QUANTUM × CRYPTO · UPDATE ONE · RE-SCORED AUGUST 6, 2026

The Tripwires: A Six-Week Re-Score

In June we closed our first quantum assessment with three pre-registered tripwires — the observations that would change our mind about the timeline. Six weeks later the period splits three ways: **the lab sped up, the tape gave up, and the chain hasn't decided.** This note re-scores our own scoreboard, line by line, and then re-arms it.

August 6, 2026 · IONQ $39.72 · QBTS $19.41 · RGTI $16.53 · QUBT $8.78 (Aug 6 closes) · BTC ≈$64,650 · Follow-up to How Close Is Quantum Computing to Being Useful? (v1.0, Jun 28, 2026)

1 of 3	**70**	**−17%**	**$15M**
v1.0 tripwires now partially triggered — the logical-qubit count	logical qubits in IBM's Jul 30 advantage claim, vs. 48 at the June frontier; ~1,200 breaks Bitcoin's curve	the pure-play quantum basket since Jun 24 — through the field's best-ever news cycle	pledged for Bitcoin's quantum defense, against ~$446B of exposed coins

Answer first: **the timeline is accelerating — modestly, measurably, and on exactly the axes we said to watch — and the equity market is not signaling imminence; if anything it is signaling a hangover.** Those two sentences point in opposite directions, and the distance between them is this note's

subject: the lab and the tape have decoupled, and only one of them is evidence about Q-Day.

Six weeks ago we wrote that quantum computing had produced zero economically useful results, that the error-correction threshold crossing of 2024–25 was real and foundational, and that the honest pacing ran on three clocks — narrow scientific value now, commercial value in the 2030s, code-breaking at roughly one-in-three odds within ten years. We closed with three tripwires that would change our mind. This note goes back and scores them.

The re-score: **one tripwire is partially triggered.** On July 30, IBM and University of Chicago researchers claimed a verified quantum advantage on a computation using 70 logical qubits — a step from “a few dozen” toward the “hundreds” we flagged as the threshold of concern, and the strongest of three quantum-advantage claims IBM coordinated that day. The claim is fresh, the classical-computing community has not yet had its usual turn at bat, and the history of this field is that supremacy claims get caught. But our v1.0 marquee stat — “zero economically useful results” — now carries an asterisk for the first time, and intellectual honesty requires us to print that.

Meanwhile, the market answered the question “are quantum stocks signaling an imminent breakthrough?” with an emphatic no. Since v1.0’s pricing date the four pure-plays fell between 9.5% (QUBT) and 26% (IonQ); from their early-June highs, 29–43%. D-Wave announced a genuine *Nature* hardware result on August 5 and its stock fell 9% the following session. Even Arqit, the post-quantum-security name we flagged as the “defensive side” of the trade, fell 25%. The pure-play basket is worth roughly $30 billion today against trailing revenue still in the low hundreds of millions — revenue that is, to be fair, growing fast: IonQ alone printed $80 million in Q2, so the sales multiple is compressing from both directions. The de-rating tells you the December–January speculative episode is unwinding; it tells you nothing about Q-Day.

On Bitcoin, the story of the summer is that the fix moved from “monitorable” to **funded and specified — but not activated.** BIP-360 was merged into the official BIPs repository in February and redesigned into a cleaner construction (Pay-to-Merkle-Root). On July 23, nine institutions — BlackRock, Coinbase, Strategy, Fidelity Digital Assets, Galaxy, Block, Blockstream, Anchorage, and ARK Invest —

pledged $15 million over three years to fund the open-source work, explicitly disclaiming any governance role. Coinbase's independent Quantum Advisory Council put the canonical exposure at roughly 6.9 million BTC with exposed public keys, 1.7 million of them Satoshi-era. A functioning post-quantum testnet exists. What does not exist is a single activated consensus change — and the freeze-the-vulnerable-coins debate (BIP-361) is where the real fight will happen. **The binding constraint on Bitcoin's quantum defense has migrated from physics to governance.**

Our posture: we nudge the code-breaking odds modestly upward — from roughly one-in-three within ten years toward something closer to 35–40% — and we are explicit below about why we are not jumping further. The three-clocks framework survives intact. The cryptography clock ticks slightly faster. And because a re-score is only worth publishing if it re-arms, the note ends the way v1.0 did: four new tripwires, written down before the evidence arrives.

SECTION 01 · THE DOCKET

The re-score, line by line

Pre-registration only means something if the scoring is mechanical — no reinterpreting the language after the fact. v1.0 §7 said, verbatim: *"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."* Score each.

01	A sustained jump in logical-qubit counts, from dozens toward hundreds	PARTIALLY TRIGGERED	Frontier moved 48 → 70 (claimed) on Jul 30. One paper, one week old, classical-review cycle not yet run. Warm — not conclusive.
02	A fault-tolerant machine shipping ahead of the 2029 vendor targets	NOT TRIGGERED	No machine early. But the 2029 roadmaps are holding, not slipping — and slippage was our base case.
03	A credible Shor demonstration against a non-trivial key	NOT TRIGGERED	State of the art remains a 15-bit key (Apr 2026). Bitcoin's keys: 256 bits. The distance is exponential, not linear.

Why "partially" and not "triggered"

The IBM / University of Chicago paper is the strong one of the three July 30 claims: an explicit quantum-advantage argument backed by complexity-theoretic hardness, a computation completed in about 15 minutes that the authors argue would take prohibitive runtimes classically — executed on 70 logical qubits with a new error-

correction method (arXiv:2607.25941). The Algorithmiq and Qedma results make more empirical claims — “the classical methods we tested became unreliable in this regime” — a weaker standard, and the fact-checking literature has already split the three claims along exactly this line.

