2025" (26 experts; 10-yr 28–49%, 15-yr 51–70%).
10. Bitcoin: Deloitte (vulnerable-supply estimate ~25%); BIP-360 (P2QRH) and BIP-361 (freeze proposal, Apr 2026); Project Eleven "Q-Day Prize" (15-bit key, Apr 2026); BlackRock IBIT quantum-risk disclosure (2025).
11. Market data & valuations: Schwab market-data API (prices, 24 Jun 2026); company fiscal-year filings and secondary earnings roundups for revenue/loss; Quantinuum IPO (Jun 2026); McKinsey Quantum Technology Monitor 2026 and BCG quantum forecasts for market-size context.