PRESIDIO BITCOIN

Bitcoin's Quantum Readiness

Exposure, Mitigations, and Upgrade Paths

Version 1 ▪ April 2026

Key Takeaways

01

The viability and timeline of cryptographically relevant quantum computers (CRQCs) are still uncertain. While there has been recent progress in error correction and resource estimates, scaling the actual operations required for cryptographically relevant machines remains an unproven bottleneck.

02

Cryptographically relevant quantum computers (CRQCs) are not a structural threat to bitcoin consensus or its monetary foundations (e.g. issuance schedule, supply cap, etc.). The main risk to bitcoin is targeted theft of coins with exposed public keys.

03

Most bitcoin are not vulnerable at rest because their public keys are hidden until spent. However, if a CRQC existed today, an estimated ~6.5M BTC (~⅓ of supply) would be immediately vulnerable to theft due to long-exposed public keys. Notably, over two-thirds of that vulnerable supply (~4.5M BTC) is attributed to "address reuse," with much of it concentrated in a small number of large custodians, often for operational simplicity, and reducible without any protocol change.

04

The core protocol upgrade to enable quantum-safe spending is post-quantum signatures, which are already technically feasible today but come with meaningful trade-offs. Bitcoin will likely start with a conservative, optional quantum-secure spending path (e.g., SHRINCS and SHRIMPS in combination), then iterate, refine, and potentially upgrade as more efficient post-quantum schemes become available.

05

Migration capacity itself is unlikely to be a bottleneck. On-chain estimates suggest that if ~25% of block space were dedicated to migration, ~90% of bitcoin's value could move in ~4 days, and ~98% in ~3 weeks.

06

If CRQCs arrive abruptly, bitcoin has a wide range of potential interim defense options that can reduce theft risk, make migration safer, and buy time for a full upgrade (e.g., hiding public keys during spends and temporarily freezing vulnerable coins with quantum-safe recovery paths). There's also a playbook, made in conjunction with bitcoin developers, that outlines a concrete plan for rapid-response.

07

What to do with unmigrated legacy coins is an open question and the likeliest fault line for a fork, which will ultimately be resolved by the market.

08

Given bitcoin's conservatism and distributed governance, coordination and speed—not technical feasibility—are the main risks to a successful transition to post-quantum bitcoin. However, work is already underway, and developer attention is visible: for example, in bitcoin's main protocol development forum, the share of messages discussing quantum-related topics has risen steadily from ~5% in 2024 to ~50% in 2026 (through March).

Contact
Alexander Zukauskas · Presidio Bitcoin Lead
alex@presidiobitcoin.org · pbquantum.com

Context and Scope

One of bitcoin's core strengths is that it is software. This makes bitcoin easy to move, verify, and hold, but also exposes it to digital risk.

Foremost among these digital risks is a vulnerability that has been discussed since bitcoin's inception: cryptographically relevant quantum computers (CRQCs). In theory, CRQCs could break the elliptic curve cryptography that secures bitcoin, deriving private keys from public keys. This means the scope of the risk is primarily quantum-enabled theft of coins with exposed public keys. However, this does not threaten the basic functioning of the network as blocks can still be produced and core monetary rules remain intact.

Bitcoin's mitigation toolkit is broad and technically feasible today, but the path is less straightforward than in more centralized systems, as coordination can be more difficult, premature changes carry risk, and current post-quantum schemes involve meaningful trade-offs. To that end, this report maps where quantum computing stands today, bitcoin's exposure, mitigations, and how the ecosystem could coordinate an upgrade and migration under both orderly and rapid-response scenarios.

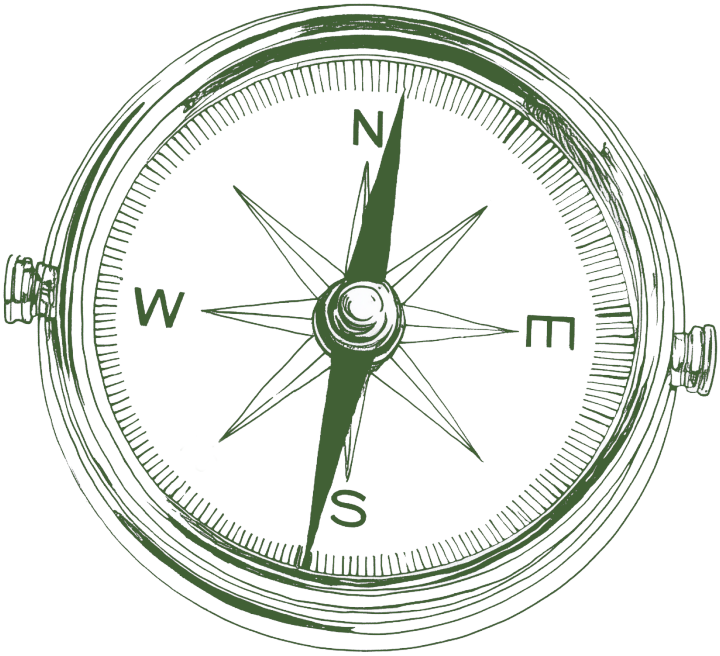

State of Quantum Computing

Today's quantum computers are still firmly in an early, experimental phase. These machines use hundreds of fragile, physical qubits to perform short, error-prone calculations. They're useful as experimental testbeds, but not as general-purpose computers capable of solving cryptographically relevant problems.