But the claim is one week old at this writing; the pattern since 2019 is that classical simulation teams take weeks-to-months to respond, and they have caught every previous claim at least in part; and 70 is still “dozens,” not “hundreds.” The tripwire said *sustained*. One paper is not sustained. But 48 → 70 at the frontier, inside a quarter, with an advantage claim attached, is precisely the shape of movement we said to watch.

For calibration: Google’s March 2026 resource estimate puts the break of ECC-256 — Bitcoin’s curve — at fewer than ~1,200 logical qubits. The frontier just moved from 48 to a claimed 70. The distance is a factor of ~17, down from ~25 in June. That is what “partially triggered” means in numbers.

On tripwire 2, the subtler change is worth recording: the 2029-vintage roadmaps are so far *holding rather than slipping*. IBM’s intermediate module (Kookaburra) remains scheduled for 2026 on the path to Starling in 2029; PsiQuantum, the most aggressive roadmap in the field, signed a $100 million CHIPS Act letter of intent in May. Government money arriving on schedule is not proof of technical schedule, but it is the opposite of a funding stall — and our v1.0 prior leaned on the slippage history. On tripwire 3, nothing: 15 bits is a universe away from 256, exactly as it was in June.

0*

Economically useful results a quantum computer has produced that a classical computer cannot — v1.0's marquee stat, now carrying its first serious asterisk. The asterisk is IBM's July 30 claim. It resolves when the classical-simulation community takes its turn — and we will score it publicly either way.

SECTION 02 · THE LAB

Six weeks of science: the densest run since Willow

Between v1.0's publication and this note — 39 days — the field produced four *Nature*-tier results and one coordinated advantage announcement. For pacing purposes the cadence is the datum; no single result decides anything, but the clustering is unlike any six-week window since December 2024.

JUL 15

Quantinuum: first universal topological gate set

Non-Abelian anyons on the H2 processor — a 54-qubit entangled state realizing S3 symmetry, in *Nature*. Topological encodings are the "protected by geometry" route to error resistance; a universal gate set is the checkbox that matters.

JUL 29

HRL Laboratories: an autonomous silicon QPU at 4 kelvin

18 silicon spin qubits run by a custom cryogenic CMOS control chip, in *Nature*. Small qubit count; large implication — control electronics living at cryogenic temperature next to the qubits is one of the two or three engineering walls between hundreds of qubits and millions. A concentration disclosure a careful reader deserves: IBM signed a definitive agreement to acquire HRL on July 23, six days before this paper — meaning IBM has a hand in four of the five headline results of the window. The cadence is real; its independence is thinner than the mastheads suggest.

JUL 30

IBM × 3: "the quantum advantage era"

Three coordinated claims with Algorithmiq, Qedma, and the University of Chicago — the docket's tripwire 1, discussed above.

AUG 5

D-Wave: a gate-model credential in *Nature*

A fast (~500 ns), high-fidelity (~99.9%) two-qubit entangling gate on a superconducting dual-rail architecture with native hardware-level error detection — with simulations (D-Wave's own, labeled as such) suggesting ~10× logical-error reduction per error-correction increment.

Read against v1.0's framework: none of this produces commercial value today, and all of it is exactly the kind of below-the-threshold plumbing — gates, control, encodings, verification — that a field crosses on the way from "threshold crossed" to "machine exists." In June we wrote that the foundation was finally real. The update is that people are now building on the foundation, quickly, in public, in the peer-reviewed literature rather than the press-release channel. That is what acceleration looks like from the outside.

And while all of this was printing, the stocks that supposedly price this future were doing precisely the opposite.

Bitcoin's real countdown metric: logical qubits. 70 claimed. ~1,200 needed.

Error-corrected (logical) qubits demonstrated at the frontier, vs. the estimated requirement to break Bitcoin's elliptic-curve cryptography — the qubit-count headlines you read are physical qubits, which do not count

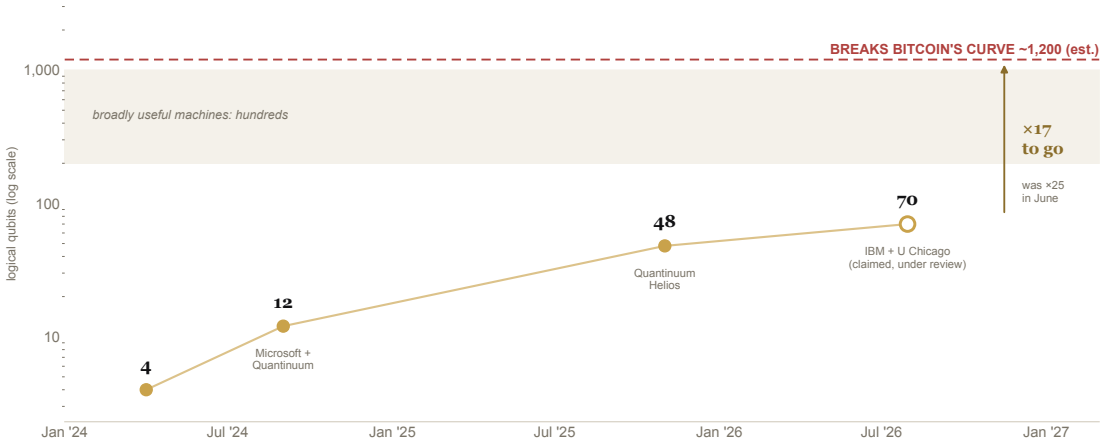

Frontier demonstrations of fully error-corrected logical qubits; IBM figure as claimed Jul 30, 2026, pending classical-simulation review. Requirement estimate per 2026 research; sourcing in the note appendix.