Recent excitement centers around results like those from Google's **Willow** (December 2024), which demonstrate breakthroughs in error-correction scalability. These experiments show, for the first time, that if you use more physical qubits to protect one "logical" qubit (an error-corrected qubit), that logical qubit can get more reliable, not less. While necessary, this is still insufficient because the more challenging, unproven bottleneck is performing operations between multiple logical qubits. Some researchers note that this bottleneck may reflect hard physical obstacles: how precisely matter can be controlled, how cold systems can be kept, and how much noise can be eliminated. Since these constraints may not budge with investment or iteration, progress could be materially slower or even plateau far below the thresholds required for cryptographically relevant machines.

Nevertheless, there are enough credible signals to treat CRQCs as technically feasible and a potential risk at some point. Others include governments planning to move systems to post-quantum cryptography by the 2030s, **recent surveys** showing that most experts assign a 50% probability to CRQCs emerging by 2030–2035, and the compression of CRQC **resource estimates**.

Breakthroughs that would warrant revising CRQC threat timelines:

- Demonstrations of low-error operations between multiple logical qubits, especially where performance holds or improves as the system grows.
- Minute-scale coherence (qubits stay stable long enough for real computations).
- Early commercial systems are leaving the lab (being used outside of rudimentary research, for more intensive, tangible use cases).

How Quantum Affects Bitcoin

A sufficiently advanced CRQC would break bitcoin's elliptic curve cryptography, enabling attackers to derive private keys from exposed public keys and steal funds. It is important to note, however, that a CRQC will not stop the bitcoin network from operating as intended, as blocks can still be produced and validated under the same consensus rules (such as the 21 million supply cap).

Within quantum-enabled theft, there are two attack modes: long-range and short-range.

Long-range theft would target coins whose public keys have been visible on-chain for a long time, either by design or by behavior. The former, structurally exposed coins, have public keys visible on-chain as a built-in feature of their design: for example, early-era Pay-to-Public-Key (P2PK) outputs and more recently introduced Pay-to-Taproot (P2TR) outputs. The latter, operationally exposed coins, are those made vulnerable by user behavior, namely, address reuse; once a public key is revealed in a spend, any remaining funds tied to that same address become targets. Long-exposed public keys would likely be the first targets of CRQCs, since attackers have no time constraints when deriving the private key from the public key, making the performance requirements and economics far more forgiving.

A recent analysis (June 2025) suggests that roughly 6.5 million BTC, about one-third of the total supply, would be vulnerable to long-range theft if a CRQC existed today.

Coins whose public keys are not visible on-chain (i.e., hashed formats, the most commonly used address types today) would only become vulnerable to short-range theft during a spend window, when the transaction reveals the public key. Whether early CRQCs would be powerful enough to exploit that brief window (usually a block, roughly 10 minutes) remains uncertain, but it is a possibility.

CRQCs also pose a potential threat to Bitcoin mining, as Grover's algorithm theoretically enables a CRQC to search for a block header whose hash is below the difficulty target quadratically faster than classical computers. That said, quantum mining is not easily parallelizable, making it difficult to compete at scale with large-scale classical mining operations. For a deeper dive into quantum's potential impact on Bitcoin mining, we recommend **Chaincode Labs' report**.

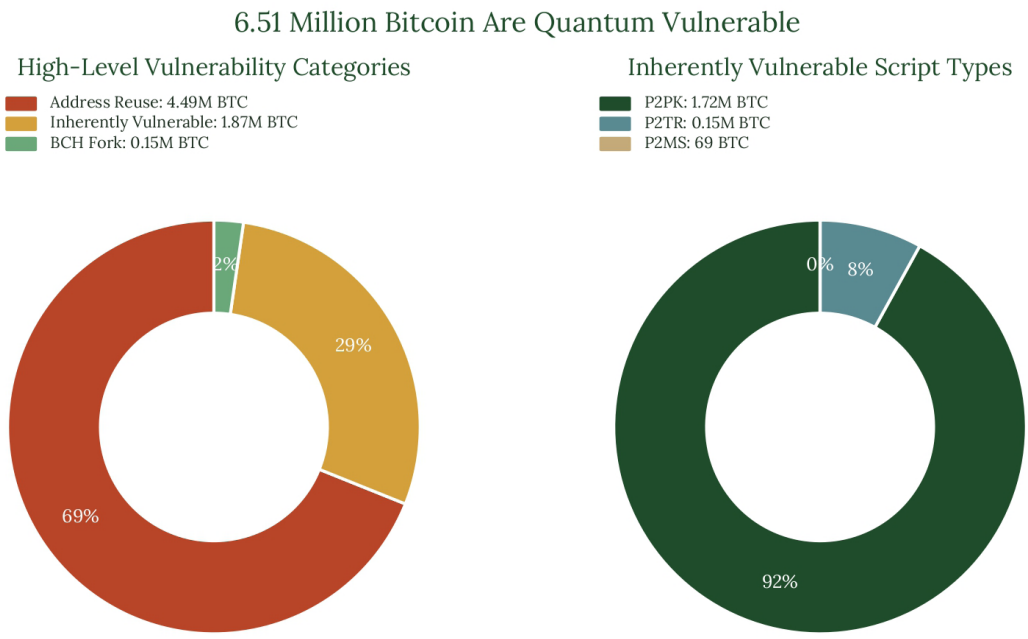

Source: **Analysis of Quantum Vulnerable Bitcoin** **As of block height 900,000 (June 6, 2025)*

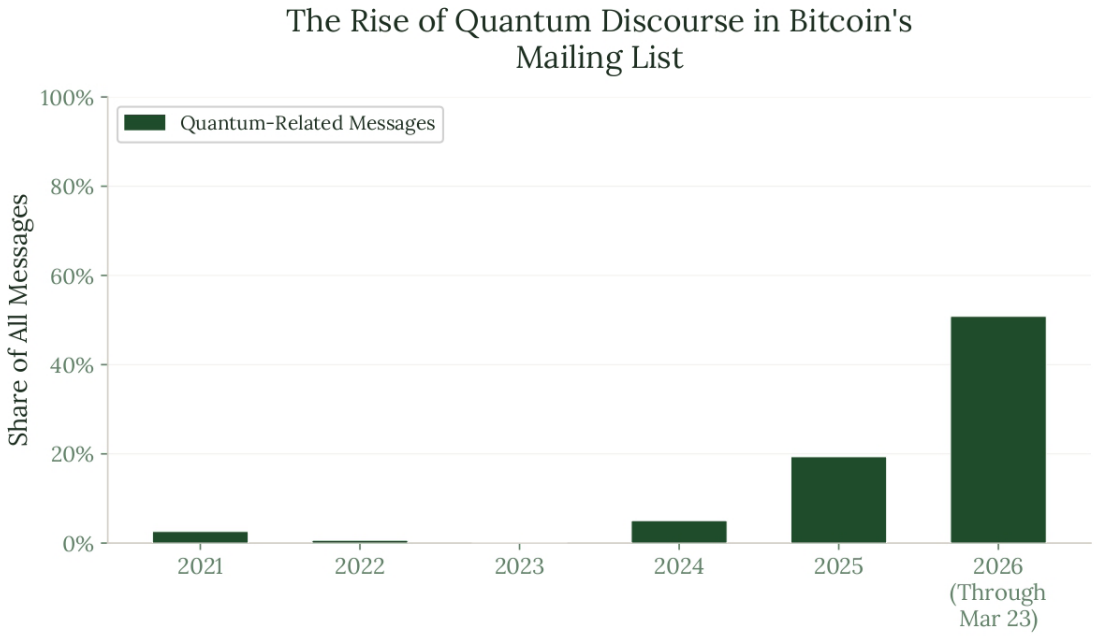

**See Appendix for methodology.*

Mitigations

Since 2023, the share of posts and replies on the Bitcoin Development Mailing List that mention quantum has increased dramatically. Developers and researchers are engaging the threat directly and now advancing a set of complementary mitigations that address different stages and layers of quantum risk to Bitcoin, collectively outlining a path toward post-quantum Bitcoin. Several proposals are in draft, but much of the work is still in research and discussion, with no final standards or activation decisions as of yet.

Custody Practices

Address reuse accounts for over two-thirds of long-range exposure. Much of it comes from a relatively small number of large custodians and institutions, often for operational simplicity. Because these holdings are under active control, they can be made safer against long-range theft risk without any protocol change, by moving funds to new addresses and rotating keys.