SECTION 03 · THE TAPE

The stocks: not a breakthrough signal — a hangover

If a breakthrough were truly close, the market should smell it first — that is the efficient-markets intuition, and it deserves a test rather than an assumption. Here is the test: **the quantum stocks are not signaling an imminent breakthrough; they are signaling the unwind of a speculative episode, and they de-rated straight through every result in the log above.**

The best six weeks of quantum news ever printed. The stocks sold off anyway.

Listed pure-play quantum computing companies, indexed to 100 on Jun 24, 2026 · through Aug 6 close

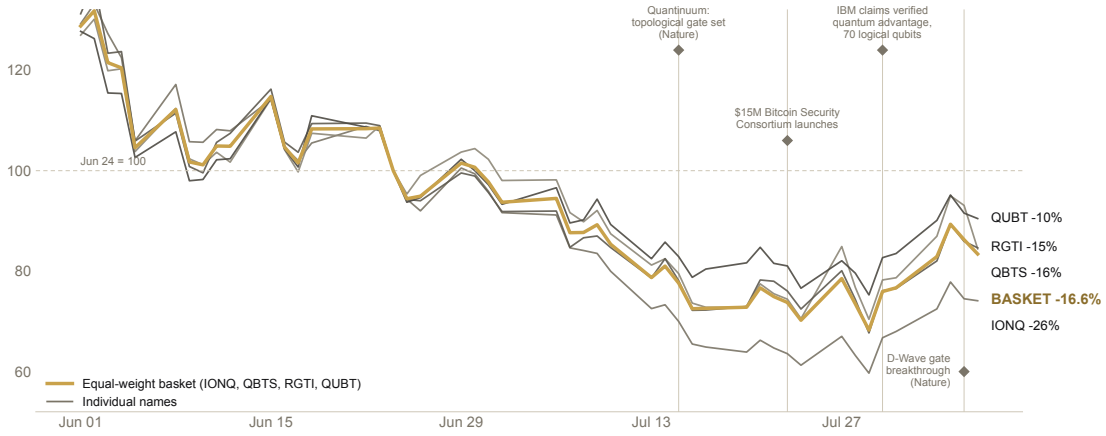

Indexed daily closes. Diamonds mark peer-reviewed results and industry events; details and sourcing in the note appendix.

Ticker	Company	Jun 24	Aug 6	Change	Mkt cap
IONQ	IonQ	$53.60	$39.72	**−25.9%**	$14.8B
QBTS	D-Wave Quantum	$22.98	$19.41	**−15.6%**	$7.2B
RGTI	Rigetti Computing	$19.53	$16.53	**−15.4%**	$5.5B
QUBT	Quantum Computing Inc.	$9.70	$8.78	**−9.5%**	$2.0B
—	Equal-weight basket	100	≈83.4	**−16.6%**	~$29.5B
ARQQ	Arqit (post-quantum security)	$28.64	$21.47	**−25.0%**	—

From the early-June highs the drawdowns are deeper still: IonQ −43% from its June 1 close, Rigetti −36%, D-Wave −33%, QUBT −29%. And the single most instructive tape print of the window: **D-Wave announced a genuine *Nature* hardware breakthrough on August 5 and fell roughly 9% the following session** — a sell-the-news reaction in which coverage framed IBM's and Google's gate-model progress as a competitive *threat* to D-Wave rather than a rising tide. When a *Nature* paper is a down-9% event, the marginal buyer is not doing physics.

Three observations for the permanent record

The de-rating and the news flow are causally disconnected. The slide began in early June, weeks before the July results, and did not pause for any of them. This is the December–January speculative episode unwinding on its own financial-market clock — rate expectations, risk appetite, supply of new shares — not a verdict on the science.

Even the defensive side fell. v1.0 suggested the durable trade might be the post-quantum-security vendors, whose revenue is driven by government migration deadlines (2030–2035) far more certain than Q-Day. Arqit fell 25% anyway. Deadline-driven revenue is a fundamentals argument, and this tape was not trading fundamentals in either direction.

Therefore: do not read equity prices as a Q-Day sensor — in either direction. The same people citing the 2025 melt-up as evidence the breakthrough was near would now have to cite this drawdown as evidence it receded. Neither inference was ever valid. The sensor that works is the docket's: logical-qubit counts in the peer-reviewed literature.

Valuation postscript: the multiple is compressing from both directions — prices down ~17% while revenue accelerates (IonQ's Q2 print: $80.0M, +287% year over year, nearly half of v1.0's combined four-company trailing figure in a single company-quarter; the other three report through August). The basket remains priced off a 2035 outcome — cheaper than June, not cheap — and the honest multiple is now severalfold lower than the 268× sales we printed in June. The v1.0 barbell framing — platform incumbents and suppliers as the core, pure-plays as deliberately-sized lottery tickets, the post-quantum migration trade as the sleeper — survives unchanged.

SECTION 04 · THE CLOCK

Is the timeline accelerating? Three lines of evidence, one answer

Strip the tape out, then. What remains is the evidence that binds: the published resource estimates, the demonstrated hardware, and the behavior of the people

whose job is to be right about this. Three lines, then the synthesis.

Line 1 — the resource estimates keep falling. In 2019 the canonical estimate for breaking RSA-2048 was ~20 million physical qubits; Gidney's May 2025 revision cut it to under 1 million — 20× from algorithms alone. The new datum is Bitcoin-specific: Google's March 2026 work puts an attack on ECC-256 at fewer than ~1,200 logical qubits and under 500,000 physical qubits, with runtimes measured in minutes on a future fault-tolerant machine. For reference, the 2022 academic estimate (Webber et al.) for breaking a Bitcoin key inside a day was ~13 million physical qubits. The target is walking toward the hardware at roughly an order of magnitude every three to four years, with no hardware improvement required.

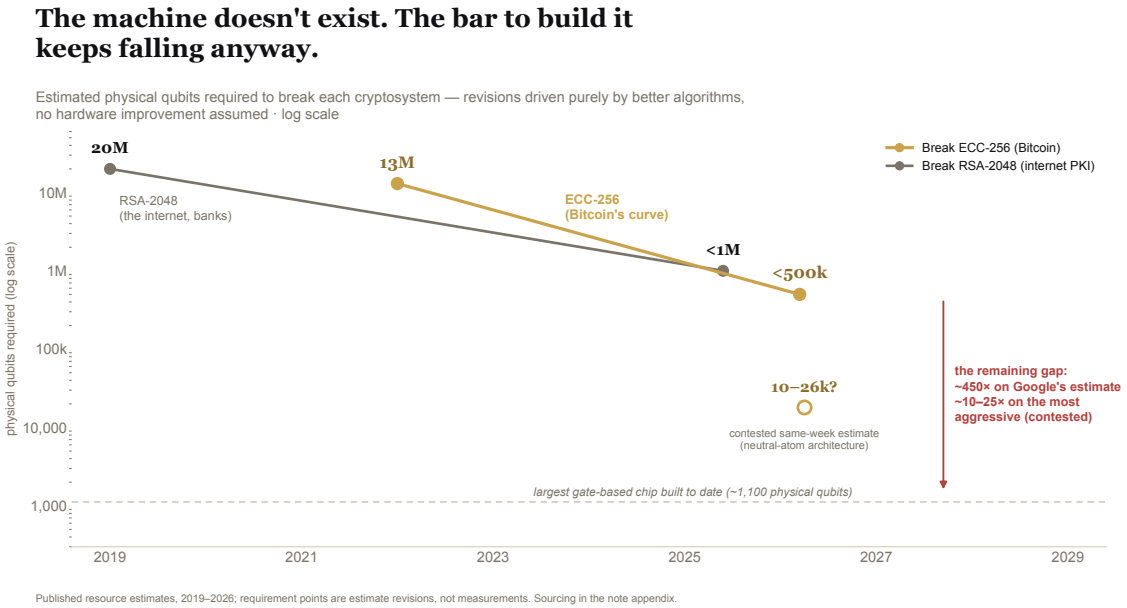

An honest complication, flagged before a reader flags it for us: Google's number is not even the aggressive end of the same-vintage literature. A Caltech/Oratomic estimate published the same week (March 2026, neutral-atom architecture, different and contested assumptions) puts an ECC-256 break at 10,000–26,000 physical qubits — twenty to fifty times below Google's figure. We treat Google's as the working estimate because its assumptions sit closer to the architectures actually being scaled today, but the spread is the point: the live range of expert estimates now runs from ~10,000 to ~500,000 physical qubits, and even its conservative end keeps falling. The direction of travel is not in dispute; only the speed is.

Line 2 — the builders moved their language, then moved their claims. IBM's CEO predicted the first real-world quantum advantage would land in 2026 (April, on the record); IBM then claimed it (July 30). The same CEO, in the same July 30 CNBC appearance, put crypto-relevant capability at three to four years — the most aggressive timeline any major-vendor executive has attached to the cryptography milestone. Nobel laureate John Martinis — who built Google's early superconducting program — spent the spring arguing Bitcoin will be among the *first* real-world targets, because a derived private key monetizes instantly and irreversibly. Discount executive timelines as marketing, as we always do; but note that in June the marketing said "advantage soon," and by August the papers said "advantage claimed." The gap between talk and print shrank, and that gap is itself a pacing indicator.

Line 3 — the deadline-setters are not waiting. Google's internal post-quantum migration targets 2029. NIST's horizon remains 2035. Coinbase's CEO works to 2029. The U.S. government PQC deadlines from v1.0 stand. Nobody with fiduciary exposure to being wrong has relaxed a deadline in six months; several have tightened operational work against them.

The synthesis. v1.0 put a code-breaking machine at roughly one-in-three within ten years, up from one-in-five a year earlier. On the evidence above we mark that to **~35–40% within ten years** — a nudge, not a jump, and the reasons for restraint are load-bearing: (a) 70 logical qubits is a claim one week old from one group, not a sustained frontier; (b) every previous advantage claim has been at least partially matched classically, and the response cycle has not run; (c) the remaining engineering wall — from tens of logical qubits to the ~1,200 that threaten ECC-256, with the required gate depth and magic-state throughput — is precisely the part of the problem that has never been demonstrated at any scale; and (d) vendor roadmaps holding for six months is weak evidence against a slippage base rate built over fifteen years. The three-clocks framework from v1.0 is unchanged: narrow scientific value is arriving on schedule (arguably early), commercial value remains a 2030s story, and the cryptography clock — the only one Bitcoin cares about — now ticks slightly faster than it did in June.