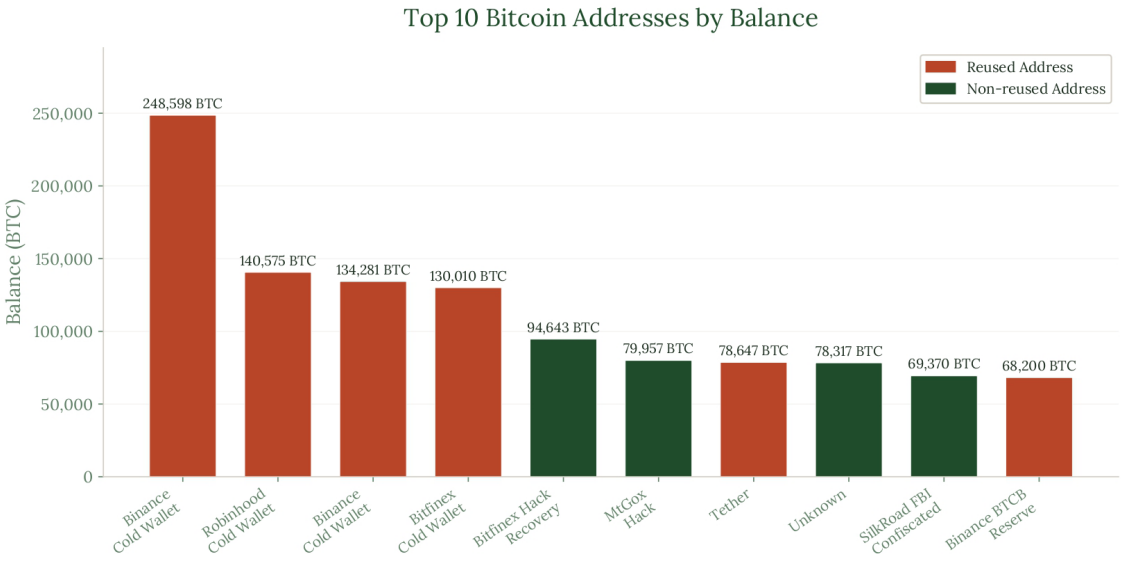

Post-Quantum Cryptography

The core mitigation, however, is to replace bitcoin's elliptic curve cryptography with post-quantum cryptographic algorithms. This protects bitcoin holders from both short-range and long-range theft. There is also recent precedent for bitcoin adopting new signature schemes, as the 2021 Taproot soft fork added Schnorr signatures alongside ECDSA.

There are two leading classes of post-quantum signatures: lattice-based and hash-based. Lattice-based signatures are better in terms of throughput and fees because they're relatively compact. However, lattice-based cryptography is still quite new, so there is a risk it ultimately turns out to be weaker than expected. Hash-based signatures, on the other hand, are simpler and more conservative, introducing no novel cryptographic assumptions since they rely solely on hash functions, which are already widely used in bitcoin (e.g., in mining, scripts, commitments, etc.). However, they are much larger on-chain, meaning lower throughput and higher fees.

Most of the bitcoin-specific work to date has been on the hash-based signature scheme SPHINCS+. **Recent research** shows that SPHINCS+ can be reconfigured for bitcoin's usage patterns, materially reducing its size.

Given that implementing new cryptography in bitcoin is a complex, network-wide upgrade, bitcoin must walk the fine line between waiting for the right cryptographic solution to mature but still deploying in a timely manner, before quantum computers become a real threat.

Output Types

Taproot-like output types can support optional spend paths in conjunction with a new post-quantum signature scheme.

In today's Pay-To-Taproot (P2TR) and in proposed Taproot-like successors such as **BIP-360**'s Pay-To-Merkle-Root (P2MR) and **Pay-to-Quantum-safe** (P2Q), wallets could create one output with two spending paths: the signature path used today and a post-quantum path used later if needed. This gives holders flexibility to spend in a quantum-safe way without migrating again. The main tradeoff is that P2TR retains long-range theft risk through its public-key-exposing key-path spend, unless the key-path spend is later disabled or frozen. BIP-360's P2MR removes that public key exposure upfront by eliminating the key path. P2Q takes a middle approach, preserving Taproot behavior today in a separate output type whose key path could later be disabled via soft fork.

PQC Precommitment for Post-Quantum Migration

PQC precommitment would allow users to pre-position funds for a future post-quantum signature scheme before bitcoin activates post-quantum signature verification in consensus.

By committing to a placeholder spend path that is valid under current rules, users can embed a future post-quantum path without any consensus change today. If a later soft fork assigns post-quantum signature verification to that path, those outputs would automatically gain a quantum-safe spend option without requiring funds to be moved again.

Quantum-Safe Bitcoin (QSB)

Quantum Safe Bitcoin (QSB) is a mechanism that enables quantum-resistant transactions using existing bitcoin consensus rules, meaning no protocol changes are required. This is accomplished by replacing reliance on elliptic curve cryptography with a hash-based construction (Binohash) embedded in bitcoin script, creating a hash-to-signature puzzle resistant to quantum attack.

While QSB does not require a soft fork to implement, it does come with its fair share of trade-offs. For instance, generating a valid transaction requires searching through billions of candidates, a GPU-intensive process that can cost $75 - $200 per transaction in compute. As a result, it is better understood as a protective fallback or specialized mitigation than as a likely replacement for protocol-level post-quantum upgrades.

Testing Environments

Protocol upgrades can be tested in production environments, but with lower-stakes than bitcoin mainnet, allowing developers to gather real performance data and build familiarity with post-quantum signatures before proposing base-layer changes.

In fact, post-quantum signatures are already being tested on live bitcoin infrastructure. In March 2026, Blockstream Research **deployed** a SHRINCS verifier on the Liquid sidechain. Other sidechains like **Anduro**, as well as quantum-resistant roll-ups, can serve a similar purpose.

Migration Safeguards

In the unlikely event of a surprise, live CRQC capable of short-range theft, the network can add an extra layer of safety for the migration period.

If a post-quantum signature scheme is ready, it could be deployed with a commit/reveal scheme. The user first publishes a post-quantum-secured commitment to the hash of a later vulnerable spend. After waiting for a set number of confirmations, they broadcast that spend. Because it is locked to the earlier commitment, a quantum attacker cannot substitute its own transaction.

If a post-quantum signature scheme has not yet been deployed, the commit transaction itself becomes vulnerable during broadcast. In this scenario, private mempools such as **MARA Slipstream**, where transactions are submitted directly to a miner without public broadcast, could shield the public key until confirmation. This is functional but undesirable, as it introduces a trust dependency on miners and is centralizing in the interim.

Each of the measures above protects holders while leaving bitcoin's current issuance, supply verifiability, ledger, and monetary foundations intact.

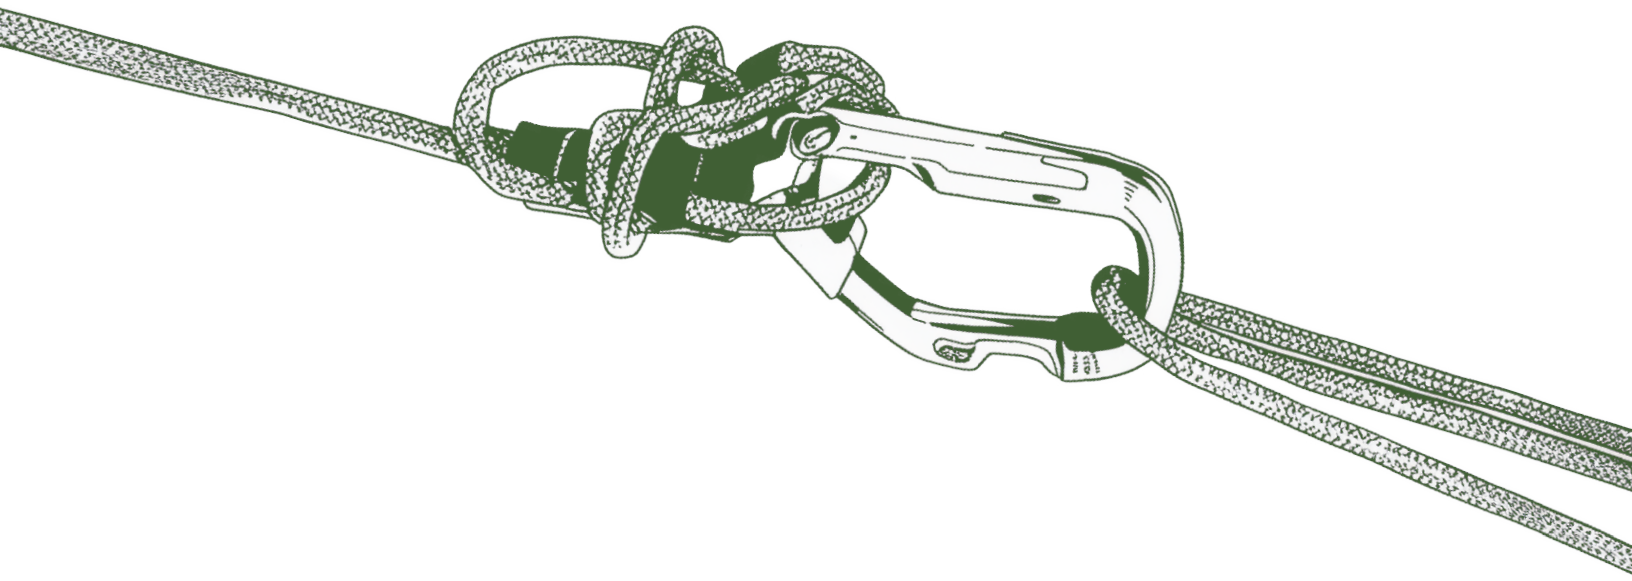

Legacy Coin Policy

Assuming a post-quantum destination exists and migration is meaningfully underway, the ecosystem still has to decide what to do with coins that remain spendable under today's signatures (and are vulnerable to quantum theft). Three possible approaches are:

Leave legacy spends valid ("steal")

Keep allowing legacy signature types, and accept that a quantum-capable actor could eventually take long-range-theft-exposed coins that never migrate.

Slow legacy spends down ("throttle")

Limit how fast vulnerable outputs can move once risk is high (e.g., 'Hourglass' proposes rate-limiting the movement of P2PK outputs to 1-per-block).

Freeze coins in legacy addresses, with a recovery path ("soft-freeze")

Recent research by **BitMEX** outlines one path by which all legacy spends are turned off, while owners can recover coins by proving they control the original wallet's seed phrase.

Most coins with exposed public keys can be reclaimed using zero-knowledge seed phrase recovery, demonstrating that they know the wallet's seed phrase without revealing it. A majority of coins are held in wallets based on seed phrases, and seed phrases are quantum-resistant: a CRQC could not derive an ordered 12-24 word seed phrase from an exposed public key. A recent **proof-of-concept** demonstrated the viability of this approach, with sub-one-minute runtime on consumer hardware.

Coins without exposed public keys could also be recovered using a variant of the commit/reveal scheme covered in Migration Safeguards.

The remaining P2PK coins (which have exposed public keys that are not derived from seed phrases) are presumed lost, but can be recovered if an owner posted a "pre-QDay" hash commitment (an on-chain commitment before quantum risk was live).

To frame the stakes from a market perspective, liquidating the core assumed-lost legacy supply, 1.72 million P2PK coins spread across almost 36,000 addresses, is roughly equivalent to about one year of "peak bull-market" long-term holder selling **per Checkonchain**.