What Bitcoin is actually doing about it: funded, specified, not activated

v1.0 treated Bitcoin's quantum exposure as a monitorable tail risk and listed the BIP-360/361 process as the live signal. Six weeks — and one busy summer — later, here is the tangible ledger. It is more substantial than we expected, and it has a hole in the middle.

One coin in three sits on an exposed public key.

Bitcoin circulating supply by quantum exposure — an exposed public key is what a future quantum machine attacks; a hashed, never-reused address is not directly exposed

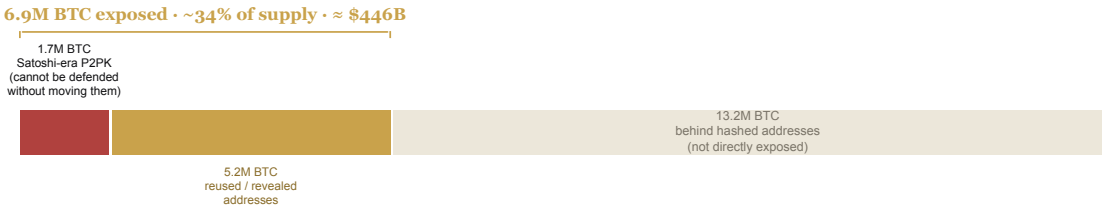

The fix exists on paper — BIP-360 was merged into Bitcoin's official proposal repository in February; a $15M industry consortium is funding the work. What does not exist: a single activated protocol change. The binding constraint has migrated from physics to governance.

The specification is real. BIP-360 — Bitcoin's first formal post-quantum proposal — was merged into the official BIPs repository on February 11, 2026. Along the way it was redesigned from the original Pay-to-Quantum-Resistant-Hash (P2QRH, June 2024) into Pay-to-Merkle-Root (P2MR): the output commits directly to a script-tree Merkle root with no internal key at all, removing the key-path exposure entirely and giving future post-quantum signature schemes a clean structural base to plug into. Merge into the repo signals documentation standards met and formal discussion open — not endorsement, not activation. But a real, reviewed, evolving specification now exists where two years ago there was a mailing-list thread.

The money is real, and correctly humble. The July 23 Bitcoin Security Consortium — nine institutions, $15 million over three years, members directing

funds independently, the consortium taking no role in governance or protocol decisions. Read both halves. The signatory list is the institutional Bitcoin complex essentially in full — the ETF issuer, the largest corporate holder, the largest exchange, the custody banks. And $15 million against a ~$446 billion exposed-coin problem is a signal, not a solution: roughly what one of these firms spends on a Super Bowl ad, aimed at a problem they all describe as existential-if-mishandled. The correct reading is that the consortium exists to fund engineers without touching governance — because governance is the part nobody can buy.

The engineering is real. Coinbase has dedicated engineering headcount to BIP-360 and the broader migration path, and stood up an independent Quantum Advisory Council whose report is now the canonical exposure census: approximately 6.9 million BTC in UTXOs with exposed public keys, including about 1.7 million BTC in Satoshi-era pay-to-public-key outputs — coins that cannot be defended by any protocol change short of freezing or moving them, because their public keys have been on-chain for fifteen years. BTQ Technologies' Bitcoin Quantum testnet (v0.3.0, March 2026) provides the first functional validation of the P2MR design running end-to-end. A phased migration proposal (new-output support → deprecation of vulnerable spends → the freeze question) is under active discussion on bitcoindev.

The hole in the middle: activation. Zero consensus changes have activated. None is scheduled. And the hard question — BIP-361's territory, what to do about coins that will never migrate, including Satoshi's — is not an engineering question at all. Freeze them and you have retroactively confiscated property to defend a principle; leave them and you have accepted that roughly a third of the supply becomes a bug bounty the day a machine exists, with the fire-sale and confidence dynamics that implies. Ethereum, with a living founder, a research culture organized around hard forks, and a coordinated migration already underway, does not face this shape of problem. This was Citi's May argument for why Bitcoin's quantum exposure exceeds Ethereum's despite identical curve cryptography — the binding constraint is governance, not physics — and everything that happened this summer is consistent with it. **The money and the code arrived. The decision has not.**

The Street picks a side

While the engineers built, the Street argued. v1.0 noted only that BlackRock had added a quantum risk factor to the IBIT filing; the intervening months produced a full-blown allocator debate, and it sorts almost perfectly by book.

ACTED

Christopher Wood (Jefferies, global head of equity strategy) removed his entire 10% bitcoin model-portfolio allocation in January, splitting it into physical gold and gold miners — calling quantum an existential threat to the store-of-value thesis and citing a 4–10M BTC exposure range. The single most significant datum in this section: an early institutional bull (allocated since 2020) who reversed on this specific risk.

WARNED

Citi (Alex Saunders, May): bitcoin more exposed than Ethereum on governance grounds; ~6.7–7M BTC exposed. **Chamath Palihapitiya**: bitcoin as the “honeypot” — the first target of non-state actors because a broken key monetizes instantly. **John Martinis** (Nobel laureate): an early target, on the record repeatedly this spring.

BOUNDED

Bernstein (Gautam Chhugani, April): “neither existential nor novel” — a manageable 3–5-year upgrade cycle; mining safe; the 1.7M Satoshi-era coins the real exposure. **Benchmark** (Mark Palmer, January): “real but distant.” **ARK Invest** (March): long-term risk, not imminent. **BlackRock**: risk factor filed, but research framing quantum as one of the last “walls of worry,” with successful migration a potential *strengthening* event.