Orderly Transition Scenario

Bitcoin is deliberately hard to change. There is **no board vote or central committee**. Bitcoin's rules are whatever software the network chooses to run. Upgrades usually start as informal, distributed research and public discussion, then get written down as a Bitcoin Improvement Proposal (BIP). In practice, a set of Bitcoin Core developers and researchers heavily influence what advances, since Bitcoin Core is the reference implementation currently run by the large majority of nodes, so even strong proposals can stall without their attention. After extensive review and testing, changes are implemented in code, but they only take effect if economically significant parts of the network, such as exchanges, custodians, payment processors, and large wallets, choose to adopt and enforce them.

So, at a high level, an orderly transition would move through three phases: research and BIP drafting, implementation and activation, and, finally, migration. Currently, the ecosystem remains firmly in the first phase. Leading developers and researchers (notably **Chaincode Labs** and **Blockstream Research**) are making active progress, but there is no agreed-upon specification and no timeline for deployment.

How an orderly transition will likely work

Choose a post-quantum signature and output type

The likely first step will be for the community to reach consensus and add a safe, conservative, bitcoin-tailored signature scheme. Today, the most likely options are the SPHINCS+ variants **SHRINCS** and **SHRIMPS** in combination.

In conjunction with this, the community would also need to decide which output type should carry that spend path: whether to extend existing P2TR outputs, despite their public-key exposure and the likelihood that key-path spends would later need to be disabled or frozen, or introduce a new output type such as P2MR that removes the public-key exposure but sacrifices some of P2TR's efficiency, privacy, and simplicity while adding yet another spend type that could make users easier to fingerprint.

Merge BIP(s) via soft fork

If the path forward involves a new output type rather than extending P2TR, it will likely be bundled into a larger soft-fork proposal. While consensus on the signature scheme and output type does not need to occur together, bundling them makes the upgrade easier to review, implement, and adopt. There is precedent for this: the Taproot upgrade, activated in November 2021, bundled three separate BIPs into a single soft fork: BIP 340 (Schnorr signatures), BIP

341 (Taproot, the new P2TR output type), and BIP 342 (Tapscript, updated script validation rules for the new output).

A post-quantum soft fork may follow a similar pattern, pairing a new signature verification opcode (e.g., OP_SHRINCSVERIFY) with a new output type like BIP-360's P2MR.

Given that current state-of-the-art post-quantum signature schemes produce significantly larger signatures, a block size increase may be proposed to maintain practical transaction throughput.

Enable it as an optional spend path

Once bitcoin supports the new post-quantum signature scheme, wallets could begin creating outputs with two ways to spend:

- a normal path that uses today's quantum-vulnerable signatures, and
- a post-quantum path that spends using the new signature's corresponding opcode (e.g., OP_SHRINCSVERIFY for SHRINCS) but stays hidden unless used.

This preserves normal spending efficiency while keeping a post-quantum option ready if a CRQC appears imminent.

Migration and adoption

At that point, users could begin moving funds into the output(s) that support the new post-quantum spend path. On-chain capacity will almost certainly not be the limiting factor: based on an estimate using **recent UTXO-set data**, if ~25% of block space were dedicated to migration, ~90% of bitcoin's value could move in ~4.4 days (~956k UTXOs), and ~98% in ~3.5 weeks (~5.3M UTXOs). The more likely constraint would be behavioral, since bitcoin wallet upgrades tend to be slow unless there is a clear forcing function.

Keep improving post-quantum signatures

Meanwhile, researchers and developers would keep working on lighter, more performant post-quantum signature schemes while tracking improvements within the broader cryptography community. If/when CRQCs are imminent, another, more mature post-quantum signature scheme could also be deployed.

Legacy outputs decision

After a safe destination exists and migration is underway, the ecosystem still must decide how to handle unmigrated, quantum-vulnerable coins.

Because this choice might divide the community, it is the clearest plausible fault line for a contentious fork. If there is a split, the winner (which chain is "bitcoin") will likely be chosen by the market. Once one version consistently trades at a premium, liquidity and users tend to follow, and the other side fades quickly, similar to how the market resolved the 2017 Bitcoin Cash split.

Rapid Response Playbook

The scenario above assumes bitcoin gets a clear warning that quantum computers are getting close to cryptographic relevance, and has time to pick a post-quantum signature standard and migrate funds in an orderly way. However, it's also important to plan for an unexpected breakthrough that requires a quick response.

In terms of precedence, bitcoin has coordinated rapid emergency fixes before (**2010 value overflow**, **2013 chain split**, **2018 inflation bug**). Those were narrower in scope than a full post-quantum migration, but they show the network can ship and adopt patches under immediate threat.

However, the first major, potential hurdle is reaching agreement that quantum-capable theft is even occurring. Quantum computing breakthroughs could be kept quiet, and early quantum thefts would likely be sporadic, making them hard to distinguish from genuine recovery and custody operations.

Activation Triggers

To reduce delays, the ecosystem can develop candidate benchmarks in advance to help guide deployment and migration decisions if quantum risk becomes more tangible.

Candidate Benchmarks: Rather than relying on a single qubit-count threshold, which **recent research** has shown can shift materially, activation triggers should track evidence that quantum computing is actually scaling low-error operations between multiple logical qubits, coherence, and gate speeds. If meaningful progress were to be made on any of these in a manner that demonstrated a repeatable scaling trajectory, that would offer a clearer warning that taking action may be prudent.

Bridge Defenses

Even if CRQCs, or credible quantum-enabled theft, surface sooner than expected, it would be imprudent to rush an untested post-quantum signature. The near-term goal, instead, would be to add safeguards that limit theft, protect transactions during migration, and reduce disorder while a long-term post-quantum upgrade is finalized and activated.

Chaincode Labs released a **CRQC Response Playbook**, developed through discussions with Bitcoin Core developers, that lays out a concrete plan and timeline for rapid response.

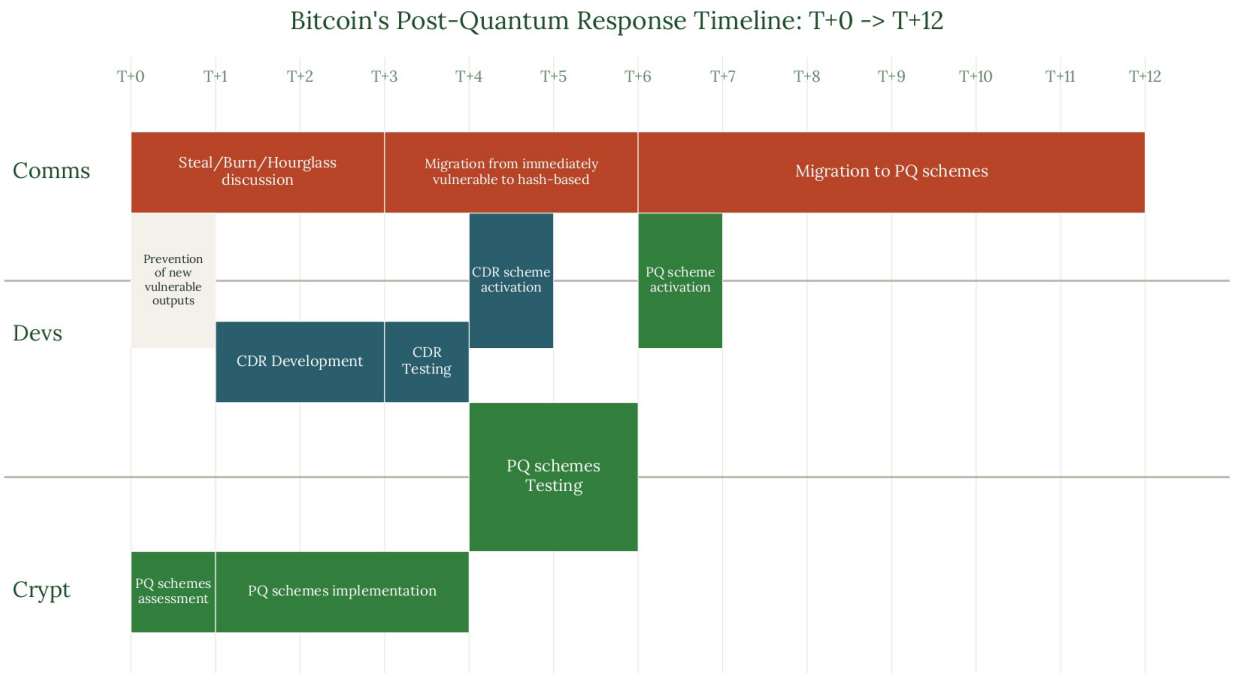

Source: **Bitcoin's CRQC Response Playbook** *(Chaincode Labs)*

In the playbook's timeline, the priority (roughly month one) is to stop the problem from getting bigger by prohibiting new quantum-vulnerable output types. In parallel, the ecosystem would need to decide how to handle legacy coins.

An unexpected live CRQC would make a "soft-freeze" more likely. Under this policy, all legacy coins are frozen, but legitimate owners can still unlock and move funds. On one hand, this intervention stops theft and buys time for a post-quantum upgrade; on the other hand, it's complex, cumbersome to roll out, and still has many edge cases.

Finally, and most importantly, is the activation of a post-quantum signature scheme. While the ecosystem will most likely have completed this well before a CRQC threat materializes, Chaincode's playbook places activation around month 6–7 if it does not. After that would come migration, in which, as mentioned above, most value could be moved in fairly short order. With 100% of block space dedicated to migration, it would take ~1.1 days to migrate ~90% of bitcoin's value (~956,000 UTXOs) and ~6.2 days to migrate ~98% of bitcoin's value (~5,267,000 UTXOs).

These bridge measures can be delivered as a soft fork: a change of ancillary rules that preserves bitcoin's core properties (blocks continue, the supply cap remains unchanged, etc.).

Emergency Safeguard in a Live CRQC Theft Environment

The discussion above assumes quantum risk is credible and imminent, but ordinary bitcoin spending remains usable. A different, though highly unlikely scenario is one in which an unexpected live CRQC is already enabling short-range theft, making ordinary spending unsafe.

In that case, developers may also work to enable a safe migration by implementing a commit/reveal mechanism. The aim would be to support safe transactions both out of already-exposed keys and, once in place, to the post-quantum scheme. While this introduces meaningful trade-offs, including reorganization risk and centralization, it would likely be preferable to widespread theft and an inability to migrate coins to post-quantum outputs safely.