DISMISSED

Mike Novogratz (Galaxy): the network adapts; the real risk is developer discord over the fix — the governance argument wearing an optimist’s coat. **Michael Saylor** (Strategy): a decade-plus away; warns against premature protocol changes. Note that Strategy signed the consortium anyway — watch what they fund, not what they say.

The pattern to flag: with the single exception of Wood, the urgency of every participant’s assessment predicts their positioning almost perfectly. The longs find it manageable; the unexposed find it existential; the gold-adjacent find it terminal. Wood is the exception that proves the diagnostic value of the exercise — he changed his book to match his assessment rather than his assessment to match his

book. That is what makes his January move the one worth studying, whatever one thinks of its conclusion.

SECTION 07 · THE DEFENSE

The steelman: Bitcoin's case, at full strength

The section above ends on an uncomfortable observation: nearly everyone's quantum verdict predicts their book. We are a Bitcoin-benchmarked manager writing about a Bitcoin risk, which means the same critique aims at us — from both directions. The only defense is to argue the strongest version of the case we did *not* lead with, and let the tripwires adjudicate. Here is Bitcoin's defense with nothing held back.

1 Most of the exposure can defend itself today, with no fork at all.

The 6.9M BTC figure is a stock, not a fate. Roughly 5.2M of it sits in reused or otherwise revealed addresses whose owners can move to a fresh, hashed address this afternoon — an ordinary transaction, no protocol change, no governance fight, no consortium. The exposure number is therefore partially self-liquidating as the threat becomes salient. What cannot be defended this way is the ~1.7M BTC in Satoshi-era outputs — the coins with no living owner. The steelman shrinks the problem from "a third of the supply" to "the dead coins"; it does not dissolve it.

2 Hashed addresses stay safe until the moment they spend.

A never-reused modern address exposes its public key only in the seconds-to-minutes between broadcast and confirmation. Attacking that window requires a machine that runs Shor in minutes — a far taller order than the years-long at-rest attack on already-exposed keys. The practical consequence is underappreciated: Bitcoin's migration can succeed even if it starts *late*, because the majority of supply is not racing the first cryptographically relevant machine; it is racing the much stronger machine that comes years after.

3 Governance latency is not governance inability — and slowness cuts both ways.

The network that supposedly cannot upgrade shipped P2SH, SegWit (2017), and Taproot (2021): a consensus change roughly every four years, which is precisely the cadence Bernstein's 3–5-year window requires. Ossification is not a bug the quantum threat exposes; it is the security model that makes a $2 trillion bearer asset possible at all — and it protects the migration itself. A chain that could be rushed into a cryptography swap by a news cycle could also be rushed into a bad one. The governance-constraint critique is real, but it describes the *freeze decision*, not the upgrade: adding a post-quantum output type is additive and uncontroversial; only the fate of the dead coins is a fight.

4 On Q-Day, everything breaks — and Bitcoin is among the few things that can be rebuilt in place.

A cryptographically relevant quantum computer does not selectively attack secp256k1. It breaks TLS, code-signing, interbank messaging, custody infrastructure, and every archived signature on earth. Bitcoin can hard-fork its cryptography; a land registry, a signed treaty, or a decade of recorded bank traffic cannot retroactively re-sign itself. BlackRock's framing is the institutional version of this point: quantum is one of the last "walls of worry," and a completed migration would remove it — the asset that crossed the post-quantum bridge first, in public, gains a property nothing else in the legacy system can demonstrate.

5 The threat is unusually observable — Bitcoin must beat the machine, not the headline.

The countdown metric (logical qubits in the peer-reviewed literature) is public, vendor roadmaps are marketing documents published years ahead, and the national-security world's own migration deadlines (2029–2035) bracket the consensus timeline. Very few tail risks announce themselves this loudly. The window in which acting is possible is wide open and well-lit.

6 Attacker economics are, at minimum, contested.

Stealing a million coins collapses the value of the loot as it moves; a rational state actor with the first machine has quieter and richer targets than crashing a transparent ledger. We flag the honest counter in the same breath: the honeypot argument says irreversibility and instant monetization make Bitcoin uniquely attractive to *non-state* actors who don't care about preserving the system — and a public quantum theft anywhere collapses confidence everywhere, whatever the attacker's P&L. This argument reduces the probability of the worst path; it does not retire it.

> *Taken together, the defense bounds the blast radius — from "a third of supply" to the dead coins, and from "race the first machine" to "race the stronger machine" — and it establishes that the upgrade is well inside*

> *Bitcoin's demonstrated governance capacity. What it cannot do is answer the freeze question, because that question is not technical, and no amount of engineering funding buys the answer. Our assessment stands, both halves: the risk is manageable on the evidence, and the remaining constraint is governance. The tripwires below are how we avoid arguing either side of that on vibes.*

SECTION 08 · THE FUND

Investment relevance: the monitoring changed, not the position

For the fund, six weeks of evidence changed the monitoring, not the position. **The barbell survives the drawdown:** platform incumbents and their suppliers remain the core way to own the theme; the pure-plays remain lottery tickets that just got ~17% cheaper and no more certain; the post-quantum migration vendors remain the fundamentals trade, now also cheaper, with their deadline-driven revenue thesis untouched by a tape that was not trading fundamentals. **The asymmetry still cuts the right way:** a credible acceleration damages every RSA/ECC-dependent incumbent — the entire banking, government, and internet stack — at least as much as it damages a Bitcoin network that now has a specification, a testnet, and a funded engineering pipeline. What Bitcoin uniquely owns is the governance risk, and that is the correct thing for a Bitcoin-benchmarked manager to monitor. The live monitoring list: the logical-qubit frontier (currently 48 sustained / 70 claimed; alarm at ~200 sustained), the classical response to the IBM/Chicago claim, BIP-360 reference-implementation progress and any activation signaling from Core maintainers (currently none), the BIP-361 freeze debate, the share of supply still exposed (currently ~6.9M BTC; falling exposure is the system working), the Q-Day Prize key size (currently 15 bits), and consortium funding actually deployed versus pledged (currently $0 visibly deployed; the first grants are the tell). The next re-score publishes when either the classical-response verdict is in, or a tripwire triggers, whichever is first. This note states no price target and no expected return on any

asset; the quantum question enters the fund's process as a monitored tail risk and a tripwire list, nothing more.