Forward Outlook

The timeline for the quantum risk to bitcoin is still unclear. What is clear, however, is the risk surface: coins with exposed public keys, and the core mitigations: ending address reuse now, and deploying a post-quantum cryptography standard.

As this report shows, developers, researchers, and the broader community are already doing serious work: research, proposals, and planning. In the future, it will be important to pair that work with clear communication to stakeholders, alongside tangible timelines for testing and implementation. Concrete benchmarks for CRQC progress and legacy-coin policy are open items that still need to be formalized. The ideal end state of all these efforts is optionality, so if quantum arrives gradually, bitcoin can execute an orderly transition; if it arrives abruptly, predefined triggers and safeguards can buy time while a resilient post-quantum signature is deployed.

A careful yet proactive approach will help protect the network's functioning and ownership guarantees, further cementing bitcoin as a reliable, long-term store of value.

Have questions or comments on this report? Get in contact with Presidio Bitcoin at ***hello@presidiobitcoin.org****.*

Appendix

Bitcoin Development Mailing List Analysis

Data was sourced from the bitcoin-dev mailing list public-inbox git mirror (~24,073 emails, 2011–2026). Messages were classified as quantum-related using keyword matching (e.g., “quantum,” “lamport,” “p2qrh,” “shor,” “hourglass,” “bip360,” etc.). After automated classification, results were analyzed to remove any false positives. Every individual email sent to the mailing list counts as a message, whether it’s an original post starting a new thread or a reply within an existing thread.

Quantum Migration Timeline Methodology

Once a quantum-safe mechanism is available, exchanges, ETFs, custodians, and many users will likely seek to migrate funds, but the commonly cited 76-305+ day estimates should not be treated as realistic forecasts of how migration would unfold in practice. As a forecast of real-world migration, that analysis is poorly specified because it treats uneconomic dust UTXOs and systemically important UTXOs as equivalent units of analysis. The May 18, 2025, mempool.space UTXO report shows why that distinction matters: 49.12% of all UTXOs contained less than 1,000 sats, while 97.75% of bitcoin’s value was concentrated in just 5,267,653 UTXOs, or about 3.04% of the 173.19 million UTXOs in the set. That means count-based migration timelines materially overstate the time required to secure the economically dominant share of BTC value.

Methodologically, our estimate begins with the April 2025 distribution of BTC value across script types and uses that value distribution as a proxy for the migrating subset. For each major script family, we estimate the weight of a representative 1 input -> 1 output migration transaction input, and then compute a value-weighted average.

Table 1. Representative Input-Weight Assumptions

Script family	April 2025 value share	Representative spend assumption	Input weight
p2pkh	32.87%	legacy single-sig	592 WU
p2wpkh	31.02%	native segwit single-sig	272 WU
p2sh	21.25%	wrapped 2-of-3 multisig	560 WU
p2pk	8.70%	pay-to-pubkey	460 WU
p2wsh	5.64%	native 2-of-3 multisig	420 WU

Using those assumptions, the representative value-weighted input size is:

`[(32.87 x 592) + (31.02 x 272) + (21.25 x 560) + (8.70 x 460) + (5.64 x 420)] / 99.48 ≈ 464 WU`

We then add the fixed, non-input portion of a one-input, one-output migration transaction. Using a Taproot-sized destination output as the proxy gives the following overhead:

Table 2. Fixed Transaction Overhead

Fixed transaction field	Weight
version	16 WU
input count	4 WU
output count	4 WU
locktime	16 WU
one destination output	172 WU
legacy subtotal	212 WU
segwit marker + flag, when applicable	2 WU
representative fixed overhead	212-214 WU

That yields an average no-batching migration transaction weight of about 678 WU.

With full blocks, usable capacity is approximately 4,000,000 − 900 = 3,999,100 WU after accounting for block and coinbase overhead. Dividing by 678 WU implies roughly 5,898 migration transactions per block, or about 5,900 when rounded. The resulting throughput and timelines are summarized below.

Table 3. Migration Throughput And Timeline Estimates

Metric	Estimate
usable block capacity	3,999,100 WU
average migration transaction	678 WU
migrations per block	~5,898
UTXOs for ~90% of BTC value	~956,830
time for ~90% of BTC value	~1.13 days
UTXOs for 97.75% of BTC value	5,267,653
time for 97.75% of BTC value	~6.20 days

These estimates are stress-test capacity estimates. They are conservative regarding transaction efficiency, assuming no batching, but optimistic about available blockspace, assuming migration can consume all block capacity essentially. That assumption is not guaranteed, but it is a defensible emergency scenario because holders of high-value UTXOs could plausibly pay fee premiums large enough to crowd out most routine transaction demand.

The ~956,830 UTXO estimate for ~90% of BTC value should also be understood as a linear interpolation within the 1-5 BTC bucket (see mempool UTXO Report) rather than a directly observed count. Because value is unlikely to be uniformly distributed within that bucket, this estimate is mildly optimistic. It may slightly understate the number of UTXOs, and therefore the time required to reach the 90% threshold. That caveat does not materially affect the headline conclusion.