SECTION 09 · RE-ARMED

The new tripwires — set August 2026

v1.0 closed with three tripwires and this note exists because we went back and scored them. Same discipline, re-armed. Four this time — three on the physics, and one on Bitcoin itself, because the conclusion of Section 05 (the binding constraint is governance) obliges us to watch a governance metric and not only a laboratory one. Each is written to be scoreable by a stranger.

T1	The refutation clock	ARMED *resolves by Jul 2027*	The IBM/Chicago claim (70 logical qubits) either survives twelve months of classical-simulation attack or it does not. Survives materially intact → the scoreboard's zero finally flips, and the next update is a jump, not a nudge. Substantially matched classically → the asterisk is deleted, the zero is restored, and we print that with equal prominence.

T2	A sustained logical-qubit frontier of ~200+	ARMED	Not one paper: at least two independent groups demonstrating on the order of 200 or more logical qubits with error rates consistent with deep circuits — the “dozens toward hundreds” language of v1.0 with the numbers filled in, set roughly at the geometric midpoint between today’s ~70 (claimed) and the ~1,200 estimated to threaten ECC-256.
T3	A Shor demonstration that scales	ARMED	A public break of an elliptic-curve key of 64 bits or more. Today’s 15-bit prize-winners are toy instances a laptop can simulate; a 64-bit break is the approximate threshold at which the remaining distance to 256 bits becomes an engineering roadmap rather than undemonstrated physics. The specific number is our choice, and we state it now so we cannot move it later.

T4	Bitcoin’s activation signal — two-sided	ARMED	Positive trigger: any post-quantum output type reaches a concrete activation mechanism — a proposed activation path with maintainer support and ecosystem signaling — or the measured exposed-supply series (start: ~6.9M BTC) begins falling at a rate consistent with deliberate migration. Negative trigger: an activation attempt is made and visibly *fails* — the first real evidence that the governance constraint binds in practice, not just in argument. Either direction re-scores Section 05; silence scores as silence.

The standing commitment is unchanged from v1.0: write it down before the evidence arrives. The lab will keep doing what it does; the tape will keep doing what it does; the chain will decide, or it won’t. When any of them crosses a line above, the next update will say so the way this one did — as a score, not a story.

APPENDIX

Sources & Method

Market data. Equity closes and daily aggregates from the Massive market-data REST API (*/v2/aggs/ticker/<T>/range/1/day/2026-06-01/2026-08-05* plus the Aug 6 daily bar via */v2/aggs/ticker/<T>/prev*), tickers IONQ, RGTI, QBTS, QUBT, ARQQ; observation vintage June 1 – August 6, 2026 (regular-session closes); pulled August 6, 2026. Basket = unweighted mean of the four pure-play closes, each indexed to 100 at the June 24, 2026 close (the pricing date of the v1.0 note, chosen for comparability; the chart baseline is labeled by date only). Market caps = August 6 close × shares outstanding from the same vendor’s reference endpoint (*/v3/reference/tickers/<T>*), pulled August 6: IONQ 373.3M, RGTI 332.4M, QBTS 370.4M (share-class), QUBT 225.5M shares. Every close, percentage change, market cap, and the basket figure were independently reproduced against a second market-data source (Schwab market-data API) on August 6 before publication. Note on comparability with v1.0: v1.0’s ~$52B combined cap was computed on June-24 share counts; today’s ~$29.5B reflects both lower prices

and subsequent share issuance — dilution is part of the story, and the two figures are not a pure price comparison.

Bitcoin figures. Exposed-supply census (≈6.9M BTC in exposed-pubkey UTXOs, ≈1.7M BTC in Satoshi-era P2PK) from the Coinbase Quantum Advisory Council report (2026). Circulating supply ≈20.05M BTC (CoinMetrics, early Aug 2026); 6.9/20.05 ≈ 34%. Exposed-coin market value ≈$446B = 6.9M × BTC ≈$64,650 (composite close-of-day quotes, August 6, 2026; approximate midpoint). These are estimates built on the Council's UTXO classification; classification methods vary across vendors and the ~25–35%-of-supply conclusion survives that variance.

Science claims, by primary source. IBM/University of Chicago verified-advantage claim: arXiv:2607.25941 (70 logical qubits, 2,415 logical two-qubit operations, 468 logical T-gates, ~15-minute runtime), announced July 30, 2026, with the Algorithmiq and Qedma companion claims (IBM newsroom, same day); the three claims' differing evidentiary weight follows the published fact-check literature. Quantinuum universal topological gate set: *Nature*, DOI 10.1038/s41586-026-10709-y, published July 15, 2026 (press coverage July 17). HRL Laboratories silicon QPU: *Nature*, July 30, 2026 issue (HRL release July 29); IBM's definitive agreement to acquire HRL was announced July 23, 2026 (IBM newsroom). D-Wave dual-rail gate result: D-Wave release / *Nature*, August 5, 2026; the ~10× logical-error projection is D-Wave's own simulation and is labeled as such. Resource estimates: Gidney & Ekerå 2019 (arXiv:1905.09749, ~20M physical qubits for RSA-2048); Gidney 2025 (arXiv:2505.15917, <1M); Webber et al. 2022 (*AVS Quantum Science*, ~13M physical qubits for a 24-hour ECC-256 break); Google Quantum AI, March 31, 2026 (arXiv:2603.28846, <1,200 logical / <500k physical qubits, minutes-scale); Caltech/Oratomic March 2026 neutral-atom estimate (10,000–26,000 physical qubits; contested assumptions, flagged as such in the text). Logical-qubit frontier series: Microsoft+Quantinuum 4 (Apr 2024) and 12 (Sep 2024, Azure Quantum blog); Quantinuum Helios 48 (Nov 2025 release); IBM/Chicago 70 claimed (Jul 2026, above). Largest gate-based chip: IBM Condor, 1,121 physical qubits (2023) — still the gate-based record as of August 2026. Q-Day Prize state of the art: 15-bit ECC key (Project Eleven, April 2026); no larger public break found in coverage through August 6, 2026.

Bitcoin workstream. BIP-360 merge into the bitcoin/bips repository: February 11, 2026 (PR #1670); P2QRH→P2MR redesign per BIP-360 v0.11.0 (February 2026). Bitcoin Security Consortium: July 23, 2026 announcement (nine members; $15M over three years; no governance role) — CoinDesk and member releases. BTQ Technologies Bitcoin Quantum testnet v0.3.0: March 19, 2026 release. BIP-361 / phased-migration discussion: bitcoindev mailing list, 2026.

Allocator record. Wood/Jefferies reallocation: January 16, 2026 (Bloomberg, The Block, CoinDesk). Citi digital-assets note (Alex Saunders): May 18, 2026. Bernstein (Gautam Chhugani): April 8, 2026. Benchmark (Mark Palmer): January 29, 2026. ARK Invest: March 12, 2026. Novogratz: February 3, 2026 (CoinDesk). Martinis: April 7, 2026 (CoinDesk). Krishna 2026-advantage prediction: April 30, 2026 (The Quantum Insider); the 3–4-year crypto-capability quote is from his July 30, 2026 CNBC appearance. IonQ Q2 2026 revenue $80.05M (+287% YoY): company report, August 2026; the other three pure-plays report through August, so no combined trailing multiple is restated here. Sell-side notes are paywalled; their figures are as reported in the cited coverage and are approximate.

Judgments, labeled as such. The ~35–40%-in-ten-years code-breaking odds, the tripwire thresholds (200 sustained logical qubits; a 64-bit key), the "partially triggered" score, and all forward-looking

characterizations are TON618 Capital opinion, not data. The 64-bit and 200-qubit thresholds are our own pre-registered choices, stated so they cannot be moved later.

Prior TON618 notes relied upon. How Close Is Quantum Computing to Being Useful? (v1.0, June 28, 2026) — the assessment this note re-scores; its §7 tripwire language is quoted verbatim in Section 01. An adversarial claim-by-claim verification pass was run on this note on August 6, 2026 before publication; the corrections it produced are incorporated.

Disclosures

As-of date. All hardware milestones, resource estimates, and prices are as of August 6, 2026 and will drift. Quantum hardware and the related equities move quickly; this is a point-in-time re-score of a June 28, 2026 baseline.

1 · Information only. TON618 Capital. This note is for information purposes only. Nothing here is an offer to sell or a solicitation of an offer to buy any security, fund interest, or digital asset, and nothing here is personalized investment advice or a recommendation regarding any instrument.

2 · Publisher's exclusion. All research is published solely as general, impersonal information of regular circulation. It is not tailored to the objectives or circumstances of any individual and is not issued in connection with compensation from any client. The Fund has no clients and distributes all research free of charge. On that basis it publishes in reliance on the publisher's exclusion from the definition of "investment adviser" under the Investment Advisers Act of 1940 (§202(a)(11)(D); cf. *Lowe v. SEC*, 472 U.S. 181 (1985)).

3 · Registration & conflicts. TON618 Capital is not registered as an investment adviser or broker-dealer in any capacity. The Fund is a Bitcoin fund and may hold or transact in the securities or digital assets it discusses; it holds long digital-asset exposure, and a slower or faster quantum timeline could affect the value of that exposure — which is precisely why the note pre-registers its tripwires rather than asking to be trusted. The Fund holds no position in IONQ, RGTI, QBTS, QUBT, or ARQQ as of publication. The Fund receives no compensation from any party in connection with its research.

4 · Use of AI. Artificial intelligence is used in the creation of this research. All methodology and data integrity are reviewed and approved before publication by TON618 Capital's Chief Investment Officer, Keyth Beck; errors may nonetheless occur, and readers should verify independently.

5 · CFA. This note was prepared to align with CFA Institute analytical standards (methodology only). CFA® and Chartered Financial Analyst® are registered trademarks owned by CFA Institute. That reference describes the analytical framework applied; it does not imply the note was prepared, reviewed, or authored by a CFA charterholder, and the note is not issued, reviewed, endorsed, certified, or approved by — nor affiliated with — CFA Institute.

6 · Risk & feedback. 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. Corrections and feedback are welcome — please direct them to CIO Keyth Beck at keyth@ton618capital.com.